UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

An atom interferometer for a precision measurement of the cesium recoil frequency

Permalink

https://escholarship.org/uc/item/0f28c5nh

ISBN

9798173787446

Author

Roth, Jack Clark

Publication Date

2026-08-10

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

An atom interferometer for a precision measurement of the cesium recoil frequency

by

Jack Clark Roth

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Holger Mueller, Chair
Assistant Professor Alp Sipahigil
Professor Hartmut Haeffner

Summer 2026

An atom interferometer for a precision measurement of the cesium recoil frequency

Copyright 2026
by
Jack Clark Roth

Abstract

An atom interferometer for a precision measurement of the cesium recoil frequency

by

Jack Clark Roth

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Holger Mueller, Chair

Measurements of atomic recoil frequency produce the most precise direct determinations of the fine-structure constant. Tension among recent determinations of the fine-structure constant motivates improved experiments with better understood systematic effects. This thesis presents the operation and characterization of an atom interferometer based on Bragg diffraction designed to measure the cesium recoil frequency to high precision by suppressing systematic effects present in a previous generation experiment.

We describe the experiment apparatus, state preparation procedure, and mechanism by which an atom interferometer is formed and the cesium recoil frequency is measured. The performance of the experiment is characterized, revealing a 24 hour precision of 400 ppt in the cesium recoil frequency. We describe newly identified systematic effects, and discuss how these effects and previously known ones can be suppressed or measured. Several simulation methods are reviewed or developed, and verified against experiment results. Simulations are found to reproduce much of the observed interferometer behavior, but do not yet reproduce experimental results at the required level of precision. Finally, we review several new interferometry techniques, including new interferometer geometries and a novel beamsplitter method, that will find use in measuring the cesium recoil frequency and beyond.

Together, these results improve understanding of systematic effects in Bragg beamsplitter based atom interferometers through well characterized experiment-simulation comparisons, and provide a foundation for a next-generation measurement of the cesium recoil frequency.

Contents

Contents **i**

List of Figures **iv**

List of Tables **xvii**

1 Introduction **1**
1.1 Atom interferometry . . . 1
1.2 The fine-structure constant . . . 3
1.3 Outline of thesis . . . 6

2 Matterwave Optics **8**
2.1 Two-level system in an off-resonant bichromatic field . . . 9
2.2 Bragg diffraction . . . 12
2.2.1 Equations of motion in the Bragg regime . . . 17
2.2.2 Multifrequency Bragg diffraction . . . 21
2.3 Bloch oscillations and adiabatic Bragg diffraction . . . 22
2.4 Raman transitions . . . 24
2.5 Laser phase . . . 28

3 Theory of Atom Interferometers **29**
3.1 Ramsey interferometry . . . 29
3.2 Interferometer geometries . . . 32
3.2.1 Free evolution and laser phase . . . 33
3.2.2 Ramsey-Bordé interferometer . . . 33
3.2.3 Simultaneous Conjugate Ramsey-Bordé interferometer . . . 37
3.2.4 Mach-Zehnder interferometer . . . 41
3.3 Separation phase . . . 42
3.4 Interferometer contrast . . . 43

4 Experiment Apparatus **44**
4.1 Experiment vacuum chamber, optics, and coils . . . 45

4.1.1 2D MOT chamber . . . 45
4.1.2 3D MOT chamber . . . 47
4.1.3 Detection chamber . . . 48
4.1.4 Interferometry chamber . . . 51
4.1.5 Retroreflection chamber . . . 51
4.1.6 Compensation coils . . . 51
4.1.7 Vacuum pumping . . . 51
4.2 Laser systems . . . 53
4.2.1 Repump laser system . . . 53
4.2.2 Cooling laser system . . . 56
4.2.3 Tapered amplifier systems . . . 58
4.2.4 Switchyard optical system . . . 59
4.2.5 Interferometry laser system . . . 61
4.2.6 Phase-locking and automatic relocking . . . 66
4.2.7 Frequency comb and wavemeter . . . 72
4.2.8 Interferometry beam alignment . . . 72
4.3 Control system . . . 72
4.4 Atom cooling, trapping, and state preparation . . . 73
4.4.1 Magneto-optical traps . . . 74
4.4.2 Molasses launch . . . 77
4.4.3 Raman sideband cooling . . . 80
4.4.4 Adiabatic passage . . . 85
4.4.5 Blowaway light . . . 87
4.4.6 Velocity selection . . . 87
4.5 Interferometry sequences . . . 90
4.5.1 Bragg diffraction . . . 90
4.5.2 Bloch oscillations . . . 92
4.6 Developing a robust experiment . . . 93

5 Experiment Characterization **95**
5.1 Extracting ω_r from interferometer output . . . 95
5.2 Projection noise limit . . . 98
5.3 Atom number measurements . . . 99
5.3.1 Absorption measurements . . . 99
5.3.2 Fluorescence measurements . . . 102
5.4 Time-averaging . . . 103
5.5 Current experiment precision . . . 105

6 Simulations **107**
6.1 ODE integration methods . . . 109
6.1.1 Lookup tables . . . 109
6.1.2 Parallelized ODE method . . . 111

6.2 Full interferometer simulations . . . 112
6.2.1 Tree approach . . . 112
6.2.2 Wavepacket integration approach . . . 116
6.3 Spatial separation of wavepackets . . . 121

7 Investigations **124**
7.1 Systematics . . . 124
7.1.1 Velocity selection magic wavelength . . . 127
7.1.2 Bragg vibration systematic . . . 130
7.1.3 Wavefunction width . . . 132
7.1.4 Parasitic interferometers . . . 133
7.1.5 Bragg pulse matching . . . 135
7.1.6 Bloch matching . . . 143
7.1.7 Ellipse matching . . . 144
7.2 Methods . . . 147
7.2.1 Novel interferometer geometries . . . 147
7.2.2 Serrodyne modulated matterwave optics . . . 157
7.3 Measuring wavefunction spatial extent . . . 164

8 Summary and outlook **167**
8.1 Summary . . . 167
8.2 Outlook . . . 167

A Reference math **171**
A.1 Unitary transformations . . . 171
A.2 Time-evolution operator . . . 172
A.3 Atom-field interaction . . . 172
A.4 Position and momentum space identities . . . 173
A.5 Supporting proofs for adiabatic elimination . . . 174
A.6 Two-level system in position space . . . 177
A.7 Serrodyne phase modulation equations of motion . . . 178

B Coil switching electronics **181**

Bibliography **184**

List of Figures

1.1 Illustration of two sinusoids with phase difference θ. When the two sinusoids are summed along the x-axis, the resulting waveform reveals mostly destructive interference. 2
1.2 Current determinations of the fine-structure constant. The partially transparent data points for the Berkeley-18 and LKB-20 measurements are the values of those measurements when beam-related effects are ignored. Data sources are cited in the text. 5

2.1 Schematic of stimulated absorption followed by stimulated emission of a photon by an atom. Initially the atom is in the ground state with zero momentum $|g, 0\hbar k\rangle$. A right-moving photon (displayed in blue) excites the atom and imparts a momentum impulse, putting it in state $|e, 1\hbar k\rangle$. Next, a left-moving photon stimulates emission of a photon, causing the atom to enter the ground state again but with the momentum kick from two photons, putting it in state $|g, 2\hbar k\rangle$. Note that the left-moving photons are slightly red-detuned relative to the right-moving photon, as the atom motion in the $|e, 1\hbar k\rangle$ state Doppler shifts the atomic resonance. 9
2.2 The level diagram of a two-level system with an external degree of freedom, with Bragg diffraction shown. The kinetic energy associated with the external degree of freedom is described by the gray dashed parabolas seen in the figure. Because the off-resonant bichromatic field only couples momentum states in discrete spacings of $\hbar k$, the continuous set of states defined by the parabola turns into a discrete set of states represented by the black horizontal bars. The parameter Δ specifies the detuning of the laser fields from the excited state, and is referred to as the single-photon detuning. The parameter δ specifies the difference in frequency between the two laser fields, and determines whether the Bragg diffraction process is resonant. Here 2nd-order Bragg diffraction is driven, as $\omega_1 - \omega_2$ is equal to half the kinetic energy difference between the initial and final state. 13

2.3 Simulation of a 4th-order Bragg diffraction process in both position and momentum space as a function of time. *Top:* The effective Rabi frequency of the Bragg pulse as a function of time. This is described by the truncated Gaussian pulse given in Eq. (2.20) where $\sigma = 0.259\omega_{\mathrm{r}}^{-1}$, $t_0 = 3\sigma$, and $\Omega_{\mathrm{eff},0} = 19.4\omega_{\mathrm{r}}$. *Upper middle:* The population in momentum space, computed by integrating Eq. (6.2). *Lower middle:* The population in position space, computed by integrating Eq. (6.14) using the split-step method. The wavefunction is initialized as a Gaussian wavepacket with amplitude width $\sigma_x = 5\sqrt{2}/k$, where k is the laser wavenumber. *Lower:* The absolute value of the potential that the position space wavefunction is subject to as a function of time, given by Eq. (A.47). 14

2.4 Bragg dynamics in momentum space for several Bragg pulse durations. All dynamics show 4th-order Bragg diffraction. For longer pulse durations we see fewer intermediate momentum states are populated. For each value of σ the value of $\Omega_{\mathrm{eff},0}$ that produces a $\pi/2$ pulse is used. 16

2.5 Momentum state populations at the end of a $\sigma = 0.259\omega_{\mathrm{r}}^{-1}$, 4th-order Bragg pulse. We see that, while suppressed, intermediate momentum states are populated at the end of the pulse. From the plot we see that a $\pi/2$ pulse (a beamsplitter) occurs slightly below $\Omega_{\mathrm{eff},0} = 20\omega_{\mathrm{r}}$ and a π pulse (a mirror) occurs slightly above $\Omega_{\mathrm{eff},0} = 25\omega_{\mathrm{r}}$. 16

2.6 *Upper:* The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of the Bragg pulse width σ. The inset zooms in on the rightmost five points in the plot to show their deviation from $-\pi/2$. *Lower:* The required $\Omega_{\mathrm{eff},0}$ to achieve a $\pi/2$ pulse for each σ value used in the upper plot. These generally decrease as σ gets larger. Insets on this plot show the momentum state population dynamics as a function of time for two different σ values. All processes in this figure are 4th-order Bragg diffraction pulses. . . . 17

2.7 The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of $\Omega_{\mathrm{eff},0}$ for three different values of σ. The vertical dashed lines display the values of $\Omega_{\mathrm{eff},0}$ that produce a $\pi/2$ pulse. We see that the phase difference is highly dependent on intensity for shorter Bragg pulses, and far less dependent on intensity for longer Bragg pulses. All processes in this figure are 4th-order Bragg diffraction pulses. 18

2.8 The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of δ_{offset} (the offset of the two-photon detuning from resonance with 4th-order Bragg diffraction) for three different values of σ. 18

2.9 The level diagram describing Bloch oscillations. This is similar to the Bragg level diagram shown in Fig. 2.2, but the laser fields do not resonantly couple two states. Instead one of the laser fields (in this case ω_2) is swept in time such that $\delta = \omega_1 - \omega_2$ sweeps over the energy difference between $|g, 0\hbar k\rangle$ and $|g, 2\hbar k\rangle$. This causes adiabatic population transfer. As δ continues to sweep in time, it will eventually sweep over the energy difference between $|g, 2\hbar k\rangle$ and $|g, 4\hbar k\rangle$, causing a second adiabatic transfer of population. This continues to happen until the sweep of δ stops. 22
2.10 Population dynamics in momentum space during 5th-order Bloch oscillations. We see that population is transferred from $p = 0\hbar k$ to $p = 10\hbar k$ at the end of the Bloch oscillation process. Here $\Omega_{\text{eff}} = 2\omega_{\text{r}}$ (meaning it is constant in time) and $\tau = 40\omega_{\text{r}}^{-1}$. 23
2.11 The final population in each momentum state for 5th-order Bloch oscillations as a function of τ. $\Omega_{\text{eff}} = 2\omega_{\text{r}}$ in this plot. 24
2.12 Population dynamics in momentum space during a 4th-order adiabatic Bragg diffraction pulse. Here $\Omega_{\text{eff},0} = 14\omega_{\text{r}}$, $\sigma = 2\omega_{\text{r}}^{-1}$ (the pulse shape defined in Eq. (2.20) is used, but with the truncation occurring at 4σ rather than 3σ), and δ is swept from $4(3+4)\omega_{\text{r}}/2$ to $4(5+4)\omega_{\text{r}}/2$ over the duration that the pulse has a nonzero intensity. 25
2.13 The level diagram describing Raman transitions. This is similar to the level diagrams considered in Figs. 2.2 and 2.9, but includes an additional ground state. In the diagram ω_1 and ω_2 are selected to drive a transition between the two ground states. Note that the field directions were chosen so that the transitioned state has momentum $-2\hbar k$ relative to the initial state. 26

3.1 Ramsey interferometry sequence applied to a standard two-level system. Initially the system is in the ground state $|g\rangle$, as is shown in the leftmost diagram. Then a $\pi/2$ pulse is applied by pulsing on a coupling field, leaving the system in state $(|g\rangle - i|e\rangle)/\sqrt{2}$. Afterwards the system is allowed to evolve for time T. The clocks above each state represent the phase accumulated by that state. After evolving, the excited state clock has accumulated some phase, as states accumulate phase at a rate proportional to their energy. Finally, the second $\pi/2$ pulse is applied, which transfers amplitude from the state $|g\rangle$ to $|e\rangle$, and vice versa. Note that the transferred amplitude retains whatever phase it accumulated, which leads to interference. 30
3.2 Fringes produced by measuring the excited state at the end of a Ramsey sequence, given by Eq. (3.3). For short free evolution times T, the fringe oscillates slowly as the drive frequency ω is scanned. As T increases the fringe oscillates more rapidly as ω is scanned. When running a Ramsey interferometry experiment, the goal is to determine the drive frequency where $\omega = \omega_e$. This equality can be measured to higher precision at longer T, because the fringe becomes much more sensitive to the drive frequency ω. 31

3.3 Position-space simulation of a Ramsey-Bordé interferometer, computed by integrating Eq. (6.14) using the split-step method. Four Bragg beamsplitters are applied at times $t = 0$, $t = T$, $t = T + T'$, and $t = 2T + T'$. Between beamsplitter pulses classical trajectories are followed, forming a trapezoidal interferometer geometry. Second-order Bragg beamsplitters are used with $\Omega_{\text{eff},0} = 5.1\omega_{\text{r}}$, $\sigma = 0.32\omega_{\text{r}}^{-1}$, $T = 10\omega_{\text{r}}^{-1}$, $T' = 5\omega_{\text{r}}^{-1}$. Note that δ is shifted from $\delta = 4n\omega_{\text{r}}$ to $\delta = 4n\omega_{\text{r}} - \omega_{\text{m}}$, where $\omega_{\text{m}} = 8n\omega_{\text{r}} - \Phi/nT$, where the target phase $\Phi = \pi/4$ of the Ramsey-Bordé interferometer. The wavefunction is initialized as a Gaussian wavepacket with amplitude width $\sigma_x = 5\sqrt{2}/k$, where k is the laser wavenumber. 34

3.4 The classical trajectories followed in a Ramsey-Bordé interferometer geometry. Light vertical lines denote the times at which beamsplitters are applied. Note that these trajectories match those from the simulation in Fig. 3.3, although we have omitted trajectories that do not interfere at the end of the interferometer. The non-interfering trajectories are referred to as junk ports. 35

3.5 Trajectories followed in an SCI geometry in a linear gravitational potential. Light vertical lines denote the times at which beamsplitters are applied. The lower interferometer is the same as that in Fig. 3.4, but in a gravitational potential. The upper interferometer looks like an upside-down version of the lower interferometer, which leads to the cancellation of any phase due to the linear gravitational potential. 38

3.6 Trajectories followed in an SCI geometry with Bloch oscillations (BO, in orange) occurring halfway between the second and third beamsplitter pulses. The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters are applied. 40

3.7 Trajectories followed in a Mach-Zehnder interferometer (MZI). The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters (or optionally a mirror in the case of the intermediate pulse) are applied. 41

4.1 Renderings of the experiment chamber. *a:* The 3D MOT, detection, interferometry, and retroreflection chambers (labeled in figure). *b:* The 2D MOT chamber. *c:* The 3D MOT telescope. *d:* The 3D MOT chamber with the 2D MOT chamber attached. 46

4.2 *a:* The 2D MOT. *b:* A 3D MOT beam telescope mounted upside down. *c:* The RSC X telescope, with the top of a 3D MOT beam telescope visible. *d:* Optics underneath the 3D MOT chamber used to align the interferometry beam and one of the blowaway beams. The microwave horn is also visible. *e:* The detection optics. Usually a light-blocking sheet is draped over the setup to reduce the detection background. *f:* The top of the interferometry chamber. 49

4.3 Schematic of the detection chamber and optics. The orange cloud represents the atoms at the time of detection. A relay lens system recreates an image of the atoms on a photodiode. The resonant detection light is expanded using a cylindrical lens and enters the chamber perpendicular to the relay lens system optical axis, illuminating the cloud. The width of the detection light, positioning, and active area of the photodiode set the region of the cloud that is detected. . 50
4.4 Recorded vacuum pressure over time. Spikes correspond to times the titanium sublimation pump was fired. 52
4.5 Typical optics used before every polarization-maintaining fiber in the experiment. Both the mirror and fiberport should be placed on a kinematic mount, which enables the beam location and angle relative to the fibercoupler to be adjusted properly. Ideally the fibercoupler also has an adjustable position aspheric lens that can be moved to maximize fibercoupling efficiency. If polarized light is being fibercoupled, a $\lambda/2$ and $\lambda/4$ waveplate should be placed before the fiber. These waveplates should be aligned to make the polarization linear on the fiber output and to minimize any polarization instability. The polarization drift can be recorded using a polarimeter such as the Thorlabs PAX1000IR1. If the fiber is undamaged, it should be possible to use the waveplates to keep the output polarization stable to within 2° in both azimuth and ellipticity on the Poincaré sphere. This can be tested by straining or heating the fiber (easily done by compressing the fiber in your hand), which changes the birefringence of the fiber and leads to a polarization shift if the input light is not perfectly coupled to a single fiber axis. It is also good practice to add a PBS cube to the fiber output to further ensure polarization stability. 54
4.6 Cesium level structure with most laser frequencies used in the experiment shown. 55
4.7 Optical schematic of the ECDL repump laser and cesium spectroscopy. 56
4.8 Optical schematic of the DBR cooling laser. 57
4.9 Optical schematic of the tapered amplifier systems. 59
4.10 Optical schematic of the switchyard used to generate light for the 3D MOT, RSC, blowaway, and detection. 60
4.11 Optical schematic of the interferometry laser and the optical chain used to generate light for Raman transitions, Bragg diffraction, and Bloch oscillations. The diagram also displays how RSC Z and Bloch compensation light are injected into the interferometry fiber. 62
4.12 Ti:Sa laser intensity and cavity element positions over time after the Ti:Sa was turned on on 2 January 2026. All element positions are set by active feedback loops that stabilize the laser frequency. The intensity and cavity elements take around 12 h to stabilize. BiFi refers to the birefringent filter, TE refers to the thin etalon, and PE refers to the piezo etalon in the cavity. 63
4.13 The circuit used to stabilize laser intensity to the required pulse shape. The dashed box is the original final op-amp stage. Yuno Iwasaki developed an upgraded output stage shown in the solid boxed region. 64

4.14 The circuit used to produce an error signal from the beatnote through use of the AD9901 phase-frequency detector. 67
4.15 The circuit used to lock the DBR laser using the error signal produced by the circuit in Fig. 4.14. 68
4.16 Measured spectrum of the beatnote between the ECDL and DBR lasers when locked. Inset displays the same signal measured with a lower resonant bandwidth. 69
4.17 Measured distribution of the phase between the ECDL and DBR frequencies when locked. The phase difference is measured after the frequency divider. Distribution is recorded over an 8 ms period at a 0.5 GS/s sample rate. The distribution is normalized and reported as a probability distribution. 70
4.18 Short-time Fourier transforms of the beatnote measured between the ECDL and DBR while the lock is activated and the reference beatnote frequency is jumped. The reference beatnote frequency is given by the dashed blue line. 71
4.19 Falling cloud time-of-flight traces observed after a molasses launch, RSC, and VS, reported in units of optical power. A basic model of the detection optical system and the number of photons emitted per atom allows the total detected energy to be converted into an atom number. This figure is further discussed in Section 5.3.2. 75
4.20 Schematic of the experiment chamber during different stages of the experiment. The 3D MOT coils are shown in blue, lasers in red, and viewports in gray. Atom clouds are orange, and their trajectory is shown in green. The three layers of Mu-metal shields 2.34 m in length [24], solenoid, and baffles are represented by the cylinder in the interferometry chamber. The position degree of freedom in Eq. (2.1) and the position axis in Fig. 3.3 correspond to the vertical dimension in the chamber. *Left:* After the 3D MOT is formed the molasses launch and PGC occur. Shortly afterwards the atom cloud passes through the RSC region (the retroreflected RSC X beam is shown) and is further cooled. *Middle:* After the VS pulses the atoms enter the magnetically shielded region of the chamber and the interferometry beam is pulsed to form Bragg beamsplitters and Bloch oscillations. *Right:* After the interferometer has completed, the spatially separated output ports of the interferometer fall through resonant light and fluoresce. The fluorescence is detected by a photodiode. 76
4.21 Fluorescence observed in the 3D MOT region by a camera when all but one pair of 3D MOT beams is turned off. The fluorescent region reveals the volume that the 3D MOT beam pair illuminates, and can be used to center the beam on the 3D MOT location. 77
4.22 Time-of-flight traces of the rising and falling molasses clouds measured in the detection region. The dashed line on the falling cloud represents the point at which the full width at half maximum is measured. 79

4.23 Height of the atom cloud as a function of time for several launch velocities. Horizontal dashed lines denote various relevant heights in the interferometry chamber. A launch velocity of 6.2 m/s is chosen as it balances sufficient time in the shielded region and interferometer duration with signal size, which goes down with larger launch velocities. 80
4.24 Velocity of atoms in the experiment chamber as a function of time during an SCI with Bloch oscillations. The interferometer starts 225 ms after launch, has a launch velocity of 6.2 m/s, 4th-order Bragg diffraction, 100th-order Bloch oscillations, a Bloch oscillation lattice acceleration of $7kg$, $T = 150$ ms, and $T' = 8$ ms. The velocities of other lattices formed in the interferometer during Bloch oscillations are shown. Care must be taken to ensure the velocity of these lattices does not intersect the velocity of the atoms, as otherwise the atoms will be ejected from their Bloch lattice as the two lattices accelerate past each other. We see that the Bloch lattice velocities do not intersect while Bloch is on, although they become relatively close (this is clear from the inset). After Bloch oscillations the third Bragg pulse occurs and the atom velocity lines bifurcate. 81
4.25 The potential formed by the RSC lattice, and its derivative as a function of displacement from the potential minimum for different displacement directions. 82
4.26 Time-of-flight trace of the falling RSC clouds measured in the detection region. The dashed line on the falling cloud represents the point at which the full width at half maximum is measured. 85
4.27 m_F states in the $F = 3$ and $F = 4$ hyperfine ground states of cesium, along with their allowed transitions. During the adiabatic passage stage, population is transferred from $F = 3, m_F = +3$ to $F = 3, m_F = 0$ along the dashed blue line. 86
4.28 Wavefunction in each of the two ground states computed at various points of the VS process. Each panel is computed with a different Raman pulse duration τ (both panels are square pulses in time). In the experiment, $\tau = 5.2\omega_r^{-1}$ (the panel on the right) is used. 88
4.29 Population in each of the two ground states computed at various points of the VS process when a Raman pulse with a Gaussian profile in time is used. The Gaussian has $\sigma = 0.5\omega_{\mathrm{r}}^{-1}$. 89
4.30 The efficiency and selected population width of the VS process computed as a function of pulse duration. A square pulse is used in time, with Ω_{eff} adjusted to maintain a π pulse. The efficiency and selected population width are shown both after a single VS pulse (g_2), and after two VS pulses (g_1). 90
4.31 Measured time-of-flight trace of the SCI signal. For this interferometer, $T = 100$ ms, $n = 4$, and $N = 100$. The pair of peaks around 50 ms are the outputs of the lower interferometer, and the pair of peaks around 235 ms are the outputs of the upper interferometer. The first pair of peaks have FWHM $\sim$2 ms, which is limited by the cloud size rather than the cloud expansion rate. 91

4.32 The measured population inversion as a function of the Bragg pulse amplitude and the Bragg two-photon detuning δ. Note the inversion here is defined as $w = (p_n - p_0)/(p_0 + p_n)$, where $p_0 = |\langle 0\hbar k|\phi\rangle|^2$ and $p_n = |\langle 2n\hbar k|\phi\rangle|^2$ in whatever frame puts the initial state at $n = 0$. The Bragg pulse amplitude is used to set the reference waveform amplitude for the Gaussian Bragg diffraction pulse, as described in Section 4.2.5. This is proportional to the effective Rabi frequency. Note that a downward 4th-order Bragg pulse is used, for which $\delta = -16\omega_{\mathrm{r}}$. . . . 92

5.1 An experimentally recorded ellipse. 95
5.2 Optical power of the probe beam detected on a Thorlabs PDA100A2 photodiode as a function of time. When a MOT has been loaded, there is a dip in the detected power due to absorption. By taking the difference of the cases without and with the MOT and integrating, the total optical energy absorbed by atoms in the MOT is measured. 101
5.3 Images recorded by a Thorlabs Zelux CMOS camera by exposing the probe beam with and without a 3D MOT. The total optical power absorbed by the 3D MOT is computed by integrating over their difference. The difference also reveals the size of the 3D MOT projected onto a plane perpendicular to the propagation of the probe beam. The position units specify pixel positions on the camera. . . . 102
5.4 *Left:* Ellipses recorded by an SCI for several values of T. *Right:* An experiment run in an identical configuration to the left panel, but with the final beamsplitter pulse delayed so that interference does not occur. Taken together, these datasets help determine which noise sources are due to phase-sensitive processes and which are phase-insensitive. 104
5.5 A synthetically generated ellipse overlaid on top of the $T = 25\,\mathrm{ms}$ ellipse in Fig. 5.4. Gaussian noise is injected into the synthetic ellipse to mimic the shot-to-shot variation in the single Bragg pulse inversion. 104

6.1 An SCI composed of 4th-order Bragg pulses. Parasitic trajectory paths propagate through the interferometer and interfere at the time of the final beamsplitter pulse. These are shown as partially transparent lines. These parasitic momentum states can be produced during Bragg diffraction pulses. 108
6.2 Histogram of the error on interpolated Bragg beamsplitter integrations with respect to direct Bragg beamsplitter integrations. Phase error is computed as $\arg(\phi_{\mathrm{integrated}}/\phi_{\mathrm{interpolated}})$, and population error is computed as $|\phi_{\mathrm{interpolated}}|^2 - |\phi_{\mathrm{integrated}}|^2$. 110

6.3 Diagram of the nodes used to represent an SCI when using the tree-based simulation method. Levels denote the interferometer state after different operations. For example, a beamsplitter is applied to go from level 0 to 1. At the end of the interferometer, nodes with equal positions and momenta interfere. These are represented by the nodes connected with black lines. Other nodes do not interfere, and are so-called "junk" outputs. These are represented in gray. Not all junk outputs are shown here, as we only include paths that end the interferometer with a population greater than 0.001 for the chosen Bragg parameters (parameters were chosen to minimize parasitic output ports). If parasitic interferometers are included in the tree, additional paths will appear in the tree that will ultimately interfere. 114
6.4 Simulated ellipses generated using a tree-based simulation method of an SCI. Several values of T are simulated, and, as expected, a decay of contrast is observed as T increases. 116
6.5 The discretization scheme used when setting $\phi_{j,s} = \phi(2\hbar kj + \hbar ks/S)$ for the case where $S = 3$. 119
6.6 Difference between computed interferometer phases for the stacked ODE method, the split-step method (ss), and the Crank-Nicolson method (cn). 121
6.7 Error in the computed interferometer phases as states are removed from the stacked ODE method. The left-hand-side axis displays the error between the truncated and untruncated simulations. The right-hand-side axis displays the remaining fraction of s states for a particular truncation level. 122

7.1 Calculated dependence of the Δ_{AC} on the single photon detuning Δ for the cesium $6\,^2\mathrm{S}_{1/2}$ to $6\,^2\mathrm{P}_{3/2}$ transition. We see that Δ_{AC} goes to zero around $\Delta = 6000\,\mathrm{MHz}$. Δ_{AC} is computed via Eq. (7.5) for the configuration used in the experiment. Note that the calculation result depends strongly on the relative beam intensities and EOM modulation depth, and so this result is not expected to describe the actual experiment. A similar calculation is performed in [51]. . . 129
7.2 The differential phase ϕ_d measured as a function of the time t_0 of the first Bragg beamsplitter pulse after the molasses launch. The differential phase is shown in three different conditions: with the new 3D MOT coils installed on the experiment, with the original 3D MOT coils installed on the experiment, and with damping sandbags placed on top of the retro mirror chamber while the original 3D MOT coils were installed on the experiment. The $\Phi = 2\phi_d = \pi/2$ condition is also shown. It is clear that the new 3D MOT coils remove much of the observed oscillation. 130
7.3 The position and velocity of the retro mirror as a function of time after the 3D MOT coil turn off. The position is recorded by an optical Michelson interferometer placed underneath the 3D MOT chamber, and the velocity is determined by numerically taking the derivative of the position. 132

7.4 Simulated interferometer phase and contrast as a function of wavefunction width for a single atom and cloud of atoms. "Masked" denotes an interferometer where parasitic momentum states are removed from the simulation after each beamsplitter pulse. The phase is reported as $\phi_{\text{systematic}}$, which denotes the deviation of the interferometer phase from the expected phase for a Ramsey-Bordé interferometer. 134
7.5 Simulated population in the $2\hbar k$ momentum state for a 5th-order Bragg diffraction pulse as a function of the pulse width σ. The loss is any population not in $0\hbar k$ or $10\hbar k$. The figure is intended to replicate Fig. (12) of [100]. 134
7.6 Simulated population in the $2\hbar k$ momentum state for a 5th-order Bragg diffraction pulse with $\sigma = 0.22\omega_{\text{r}}^{-1}$ (which suppressed parasitic populations in Fig. 7.5) as a function of the peak effective Rabi frequency. The loss is any population not in $0\hbar k$ or $10\hbar k$. 135
7.7 Method by which momentum state populations after a Bragg pulse are measured. The PD trace is fit to a sum of Gaussians centered at each of the vertical dashed lines. Each line corresponds to a particular momentum state (labeled at the top of the figure). The extracted Gaussian peak height is taken to be the population in that state. 136
7.8 Experimentally measured momentum state populations as a function of $\Omega_{\text{eff},0}$, along with a simulation of an atomic cloud undergoing Bragg diffraction (solid dashed lines). A basic simulation of a single atom undergoing Bragg diffraction in a uniform beam is also shown (transparent dashed lines). 138
7.9 Measured Bragg pulse waveforms for two different AWG pulse shape amplitudes. The Gaussian waveform sits on a pedestal, which must be considered for the simulated Bragg dynamics to match the experimentally measured dynamics. The left y-axis shows the measured photodiode voltage for the 50 mV amplitude AWG pulse, and the right y-axis shows the measured photodiode voltage for the 775 mV amplitude AWG pulse. 138
7.10 Measured population inversion resulting from 4th-order Bragg diffraction as a function of $\Omega_{\text{eff},0}$ and the Bragg pulse delay. 140
7.11 Residuals between the measured and simulated Bragg diffraction inversion as a function of $\Omega_{\text{eff},0}$ and the Bragg pulse delay. The measured inversion is displayed in Fig. 7.10. 141
7.12 *Left:* Measured population inversion from 4th-order Bragg diffraction as a function of $\Omega_{\text{eff},0}$ and δ. *Right:* Difference between the experimentally measured and simulated population inversions. 141
7.13 Measured populations after an SCI beamsplitter sequence, where the last beamsplitter pulse is delayed in time to prevent interference. Populations are shown as a function of the peak Bragg pulse intensity for the first two ($\Omega_{\text{eff},0}^{(1,2)}$) and last two ($\Omega_{\text{eff},0}^{(3,4)}$) Bragg pulses. The simulation parameters are determined by fitting a dataset with the first two Bragg beamsplitters to a simulation, and then using those parameters to simulate the more complex four beamsplitter case shown here. 143

7.14 Measured and simulated populations after a single $N = 10$ Bloch oscillation pulse with $\tau = 7.2\,\text{ms}$. 144
7.15 Experimentally measured and simulated SCI ellipses. 145
7.16 Experimentally measured and simulated SCI ellipse parameters extracted via a fit. 146
7.17 The simultaneous conjugate Mach-Zehnder interferometer geometry overlaid on a partially transparent simultaneous conjugate Ramsey-Bordé interferometer geometry. Bloch oscillations (BO) are shown in orange. The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters are applied. 148
7.18 An SCI implemented with multiple 1st-order Bragg diffraction pulses. Trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters or mirrors are applied. The pulse type ($\pi/2$ beamsplitter or π mirror) is shown underneath each vertical line. 149
7.19 Experimentally measured contrast and differential phase ϕ_d for several $\pi/2$ and π pulse intensities used for the multiple 1st-order Bragg pulse SCI geometry for $n_{\text{target}} = 4$. The interferometer uses $T = 10\,\text{ms}$, $T' = 2\,\text{ms}$, and $T_s = 250\,\mu\text{s}$. All pulses are described by Eq. (2.20) with $\sigma = 30\,\mu\text{s}$. 151
7.20 Experimentally measured ellipse contrast as a function of n_{target}. Here $\sigma = 30\,\mu\text{s}$, $T = 10\,\text{ms}$, $T' = 2\,\text{ms}$, $T_s = 250\,\mu\text{s}$, $\Omega_{\text{eff},0}^{(\pi/2)}$ uses a 39 mV AWG peak pulse amplitude and $\Omega_{\text{eff},0}^{(\pi)}$ uses a 66 mV AWG peak pulse amplitude. 151
7.21 Experimentally measured ellipse contrast as a function of T for the SCI composed of multiple 1st-order Bragg pulses for several different n_{target} and interferometry beam fiber outcoupler focal lengths. Here $\sigma = 30\,\mu\text{s}$, $T' = 2\,\text{ms}$, and $T_s = 250\,\mu\text{s}$. $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ are set by performing the scan in Fig. 7.19 and choosing the values that maximize the contrast. 152
7.22 Experimentally measured interferometer per-shot sensitivity to the recoil frequency as a function of T for the SCI composed of multiple 1st-order Bragg pulses for several different n_{target} and interferometry beam fiber outcoupler focal lengths. Here $\sigma = 30\,\mu\text{s}$, $T' = 2\,\text{ms}$, and $T_s = 250\,\mu\text{s}$. $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ are set by performing the scan in Fig. 7.19 and choosing the values that maximize the contrast. Note that no Bloch oscillations are used here. 153
7.23 Simulated interferometer contrast as a function of the peak effective Rabi frequency $\Omega_{\text{eff},0}$ for several different Gaussian pulse widths σ. As σ is varied, different values of $\Omega_{\text{eff},0}$ are needed to produce a $\pi/2$ pulse. The x-axis is shown relative to the optimal $\Omega_{\text{eff},0}^{\text{opt}}$. Several types of interferometers are shown: SCRBI, SCMZI, and an SCRBI composed of multiple $n = 1$ Bragg pulses. 154

7.24 Simulated interferometer phase as a function of the peak effective Rabi frequency $\Omega_{\text{eff},0}$ for several different Gaussian pulse widths σ. As σ is varied, different values of $\Omega_{\text{eff},0}$ are needed to produce a $\pi/2$ pulse. The x-axis is shown relative to the optimal $\Omega_{\text{eff},0}^{\text{opt}}$. Several types of interferometers are shown: SCRBI, SCMZI, and an SCRBI composed of multiple $n = 1$ Bragg pulses. 155
7.25 Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of δ_{offset}, which is the offset of the Bragg two-photon detuning, δ, from resonance. Several Bloch orders, N, are shown. Note that the $N = 20$ and $N = 50$ cases overlap for both the SCMZI and SCRBI. 155
7.26 Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of t_0, the time of the first Bragg pulse, for $N = 10$. 156
7.27 Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of T. Several Bloch orders, N, are shown. 156
7.28 Simulated dependence of the interferometer phase on T for SCRBI and SCMZI interferometers with $N = 50$ Bloch oscillations. The phase is the same as that reported in Fig. 7.27, but ϕ_{offset} has been subtracted, where ϕ_{offset} is the interferometer phase for either the SCRBI or SCMZI at $T = 5\,\text{ms}$. 157
7.29 Simulation of a 4th-order Bragg diffraction process where the two-photon detuning is generated through serrodyne modulation. The simulation is made by integrating Eq. (A.51). *Top:* The effective Rabi frequency of the Bragg pulse as a function of time. *Middle:* The phase modulation applied to the Bragg process as a function of time. *Bottom:* The population in momentum space. 159
7.30 Simulated population in the initial $0\hbar k$ state and final $8\hbar k$ state for a 4th-order Bragg pulse as a function of the number of Fourier modes included in the phase modulation waveform. Leakage to unwanted states is also plotted. 159
7.31 Simulated population in the initial $0\hbar k$ state and final $8\hbar k$ state for a 4th-order Bragg pulse as a function of the amplitude of the phase modulation. Leakage to unwanted states is also plotted. 160
7.32 Simulated population inversion between the $0\hbar k$ and $8\hbar k$ momentum states as a function of $\Omega_{\text{eff},0}$ and the phase of the sawtooth waveform θ_{serr}. 160
7.33 Experimentally measured Bragg inversion as a function of $\Omega_{\text{eff},0}$ for several falling ramp durations and Bragg orders. The normal line denotes a Bragg diffraction pulse performed without serrodyne-modulation. 161
7.34 Experimentally measured population leakage to unwanted momentum states as a function of $\Omega_{\text{eff},0}$ for several falling ramp durations and Bragg orders. The normal line denotes a Bragg diffraction pulse performed without serrodyne-modulation. 162
7.35 Simulated population in the initial $0\hbar k$ state and final $\pm 8\hbar k$ states for 4th-order double Bragg diffraction. 162
7.36 Experimentally measured MZI fringes composed of normal and serrodyne-modulated Bragg pulses for different Bragg orders. Dots are experimentally recorded data points. Dashed lines are fits to those points. 163

7.37 Experimentally measured MZI fringe contrast as a function of the falling ramp duration for different Bragg orders. Contrasts are measured by fitting a sinusoid to MZI fringes as in Fig. 7.36 and extracting the amplitude. The theoretical maximum contrast of the interferometer is 1, and in practice is limited to a maximum of approximately 0.6 due to the variation of Ω_{eff} across the atom cloud. 164
7.38 Experimentally measured ellipse contrast as a function of the delay, T_{delay}, of the fourth beamsplitter pulse. As the Raman pulse duration, τ, increases and the double VS pulse process becomes more velocity selective, the position space wavefunction becomes wider and the contrast decays more slowly as a function of T_{delay}. 165
7.39 Experimentally measured ellipse contrast as a function of the delay, T_{delay}, of the fourth beamsplitter pulse. As the Raman pulse duration, τ, increases and the VS pulse becomes more velocity selective, the position space wavefunction becomes wider and the contrast decays more slowly as a function of T_{delay}. Two VS sequences are shown, one with the standard two $\tau = 400$µs pulses, and another with the standard two pulses followed by a single 1200µs square Raman pulse. . 166

B.1 The 3D MOT coil switching circuit. 182
B.2 The 3D MOT coil switching circuit turn off. The lower dashed line is at $-15\,\text{V}$, which we expect the MOSFET drain voltage to be close to before the switch. The upper dashed line is at $0\,\text{V}$, which we expect the probe voltage to decay to once no current flows through the MOT coil. 182
B.3 Basic coil switching circuit. 183
B.4 Coil driver circuit. 183

List of Tables

2.1 Parameters used for all numerical integrations in this chapter, unless specified otherwise. 13

5.1 Comparison of current experiment performance with possible (but optimistic) and idealized performance. 106

6.1 Per-atom integration times for the DOP853, parallelized ODE, and lookup table methods. 112

7.1 Error budget of the 2018 Berkeley fine-structure constant measurement [3]. The thermal motion of atoms and speckle phase shift systematics produce systematic shifts via the Gouy phase and diffraction phase mechanisms. For this reason, the shifts they produce are included in the Gouy phase shift (and their shifts are listed as zero), but their errors are listed separately. The thermal motion of atoms error, for example, describes the uncertainty in the Gouy phase shift (and therefore α) resulting from the uncertainty surrounding atom positions in the cloud over the interferometer duration. This is true for the speckle phase shift and parasitic interferometer systematics as well. 125

7.2 Pulse sequence for the stacked first-order Bragg SCI geometry. 150

Acknowledgments

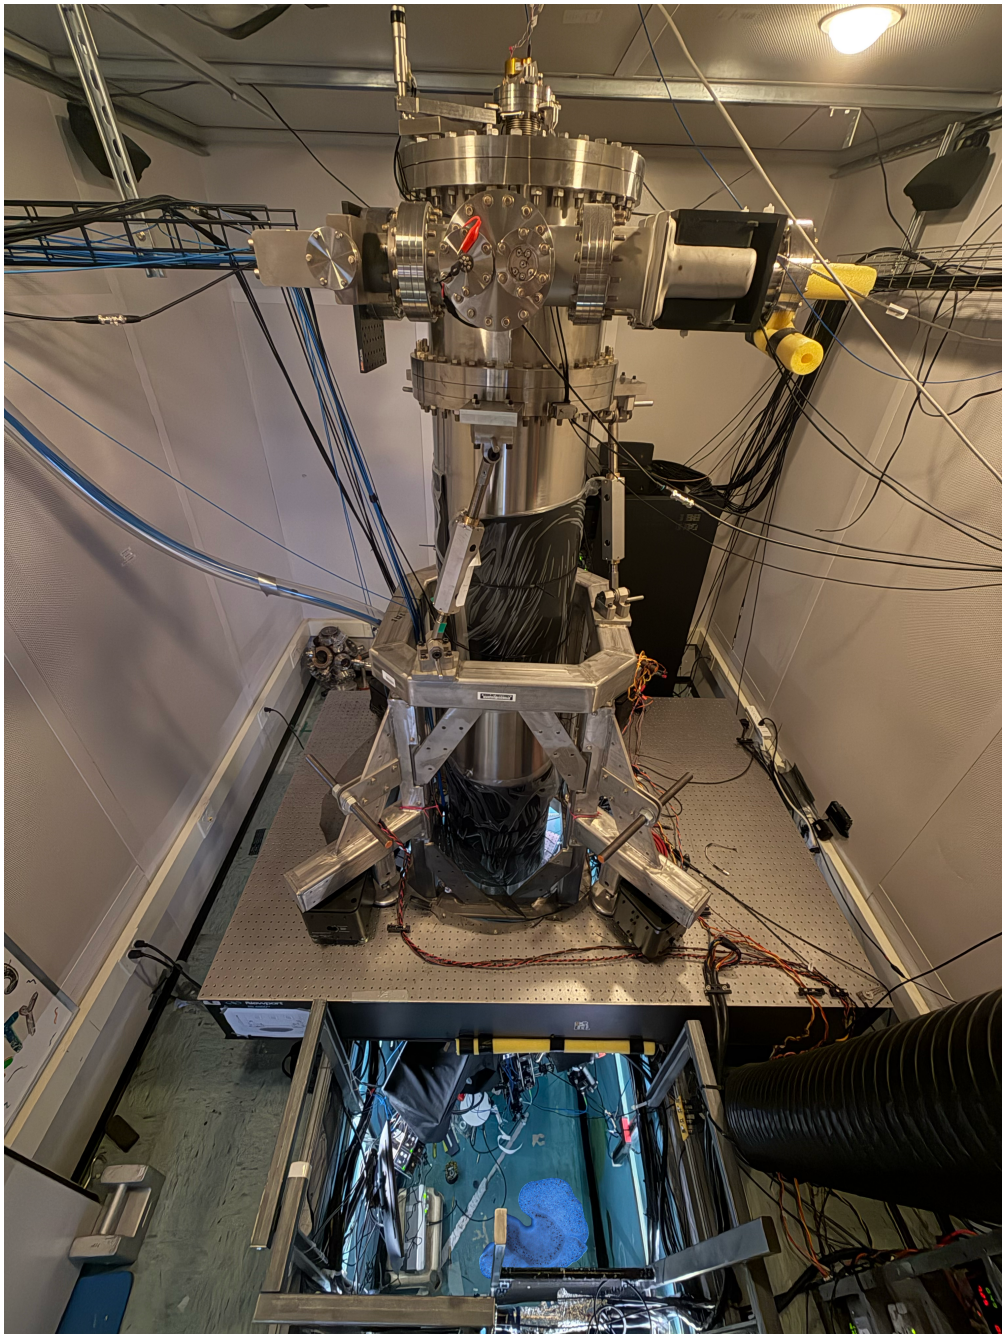
I would like to thank all those who supported me during my PhD and throughout my scientific training. First, I would like to thank the Physics Department administrative and support staff. I would especially like to thank Anthony Vitan, Espe Munoz-Riddle, Carlos Bustamante, and Jon Deasy. This work would not have been possible without their support and the support of many others at Berkeley.

I thank all members of the Mueller group for creating a welcoming and collaborative research group. In particular, I would like to thank Ashwin Singh, Eric Cooper, Garrett Louie, Isaac Pope, James Egelhoff, Jeremy Axelrod, Jessie Zhang, Lothar Maisenbacher, Matt Tao, Peter Petrov, Peter Stromberger, and Storm Weiner. I am especially grateful to the other members of the alpha experiment, who often spent many long hours in the pit in exchange for a functional experiment: Zachary Pagel, Yair Segev, Madeline Bernstein, Andrew Christensen, and Yuno Iwasaki.

Next, I would like to thank my advisor, Holger Mueller, for his insight and guidance throughout the process of building, debugging, and running the alpha experiment. I would also like to thank Holger for passing along his knowledge of analog electronics, optics, and laser systems.

Before graduate school, I had many advisors and mentors as a physics undergraduate. In particular, I would like to thank Hong Tang for giving me the opportunity to work in a research lab for the first time, Benjamin Ofori-Okai for advising me through my first real simulation project, and David DeMille for giving me the opportunity to join his lab and for introducing me to experimental atomic, molecular, and optical physics. I would also like to thank Maximilian Beyer for his mentorship and friendship, and for pushing me to approach research in an efficient, effective, and enjoyable way.

I am very grateful to my friends and family for their support over the past six years. I thank my siblings, Max and Helen, for growing up alongside me and for being my very first teammates. I thank my parents, Alex and Erin, for supporting me in whatever I choose to do. I especially thank my wife, Charlotte, who, despite having her hands full with her own PhD work, found time to support me through mine and to have many adventures along the way.



Chapter 1

Introduction

This thesis describes the atom interferometer that I have spent my PhD building, operating, and simulating. The atom interferometer is designed to measure the cesium recoil frequency, from which the fine-structure constant can be determined. This determination amounts to a precision test of the Standard Model. In this chapter I will provide an overview of the mechanism behind atom interferometry, a description of what the cesium recoil frequency is, a description of the fine-structure constant, and an explanation of why measuring the cesium recoil frequency tells you the fine-structure constant.

1.1 Atom interferometry

We will begin by reviewing the concept of interference. In Fig. 1.1 we have drawn a red wave and a blue wave of the same frequency and with phase difference θ. Since the waves are overlapping in space, they will add to form the green wave seen in the figure. As θ is scanned the green wave will change adding destructively to constructively. This is called interference. An interferometer is any device that measures the amplitude of the green wave in order to determine the phase difference θ between the red and blue waves.

Interferometers are useful because the phase difference between the two interfering waves can contain interesting information. For example, an optical interferometer can be constructed by sending two waves derived from the same light source through materials with differing indices of refraction. Because the wavelength of light is modified by the index of refraction, the two waves will accumulate different phases as they pass through the material. When the two waves add, the amplitude of the resulting wave provides information about the phase difference between the two initial waves. This phase difference encodes information about the difference in refractive index between the two materials.

In theory, one could build an interferometer using any wave. To do this you need: 1) a way to generate the wave; 2) a way to split the wave along different paths; 3) a way to measure the amplitude of the wave. For an optical interferometer, the wave is an electromagnetic wave generated by any light source (often a laser), the light-wave is split using a beamsplitter

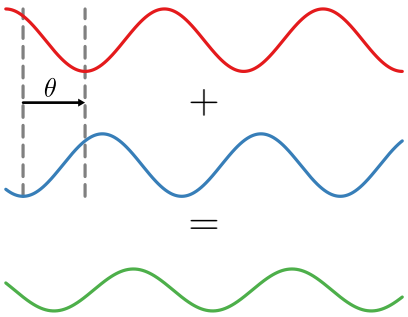


Figure 1.1: Illustration of two sinusoids with phase difference θ. When the two sinusoids are summed along the x-axis, the resulting waveform reveals mostly destructive interference.

(often a partially transparent mirror), and the wave intensity is measured by a photodiode (or the human eye).

In an atom interferometer, the wave is the quantum mechanical wavefunction of an atom, as opposed to the electromagnetic wave of an optical interferometer[1]. The propagation of the wave in time and space is described by the Schrödinger equation, as opposed to Maxwell's equations in the case of an optical interferometer. It is nontrivial to prepare an atom such that its quantum mechanical wavefunction can be used for an atom interferometer, as we will discuss in Section 4.4. The atom wavefunction can be split along different paths using a variety of different methods. In this experiment we use laser-based methods, which are described in Chapter 2.

To measure the output of an atom interferometer, light resonant with the atom is turned on after interferometry is completed, causing the atom to fluoresce and reveal its position. The fluorescence collapses the wavefunction and localizes the atom to a single interferometer output. By running the interferometer many times, we can determine the probability of the atom exiting any given output port, which gives us something analogous to the intensity measured in a light interferometer described using classical electrodynamics.

To compute the phase accumulated in an atom interferometer we must integrate the time-dependent 1D Schrödinger equation:

$$i\hbar\frac{\partial}{\partial t}\psi(t,z) = \left[-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2} + V(t,z)\right]\psi(t,z) \tag{1.1}$$

Most of the time, $V(t,z) = 0$ in an atom interferometer, and the Schrödinger equation

[1]Throughout this thesis I primarily use the name "atom interferometer". The name "matterwave interferometer" is another commonly used term when atoms are not used [1, 2], and is generally interchangeable with "atom interferometer" for the purpose of this thesis.

describes a free particle

$$i\hbar\frac{\partial}{\partial t}\psi(t,z) = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\psi(t,z) \tag{1.2}$$

In this case, it is useful to represent the wavefunction as a wavepacket

$$\psi(t,z) = \frac{1}{\sqrt{2\pi\hbar}}\int_{-\infty}^{\infty}\phi(t,p)\exp(ipz/\hbar)\,\mathrm{d}p \tag{1.3}$$

Plugging this form for $\psi(t,z)$ into the Schrödinger equation and solving we find

$$\phi(t,p) = \phi(0,p)\exp\left(-i\frac{p^2}{2m\hbar}t\right) \tag{1.4}$$

And so we see the part of the wavepacket with momentum p accumulates phase at a rate of $p^2/2m\hbar$ relative to the part of the wavepacket with momentum 0. In the atom interferometer used for this experiment, the atom wavefunction is split into two parts, one of which has $p = 0$ and the other of which has $p = 2n\hbar k$, where k is the laser wavenumber and n is an integer called the diffraction order. k can be precisely measured, and n can be chosen in the lab. Therefore the phase difference between the two parts of the wavefunction in the interferometer will be $2\hbar n^2k^2t/m$ for an interferometer of duration t. We define the rate at which phase is accumulated from a single photon recoil of $\hbar k$ as the recoil frequency $\omega_\mathrm{r} = \hbar k^2/2m$. The accumulated phase then becomes $4n^2\omega_\mathrm{r}t$. Since n and t are set in the lab, a measurement of this phase amounts to a measurement of the recoil frequency ω_r. By measuring the laser frequency precisely, k can be determined as well, which allows a precise value of the particle mass, m, to be determined from ω_r! The interferometer described in this thesis creates interferometers using cesium atoms, and so the experiment measures m_Cs. As we will see in the next section, a measurement of m_Cs can be used to determine the value of the fine-structure constant.

To split the atom wavefunction, lasers are used to create a particular form for $V(t,z)$ that splits the atom wavefunction into two momentum states. This is described in Chapter 2. In practice, atom interferometers are susceptible to many additional phases in addition to the phase proportional to ω_r. These are described, along with additional considerations necessary for understanding atom interferometers, in Chapter 3.

1.2 The fine-structure constant

The fine-structure constant α quantifies the strength of the electromagnetic interaction in the Standard Model [3]. It is defined as

$$\alpha = \frac{e^2}{4\pi\epsilon_0\hbar c} \tag{1.5}$$

The fine-structure constant appears in many areas of physics. The first experimental hint of the fine-structure constant was identified by Michelson and Morley in their 1887 paper *On a Method of making the Wave-length of Sodium Light the actual and practical Standard of Length* [4, 5]. They were attempting to develop a length standard with sodium light using the same Michelson interferometer they built for the Michelson-Morley experiment. Using a sodium lamp as the light source for their interferometer, they counted the number of observed interference fringes as they scanned the position of a mirror on one of the interferometer arms. While scanning the mirror position, they noticed that at certain positions the interference would completely disappear. It was already known, as Michelson and Morley put it, that "there are two series of waves in sodium-light" from the astronomical observations of Fraunhofer [6]. They understood that the two wavelengths emitted by the sodium lamp were beating together to cause the fringe to disappear at certain locations. As part of their investigation, they attempted the same measurement with hydrogen, finding that it also emitted two wavelengths [4]. Thus, they had determined that one of the spectral lines in hydrogen was actually two lines very close in wavelength. The paper is very short, and is well worth a read.

In 1915 Bohr showed that his model of the atom reproduced most of the then observed spectral lines in hydrogen. However, it only predicted a single line where Michelson and Morley had observed two [7]. In 1916 Sommerfeld extended Bohr's model, showing that Bohr's single line split in two if a relativistic correction was included in the model. In computing the splitting he wrote down the ratio p_0/p_1, where p_1 was the angular momentum of an electron in hydrogen derived by Bohr and p_0 was the relativistic correction derived by Sommerfeld. Sommerfeld named the ratio the fine-structure constant and defined it in Eq. (12a) of his paper [8]. His definition is equivalent to our definition in Eq. (1.5).

Sommerfeld computed α to be around 7×10^{-3}. Now, 110 years later, α is known to much higher precision. Note that α is unitless. It is unknown why α takes on this value [9]; it seems to be an intrinsic property of the universe. A more complete history of the fine-structure constant is provided in [10].

To compute the value of α we require a measurement of ϵ_0, as the SI unit system fixes values for e, $\hbar$, and c. In practice this is done by deriving ϵ_0 from several other well-measured values. Here we rewrite α in terms of the Rydberg constant R_∞, the mass ratio $m_{\mathrm{Cs}}/m_{\mathrm{e}}$, the cesium mass:

$$\alpha^2 = \frac{2R_\infty}{c}\frac{h}{m_{\mathrm{e}}} = \frac{2R_\infty}{c}\frac{m_{\mathrm{Cs}}}{m_{\mathrm{e}}}\frac{h}{m_{\mathrm{Cs}}} \tag{1.6}$$

The Rydberg constant has been measured to 1 ppt [11] by hydrogen spectroscopy [12]. The mass ratio $m_{\mathrm{Cs}}/m_{\mathrm{e}}$ is determined by combining measurements of the electron mass in atomic mass units m_{e}/u and the cesium mass in atomic mass units m_{Cs}/u made in an ion trap. These contribute uncertainties 20 ppt [13] and 65 ppt [14] to a cesium recoil determination of the fine-structure constant, respectively. The remaining undetermined value is m_{Cs}, which can be measured using an atom interferometer as we found in the previous section. Other

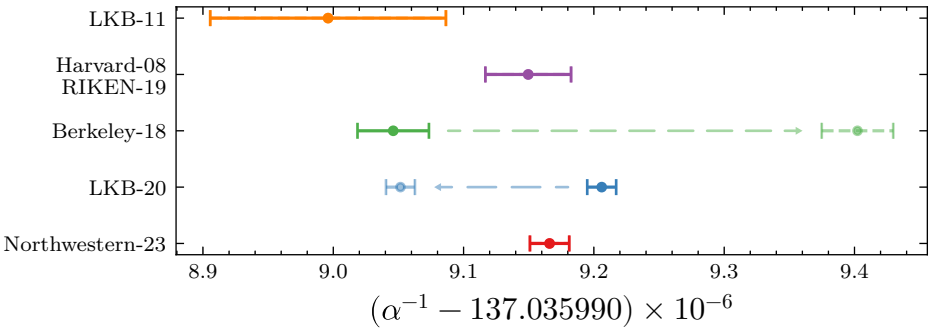


Figure 1.2: Current determinations of the fine-structure constant. The partially transparent data points for the Berkeley-18 and LKB-20 measurements are the values of those measurements when beam-related effects are ignored. Data sources are cited in the text.

methods for measuring the fine-structure constant are discussed in [15], although some of these are no longer valid measurement approaches after the 2019 SI unit redefinition.

We are interested in determining α to even higher precision because it will enable a precise test of the Standard Model. The Standard Model can be used to calculate the electron magnetic moment $g/2$, using α as an input. Thus, the precision to which $g/2$ can be computed is limited by the precision at which α is known. Currently calculations of $g/2$ have comparable precision to direct measurements of $g/2$, and so improving measurements of α and $g/2$ amounts to a precise test of the Standard Model. $g/2$ is measured to high precision in a Penning trap at Northwestern [16].

The calculation of $g/2$ is a power series in α with additional contributions coming from the muon, tauon, hadronic, and weak interaction terms [16]

$$\frac{g}{2} = 1 + \sum_n a_n \left(\frac{\alpha}{\pi}\right)^n + a_{\mu\tau} + a_{\text{hadronic}} + a_{\text{weak}} \tag{1.7}$$

The a_n terms are computed via quantum electrodynamics. Performing this comparison at higher and higher precisions makes it sensitive to smaller and smaller effects. A deviation between the computed and measured values of $g/2$ could be evidence for new physics.

The current state of measurements of α and $g/2$ is shown in Fig. 1.2. The first displayed recoil measurement was made in 2011 at Laboratoire Kastler Brossel (LKB) [17], followed by a measurement made in our group at Berkeley in 2018 [3], and then another LKB measurement in 2020 [18]. The reported values are the solid points in the figure. The LKB interferometer utilizes rubidium atoms, as opposed to the cesium atoms used in the Berkeley measurement. All determinations of α from measurements of $g/2$ (computed by inverting Eq. (1.7)) were made by Gerald Gabrielse's group, first at Harvard in 2008 [19] and more recently at Northwestern in 2023 [16]. A more precise theoretical calculation of α from the

2008 Harvard $g/2$ measurement was made by RIKEN in 2019 [20]. This result is displayed on the plot. The 2018 Berkeley, 2020 LKB, and 2023 Northwestern measurements contribute to the CODATA 2022 α value of $\alpha^{-1} = 137.035999177(21)$ [11]. The 2011 LKB measurement did not contribute to this value.

The Berkeley measurement is consistent with the 2011 LKB recoil measurement and the Harvard/RIKEN $g/2$ measurement. However, it is in tension with the 2020 LKB recoil measurement, and, since 2023, the latest $g/2$ measurement. It seems unlikely that this discrepancy is due to new physics or due to errors in the mass ratio measurements made for cesium and/or rubidium. It is more likely that the observed discrepancy comes from mischaracterized systematic shifts related to the atom-laser interaction. The Berkeley measurement and the 2020 LKB measurement report large systematic shifts due to the interaction of atoms in the interferometer with the laser responsible for splitting the atom wavefunction. Removing these shifts leads to the partially transparent points in Fig. 1.2. It is apparent that these shifts are significantly larger than the reported error, and so it is critical that they be precisely understood. These systematics heavily influenced the design of the second-generation experiment apparatus built at Berkeley and described in this thesis. The design should mitigate many of the systematic effects identified in the 2018 measurement, as described in Section 7.1. Substantial effort has also been put into developing models of these systematics, as is described in this thesis.

1.3 Outline of thesis

The thesis is organized as follows. First the concept of matterwave optics is introduced in Chapter 2. This discusses how lasers can be used to create beamsplitters for atom wavefunctions, as well as a few other related techniques used in the experiment. Chapter 3 discusses how beamsplitters for atoms can be used to create atom interferometers sensitive to the recoil frequency ω_{r}. Several different kinds of interferometers are discussed. Both chapters review information that (mostly) already exists in the literature, but present it in a way that reflects how I think about atom interferometry. For a description of the 2018 Berkeley measurement see the theses of Brian Estey [21], Chenghui Yu [22], or Weicheng Zhong [23]. For an earlier description of the second-generation Berkeley measurement, see Zachary Pagel's thesis [24].

Chapter 4 onwards describes novel work carried out over the course of this thesis. In Chapter 4, I describe the design of the experimental apparatus used for running the Berkeley atom interferometer. I review the vacuum chamber, magnetic field coils, laser systems, the laser frequency stabilization technique (resulting in a publication [25]), several custom electrical circuits, the control system, the state preparation sequence, the interferometry sequence, and optical alignment procedures. I conclude with a brief discussion of the lessons learned for designing a robust experiment. In Chapter 5, I discuss how the number of atoms propagating through the interferometer was measured, and sources of noise in the experiment. In Chapter 6, I discuss the simulation methods developed to characterize the systematic effects present in the experiment, and discuss how these methods are related to

each other. I also discuss a novel simulation technique that resulted in a paper [26]. In Chapter 7, I discuss several novel investigations made using the experiment apparatus. This chapter includes extensive comparisons of experimentally measured interferometer output with simulated interferometer output, finding agreement and grounding future simulations of systematic effects in a verified simulation procedure. There is also a discussion of some simulated results that have not been experimentally tested. This chapter also includes a description of a novel beamsplitter technique that resulted in a paper [27]. The thesis ends with Chapter 8, which describes improvements that might be made to future generations of the experiment to further reduce systematic effects or otherwise improve its performance. Finally, Appendix A contains math supporting derivations made elsewhere in the thesis, and Appendix B contains several circuit diagrams that are not included earlier in the thesis.

Chapter 2

Matterwave Optics

In this chapter we will discuss the matterwave optics used to split atom wavefunctions into a superposition of momentum states. Some matterwave optics, such as mirrors and beamsplitters, are built from matter, such as nanostructured diffraction gratings [28]. Others are based on the interaction between laser light and an atom wavefunction. All of the optics discussed in this thesis are light-based. In light-based matterwave optics, the basic interaction at play is the exchange of momentum between the atom wavefunction and a light field through the absorption and emission of photons. This interaction can happen through the exchange of a single photon [29–31], or multiple photons.

Since the atom interferometer developed in this thesis uses cesium atoms, single-photon transitions cannot be used as the excited state would decay before an atom interferometer could be formed (the excited state lifetime is only $\sim$30 ns [32]). Instead we use several types of multiphoton transitions that leave the atom in a ground state: Bragg diffraction [33–35], Bloch oscillations [36–39], and Raman transitions [40, 41]. Each transition occurs in the presence of two laser fields–one propagating to the right and the other to the left, as in Fig. 2.1–and can be explained as the absorption of a photon from one of the beams (resulting in a kick of $+\hbar k$) followed by the stimulated emission of a photon into the other beam (resulting in another kick of $+\hbar k$) for a total momentum transfer of $+2\hbar k$. To satisfy energy conservation the frequency difference between the two counterpropagating beams must be $\delta = \pm 4\omega_{\mathrm{r}}$ (recall that we defined $\omega_{\mathrm{r}} = \hbar k^2/2m$ in Section 1.1). The sign of δ determines the direction of momentum transfer.

In an atom interferometer, matterwave optics play roles analogous to optical beamsplitters and mirrors by directing atom wavefunctions along different paths in time and one-dimensional space, similar to how optical beamsplitters and mirrors reflect light waves along different paths in two-dimensional space. For this reason we often refer to matterwave optics as beamsplitters or mirrors, where a beamsplitter creates a 50-50 superposition of momentum states and a mirror transitions one momentum state to another. I will often refer to the state produced by a beamsplitter or mirror as the diffracted state. In Chapter 3 we will discuss how matterwave optics can be applied repeatedly to form an atom interferometer.

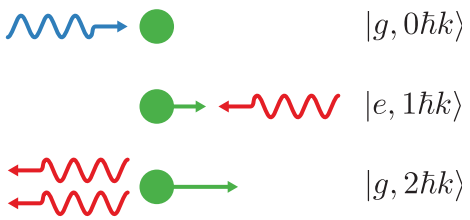


Figure 2.1: Schematic of stimulated absorption followed by stimulated emission of a photon by an atom. Initially the atom is in the ground state with zero momentum $|g, 0\hbar k\rangle$. A right-moving photon (displayed in blue) excites the atom and imparts a momentum impulse, putting it in state $|e, 1\hbar k\rangle$. Next, a left-moving photon stimulates emission of a photon, causing the atom to enter the ground state again but with the momentum kick from two photons, putting it in state $|g, 2\hbar k\rangle$. Note that the left-moving photons are slightly red-detuned relative to the right-moving photon, as the atom motion in the $|e, 1\hbar k\rangle$ state Doppler shifts the atomic resonance.

2.1 Two-level system in an off-resonant bichromatic field

Bragg diffraction and Bloch oscillations use the potential resulting from an optical lattice to diffract an atomic wavefunction [33, 36]. The optical lattice is formed by counterpropagating off-resonant laser beams. Their frequency difference sets the lattice velocity. The equations of motion describing Bragg diffraction and Bloch oscillations can be derived by considering a two-level system with a one-dimensional external degree of freedom in an off-resonant bichromatic field

$$\hat{H}_{\text{Bragg}} = \frac{\hat{p}^2}{2m} + \hbar\omega_e|e\rangle\langle e| + \hbar\sum_\ell \left[\Omega_\ell(t)\cos(k_\ell z - \omega_\ell t)|e\rangle\langle g| + \text{h.c.}\right] \tag{2.1}$$

where ω_e is the excited state frequency, Ω_ℓ is the Rabi frequency as defined by Eq. (A.14), k_ℓ is the wavevector, and ω_ℓ is the field frequency. For a bichromatic field the sum is over two fields. The same Hamiltonian describes Bragg diffraction and Bloch oscillations.

The first step is to transform into the interaction picture via

$$U = 1|g\rangle\langle g| + \exp(i\omega_e t)|e\rangle\langle e| \tag{2.2}$$

Applying the transformation as usual (see Theorem A.1) and applying the rotating wave approximation leads to

$$\widetilde{\hat{H}}_{\text{Bragg}} = \frac{\hat{p}^2}{2m} + \frac{\hbar}{2}\sum_\ell \left[\Omega_\ell(t)\exp(i(k_\ell z - \omega_\ell t + \omega_e t))|e\rangle\langle g| + \text{h.c.}\right] \tag{2.3}$$

Plugging into the Schrödinger equation we find

$$i\hbar\frac{\partial\big|\widetilde{\psi}\big\rangle}{\partial t} = \widehat{\widetilde{H}}_{\text{Bragg}}\big|\widetilde{\psi}\big\rangle = \frac{\hat{p}^2}{2m}\big|\widetilde{\psi}\big\rangle + \frac{\hbar}{2}\sum_{\ell}\left[\Omega_\ell(t)\exp(i(k_\ell z - \omega_\ell t + \omega_e t))|e\rangle\langle g| + \text{h.c.}\right]\big|\widetilde{\psi}\big\rangle \tag{2.4}$$

At this point we must decide to compute the dynamics in position or momentum space. Here we will proceed in momentum space, leaving the position space derivation to Appendix A.6. We successively apply $\langle p|\otimes\langle g|$ and $\langle p|\otimes\langle e|$ to the left-hand side of Eq. (2.4). Using Theorem A.2 to evaluate the matrix elements we find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_g(t,p) =& \frac{p^2}{2m}\tilde{c}_g(t,p) + \frac{\hbar}{2}\sum_{\ell}\Omega_\ell(t)\exp(-i(-\omega_\ell+\omega_e)t)\tilde{c}_e(t,p+\hbar k_\ell) \\ i\hbar\dot{\tilde{c}}_e(t,p) =& \frac{p^2}{2m}\tilde{c}_e(t,p) + \frac{\hbar}{2}\sum_{\ell}\Omega_\ell(t)\exp(i(-\omega_\ell+\omega_e)t)\tilde{c}_g(t,p-\hbar k_\ell) \end{aligned} \tag{2.5}$$

where $\tilde{c}_g(t,p) = \langle p|\otimes\langle g|\big|\widetilde{\psi}\big\rangle$ and $\tilde{c}_e(t,p) = \langle p|\otimes\langle e|\big|\widetilde{\psi}\big\rangle$. If we assume a far-detuned field such that $\Omega_\ell \ll \omega_\ell - \omega_e$ very little amplitude should flow through the excited state[1]. Therefore we should be able to adiabatically eliminate $\tilde{c}_e(t,p)$ from our expression for $\tilde{c}_g(t,p)$. To do this, we first transform the excited state expression to a frame rotating with the kinetic energy via $\tilde{c}_e(t,p) = c'_e(t,p)\exp(-ip^2t/2m\hbar)$

$$i\hbar\dot{c}'_e(t,p) = \frac{\hbar}{2}\sum_{\ell}\Omega_\ell(t)\exp(i(p^2/2m\hbar - \omega_\ell + \omega_e)t)\tilde{c}_g(t,p-\hbar k_\ell) \tag{2.6}$$

We then directly integrate this expression to find a form for $c'_e(t,p)$. For a generic two-level system, the optical Bloch equations are

$$\begin{aligned} \dot{a}_g(t) &= \exp(i\beta t)f(t)a_e(t) \\ \dot{a}_e(t) &= \exp(i\alpha t)h(t)a_g(t) \end{aligned} \tag{2.7}$$

where f and h depend on the form of the particular Hamiltonian. Therefore we can compute $a_e(t)$ via

$$a_e(t) = \int_0^t \exp(i\alpha t')h(t')a_g(t')\,\mathrm{d}t' \tag{2.8}$$

We can rewrite this integral exactly using Theorem A.3

$$\begin{aligned} a_e(t) = \int_0^t \exp(i\alpha t')h(t')a_g(t')\,\mathrm{d}t' =& -\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1} e^{i\alpha t'}\frac{\partial^k}{\partial t^k}\left[h(t')a_g(t')\right]\Big|_0^t \\ &+ \left(\frac{i}{\alpha}\right)^{m+1}\int_0^t e^{i\alpha t'}\frac{\partial^{m+1}}{\partial t^{m+1}}\left[h(t')a_g(t')\right]\,\mathrm{d}t' \end{aligned} \tag{2.9}$$

[1] When the system is close to resonance a similar experimental setup can be used to produce the bichromatic force [42–44].

If $\alpha \gg h(t')a_g(t')$ and $\partial^k/\partial t^k h(t')a_g(t') < \partial^{k-1}/\partial t^{k-1} h(t')a_g(t')$ (these are the approximations required by the adiabatic elimination) then we can drop the integral term on the right hand side and truncate the sum after $k = 0$ to find

$$a_e(t) \approx \frac{i}{\alpha}\left[h(0)a_g(0) - h(t)a_g(t)e^{i\alpha t}\right] \tag{2.10}$$

Applying this prescription to $i\hbar\dot{c}'_e(t,p)$, and under the assumption $\Omega_\ell \ll \Delta$ for all ℓ, $c'_e(0,p) = 0$, and $\Omega_\ell(0) = 0$, we find

$$\begin{aligned} i\hbar c'_e(t,p) =& i\hbar \int_0^t \dot{c}'_e(t',p)\,\mathrm{d}t' \\ =& \frac{\hbar}{2}\sum_\ell \int_0^t \Omega_\ell(t')\exp(i(p^2/2m\hbar - \omega_\ell + \omega_e)t')\tilde{c}_g(t', p-\hbar k_\ell)\,\mathrm{d}t' \\ =& -i\frac{\hbar}{2}\sum_\ell \frac{\Omega_\ell(t)\tilde{c}_g(t,p-\hbar k_\ell)\exp(i(p^2/2m\hbar - \omega_\ell + \omega_e)t)}{p^2/2m\hbar - \omega_\ell + \omega_e} \end{aligned} \tag{2.11}$$

For the case of a bichromatic field with two counterpropagating lasers ($k = k_1 \approx -k_2$) our expression for $i\hbar\dot{\tilde{c}}_g(t,p)$ becomes

$$\begin{aligned} i\hbar\dot{\tilde{c}}_g(t,p) =& \frac{p^2}{2m}\tilde{c}_g(t,p) - \hbar\frac{\Omega_1(t)^2}{4\Delta}\tilde{c}_g(t,p) - \hbar\frac{\Omega_2(t)^2}{4\Delta}\tilde{c}_g(t,p) \\ & - \hbar\frac{\Omega_1(t)\Omega_2(t)}{4\Delta}\exp(i(\omega_1-\omega_2)t)\tilde{c}_g(t,p+2\hbar k) \\ & - \hbar\frac{\Omega_1(t)\Omega_2(t)}{4\Delta}\exp(-i(\omega_1-\omega_2)t)\tilde{c}_g(t,p-2\hbar k) \end{aligned} \tag{2.12}$$

where we made the approximation $(p+\hbar k_{1,2})^2/2m\hbar - \omega_{1,2} + \omega_e \approx \Delta$ in the denominators. In practice ω_1 and ω_2 are within a few megahertz of each other, and Δ is several gigahertz. We define

$$\Omega_1^{\mathrm{AC}}(t) = \frac{\Omega_1(t)^2}{4\Delta} \qquad \Omega_2^{\mathrm{AC}}(t) = \frac{\Omega_2(t)^2}{4\Delta} \qquad \Omega_{\mathrm{eff}}(t) = \frac{\Omega_1(t)\Omega_2(t)}{2\Delta} \qquad \delta = \omega_1 - \omega_2 \tag{2.13}$$

and switch to the notation $c_j(t) = c_g(t, 2j\hbar k)$

$$\begin{aligned} i\hbar\dot{c}_j(t) =& 4\hbar j^2\omega_{\mathrm{r}}c_j(t) - \hbar\frac{\Omega_{\mathrm{eff}}(t)}{2}\left[\exp(i\delta t)c_{j+1}(t) + \exp(-i\delta t)c_{j-1}(t)\right] \\ & - \hbar\Omega_1^{\mathrm{AC}}(t)c_j(t) - \hbar\Omega_2^{\mathrm{AC}}(t)c_j(t) \end{aligned} \tag{2.14}$$

where $\omega_{\mathrm{r}} = \hbar k^2/2m$. We see that the dynamics of an atom in a far-detuned bichromatic laser field are described by a ladder of momentum states spaced by $2\hbar k$. Note that we are not necessarily considering states with zero width in momentum space; all we have shown is

that there is a discrete coupling between momentum states. However, we will assume that initially the wavefunction has width less than $2\hbar k$.

From Eq. (2.14) we can state an effective Hamiltonian for the two-level system in a bichromatic field in the off-resonant limit

$$\begin{aligned}\hat{H}_{\text{eff}} =&\hbar \begin{pmatrix} \ddots & \ddots & 0 & 0 & 0 \\ \ddots & 4(j+1)^2\omega_{\text{r}} & -\Omega_{\text{eff}}(t)\exp(-i\delta t)/2 & 0 & 0 \\ 0 & -\Omega_{\text{eff}}(t)\exp(i\delta t)/2 & 4j^2\omega_{\text{r}} & -\Omega_{\text{eff}}(t)\exp(-i\delta t)/2 & 0 \\ 0 & 0 & -\Omega_{\text{eff}}(t)\exp(i\delta t)/2 & 4(j-1)^2\omega_{\text{r}} & \ddots \\ 0 & 0 & 0 & \ddots & \ddots \end{pmatrix} \\ =&\hbar \sum_{j=-\infty}^{\infty} \Bigg[4j^2\omega_{\text{r}} |2j\hbar k\rangle\langle 2j\hbar k| - \frac{\Omega_{\text{eff}}(t)}{2}\exp(-i\delta t)|2j\hbar k\rangle\langle 2(j-1)\hbar k| \\ &-\frac{\Omega_{\text{eff}}(t)}{2}\exp(i\delta t)|2j\hbar k\rangle\langle 2(j+1)\hbar k| \Bigg]\end{aligned} \tag{2.15}$$

with corresponding wavefunction

$$|\phi\rangle = \cdots + c_{j+1}(t)|2(j+1)\hbar k\rangle + c_j(t)|2j\hbar k\rangle + c_{j-1}(t)|2(j-1)\hbar k\rangle + ... \tag{2.16}$$

Note that the AC Stark shift terms are common to all states, and so we dropped them from the effective Hamiltonian. The eigenstates of the system are investigated in detail in [45]. It is convenient in certain cases to apply the unitary transformation

$$U = \exp\left(i\hat{H}_{\text{eff}} t\big|_{\Omega_{\text{eff}}(t)=0}/\hbar \right), \tag{2.17}$$

which stops the states from accumulating phase due to kinetic energy. This is useful when numerically integrating Bragg or Bloch dynamics.

2.2 Bragg diffraction

In the previous section we showed that a two-level system in a counterpropagating bichromatic field creates a ladder of momentum states spaced by intervals of $2\hbar k$. Bragg diffraction couples two states on this ladder through intermediate momentum states, as shown in Fig. 2.2. In the transition in the figure two photons will be absorbed from the blue laser field and two photons will be emitted into the red laser field, producing 2nd-order Bragg diffraction. The higher order the Bragg diffraction process, the higher peak effective Rabi frequency $\Omega_{\text{eff},0}$ required. The resonance condition for δ for Bragg diffraction of a particular order is easily determined from energy conservation

$$\hbar\delta = \left(\frac{p_f^2}{2m} - \frac{p_0^2}{2m} \right) / (n_f - n_0) = 4\left(n_0 + n_f\right)\hbar\omega_{\text{r}} \tag{2.18}$$

Table 2.1: Parameters used for all numerical integrations in this chapter, unless specified otherwise.

Parameter	**Description**
Bragg pulse shape	$\Omega_{\text{eff}}(t)$ is a truncated Gaussian described by Eq. (2.20) with $\sigma = 0.259\omega_{\text{r}}^{-1}$.
δ	δ is set to be resonant with the specified Bragg process. For adiabatic processes such as Bloch and adiabatic Bragg, δ is time-dependent.
n	All Bragg pulses are 4th-order pulses.
Spatial beam profile	The beam is assumed to have uniform intensity throughout space, meaning that $\Omega_{\text{eff},0}$ is constant.
Atom number	All numerical integrations are of a single atom.
Initial state	All amplitude is initially in $p = 0\hbar k$.

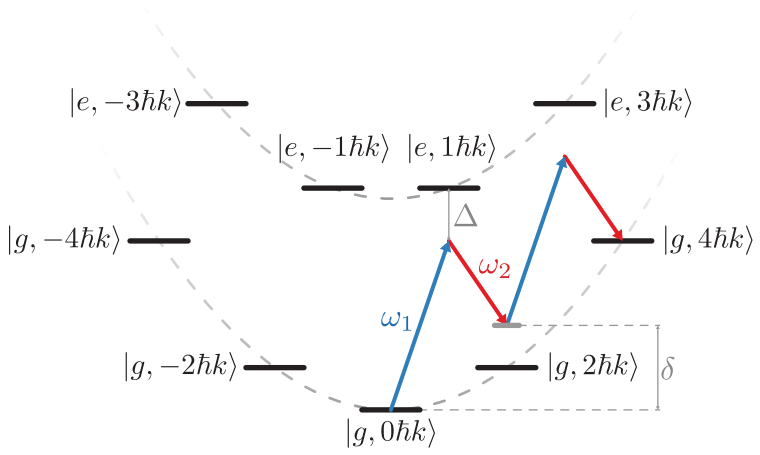


Figure 2.2: The level diagram of a two-level system with an external degree of freedom, with Bragg diffraction shown. The kinetic energy associated with the external degree of freedom is described by the gray dashed parabolas seen in the figure. Because the off-resonant bichromatic field only couples momentum states in discrete spacings of $\hbar k$, the continuous set of states defined by the parabola turns into a discrete set of states represented by the black horizontal bars. The parameter Δ specifies the detuning of the laser fields from the excited state, and is referred to as the single-photon detuning. The parameter δ specifies the difference in frequency between the two laser fields, and determines whether the Bragg diffraction process is resonant. Here 2nd-order Bragg diffraction is driven, as $\omega_1 - \omega_2$ is equal to half the kinetic energy difference between the initial and final state.

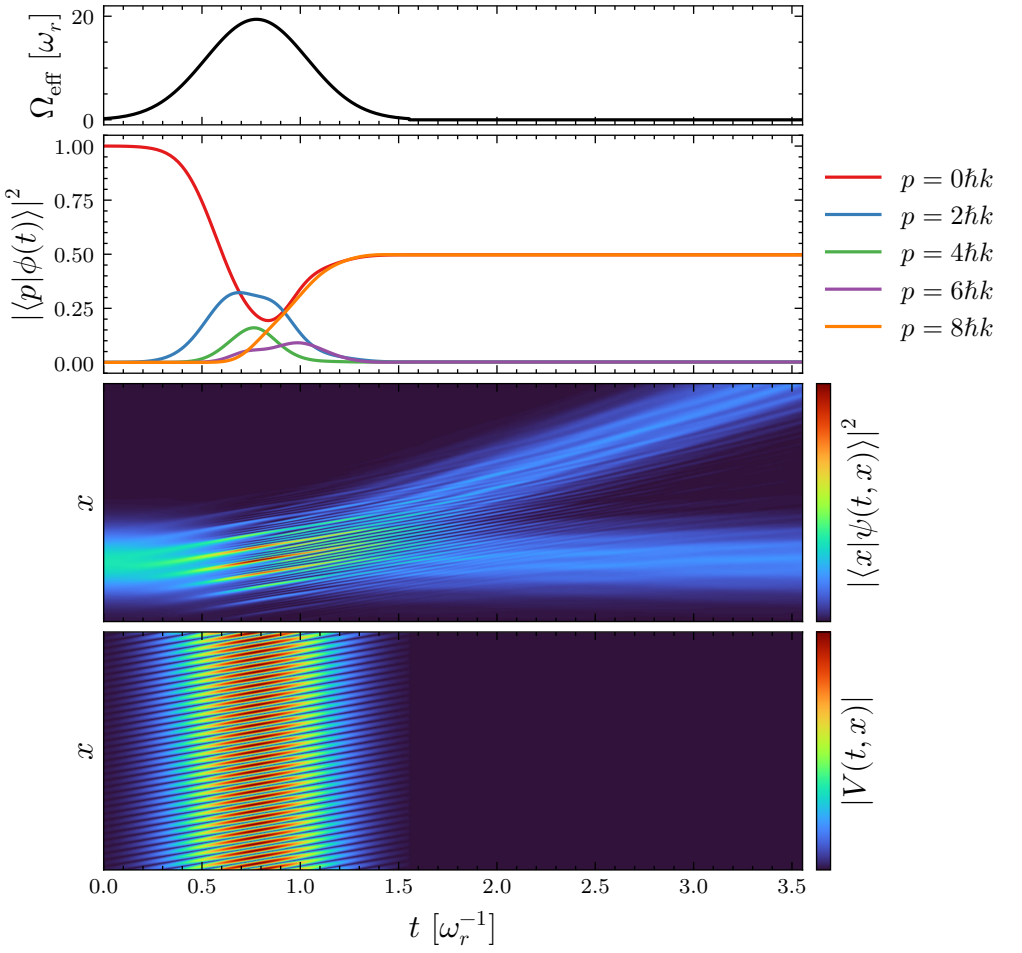


Figure 2.3: Simulation of a 4th-order Bragg diffraction process in both position and momentum space as a function of time. *Top:* The effective Rabi frequency of the Bragg pulse as a function of time. This is described by the truncated Gaussian pulse given in Eq. (2.20) where $\sigma = 0.259\omega_{\text{r}}^{-1}$, $t_0 = 3\sigma$, and $\Omega_{\text{eff},0} = 19.4\omega_{\text{r}}$. *Upper middle:* The population in momentum space, computed by integrating Eq. (6.2). *Lower middle:* The population in position space, computed by integrating Eq. (6.14) using the split-step method. The wavefunction is initialized as a Gaussian wavepacket with amplitude width $\sigma_x = 5\sqrt{2}/k$, where k is the laser wavenumber. *Lower:* The absolute value of the potential that the position space wavefunction is subject to as a function of time, given by Eq. (A.47).

where $p_0 = 2n_0\hbar k$ and $p_f = 2n_f\hbar k$. Note that the sign of δ determines the direction of momentum transfer. We refer to δ as the two-photon detuning. The order of the Bragg diffraction process is given by $|n_f - n_0|$. When Ω_{eff} is much smaller than the detuning of every intermediate momentum state from the Bragg transition, the intermediate momentum states can be adiabatically eliminated from the system, similar to how the excited state was eliminated in the previous section [33]. However, in practice, this approximation is not valid and the full ladder of momentum states in Eq. (2.14) must be considered. We rewrite Eq. (2.14) in the frame $\tilde{c}_n(t) = \exp(i4n^2\omega_{\text{r}}t)c_n(t)$ and drop the AC Stark shift terms, finding

$$i\hbar\dot{\tilde{c}}_j(t) = -\frac{\hbar\Omega_{\text{eff}}(t)}{2}\left[\exp(i(\delta - 4\omega_{\text{r}} - 8j\omega_{\text{r}})t)\tilde{c}_{j+1}(t) + \exp(i(-\delta - 4\omega_{\text{r}} + 8j\omega_{\text{r}})t)\tilde{c}_{j-1}(t)\right] \tag{2.19}$$

These equations can be numerically integrated to determine the dynamics of the Bragg diffraction process. To create a Bragg beamsplitter we tune $\Omega_{\text{eff},0}$ such that there is equal population in $|2n_0\hbar k\rangle$ and $|2n_f\hbar k\rangle$ at the end of the Bragg diffraction process. In Fig. 2.3 we show the result of such an integration in both position and momentum space for a 4th-order Bragg beamsplitter. We use a truncated Gaussian pulse in time described by

$$\Omega_{\text{eff}}(t, t_0) = \begin{cases} \Omega_{\text{eff},0}\exp\left(-\frac{(t-t_0)^2}{2\sigma^2}\right), & |t - t_0| < 3\sigma \\ 0, & \text{otherwise} \end{cases} \tag{2.20}$$

as shown in the top panel of the figure. The second panel shows the population in each momentum state as a function of time. During the pulse population flows through the intermediate momentum states. At the end of the pulse the system is in a 50-50 superposition of the $0\hbar k$ and $8\hbar k$ momentum states. The third panel shows the position space wavefunction as a function of time. During the pulse we observe that the wavefunction is localized in the lattice formed by the potential. The atom wave diffracts off of this lattice to form two wavepackets that separate as time evolves. In the fourth panel we see the absolute value of the potential plotted as a function of time and position. The numerical integration methods used to make this plot are described in Chapter 6.

As mentioned earlier, the intermediate states should not become substantially populated when Ω_{eff} is much smaller than the detuning of every intermediate state from the Bragg transition. In Fig. 2.4 we show the population during Bragg beamsplitters with several different σ values. Longer σ values achieve a 50-50 superposition for a lower $\Omega_{\text{eff},0}$ (as shown in Fig. 2.6), and exhibit less population flowing through the intermediate states. This regime is known as the Bragg regime. For extremely short values of σ, population is scattered into many states by the end of the pulse. This is known as the Raman-Nath regime [33]. The regime where σ is small enough that intermediate states are populated during the pulse but large enough that only the target momentum states are populated after the pulse is known as the quasi-Bragg regime [33, 46]. This is the regime in which our atom interferometer operates. It is useful to see how these populations vary as a function of $\Omega_{\text{eff},0}$, as shown in

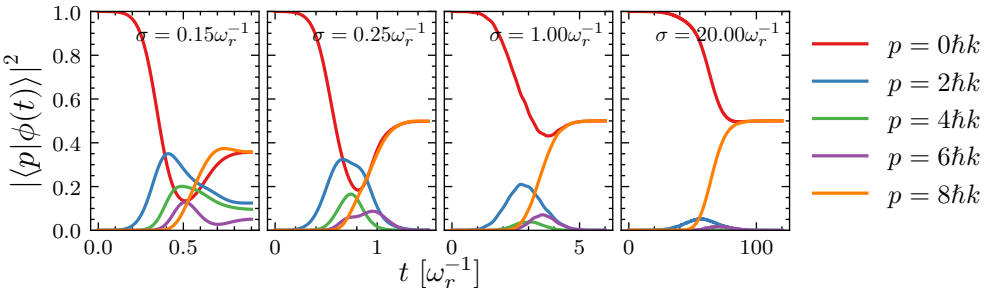


Figure 2.4: Bragg dynamics in momentum space for several Bragg pulse durations. All dynamics show 4th-order Bragg diffraction. For longer pulse durations we see fewer intermediate momentum states are populated. For each value of σ the value of $\Omega_{\text{eff},0}$ that produces a $\pi/2$ pulse is used.

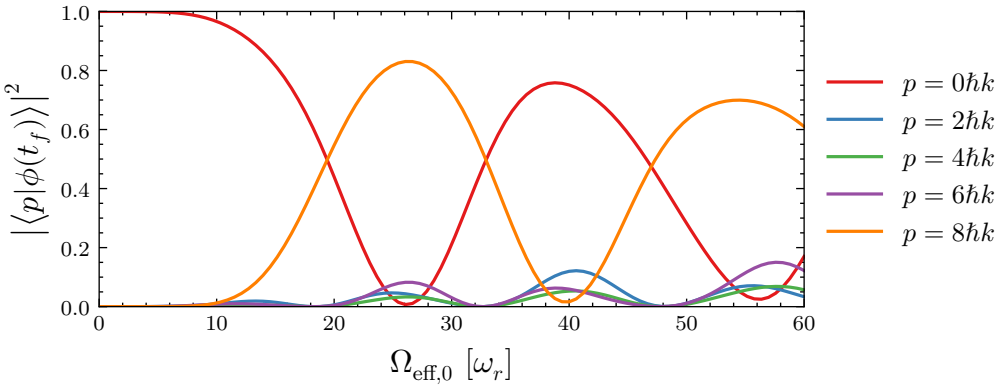


Figure 2.5: Momentum state populations at the end of a $\sigma = 0.259\omega_{\text{r}}^{-1}$, 4th-order Bragg pulse. We see that, while suppressed, intermediate momentum states are populated at the end of the pulse. From the plot we see that a $\pi/2$ pulse (a beamsplitter) occurs slightly below $\Omega_{\text{eff},0} = 20\omega_{\text{r}}$ and a π pulse (a mirror) occurs slightly above $\Omega_{\text{eff},0} = 25\omega_{\text{r}}$.

Fig. 2.5. We see that some intermediate momentum states are populated at the end of the pulse, but are for the most part suppressed. We also see that at certain values of $\Omega_{\text{eff},0}$ we obtain a mirror pulse.

In the quasi-Bragg regime, Bragg beamsplitters can produce an anomalous phase between the initial and target momentum states, as shown in Fig. 2.6. From the figure we see that in the Bragg regime the phase difference goes to $\pi/2$. For shorter values of σ the phase starts to deviate from $\pi/2$. In Fig. 2.7 we see that this phase also has a dependence on $\Omega_{\text{eff},0}$, which can be significant for shorter values of σ. We will explore why the phase goes to $\pi/2$ for long σ in the next section.

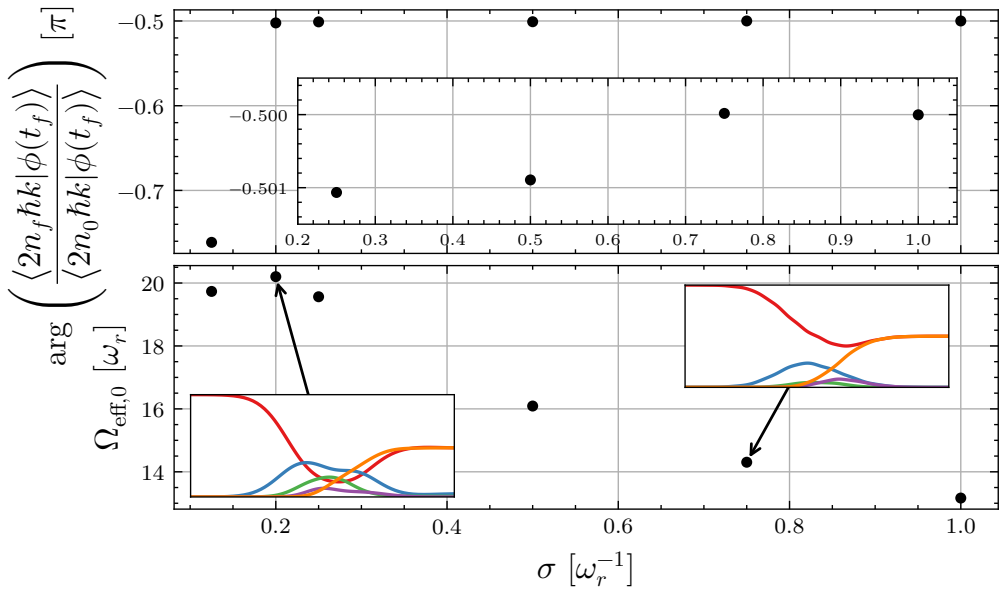


Figure 2.6: *Upper:* The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of the Bragg pulse width σ. The inset zooms in on the rightmost five points in the plot to show their deviation from $-\pi/2$. *Lower:* The required $\Omega_{\text{eff},0}$ to achieve a $\pi/2$ pulse for each σ value used in the upper plot. These generally decrease as σ gets larger. Insets on this plot show the momentum state population dynamics as a function of time for two different σ values. All processes in this figure are 4th-order Bragg diffraction pulses.

We also show the dependence of the anomalous phase on deviations of δ from resonance in Fig. 2.8. We see that on resonance the anomalous phase goes to $\pi/2$. The anomalous phase varies linearly as δ acquires an offset from resonance, and has a larger slope for longer pulse durations.

We refer to the deviation of this phase from $\pi/2$ as the diffraction phase, which is consistent with past literature [21, 47]. Note that some sources define the diffraction phase as the phase introduced to an interferometer by parasitic momentum states [48] (see Section 7.1.4), which is inconsistent with the definition used in this thesis. The dependence of this phase on the effective Rabi frequency $\Omega_{\text{eff},0}$ and the two-photon detuning δ leads to a major systematic in the measurement of the fine-structure constant, and is investigated further in Section 7.1.7.

2.2.1 Equations of motion in the Bragg regime

Here we reduce the ladder of momentum states to a two-level system, as is done in [33]. Consider a Bragg process coupling $|2n_0\hbar k\rangle$ to $|2(n_0+n)\hbar k\rangle$ where $n > 0$. Starting from

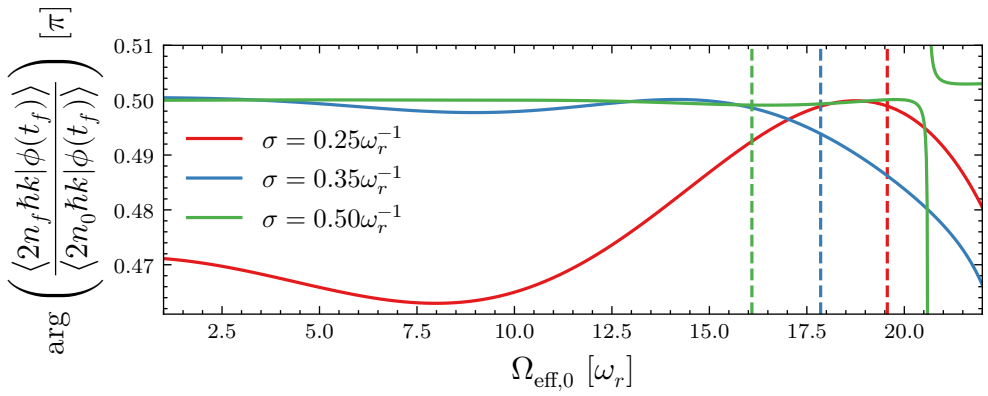


Figure 2.7: The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of $\Omega_{\text{eff},0}$ for three different values of σ. The vertical dashed lines display the values of $\Omega_{\text{eff},0}$ that produce a $\pi/2$ pulse. We see that the phase difference is highly dependent on intensity for shorter Bragg pulses, and far less dependent on intensity for longer Bragg pulses. All processes in this figure are 4th-order Bragg diffraction pulses.

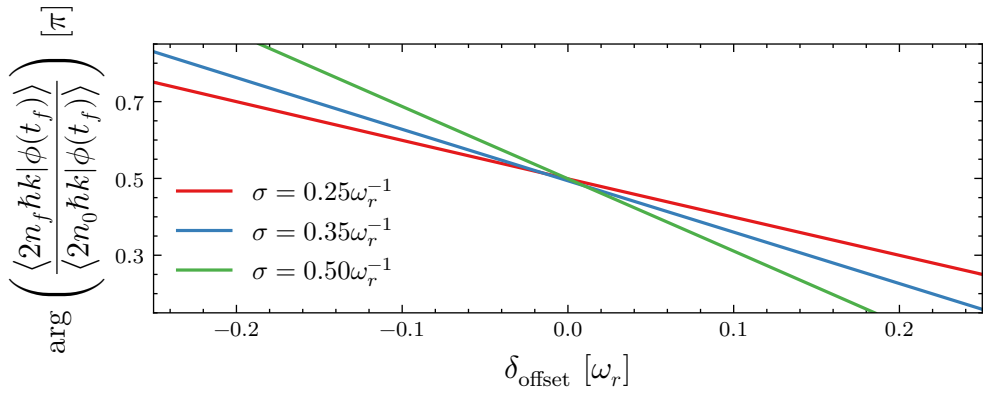


Figure 2.8: The phase difference between the undiffracted and diffracted states (i.e. the diffraction phase) as a function of δ_{offset} (the offset of the two-photon detuning from resonance with 4th-order Bragg diffraction) for three different values of σ.

Eq. (2.19) and taking $\tilde{c}_{n_0+q}(t) \to \tilde{c}_q(t)$, we find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_q(t) = &-\hbar\frac{\Omega_{\text{eff}}(t)}{2}\left[\exp(i(\delta - 4\omega_{\text{r}} - 8(n_0+q)\omega_{\text{r}})t)\tilde{c}_{q+1}(t)\right. \\ &\left. + \exp(i(-\delta - 4\omega_{\text{r}} + 8(n_0+q)\omega_{\text{r}})t)\tilde{c}_{q-1}(t)\right] \end{aligned} \tag{2.21}$$

We define

$$\Delta_q = q\delta - 4\left[(n_0+q)^2 - n_0^2\right]\omega_{\text{r}} \tag{2.22}$$

which lets us rewrite Eq. (2.21) as

$$i\hbar\dot{\tilde{c}}_q(t) = -\hbar\frac{\Omega_{\text{eff}}(t)}{2}\left[\exp(i(\Delta_{q+1} - \Delta_q)t)\tilde{c}_{q+1}(t) + \exp(i(\Delta_{q-1} - \Delta_q)t)\tilde{c}_{q-1}(t)\right] \tag{2.23}$$

For Bragg diffraction of order n we consider the most direct coupling from n_0 to $n_0 + n$. The most direct coupling path occurs when population flows from q to $q+1$, and so we drop $\tilde{c}_{q-1}(t)$-like terms in our analysis. In dropping this term, we are now only approximately describing the dynamics of the system. We define $\bar{c}_{q+1}(t)$ as

$$\bar{c}_{q+1}(t) = \tilde{c}_{q+1}(t)\exp(i(\Delta_{q+1} - \Delta_n)t) \tag{2.24}$$

We now use the adiabatic elimination developed in Section 2.1 to evaluate the integral of Eq. (2.23).

$$\begin{aligned} \tilde{c}_q(t) \approx &\frac{i}{2}\int_0^t \Omega_{\text{eff}}(t')\exp(i(\Delta_n - \Delta_q)t')\bar{c}_{q+1}(t')\,\mathrm{d}t' \\ \approx &\frac{\Omega_{\text{eff}}(t)}{2(\Delta_n - \Delta_q)}\exp(i(\Delta_n - \Delta_q)t)\bar{c}_{q+1}(t) \\ \approx &\frac{\Omega_{\text{eff}}(t)}{2(\Delta_n - \Delta_q)}\exp(i(\Delta_{q+1} - \Delta_q)t)\tilde{c}_{q+1}(t) \end{aligned} \tag{2.25}$$

where we required $\Omega_{\text{eff}}(t) \ll |\Delta_n - \Delta_q|$. We use this relation to rewrite $\tilde{c}_1(t)$ in terms of $\tilde{c}_n(t)$

$$\tilde{c}_1(t) \approx \exp(i(\Delta_n - \Delta_1)t)\tilde{c}_n(t)\prod_{\ell=1}^{n-1}\frac{\Omega_{\text{eff}}(t)}{2(\Delta_n - \Delta_\ell)} \tag{2.26}$$

By symmetry we write $\tilde{c}_{n-1}(t)$ in terms of $\tilde{c}_0(t)$

$$\tilde{c}_{n-1}(t) \approx \exp(i(\Delta_0 - \Delta_{n-1})t)\tilde{c}_0(t)\prod_{\ell=1}^{n-1}\frac{\Omega_{\text{eff}}(t)}{2(\Delta_0 - \Delta_\ell)} \tag{2.27}$$

$$\approx \exp(-i\Delta_{n-1}t)\tilde{c}_0(t)\prod_{\ell=1}^{n-1}\left(-\frac{\Omega_{\text{eff}}(t)}{2\Delta_\ell}\right) \tag{2.28}$$

Substituting these expressions into our expressions for $\dot{\tilde{c}}_0(t)$ and $\dot{\tilde{c}}_n(t)$ we find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_0(t) \approx& -\hbar\frac{\Omega_{\text{eff}}(t)}{2}\exp(i\Delta_1 t)\tilde{c}_1(t) \\ \approx& -\hbar\exp(i\Delta_n t)\tilde{c}_n(t)\frac{\Omega_{\text{eff}}(t)}{2}\prod_{\ell=1}^{n-1}\frac{\Omega_{\text{eff}}(t)}{2(\Delta_n-\Delta_\ell)} \end{aligned} \tag{2.29}$$

$$\begin{aligned} i\hbar\dot{\tilde{c}}_n(t) \approx& -\hbar\frac{\Omega_{\text{eff}}(t)}{2}\exp(i(\Delta_{n-1}-\Delta_n)t)\tilde{c}_{n-1}(t) \\ \approx& -\hbar\exp(-i\Delta_n t)\tilde{c}_0(t)\frac{\Omega_{\text{eff}}(t)}{2}\prod_{\ell=1}^{n-1}\left(-\frac{\Omega_{\text{eff}}(t)}{2\Delta_\ell}\right) \end{aligned} \tag{2.30}$$

This is the Bragg regime, where the ladder of momentum states has been reduced to a two-level system where each level is a momentum state. The resonance condition occurs when $\Delta_n = 0$, so we find the resonance condition for δ to be

$$0 = n\delta - 4\left[(n_0+n)^2 - n_0^2\right]\omega_{\text{r}} \rightarrow \delta = 4(2n_0+n)\omega_{\text{r}} \tag{2.31}$$

This is the same resonance condition we found from our energy conservation argument in Eq. (2.18). Approximating Δ_ℓ in the product to the on-resonance case, the equations of motion reduce further to

$$\begin{aligned} i\hbar\dot{\tilde{c}}_0(t) \approx& -\hbar\Omega_n(t)\exp(i(n\delta - 4((n_0+n)^2-n_0^2)\omega_{\text{r}})t)\tilde{c}_n(t) \\ i\hbar\dot{\tilde{c}}_n(t) \approx& -\hbar\Omega_n(t)\exp(-i(n\delta - 4((n_0+n)^2-n_0^2)\omega_{\text{r}})t)\tilde{c}_0(t) \end{aligned} \tag{2.32}$$

where

$$\begin{aligned} \Omega_n(t) =& \frac{\Omega_{\text{eff}}(t)^n}{2^n}\prod_{\ell=1}^{n-1}\left(-\frac{1}{\Delta_\ell}\right) \\ =& \frac{\Omega_{\text{eff}}(t)^n}{2^n}\prod_{\ell=1}^{n-1}\left(-\frac{1}{4\ell(n-\ell)\omega_{\text{r}}}\right) \\ =& 4\frac{\Omega_{\text{eff}}(t)^n}{8^n\omega_{\text{r}}^n\left((n-1)!\right)^2}(-1)^{n-1}\omega_{\text{r}} \end{aligned} \tag{2.33}$$

This result allows the estimation of the effective Rabi frequency required for a Bragg beamsplitter of order n [33]. Solving the system of equations on resonance for $\tilde{c}_0(0) = 1$, $\tilde{c}_n(0) = 0$, and constant $\Omega_n(t)$ yields

$$\tilde{c}_0(t) = \cos(\Omega_n t) \qquad \tilde{c}_n(t) = i\sin(\Omega_n t) \tag{2.34}$$

The i in front of the sinusoid reveals that the diffracted state acquires a phase shift of $\pi/2$ relative to the undiffracted state in the idealized Bragg regime limit.

2.2.2 Multifrequency Bragg diffraction

For some interferometer geometries it is useful to operate two Bragg beamsplitters simultaneously. One approach for doing this utilizes two copropagating and one counterpropagating field. These fields will form two lattices moving at velocities set by the frequency difference between each copropagating field and the counterpropagating field. Assuming $k = k_1 \approx k_2 \approx -k_3$, we eliminate the excited state from Eq. (2.5) using Eq. (2.11) in the three-field case to find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_g(t,p) =& \frac{p^2}{2m}\tilde{c}_g(t,p) - \hbar\frac{\Omega_1(t)^2}{4\Delta}\tilde{c}_g(t,p) - \hbar\frac{\Omega_2(t)^2}{4\Delta}\tilde{c}_g(t,p) - \hbar\frac{\Omega_3(t)^2}{4\Delta}\tilde{c}_g(t,p) \\ &- \hbar\frac{\Omega_1(t)\Omega_2(t)}{2\Delta}\cos(2\omega_\mathrm{m}t)\tilde{c}_g(t,p) \\ &- \hbar\frac{\Omega_1(t)\Omega_3(t)}{4\Delta}\exp(-i(\delta-\omega_\mathrm{m})t)\tilde{c}_g(t,p-2\hbar k) \\ &- \hbar\frac{\Omega_1(t)\Omega_3(t)}{4\Delta}\exp(i(\delta-\omega_\mathrm{m})t)\tilde{c}_g(t,p+2\hbar k) \\ &- \hbar\frac{\Omega_2(t)\Omega_3(t)}{4\Delta}\exp(-i(\delta+\omega_\mathrm{m})t)\tilde{c}_g(t,p-2\hbar k) \\ &- \hbar\frac{\Omega_2(t)\Omega_3(t)}{4\Delta}\exp(i(\delta+\omega_\mathrm{m})t)\tilde{c}_g(t,p+2\hbar k) \end{aligned} \tag{2.35}$$

where we set $\delta = (\omega_1+\omega_2)/2-\omega_3$ and $\omega_\mathrm{m} = (\omega_2-\omega_1)/2$. From these expressions we see that we can tune δ and ω_m such that $\delta+\omega_\mathrm{m}$ and $\delta-\omega_\mathrm{m}$ satisfy different Bragg resonance conditions, leading to two simultaneous Bragg diffraction pulses. If we set $\Omega_1 = \Omega_2 = \Omega_3 = \Omega$, then this expression simplifies to

$$\begin{aligned} i\hbar\dot{\tilde{c}}_g(t,p) =& \frac{p^2}{2m}\tilde{c}_g(t,p) - 3\hbar\frac{\Omega(t)^2}{4\Delta}\tilde{c}_g(t,p) - \hbar\frac{\Omega(t)^2}{2\Delta}\cos(2\omega_\mathrm{m}t)\tilde{c}_g(t,p) \\ &- \hbar\cos(\omega_\mathrm{m}t)\frac{\Omega(t)^2}{2\Delta}\exp(-i\delta t)\tilde{c}_g(t,p-2\hbar k) \\ &- \hbar\cos(\omega_\mathrm{m}t)\frac{\Omega(t)^2}{2\Delta}\exp(i\delta t)\tilde{c}_g(t,p+2\hbar k) \end{aligned} \tag{2.36}$$

which we see is equivalent to Eq. (2.19) up to the prefactor of $2\cos(\omega_\mathrm{m}t)$ after dropping the AC Stark shift terms and transforming to a rotating frame. Therefore, the Bragg dynamics in the case of three fields can be computed from Eq. (2.19) by including the factor $2\cos(\omega_\mathrm{m}t)$ in $\Omega_\mathrm{eff}(t)$. The resonance condition for each multifrequency lattice can be determined by plugging $\delta \to \delta \pm \omega_\mathrm{m}$ into Eq. (2.18).

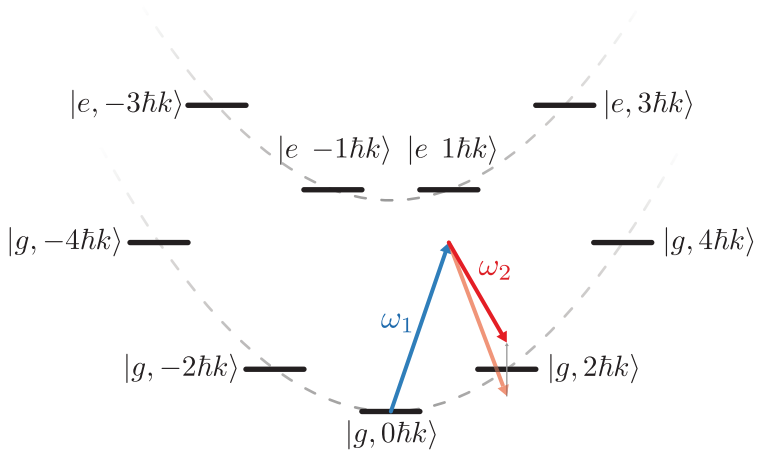


Figure 2.9: The level diagram describing Bloch oscillations. This is similar to the Bragg level diagram shown in Fig. 2.2, but the laser fields do not resonantly couple two states. Instead one of the laser fields (in this case ω_2) is swept in time such that $\delta = \omega_1 - \omega_2$ sweeps over the energy difference between $|g, 0\hbar k\rangle$ and $|g, 2\hbar k\rangle$. This causes adiabatic population transfer. As δ continues to sweep in time, it will eventually sweep over the energy difference between $|g, 2\hbar k\rangle$ and $|g, 4\hbar k\rangle$, causing a second adiabatic transfer of population. This continues to happen until the sweep of δ stops.

2.3 Bloch oscillations and adiabatic Bragg diffraction

Bloch oscillations

In the previous section we saw that Bragg diffraction can be seen as a resonant drive coupling two momentum states. Instead of resonantly driving the transition, one can sweep the drive frequency δ across the resonance condition to adiabatically transfer population between momentum states. The adiabatic transfer process is often referred to as adiabatic passage and is described by the Landau-Zener tunneling probability equation [49, 50]. There are two ways that this adiabatic sweep can be implemented to transfer population from some initial momentum state n_0 to a final state n_f. The first approach is to sweep δ across the 1st-order Bragg resonance as in Fig. 2.9. In the figure this will transfer population from $|g, 0\hbar k\rangle$ to $|g, 2\hbar k\rangle$. As δ continues to sweep another transfer will occur from $|g, 2\hbar k\rangle$ to $|g, 4\hbar k\rangle$. This will continue until the sweep is stopped. This process is referred to as Bloch oscillations [36–39]. To apply Bloch oscillations from some initial momentum state n_0 to a final momentum state n_f, we start the sweep underneath the resonance condition for $n_0 \to n_0 + 1$ and end the sweep above the resonance condition for $n_f - 1 \to n_f$

$$\delta(t) = 8n_0\omega_\mathrm{r} + 8\omega_\mathrm{r}\frac{n_f - n_0}{\tau}t \tag{2.37}$$

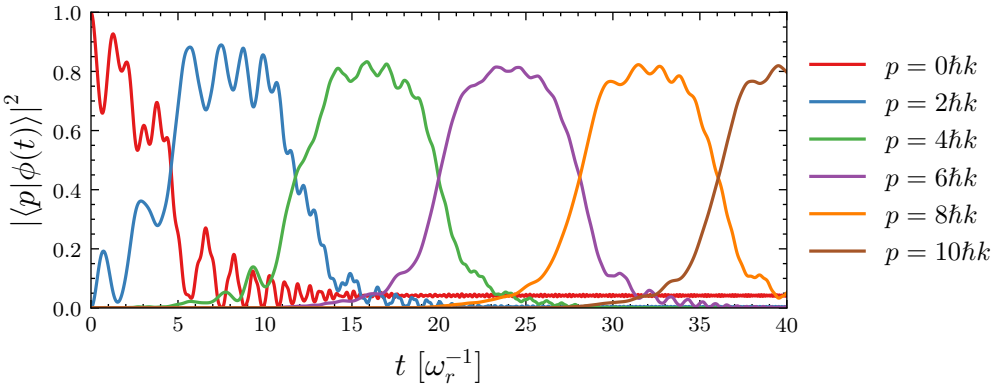


Figure 2.10: Population dynamics in momentum space during 5th-order Bloch oscillations. We see that population is transferred from $p = 0\hbar k$ to $p = 10\hbar k$ at the end of the Bloch oscillation process. Here $\Omega_{\text{eff}} = 2\omega_{\text{r}}$ (meaning it is constant in time) and $\tau = 40\omega_{\text{r}}^{-1}$.

where τ is the sweep duration. The result of a 5th-order Bloch oscillation is shown in Fig. 2.10, which shows population moving sequentially between neighboring momentum states. In general, an Nth-order Bloch oscillation pulse transfers $2N\hbar k$ momentum. The choice of τ can have a dramatic impact on the loss associated with Bloch oscillations, as seen in Fig. 2.11, with longer values of τ mostly corresponding to more efficient population transfer. In practice, even though the lasers used to produce Bloch oscillations are far-detuned from resonance, if τ is long enough single-photon scattering can occur and cause additional losses (see Section 7.1.6). Landau-Zener losses can be computed using the Landau-Zener formula [50]

$$P_{\text{loss}} = \exp\left(-\frac{\pi\Omega_{\text{eff}}^2}{2|\alpha|}\right) \tag{2.38}$$

where α is the ramp rate given by $\alpha = 8\omega_{\text{r}}(n_f - n_0)/\tau$.

In the position-space picture the wavefunction evolves according to Eq. (A.47), reproduced below

$$i\hbar\dot{\psi}(t,x) = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2}\psi(x) - \hbar\Omega_{\text{eff}}(t)\cos\left(2kx - \int_0^t \delta(t')\,\mathrm{d}t'\right)\psi(t,x) \tag{2.39}$$

The position-space lattice potential has velocity $\delta(t)/2k$. In the position-space picture we can imagine Bloch oscillations as trapping a wavefunction in an accelerating lattice. The wavefunction accelerates with the lattice and changes its momentum. There is a third picture in which Bloch oscillations are described using bandstructure dynamics [38]. The dynamics of Bloch oscillations are further affected by how the Bloch lattice is initially turned on or off (i.e. the rate at which the Bloch intensity is ramped on/off, the shape of the ramp on/off,

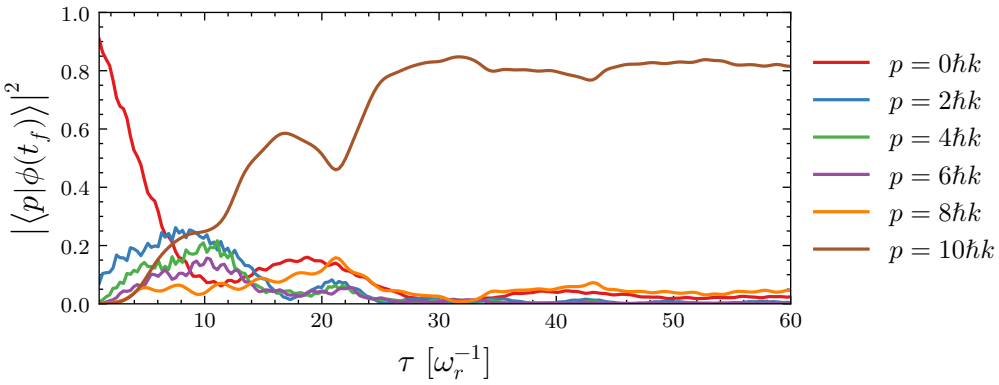


Figure 2.11: The final population in each momentum state for 5th-order Bloch oscillations as a function of τ. $\Omega_{\text{eff}} = 2\omega_{\text{r}}$ in this plot.

etc.), as discussed in [39]. In the numerical integrations performed here the intensity ramp on and ramp off are not considered.

Note that multifrequency Bloch oscillations can be performed using the same experimental setup as multifrequency Bragg diffraction. By sweeping ω_{m} in time two Bloch lattices will accelerate in opposite directions. In the experiment described in this thesis, multifrequency Bloch oscillations are used as a momentum "mirror". Previous work has developed methods to make Bloch beamsplitters as well as other matterwave optics [39].

Adiabatic Bragg diffraction

The second approach for transferring population between momentum states adiabatically is to sweep $\delta(t)$ over a single higher-order Bragg resonance instead of successively sweeping over 1st-order Bragg resonances [35]. This is known as an adiabatic Bragg transition and is shown in Fig. 2.12.

2.4 Raman transitions

Raman transitions are similar to Bragg transitions but are implemented in a three-level system instead of a two-level system [40, 41]. The three-level system in question is often a lambda system with two ground states with a microwave energy splitting and a single excited state coupled to the ground states with an optical electric dipole transition, as shown in Fig. 2.13. The excited state is far-detuned from the optical driving field(s) and is adiabatically eliminated from the system. We begin with the following Hamiltonian for a

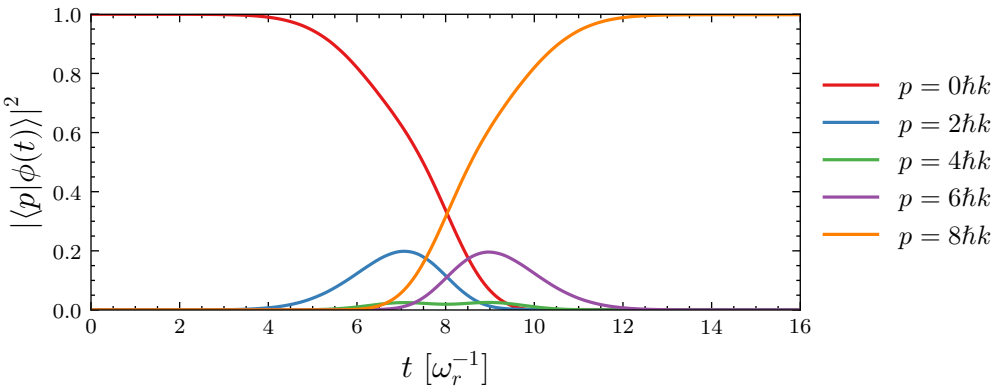


Figure 2.12: Population dynamics in momentum space during a 4th-order adiabatic Bragg diffraction pulse. Here $\Omega_{\text{eff},0} = 14\omega_\text{r}$, $\sigma = 2\omega_\text{r}^{-1}$ (the pulse shape defined in Eq. (2.20) is used, but with the truncation occurring at 4σ rather than 3σ), and δ is swept from $4(3+4)\omega_\text{r}/2$ to $4(5+4)\omega_\text{r}/2$ over the duration that the pulse has a nonzero intensity.

three-level system

$$\begin{aligned}\hat{H}_\text{Raman} =& \frac{\hat{p}^2}{2m} + \hbar\omega_e|e\rangle\langle e| + \hbar\omega_{g_2}|g_2\rangle\langle g_2| \\ &+ \hbar\sum_\ell \left[\Omega_{g_1,\ell}(t)\cos(k_\ell z - \omega_\ell t)|e\rangle\langle g_1| + \Omega_{g_2,\ell}(t)\cos(k_\ell z - \omega_\ell t)|e\rangle\langle g_2| + \text{h.c.}\right]\end{aligned} \tag{2.40}$$

where $\Omega_{g_1,\ell}(t)$ and $\Omega_{g_2,\ell}(t)$ denote the ℓth field coupling states $|g_1\rangle$ and $|g_2\rangle$ to $|e\rangle$, respectively. We follow the same steps from Section 2.1 of transforming into the interaction picture, applying the rotating wave approximation to eliminate optical frequencies, and integrating the excited state equation of motion to yield

$$\begin{aligned}i\hbar\dot{c}_{g_1}(t,p) =& \frac{p^2}{2m}c_{g_1}(t,p) + \frac{\hbar}{2}\sum_\ell \Omega_{g_1,\ell}(t)\exp(-i(\omega_e-\omega_\ell)t)c_e(t,p+\hbar k_\ell) \\ i\hbar\dot{c}_{g_2}(t,p) =& \frac{p^2}{2m}c_{g_2}(t,p) + \frac{\hbar}{2}\sum_\ell \Omega_{g_2,\ell}(t)\exp(-i(\omega_e-\omega_\ell-\omega_{g_2})t)c_e(t,p+\hbar k_\ell) \\ i\hbar\dot{c}_e(t,p) =& -i\hbar\sum_\ell\left[\frac{\Omega_{g_1,\ell}(t)}{2(\omega_e-\omega_\ell)}\exp(i(\omega_e-\omega_\ell)t)c_{g_1}(t,p-\hbar k_\ell)\right. \\ &\left.+\frac{\Omega_{g_2,\ell}(t)}{2(\omega_e-\omega_\ell-\omega_{g_2})}\exp(i(\omega_e-\omega_\ell-\omega_{g_2})t)c_{g_2}(t,p-\hbar k_\ell)\right]\end{aligned} \tag{2.41}$$

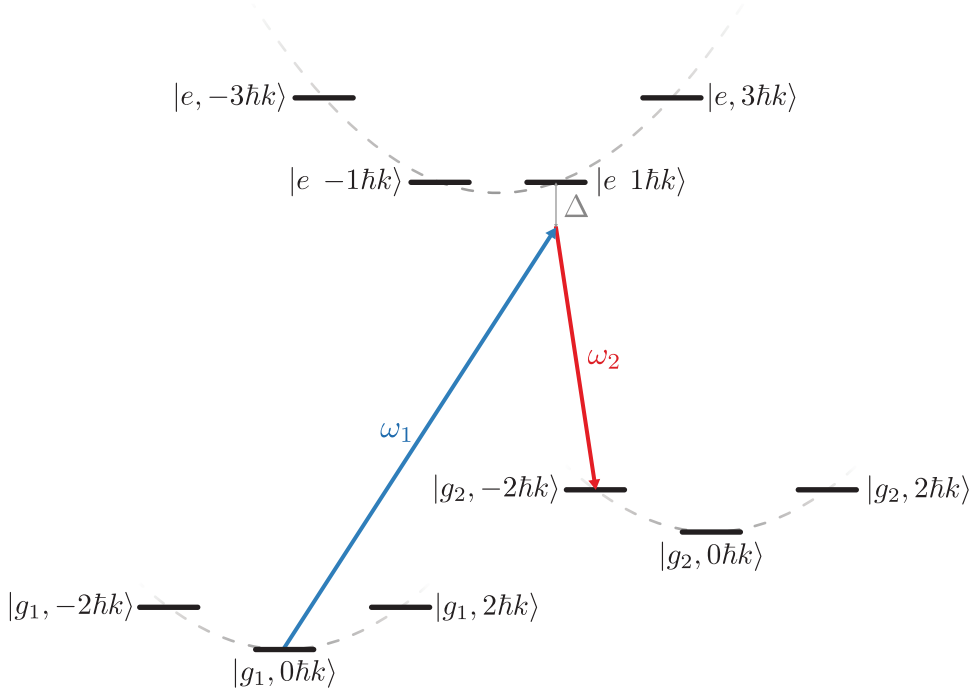


Figure 2.13: The level diagram describing Raman transitions. This is similar to the level diagrams considered in Figs. 2.2 and 2.9, but includes an additional ground state. In the diagram ω_1 and ω_2 are selected to drive a transition between the two ground states. Note that the field directions were chosen so that the transitioned state has momentum $-2\hbar k$ relative to the initial state.

For the case of two counterpropagating fields, eliminating $c_e(t,p)$ from the system reduces this to

$$\begin{aligned} i\hbar\dot{c}_{g_1}(t,p) =& \frac{p^2}{2m}c_{g_1}(t,p) - \hbar\Omega^{g_1}_{\text{AC}}c_{g_1}(t,p) \\ & - \hbar\Omega^{(1)}_{\text{eff}}\exp(i(\omega_1-\omega_2-\omega_{g_2})t)c_{g_2}(t,p+\hbar(k_1-k_2)) \\ i\hbar\dot{c}_{g_2}(t,p) =& \frac{p^2}{2m}c_{g_2}(t,p) - \hbar\Omega^{g_2}_{\text{AC}}c_{g_2}(t,p) \\ & - \hbar\Omega^{(2)}_{\text{eff}}\exp(-i(\omega_1-\omega_2-\omega_{g_2})t)c_{g_1}(t,p-\hbar(k_1-k_2)) \end{aligned} \tag{2.42}$$

where

$$\begin{aligned} \Omega^{g_2}_{\text{AC}} &= \frac{\Omega^2_{g_2,1}}{4(\omega_e-\omega_1-\omega_{g_2})} + \frac{\Omega^2_{g_2,2}}{4(\omega_e-\omega_2-\omega_{g_2})} & \Omega^{g_1}_{\text{AC}} &= \frac{\Omega^2_{g_1,1}}{4(\omega_e-\omega_1)} + \frac{\Omega^2_{g_1,2}}{4(\omega_e-\omega_2)} \\ \Omega^{(2)}_{\text{eff}} &= \frac{\Omega_{g_1,1}\Omega_{g_2,2}}{4(\omega_e-\omega_1)} & \Omega^{(1)}_{\text{eff}} &= \frac{\Omega_{g_1,1}\Omega_{g_2,2}}{4(\omega_e-\omega_2-\omega_{g_2})} \end{aligned} \tag{2.43}$$

The adiabatic elimination is valid in the case where $\Omega_{\text{eff}}^{(1)}$ and $\Omega_{\text{eff}}^{(2)}$ have the same denominators, in which case the result has an analytic solution [15]. We see that the relative direction of k_1 and k_2 will determine whether a Raman transition results in momentum transfer or not.

The equation for each ground state has three terms: an AC Stark shift term, a kinetic energy term, and an interaction term. The AC Stark shift term is not necessarily equal for both levels. We refer to $\Delta_{\text{AC}} = \Omega_{\text{AC}}^{g_1} - \Omega_{\text{AC}}^{g_2}$ as the differential AC Stark shift [51]. The differential AC Stark shift can be intensity dependent, which produces the systematic effect discussed in Section 7.1.1.

In the frame $\tilde{c}_{g_{1,2}}(t,p) = \exp(i(p^2/2m\hbar - \Omega_{\text{AC}}^{g_{1,2}})t)c_{g_{1,2}}(t,p)$, Eq. (2.42) reduces to

$$\begin{aligned} i\hbar\dot{\tilde{c}}_{g_1}(t,p) =& -\hbar\Omega_{\text{eff}}^{(1)} \exp\left(i\left[\delta - \frac{p(k_1-k_2)}{m} - \frac{\hbar(k_1-k_2)^2}{2m}\right]t\right)\tilde{c}_{g_2}(t, p+\hbar(k_1-k_2)) \\ i\hbar\dot{\tilde{c}}_{g_2}(t,p) =& -\hbar\Omega_{\text{eff}}^{(2)} \exp\left(i\left[-\delta + \frac{p(k_1-k_2)}{m} - \frac{\hbar(k_1-k_2)^2}{2m}\right]t\right)\tilde{c}_{g_1}(t, p-\hbar(k_1-k_2)) \end{aligned} \tag{2.44}$$

where

$$\delta = \omega_1 - \omega_2 - \omega_{g_2} - \Delta_{\text{AC}} \tag{2.45}$$

Resonance occurs when the argument of the exponential in the $\dot{\tilde{c}}_{g_1}(t,p)$ expression matches the argument of the exponential in the $\dot{\tilde{c}}_{g_2}(t,p)$ expression. Therefore, for the case of two counterpropagating beams with approximately equal wavenumbers, resonance occurs when

$$\omega_1 - \omega_2 = \frac{2kp}{m} + 4\omega_{\text{r}} + \omega_{g_2} + \Delta_{\text{AC}} \tag{2.46}$$

For a state initially at rest with $p = 0$, we see that the energy $\hbar(\omega_1 - \omega_2)$ must be equal to the energy of two photon recoils plus the energy of the hyperfine splitting plus the energy associated with any differential AC Stark shift.

To compute the Rabi frequencies in this expression it is necessary to consider the angular momentum of the laser field. This is considered for a Raman transition in [22].

Atom beamsplitters and mirrors can be implemented using Bragg diffraction or Raman transitions. Raman transitions provide the benefit of state-sensitive detection and are often experimentally more straightforward to implement, as they do not require as cold of an atomic sample to operate on. However, Raman transitions limit the experimenter to $2\hbar k$ photon recoils per beamsplitter or mirror pulse (unlike Bragg diffraction, which allows for higher order transitions). Additionally, Raman transitions result in interferometer arms accumulating phase due to the hyperfine energy splitting between the two Raman ground states [18]. For these reasons, we use Bragg diffraction to form atom interferometers in this thesis.

2.5 Laser phase

For both Bragg and Raman transitions we saw that the right-hand side of the equations of motion always has terms containing $\exp(i\delta t)c_{j\pm1}(t)$. This term imparts a phase of δt onto the diffracted or transitioned state. In other words, the local phase of the driving field (often a laser) at the time of the Bragg or Raman pulse is imprinted onto the diffracted or transitioned wavefunction.

When multiple Bragg or Raman transitions are applied in a sequence, the phase of the field must be referenced to the same initial position or time. For example, if a first Raman pulse is applied at time $t = t_0$ and then a second pulse at time $t = t_1$, then a phase $\delta(t_1 - t_0)$ will be imparted on the second pulse relative to the first. For Bragg diffraction of order n this phase is applied n times, and so the imparted phase is $n\delta t$.

To see this result for Bragg diffraction more explicitly, we can use the equations of motion in the Bragg regime from Eq. (2.32) by transforming from the rotating frame to the lab frame. From this result we immediately see that in the limit where Bragg diffraction can be described as a two-level system, phase $n\delta t$ is picked up by the diffracted state. Therefore

$$\phi_{\text{laser}} = n(2kz - \delta t) \tag{2.47}$$

Note that the $2kz$ term comes from transforming the equations of motion to position space. We see this by transforming the Hamiltonian derived in Eq. (2.15) to position space via $\langle z|\hat{H}_{\text{eff}}|\psi\rangle$, which yields coupling terms with exponentials of the form $\exp(i(2kz - \delta t))$. The laser phase can also be derived by treating the atom-laser interaction as a perturbation to the atom Hamiltonian and generalizing to an exchange of $2n$ photons [21].

Chapter 3

Theory of Atom Interferometers

Atom interferometry is often introduced with an analogy to optical interferometry. The analogy goes something like this: an interferometer splits a wave along two different paths. These two paths eventually recombine by interfering on a beamsplitter. Information about the relative phase accumulated along the two paths is contained in how constructive or destructive the outputs of the interferometer are. In an optical interferometer the wave is produced using a light source, like a laser. In an atom interferometer the wave is the quantum mechanical wavefunction of an atom.

The analogy makes it clear that the two ingredients you need to make an interferometer are 1) a wave and 2) a beamsplitter. In Chapter 2 we discussed how to make an atom beamsplitter using Bragg diffraction along with other matterwave optics. In Chapter 4 we will discuss the experimental apparatus required to prepare a gas of atoms in such a way that their state is well described by a simple wavefunction. In this chapter I will describe how to compute the phase accumulated in an atom interferometer and the requirements of the wavefunction that propagates through the atom interferometer.

3.1 Ramsey interferometry

An atom interferometer can be viewed as an extension of the Ramsey interferometer sequence, shown in Fig. 3.1, used for measuring the energy of a two-level system in a driving field. First the atom is prepared in its ground state with the driving field turned off. Then the experimenter pulses the driving field to produce a "beamsplitter" pulse (also called a $\pi/2$-pulse) that puts the system into an equal superposition of states $|g\rangle$ and $|e\rangle$. The system then freely evolves for time T, and each state in the superposition acquires phase at a rate proportional to its energy. Finally the experimenter triggers a second beamsplitter pulse to interfere the states and measures if the system is in the ground or excited state. The population in each state gives the experimenter information about the differential rate of phase accumulation in the two states [52–54].

We compute the probability of finding the system in the ground or excited state in the

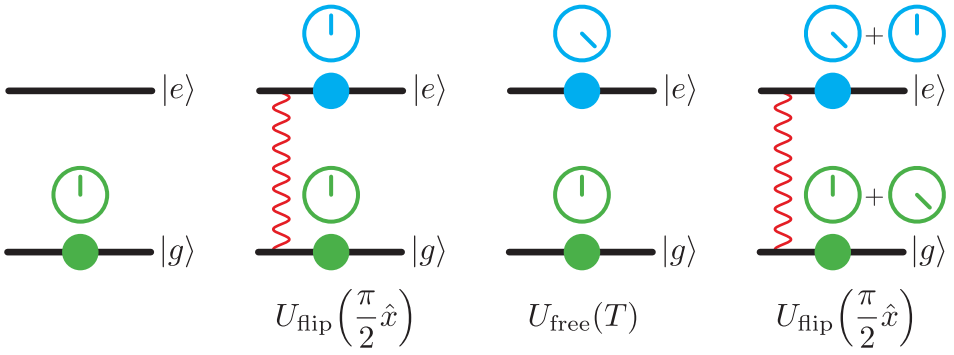


Figure 3.1: Ramsey interferometry sequence applied to a standard two-level system. Initially the system is in the ground state $|g\rangle$, as is shown in the leftmost diagram. Then a $\pi/2$ pulse is applied by pulsing on a coupling field, leaving the system in state $(|g\rangle - i|e\rangle)/\sqrt{2}$. Afterwards the system is allowed to evolve for time T. The clocks above each state represent the phase accumulated by that state. After evolving, the excited state clock has accumulated some phase, as states accumulate phase at a rate proportional to their energy. Finally, the second $\pi/2$ pulse is applied, which transfers amplitude from the state $|g\rangle$ to $|e\rangle$, and vice versa. Note that the transferred amplitude retains whatever phase it accumulated, which leads to interference.

rotating frame. The two-level system Hamiltonian without interactions is $\hat{H} = \hbar\omega_e|e\rangle\langle e|$. We transform to a frame rotating at the frequency of the driving field ω by applying the unitary transformation

$$U = \begin{pmatrix} 1 & 0 \\ 0 & \exp(i\omega t) \end{pmatrix} \tag{3.1}$$

via Theorem A.1. The transformed Hamiltonian is $\widetilde{\hat{H}} = \hbar(\omega_e - \omega)|\tilde{e}\rangle\langle\tilde{e}|$. We define the free evolution operator $\widetilde{U}_{\text{free}}(t) = \exp(-i\widetilde{\hat{H}}t/\hbar)$ and the beamsplitter operator $\widetilde{U}_{\text{BS}}(\boldsymbol{\theta}) = \exp(-i\boldsymbol{\theta}\cdot\boldsymbol{\sigma}/2)$ (taken from [55] Eq. (27)). The rotation angle $\boldsymbol{\theta}$ is set by the form of the interaction and $\boldsymbol{\sigma}$ is the Pauli vector. Computing the matrix exponentials for $t = T$ and $\boldsymbol{\theta} = \pi\hat{x}/2$ we find

$$\widetilde{U}_{\text{BS}}(\pi\hat{x}/2) = \widetilde{U}_{\text{BS}} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix} \qquad \widetilde{U}_{\text{free}}(T) = \widetilde{U}_{\text{free}} = \begin{pmatrix} 1 & 0 \\ 0 & \exp(-i\delta T) \end{pmatrix} \tag{3.2}$$

where $\delta = \omega_e - \omega$. Applying the operations in order we find the probability of measuring each state at the end of the sequence to be

$$p_g = |\langle\tilde{g}|\widetilde{U}_{\text{BS}}\widetilde{U}_{\text{free}}\widetilde{U}_{\text{BS}}|\tilde{g}\rangle|^2 = \sin^2(\delta T/2) \qquad p_e = |\langle\tilde{e}|\widetilde{U}_{\text{BS}}\widetilde{U}_{\text{free}}\widetilde{U}_{\text{BS}}|\tilde{g}\rangle|^2 = \cos^2(\delta T/2) \tag{3.3}$$

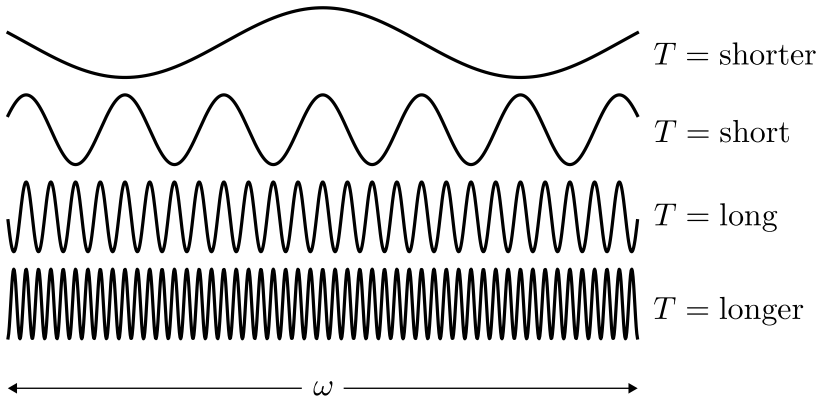


Figure 3.2: Fringes produced by measuring the excited state at the end of a Ramsey sequence, given by Eq. (3.3). For short free evolution times T, the fringe oscillates slowly as the drive frequency ω is scanned. As T increases the fringe oscillates more rapidly as ω is scanned. When running a Ramsey interferometry experiment, the goal is to determine the drive frequency where $\omega = \omega_e$. This equality can be measured to higher precision at longer T, because the fringe becomes much more sensitive to the drive frequency ω.

If the experimenter has set the driving field frequency equal to the frequency difference between the ground and excited states, then $\delta = 0$ and the system is guaranteed to be in the excited state regardless of the free evolution time T. If the experimenter sets the driving field frequency such that $\delta \neq 0$ then the probability oscillates sinusoidally as a function of T.

This provides a useful method for measuring ω_e. Initially, the Ramsey interferometer sequence is run with short T. The experimenter scans the drive frequency ω to maximize p_e (or minimize p_g). The experimenter then repeats the experiment at progressively longer T, making more and more precise adjustments to ω each time. This method is displayed in Fig. 3.2, where we see that for longer and longer T the exact value of ω must be chosen with increased precision to land on the extrema of the sinusoid at ω_e. In the context of a Ramsey interferometer sequence, these sinusoids are referred to as Ramsey fringes [56]. We will use this same method in the atom interferometer used to measure ω_{r}.

When ω_e and ω are matched, the laser phase cancels the free evolution phase and we *zero the phase of the interferometer*. The experimenter knows the phase has been zeroed because the interferometer populations do not vary as T is scanned.

Lab frame picture and the laser phase

The definition of the beamsplitter unitary operator in a frame rotating at the driving field frequency masks an important property of the beamsplitter operation. Transforming the beamsplitter operation into the lab frame, we find

$$U_{\mathrm{BS}} = U^{\dagger}\widetilde{U}_{\mathrm{BS}}U = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -i\exp(i\omega t) \\ -i\exp(-i\omega t) & 1 \end{pmatrix} \tag{3.4}$$

The exponential terms in the off-diagonal elements imprint the phase of the driving field onto the transitioned state. This phase is equivalent to the laser phase imprinted by Bragg diffraction that we discussed in Section 2.5.

In the lab frame the Ramsey interferometer sequence can be seen as a comparison of two oscillator frequencies. One oscillator is the excited state phase, which has frequency ω_e. The other oscillator is the driving field, which has frequency ω. At time $t = 0$ the first beamsplitter is applied, putting the system in state $|g\rangle - i|e\rangle$. At some later time $t = T$ we apply the second beamsplitter pulse. By this time, the excited state has accumulated phase $-\omega_e T$ from free evolution. This amplitude, when transferred from $|e\rangle$ to $|g\rangle$, acquires the driving field phase $+\omega T$. If $\omega = \omega_e$, then these phases cancel and the final state is clearly independent of T. In this situation the oscillators are synchronized and so any operation we apply becomes independent of T.

3.2 Interferometer geometries

In the previous section, we described how a Ramsey interferometer is used to measure the energy difference between states in a two-level system. Ramsey interferometry is often used to measure the energy difference between internal states. However, by using the Bragg beamsplitter introduced in Section 2.2, we could implement a Ramsey interferometer that measures the energy difference between different momentum states. In our discussion of the Bragg beamsplitter, the quantity $\delta = \omega_1 - \omega_2$ is equivalent to the drive frequency ω that we introduced when discussing Ramsey interferometry. We compute the time evolution operator for the Hamiltonian in Eq. (2.15), evaluated at $\Omega_{\mathrm{eff}}(t) = 0$,

$$U_{\mathrm{free}}(t) = \exp(-i\hat{H}_{\mathrm{eff}}t/\hbar) = \sum_{j=-\infty}^{\infty} \exp(-i4j^2\omega_{\mathrm{r}}t)|2j\hbar k\rangle\langle 2j\hbar k| \tag{3.5}$$

Therefore the free evolution phase for some state with momentum $2j\hbar k$ is $-4j^2\omega_{\mathrm{r}}t$. The phase accumulated in such a Ramsey interferometer is proportional to the recoil frequency ω_{r}, which is the quantity we are interested in measuring to determine the fine-structure constant.

To maximize the precision at which we can measure ω_{r}, we wish to make the free evolution time as long as possible and the diffraction order n as large as possible. This creates a problem. If the atom has mass m and wavepacket length ℓ, then for free evolution times

greater than $\ell m/2n\hbar k$ the excited state wavepacket will become spatially separated from the ground state wavepacket by the time of the second beamsplitter pulse. We saw this separation in Fig. 2.3. If there is no overlap between these two wavepackets at the time of the second beamsplitter pulse, then there will be no interference. To enable long free evolution times we must either create longer wavepackets, which is experimentally non-trivial, or introduce additional beamsplitter pulses that redirect the spatially separated wavefunctions so that they overlap at the time of the final beamsplitter. Depending on the direction of the momentum kick imparted by each beamsplitter, a variety of different interferometer geometries can be created. First we will consider the Ramsey-Bordé interferometer, which is used to measure ω_{r}.

3.2.1 Free evolution and laser phase

An alternate method of computing the free evolution phase is to use the Van Vleck formula for the propagator [57]

$$K(\mathbf{x}, \mathbf{x}_0, t) = \langle \mathbf{x}|U(t)|\mathbf{x}_0\rangle = A \exp\left(\frac{i}{\hbar} S(\mathbf{x}, \mathbf{x}_0, t)\right) \tag{3.6}$$

where $S(x, x_0, t)$ is the classical action along the path from $\mathbf{x}_0$ to $\mathbf{x}$ and A is some prefactor. For the case considered in Eq. (3.5) there is no potential, so $\mathcal{L} = p^2/2m$ and S is easily computed to be $S = p^2 t/2m$. Therefore the free evolution phase in the absence of any potential is

$$\phi_{\mathrm{free}} = \frac{p^2}{2m\hbar} t \tag{3.7}$$

For $p = 2j\hbar k$, $S = 4\hbar j^2 \omega_{\mathrm{r}} t$ and so the free evolution phase accumulated by a particle with this momentum is $4j^2\omega_{\mathrm{r}} t$. We note that the sign of this phase is the opposite of what was derived in Eq. (3.5). This is because the Van Vleck formula computes the free evolution phase in position space, and Eq. (3.5) computed the free evolution phase in momentum space. If we transform Eq. (3.5) to position space, we will get the same result but with the sign expected from the Van Vleck formula.

We restate the laser phase from Eq. (2.47)

$$\phi_{\mathrm{laser}} = n(2kz - \delta t) \tag{3.8}$$

We shall use these two phases in computing the phase accumulated in the interferometer geometries that follow.

3.2.2 Ramsey-Bordé interferometer

The Ramsey-Bordé interferometer uses a sequence of four beamsplitters, where the first and last beamsplitter pairs are spaced by time T and the intermediate pair is spaced by

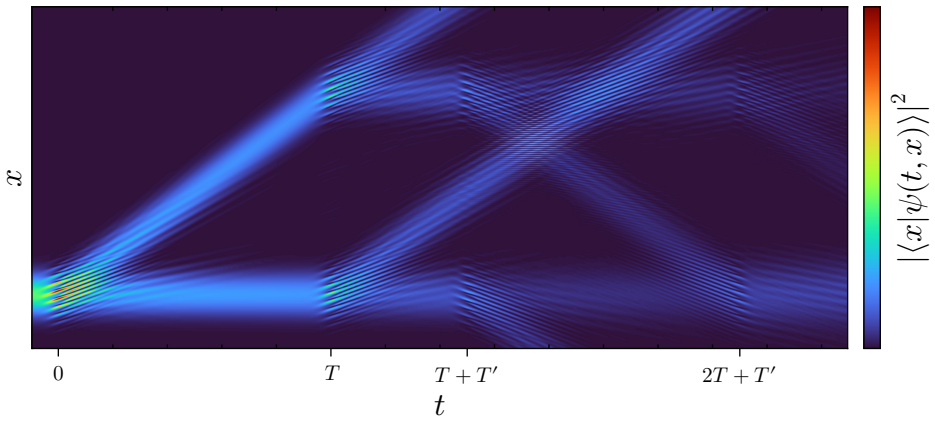


Figure 3.3: Position-space simulation of a Ramsey-Bordé interferometer, computed by integrating Eq. (6.14) using the split-step method. Four Bragg beamsplitters are applied at times $t = 0$, $t = T$, $t = T + T'$, and $t = 2T + T'$. Between beamsplitter pulses classical trajectories are followed, forming a trapezoidal interferometer geometry. Second-order Bragg beamsplitters are used with $\Omega_{\text{eff},0} = 5.1\omega_{\text{r}}$, $\sigma = 0.32\omega_{\text{r}}^{-1}$, $T = 10\omega_{\text{r}}^{-1}$, $T' = 5\omega_{\text{r}}^{-1}$. Note that δ is shifted from $\delta = 4n\omega_{\text{r}}$ to $\delta = 4n\omega_{\text{r}} - \omega_{\text{m}}$, where $\omega_{\text{m}} = 8n\omega_{\text{r}} - \Phi/nT$, where the target phase $\Phi = \pi/4$ of the Ramsey-Bordé interferometer. The wavefunction is initialized as a Gaussian wavepacket with amplitude width $\sigma_x = 5\sqrt{2}/k$, where k is the laser wavenumber.

T'. The last beamsplitter pair transfers momentum in the opposite direction from the first beamsplitter pair. This is accomplished by flipping the sign of δ to produce momentum transfer in the opposite direction. A simulation of the position-space free-particle wavefunction subjected to this sequence is shown in Fig. 3.3[1]. After the first Bragg beamsplitter, the wavefunction is described by two spatially separated wavepackets. The wavepackets travel along trajectories that we would expect from classical mechanics, which can be proven using the path integral formulation of quantum mechanics [21]. Because of this, we often display interferometer geometries using their classical trajectories, as shown in Fig. 3.4 for the case of the Ramsey-Bordé interferometer.

For a Bragg beamsplitter, recall from Eq. (2.18) that the resonance condition is $\delta = 4n\omega_{\text{r}}$ for the first beamsplitter pair and $\delta' = -4n\omega_{\text{r}}$ for the second beamsplitter pair (where n is the Bragg order). It is convenient to define the modulation frequency, ω_{m}, as the frequency by which we need to shift the value of δ for the first beamsplitter pair to get δ for the second beamsplitter pair, so that $\delta' = \delta - \omega_{\text{m}}$. To compute the phase of the interferometer, we sum

[1] At the beginning of this chapter I mentioned that atom interferometers are analogous to optical interferometers. Fig. 3.3 looks similar to an optical interferometer diagram. However, in an optical interferometer the x-axis would have units of position and not time. The position width of the wavepacket is analogous to the coherence length of light in an optical interferometer.

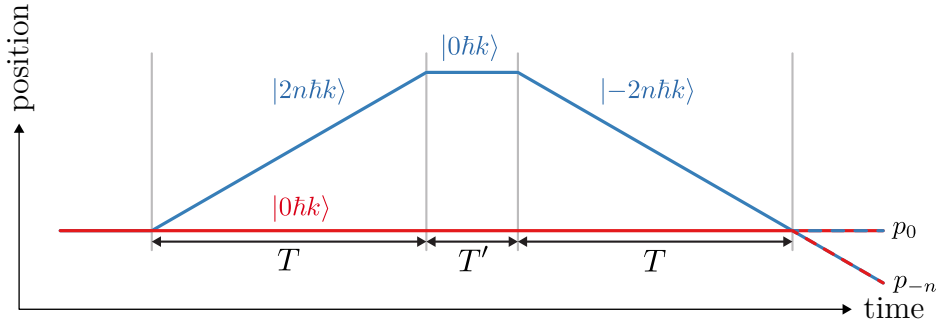


Figure 3.4: The classical trajectories followed in a Ramsey-Bordé interferometer geometry. Light vertical lines denote the times at which beamsplitters are applied. Note that these trajectories match those from the simulation in Fig. 3.3, although we have omitted trajectories that do not interfere at the end of the interferometer. The non-interfering trajectories are referred to as junk ports.

the laser phase and free evolution phase along each path exiting the upper port.

$$\begin{aligned} \phi_{\text{lower}} &= 0 \\ \phi_{\text{upper}} &= \phi_1 + 4n^2\omega_{\text{r}}T - \phi_2 - \phi_3 + 4n^2\omega_{\text{r}}T + \phi_4 \end{aligned} \tag{3.9}$$

where ϕ_1, ϕ_2, ϕ_3, and ϕ_4 are the laser phases given by Eq. (2.47). Note that because the laser phases were computed in position space, the free evolution phases must be as well. Therefore Eq. (3.5) must be expressed in position space to compute the relevant phases here. The lower path has no momentum and is never diffracted by a beamsplitter, so it has no free evolution phase or laser phase. The upper interferometer has momentum $2n\hbar k$ for duration $2T$ and is diffracted four times. The laser phase at the time and position of each beamsplitter pulse is

$$\begin{aligned} \phi_1 &= 0 \\ \phi_2 &= n(8n\omega_{\text{r}} - \delta)T \\ \phi_3 &= n(8n\omega_{\text{r}}T - (\delta - \omega_{\text{m}})(T + T')) \\ \phi_4 &= -n(2T + T')(\delta - \omega_{\text{m}}) \end{aligned} \tag{3.10}$$

Plugging everything in, we find the phase difference between the lower and upper interferometer paths to be

$$\Phi = \phi_{\text{lower}} - \phi_{\text{upper}} = 8n^2\omega_{\text{r}}T - n\omega_{\text{m}}T \tag{3.11}$$

This is the phase of a Ramsey-Bordé interferometer. From this result we see that the phase difference between the two arms is zeroed for $\omega_{\mathrm{m}} = 8n\omega_r$. In Section 2.2.1 we found that a Bragg beamsplitter imparts a phase of $\pm\pi/2$ on the diffracted state's phase in addition to the laser phase. We ignored this additional phase when computing the phase above. Writing out the wavefunction at the end of the interferometer with these additional phase shifts included yields

$$\begin{aligned}
|\phi\rangle =& \frac{1}{2}\left[\exp(i\phi_{\mathrm{lower}}) + \exp(i(\phi_{\mathrm{upper}} + 2\pi/2 - 2\pi/2))\right]|0\hbar k\rangle \\
&+ \frac{1}{2}\left[\exp(i(\phi_{\mathrm{lower}} - \pi/2)) + \exp(i(\phi_{\mathrm{upper}} + 2\pi/2 - \pi/2))\right]|-2n\hbar k\rangle \\
=& \frac{1}{2}\left[\exp(i\phi_{\mathrm{lower}}) + \exp(i\phi_{\mathrm{upper}})\right]|0\hbar k\rangle \\
&+ \frac{1}{2}\left[\exp(i(\phi_{\mathrm{lower}} - \pi/2)) + \exp(i(\phi_{\mathrm{upper}} + \pi/2))\right]|-2n\hbar k\rangle
\end{aligned} \tag{3.12}$$

where $|0\hbar k\rangle$ is the upper port state, $|-2n\hbar k\rangle$ is the lower port state, and ϕ_{lower} and ϕ_{upper} were taken from Eq. (3.9). We compute the populations to be

$$p_0 = |\langle 0\hbar k|\phi\rangle|^2 = \cos^2\left(\frac{\Phi}{2}\right) \qquad p_{-n} = |\langle -2n\hbar k|\phi\rangle|^2 = \sin^2\left(\frac{\Phi}{2}\right) \tag{3.13}$$

Note that, since the Ramsey-Bordé interferometer uses $\pi/2$ pulses for all beamsplitters, atoms exit from so-called "junk" ports at the time of the second and third Bragg pulses. We have ignored these states in the above calculation as they do not contribute to the interference signal.

This result should make some sort of intuitive sense. In the case where $\Phi = 0$ we have set ω_{m} such that the third and fourth Bragg diffraction pulses are perfectly on resonance. This means that the net momentum transfer to the atom must be zero, which can only happen if the wavefunction at the end of the interferometer is confined to the same momentum state it entered the interferometer in.

We can plot the inversion between the populations as

$$w = \frac{p_0 - p_{-n}}{p_0 + p_{-n}} = \cos(\Phi) \tag{3.14}$$

We now see that the Ramsey-Bordé interferometer produces a fringe with frequency set by the difference between the natural frequency of the system $8n^2\omega_{\mathrm{r}}$ and the drive frequency $n\omega_{\mathrm{m}}$. This is analogous to the Ramsey fringe discussed in Section 3.1, providing a method for measuring the recoil frequency.

Note that we also ignored the anomalous diffraction phase described in Section 2.2 in this calculation. We will consider the effect of this phase in Section 7.1, where we compute the interferometer phase numerically.

The phase calculation can be automated by representing the interferometer as a tree. The tree is initialized with a single node representing the wavefunction at the start of the

interferometer. Each node has a position, momentum, and time. Each operation applied to the interferometer creates a new level in the tree. Free propagation operations create a child node for each parent node. Beamsplitter operations create two child nodes for each parent node if the parent node is resonant with the beamsplitter, and one child node if not. Each point in the interferometer geometry corresponds to a node, and so the phase at any point can be found by tracing through the lineage of that node. I describe the exact algorithm for this calculation in Section 6.2.1. The remaining interferometer phases reported in this section were computed using this method and match the values found in the literature.

3.2.3 Simultaneous Conjugate Ramsey-Bordé interferometer

In practice atom interferometers are often implemented in an atomic fountain—an experimental apparatus that launches a cloud of atoms vertically in a vacuum chamber and then waits for them to fall back down—in order to achieve long duration free evolution times. In this configuration the two interferometer paths experience different gravitational potentials, which turns our previously linear trajectories into quadratic ones (for the case of a linear gravitational potential).

Since we are interested in measuring ω_{r} and not the difference in gravitational potential along the two paths traversed by the interferometer, we must find a way to remove this phase from our measurement. This can be done by operating two Ramsey-Bordé interferometers simultaneously to form a simultaneous conjugate Ramsey-Bordé interferometer (SCI) [58], as shown in Fig. 3.5.

For a linear gravitational potential with gravity g, the accumulated phases in the upper and lower interferometers are

$$\Phi_{\pm} = 2nT(T+T')kg \mp 8n^2T\omega_{\mathrm{r}} \pm nT\omega_{\mathrm{m}} \tag{3.15}$$

where Φ_+ is the upper interferometer and Φ_- is the lower interferometer. The free evolution phase in this case is computed using Eq. (3.6) where the Lagrangian is taken to have a linear gravitational potential. The differential phase between the two interferometers is therefore

$$\Phi = \Phi_- - \Phi_+ = 16n^2T\omega_{\mathrm{r}} - 2nT\omega_{\mathrm{m}} \tag{3.16}$$

We see that this phase is no longer dependent on the linear gravitational potential. However, it is sensitive to any higher-order terms in the gravitational potential, as shown later on. To implement the SCI, a multifrequency Bragg beamsplitter is used, as described in Section 2.2.2. Note that our definition of the modulation frequency ω_{m} is consistent with the definition of ω_{m} in Section 2.2.2!

This interferometer will produce two fringes, one for the upper interferometer and one for the lower. Working through the same argument as before, we find

$$\begin{aligned} p_0 &= |\langle 0\hbar k|\phi\rangle|^2 = \cos^2\left(\frac{\Phi_-}{2}\right) \qquad p_{-n} = |\langle -2n\hbar k|\phi\rangle|^2 = \sin^2\left(\frac{\Phi_-}{2}\right) \\ p_n &= |\langle 2n\hbar k|\phi\rangle|^2 = \cos^2\left(\frac{\Phi_+}{2}\right) \qquad p_{2n} = |\langle 4n\hbar k|\phi\rangle|^2 = \sin^2\left(\frac{\Phi_+}{2}\right) \end{aligned} \tag{3.17}$$

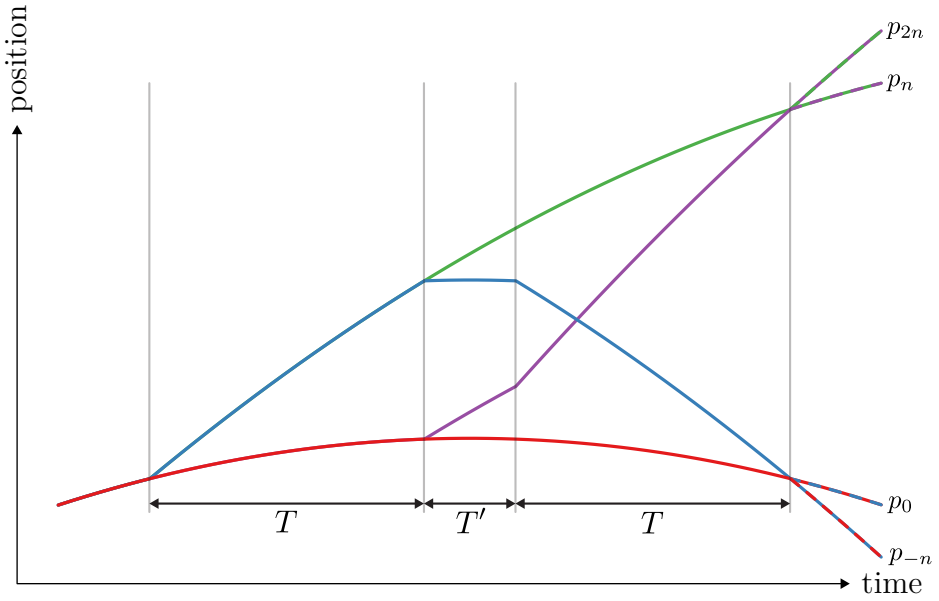


Figure 3.5: Trajectories followed in an SCI geometry in a linear gravitational potential. Light vertical lines denote the times at which beamsplitters are applied. The lower interferometer is the same as that in Fig. 3.4, but in a gravitational potential. The upper interferometer looks like an upside-down version of the lower interferometer, which leads to the cancellation of any phase due to the linear gravitational potential.

The two fringes are therefore given by

$$w_- = \frac{p_0 - p_{-n}}{p_0 + p_{-n}} = \cos(\Phi_-) = \cos(\phi_c + \phi_d) \qquad w_+ = \frac{p_n - p_{2n}}{p_n + p_{2n}} = \cos(\phi_c - \phi_d) \tag{3.18}$$

where we have defined the common and differential mode phases $\phi_c = 2nT(T + T')kg$ and $\phi_d = 8n^2T\omega_\mathrm{r} - nT\omega_\mathrm{m}$.

In the Ramsey interferometer we determined the frequency of a two-level system by matching the drive frequency to the energy difference of the two-level system. Here we do the same thing. Our drive frequency is ω_m, and the frequency of our two-level system is proportional to ω_r. By tuning ω_m to zero the phase of the interferometer, we know that $\omega_\mathrm{m} = 8n\omega_\mathrm{r}$.

The ultimate goal of the fine-structure measurement experiment is to measure ω_r to the highest possible precision. From the expression $\omega_\mathrm{m} = 8n\omega_\mathrm{r}$ we see there are two ways to do this: increase the precision with which we can determine ω_m, or increase n. The former can be done by running the experiment for longer and increasing the experiment's statistical precision, or by increasing T (this makes the interferometer fringes more sensitive to changes

in ω_{m}, as seen in Fig. 3.2). The latter can be accomplished by driving higher-order Bragg beamsplitter pulses, or by modifying the interferometer geometry.

One method of modifying the interferometer geometry is to introduce Bloch oscillations that pull the lower and upper interferometers in the SCI apart, as shown in Fig. 3.6. The figure assumes there is no gravitational potential, as we know the SCI cancels any phase stemming from the linear gravitational potential. The Bloch oscillations are performed between the second and third Bragg beamsplitter pulses, with the lower interferometer getting accelerated downwards and the upper interferometer getting accelerated upwards. The accumulated phases in the upper and lower interferometers are

$$\Phi_{\pm} = 2nT(T+T')kg \mp 8n(N+n)T\omega_{\mathrm{r}} \pm nT\omega_{\mathrm{m}} \tag{3.19}$$

where N is the number of Bloch oscillations. This yields differential phase

$$\Phi = 16n(N+n)T\omega_{\mathrm{r}} - 2nT\omega_{\mathrm{m}} \tag{3.20}$$

The common and differential mode phases ϕ_c and ϕ_d change accordingly. Note that this configuration does not lead to a larger momentum shift between arms in the same interferometer. However, it does lead to an increased kinetic energy difference between arms in the same interferometer.

If Bloch oscillations were used to pull the arms within each individual interferometer apart, then any phase shift imparted by Bloch oscillations could be different between the two arms. This phase shift would need to be understood. In the configuration we use, both arms experience approximately the same Bloch oscillation process, and so any imparted phase shift is common to the two arms and does not contribute to the interferometer phase.

Nonlinear gravitational potential

While the SCI geometry cancels the linear effect of gravity, it does not cancel any higher-order effects. The phase resulting from these effects must be measured and then removed. We now compute the free evolution phase accumulated in a one-dimensional nonlinear gravitational potential using Eq. (3.6). Its Lagrangian is

$$\mathcal{L} \approx \frac{1}{2}m\dot{z}^2 - mgz + \frac{1}{2}m\gamma z^2 \tag{3.21}$$

where we Taylor expanded the gravitational potential to second order to get the usual linear gravitational potential term dependent on g and the quadratic gravitational potential term dependent on the gravity gradient γ. Solving the Euler-Lagrange equation for this Lagrangian yields the trajectory

$$z(t) = z(0)\cosh(\sqrt{\gamma}t) + \frac{g}{\gamma}\left(1 - \cosh(\sqrt{\gamma}t)\right) + \frac{\dot{z}(0)}{\sqrt{\gamma}}\sinh(\sqrt{\gamma}t) \tag{3.22}$$

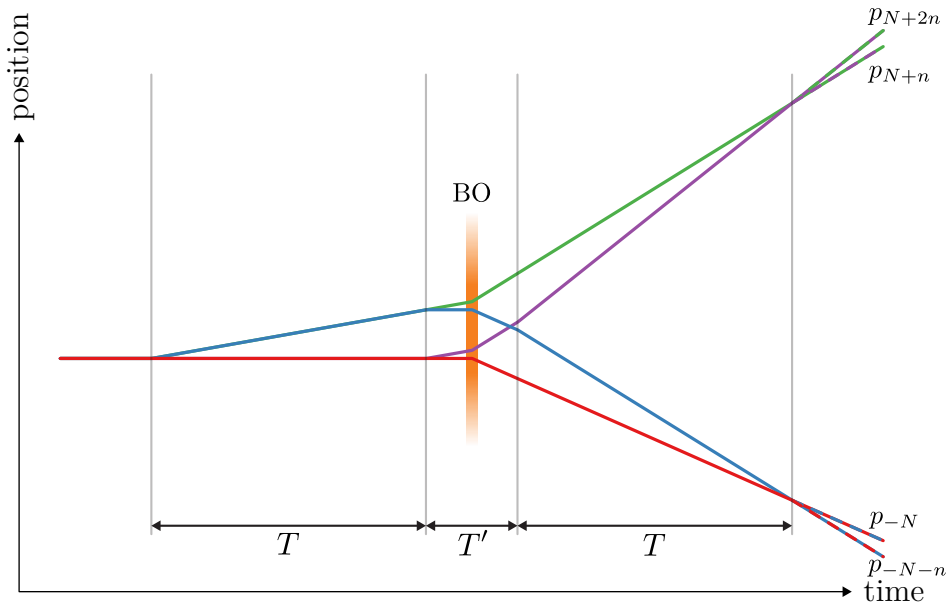


Figure 3.6: Trajectories followed in an SCI geometry with Bloch oscillations (BO, in orange) occurring halfway between the second and third beamsplitter pulses. The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters are applied.

This form for $z(t)$ can then be Taylor expanded to first order in γ, plugged back into $\mathcal{L}$, and then integrated to compute the action, i.e. $S = \int \mathcal{L}\,\mathrm{d}t$. The resulting action can then again be Taylor expanded to first order in γ to yield

$$\begin{aligned} S \approx & \frac{1}{2}m\dot{z}(0)^2T + \frac{1}{3}mgT\left(gT^2 - 3(z(0) + T\dot{z}(0))\right) \\ & + \frac{\gamma}{30}m\left[2g^2T^4 - 10gT^2(2z(0) + \dot{z}(0)T) + 5(2T^2\dot{z}(0)^2 + 6T\dot{z}(0)z(0) + 3z(0)^2)\right]T \end{aligned} \tag{3.23}$$

This approximation of the action is plugged back into Eq. (3.6) to compute the phase. Comparing this to the phase accumulated in Eq. (3.5), we see that we still have the phase resulting from kinetic energy (just plug in $\dot{z}(0) = 2j\hbar k/m$), but we also have phases due to the potential. Eq. (3.5) must be modified to include these phases in such a nonlinear potential. Including the gravity and gravity gradient changes the phase accumulated in an SCI to

$$\Phi = 16n(N+n)T\omega_\mathrm{r} - 2nT\omega_\mathrm{m} + \phi_\mathrm{gravity} \tag{3.24}$$

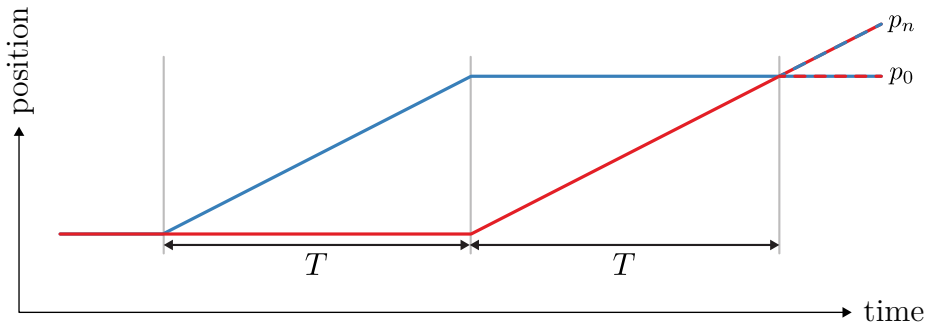


Figure 3.7: Trajectories followed in a Mach-Zehnder interferometer (MZI). The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters (or optionally a mirror in the case of the intermediate pulse) are applied.

where ϕ_{gravity}, Taylor expanded to first order in γ, is

$$\phi_{\text{gravity}} = \frac{2}{3} n\gamma\omega_{\text{r}} T \left[n(4T^2 + 6TT' + 6T'^2) + N(4T^2 + 6TT' + 3T'^2) \right] \tag{3.25}$$

Note that this assumes Bloch oscillations happen instantaneously.

3.2.4 Mach-Zehnder interferometer

The Mach-Zehnder interferometer uses a sequence of three beamsplitters. This forms the geometry seen in Fig. 3.7. It is often used to measure the gravitational potential. The accumulated phase of the interferometer is given by

$$\Phi = 2nT^2 kg \tag{3.26}$$

From this phase equation we see that there is no way to zero the phase of the interferometer. In order to zero the phase, the drive frequency of the beamsplitter can be swept in time via $\delta(t) = \delta_0 + \alpha t$, which results in an interferometer phase

$$\Phi = nT^2(2kg + \alpha) \tag{3.27}$$

By selecting the sweep rate that zeros the interferometer phase, the experimenter knows that $\alpha = -2kg$. A simultaneous Mach-Zehnder interferometer can be used to measure the gravity gradient, as described in [21] Section 5.3.

3.3 Separation phase

In certain cases interferometer arms might not close perfectly. For example, if the field wavenumber k has a spatially dependent value (this is true in a Gaussian laser beam), then the Bragg beamsplitter does not always impart the same momentum, which results in wavepackets with a slight separation at the end of the interferometer. Separation of the final wavepackets can also occur from propagating in a nonlinear potential.

This wavepacket separation produces a phase shift in the interferometer [59, 60]. For the case of two separated Gaussian wavepackets, we can derive the phase shift in a straightforward manner. Consider a wavefunction composed of two Gaussian wavepackets

$$\psi(x) = \langle x|\psi\rangle = \psi_1(x) + \psi_2(x) \tag{3.28}$$

where

$$\psi_{1,2}(x) = \frac{1}{\sqrt{\sigma\sqrt{\pi}}} \exp\left(-\frac{(x-x_{1,2})^2}{2\sigma^2} + i\frac{1}{\hbar}p_{1,2}(x-x_{1,2}) + i\theta_{1,2}\right) \tag{3.29}$$

The observed population will be

$$P = \int_{-\infty}^{\infty} \psi(x)^*\psi(x)\,\mathrm{d}x \tag{3.30}$$

$$= \int_{-\infty}^{\infty} \left[\psi_1(x)^*\psi_1(x) + \psi_2(x)^*\psi_2(x) + \psi_1(x)^*\psi_2(x) + \psi_2(x)^*\psi_1(x)\right]\,\mathrm{d}x \tag{3.31}$$

Clearly we only get interference from the last two terms in the integrand. Integrating these terms yields

$$\begin{aligned}&\int_{-\infty}^{\infty} \left[\psi_1(x)^*\psi_2(x) + \psi_2(x)^*\psi_1(x)\right]\,\mathrm{d}x \\ &\quad = 2\cos\left(\frac{p_1+p_2}{2\hbar}(x_2-x_1) + \theta_1 - \theta_2\right)\exp\left(-\frac{(x_1-x_2)^2}{4\sigma^2} - \frac{\sigma^2}{4\hbar^2}(p_1-p_2)^2\right)\end{aligned} \tag{3.32}$$

We immediately see that in addition to the expected phase $\theta_1 - \theta_2$ we have an additional term given by

$$\phi_{\text{separation}} = \frac{p_1+p_2}{2\hbar}(x_2-x_1) = \frac{\bar{p}\Delta x}{\hbar} \tag{3.33}$$

This is the separation phase. Note that computing the separation phase does not require both wavepackets to have a Gaussian profile in space. However, it does require that both wavepackets have the same real-valued profile in space. My understanding is that for two wavepackets with distinct profiles the separation phase must be computed on a case-by-case basis.

An alternative derivation of the separation phase can be made by considering the de Broglie wavelength of a wavepacket. A wavepacket with momentum p has wavelength $\lambda = h/p$. Therefore, if two wavepackets with momentum p are separated by Δx then, as expected, a separation phase of $2\pi\Delta x/\lambda = \bar{p}\Delta x/\hbar$ appears.

3.4 Interferometer contrast

In practice many atoms propagate through an atom interferometer at once, and so a population of atoms is measured in each momentum state at the end of the interferometer. This allows the populations in Eq. (3.17) to be directly read out after each experiment shot, as opposed to being measured by averaging over many experiment shots.

Each atom propagating through the interferometer will accumulate a slightly different phase along its trajectory. For example, because atoms are spread out across a finite-sized laser beam, they experience a range of Ω_{eff} values, which leads to atoms accumulating different diffraction phases, as seen in Fig. 2.7. Furthermore, the variation in Ω_{eff} means that most atoms do not experience a perfect $\pi/2$ pulse (see Fig. 2.5), reducing the amplitude of their interference fringes. The measured population in each output port of the interferometer therefore amounts to an average over sinusoids with slightly different phases and amplitudes, producing a shift in the averaged fringe phase and a suppression of the averaged fringe amplitude.

Furthermore, individual atoms can lose coherence through single-photon scattering or other effects, which can result in an atom that is detected at the interferometer output but does not contribute to interference. Additionally, the junk ports mentioned in Section 3.2.2 can also lead to a non-interfering background. These effects further suppress the fringe amplitude and can produce a fringe offset. Taken together, all of these effects require us to rewrite Eq. (3.18) as

$$w_- = b_- + A_- \cos(\phi_c + \phi_d) \qquad w_+ = b_+ + A_+ \cos(\phi_c - \phi_d) \tag{3.34}$$

where A_-, A_+, b_-, and b_+ are determined by the particular contrast loss conditions present.

The fringe amplitude is often referred to as the contrast (defined here as $\sqrt{A_+ A_-}$ for the case of a simultaneous interferometer), and any effect that reduces the fringe amplitude is referred to as a source of contrast loss. It is annoying to do all of the algebra to compute exactly how contrast loss shifts the ellipse location. The simulation methods described in Chapter 6 account for these sources of loss numerically, which enables the automatic simulation of ellipse location (see Section 7.1.7). The simulations shown up to this point assume a single atom and an infinite interferometry beam size, so are not sensitive to these contrast loss effects. Simulations that do account for these contrast loss effects are considered in Chapter 6.

Chapter 4

Experiment Apparatus

In Chapter 3 we discussed how to make an atom interferometer using the matterwave optics described in Chapter 2. To build a working interferometer, we must prepare a gas of atoms in a well-defined quantum mechanical state. To do this we need an extremely cold gas–so cold that each atom in the gas is only moving at a few millimeters per second–made up of cesium atoms. This is the responsibility of the systems described in this section.

We refer to each individual experiment run as an experiment shot. The experiment performs one shot every 2.7 s. At a high level, a single experiment shot starts with a two-dimensional magneto-optical trap (2-D MOT) loading atoms into a three-dimensional magneto-optical trap (3-D MOT) for around a second. The MOTs use radiation pressure from photon scattering to cool and trap atoms [61]. The trapped atoms are referred to as an atom cloud. After loading is completed the magnetic fields required for the MOT are turned off, meaning that the atom cloud is no longer trapped. Polarization gradient cooling (PGC) is performed to further cool the atom cloud before the cloud is launched vertically upwards into the cylindrical vacuum chamber. As the cloud moves upwards, Raman sideband cooling (RSC) is used to cool the atoms in the cloud to an even lower temperature. After RSC, atoms are in the $F = 3, m_F = 3$ state. Next a microwave pulse is used to transfer the atoms to the $F = 3, m_F = 0$ state via adiabatic passage. Two velocity selective Raman pulses are then performed to narrow the vertical velocity distribution of the atom cloud to roughly a tenth of a recoil velocity. After this, Bragg diffraction and Bloch oscillation pulses are performed to construct an atom interferometer. The interferometer ends with atoms in a series of different momentum states depending on the accumulated interferometer phase. These distinct atom clouds then fall through the detection chamber, where they are detected by fluorescence.

In this section I will describe the vacuum chamber, optical systems, laser systems, and control system used in the experiment. Each of these systems was built between 2020 and 2026 for this generation of the Berkeley fine-structure measurement, and did not exist in the apparatus used for the 2018 Berkeley measurement. I will discuss how these systems are tuned to optimize the laser cooling and trapping, state preparation, and interferometer stages of the experiment. Finally I will discuss the steps taken to make the experiment robust

and capable of continuous operation. The laser locking system described in Section 4.2.6 resulted in a publication [25].

4.1 Experiment vacuum chamber, optics, and coils

The vacuum system was designed by Zachary Pagel [24]. The main chamber is a cylinder 0.4 m in diameter and 4 m long (inner dimensions). The chamber is composed of four separate cylindrical chamber components. From the bottom to the top of the chamber, these are the 3D MOT chamber, the detection chamber, the interferometry chamber, and the retroreflection (retro) mirror chamber. A 2D MOT chamber feeds atoms into the MOT chamber. A SolidWorks rendering of the entire vacuum chamber is shown in Fig. 4.1 panel a. The large chamber diameter was chosen to accommodate a large interferometry beam, as the 2018 fine-structure constant measurement determined that increasing the interferometry beam size would reduce a variety of systematic effects [3] (this is discussed further in Section 7.1).

The chamber is held by an optical table, and passes through a circular hole bored into the middle of the table. The optical table is located over an excavated "pit" in the basement level of Campbell Hall on the UC Berkeley campus. While the pit makes it possible to fit a 4 m tall experiment in the lab, it makes accessing different parts of the experiment physically difficult, as experimenters must awkwardly crawl around the pit to access different optical elements, fibers, and sensors.

4.1.1 2D MOT chamber

A cross section of the 2D MOT chamber is shown in Fig. 4.1 panel b. It consists of a glass vacuum chamber connected to the MOT chamber through a differential pumping tube. The exact pressure differential between the MOT and 2D MOT chambers has not been measured [24]. A bellows containing a cesium ampule is attached to the glass chamber. After the vacuum system was pumped down, the ampule was broken by bending the bellows, allowing cesium vapor to fill the glass chamber. A 2D MOT is used to direct a beam of cesium atoms through the differential pumping tube to the center of the MOT chamber, where the beam is trapped in a 3D MOT. This process is referred to as loading the 3D MOT.

The differential pumping tube is 6.35 mm in inner diameter [24] and 140 mm long [21], meaning that the most extreme angle of an atom exiting the differential pumping tube is 45 mrad. The distance from the differential pumping tube to the center of the chamber is $\sim$400 mm, so the maximum diameter of the atom beam at the chamber center is $\sim$36 mm.

To create the 2D MOT, four hand-wound coils are arranged in a typical configuration (see [62]). The coils are run at 5 A, each have 120 windings, and are approximately 20 cm in diameter. We expect, and also measure, a magnetic field gradient of 4 G/cm from the coils. During experiment sequences, the coils are permanently on. This is not ideal, as the coils heat the surrounding vacuum chamber and produce a magnetic field that atoms in the main chamber interact with. In the future it would be beneficial to switch the coils on

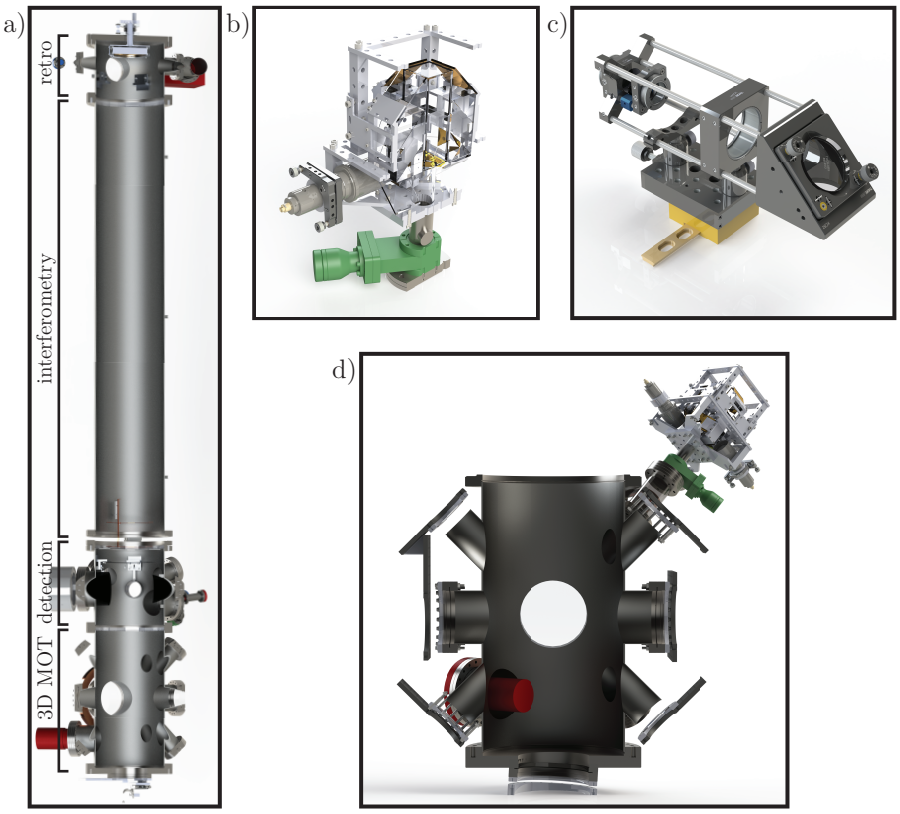


Figure 4.1: Renderings of the experiment chamber. *a:* The 3D MOT, detection, interferometry, and retroreflection chambers (labeled in figure). *b:* The 2D MOT chamber. *c:* The 3D MOT telescope. *d:* The 3D MOT chamber with the 2D MOT chamber attached.

only when the 2D MOT is used to load the 3D MOT. Combined cooling and repump light is generated from a tapered amplifier system described in Section 4.2.3, which is delivered to the 2D MOT chamber through two optical fibers. Nominally 400 mW is coupled into each fiber. Each fiber is connected to a 60FC-Q852-4-M200-37 Schäfter Kirchhoff telescope with an internal $\lambda/4$ waveplate, producing a 15.4 mm $1/e^2$ radius beam. The waveplates are aligned to produce circularly polarized light before passing through a series of non-polarizing beamsplitters designed to create three regions in the glass cell where a 2D MOT can form. A shutter before each optical fiber can be used to quickly turn the 2D MOT on and off.

This 2D MOT design provides enough atom flux to load a sufficiently large 3D MOT for the experiment, as will be discussed in Section 5.3. However, the design could be improved in the following ways: 1) the capture velocity of a 2D MOT scales with the length of the 2D MOT region [63]; and 2) there are not enough degrees of freedom in the optical system to align the direction of the 2D MOT beams generated in each of the three regions. This likely produces three somewhat independent 2D MOTs. Each 2D MOT region has a reduced

length (lowering the capture velocity), and the three regions do not necessarily produce atom beams aimed in the same direction.

4.1.2 3D MOT chamber

The 3D MOT chamber is located at the bottom of the experiment apparatus. The chamber is used to create the 3D MOT and to perform RSC (described in Section 4.4.3). The bottom of the chamber contains a 0.1 m radius viewport where the interferometry laser beam is sent into the system. The viewport is slightly angled to prevent etaloning with the retroreflection mirror at the top of the chamber. Two coils are mounted in an anti-Helmholtz configuration on either side of the vacuum chamber to produce the magnetic field required for the 3D MOT. The coils each have 36 windings, an inner radius of 0.18 m, are spaced by 0.508 m, and typically operate with 200 A current. This coil configuration produces an estimated field gradient of 7.7 G/cm at the center of the chamber. The switching circuit described in Appendix B switches the 3D MOT coils on and off.

The original coils were not tightly wound, which led the individual windings to move relative to each other due to eddy currents when the coils were switched off. It was eventually determined that the coil turn-off led to vibrations in the chamber which produced a major systematic effect. Replacement coils were designed by Madeline Bernstein in 2025 that are mounted to the pit walls and do not touch the experiment chamber. This seems to have drastically reduced the vibrations present in the chamber. The effect of these vibrations is discussed in Section 7.1.2.

The 3D MOT telescopes responsible for injecting cooling and repump light also went through several design iterations. Originally the same telescopes from the 2018 fine-structure constant measurement were used. Each telescope provided a mirror that could be angled to adjust the pointing and location of its 3D MOT beam. However, with such a large experiment chamber, it was difficult to ensure that each counterpropagating beam pair was overlapping and orthogonal to the other two beam pairs. Furthermore, as will be discussed in Section 4.4.2, the 3D MOT beams must also be aligned to create a launch aligned with gravity.

To provide both position and angle degrees of freedom to the 3D MOT beams, we used a new telescope design shown in Fig. 4.1 panel c. The fiber is outcoupled and then collimated using a Thorlabs LA1979-B $f = 200$ mm lens to create a 20 mm $1/e^2$ diameter beam. Between the fiber and the lens a $\lambda/4$ waveplate is used to circularly polarize the light (aligned using a Thorlabs PAX1000IR1 polarimeter). A shearing interferometer was used to collimate the output beam at a location roughly 1 m away from the lens. The fiberport and lenses are mounted using a cage system. At the end of the cage a kinematic mirror mount is used to reflect the beam 90° and to provide pointing degrees of freedom. The entire cage mount assembly is mounted on a linear stage and an optical rail to provide translation degrees of freedom. A gold mirror is used in the kinematic mirror mount to maintain the polarization set by the $\lambda/4$ waveplate. The laser system used for generating the 3D cooling and repump light is described in Section 4.2.4.

Approximately 10 cm above the location of the 3D MOT is the RSC region. We will discuss the mechanism for RSC and the alignment procedure used to align the RSC beams in Section 4.4.3. RSC is performed with a constant magnetic field in an optical lattice formed by two traveling waves and one standing wave. The magnetic field is formed using a pair of coils in a Helmholtz configuration wrapped around a chamber viewport and driven with 5.3 A. The coils are connected in series and switched on and off by a TTL signal using the circuit shown in Fig. B.3. The field at the center of the chamber is measured to be 140 mG using microwave spectroscopy.

The standing wave component of the optical lattice is produced by retroreflecting a laser beam. This beam is referred to as the RSC X beam. The incident RSC X beam comes from a Thorlabs F280APC-850 in a cage mount system, which is mounted using components that make the beam position and angle degrees of freedom straightforward to adjust. The outcoupled beam is 4.0 mm in $1/e^2$ diameter and has 160 mW of power. The beam is immediately passed through a polarizing beamsplitter to ensure a linear polarization, then passes through a kinematic mirror mount that allows the beam pointing to be easily adjusted, and finally passes through a $\lambda/2$ waveplate. On the opposite side of the chamber the beam passes through a $\lambda/4$ waveplate and is then retroreflected back through the waveplate and into the chamber. The double-passed $\lambda/4$ waveplate produces a retroreflected beam linearly polarized at some angle relative to the incident linearly polarized beam.

The two traveling wave components of the optical lattice are referred to as the RSC Y and RSC Z beams. The RSC Y beam comes from another Thorlabs F280APC-850 in a cage mount system with a polarizing beamsplitter and translation and pointing degrees of freedom. The outcoupled beam is 4.0 mm in $1/e^2$ diameter and has 60 mW of power. The RSC Z beam is coupled into the interferometry beam fiber as shown in Fig. 4.11, and has around 40 mW of power. Both the RSC Y and Z beams have $\lambda/2$ waveplates that can be used to adjust their linear polarizations. The RSC beams are aligned with the goal of making the RSC Z beam parallel to gravity, the RSC X beam orthogonal to gravity, and the RSC Y beam orthogonal to gravity and the RSC X beam.

The interferometry beam is sent into the chamber through the 0.1 m radius viewport at the bottom of the 3D MOT chamber. Several translation stages are used to align the beam to the atom cloud. Typically the Schäfter Kirchhoff 60FC-T-4-M75-37 $f = 75$ mm fiber collimator is used to outcouple the interferometry beam, producing a beam with $1/e^2$ radius of 7.7 mm. The beam is aligned parallel to gravity using the method described in Section 4.2.8. A microwave horn is located at the bottom of the interferometry chamber and is aimed through the bottom viewport. The horn is used to drive microwave transitions during the adiabatic rapid passage stage of the state preparation sequence.

4.1.3 Detection chamber

The detection chamber is located directly above the 3D MOT chamber and contains the lasers and optics required for fluorescence detection of atoms launched upwards from the 3D MOT chamber. To make the atoms fluoresce they are illuminated by a laser beam resonant

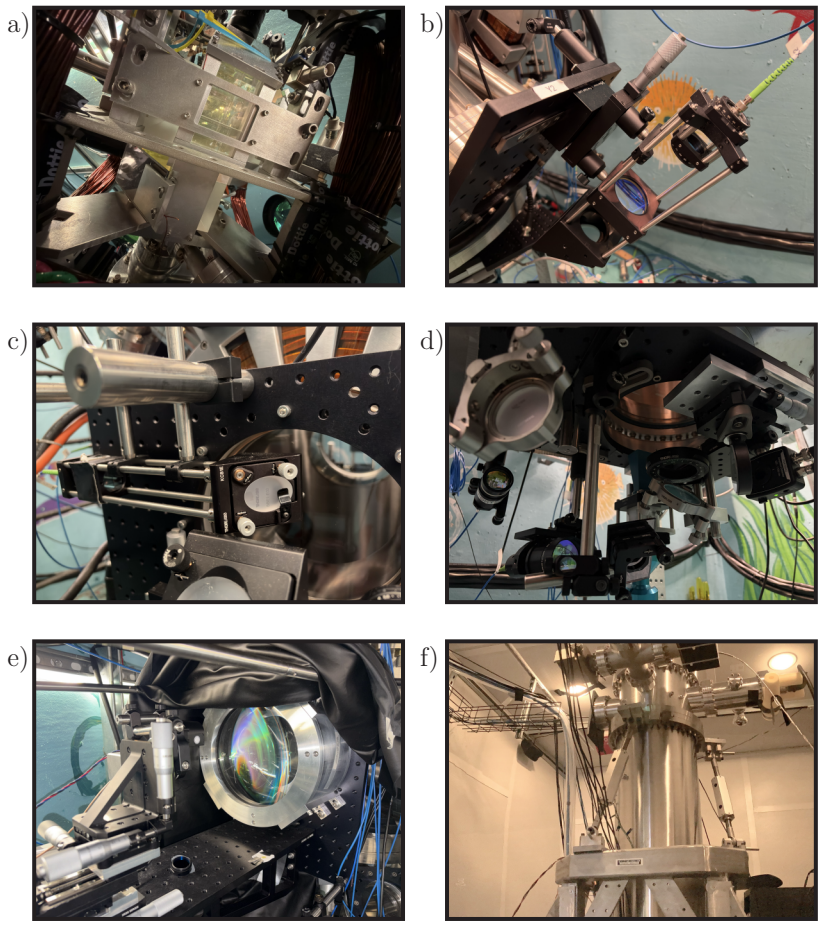


Figure 4.2: *a:* The 2D MOT. *b:* A 3D MOT beam telescope mounted upside down. *c:* The RSC X telescope, with the top of a 3D MOT beam telescope visible. *d:* Optics underneath the 3D MOT chamber used to align the interferometry beam and one of the blowaway beams. The microwave horn is also visible. *e:* The detection optics. Usually a light-blocking sheet is draped over the setup to reduce the detection background. *f:* The top of the interferometry chamber.

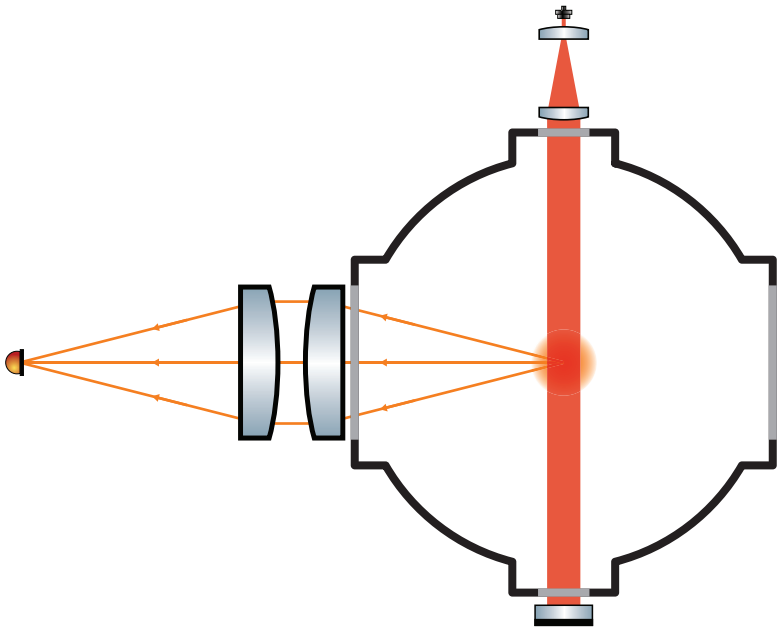


Figure 4.3: Schematic of the detection chamber and optics. The orange cloud represents the atoms at the time of detection. A relay lens system recreates an image of the atoms on a photodiode. The resonant detection light is expanded using a cylindrical lens and enters the chamber perpendicular to the relay lens system optical axis, illuminating the cloud. The width of the detection light, positioning, and active area of the photodiode set the region of the cloud that is detected.

with the $6\,^2S_{1/2}F = 4 \rightarrow 6\,^2P_{3/2}F' = 5$ transition (see Fig. 4.6). The laser beam passes through a cylindrical telescope so it expands to cover a wide region of the chamber, as shown in Fig. 4.3.

To detect the fluorescence emitted from the atoms, a pair of uncoated Edmund Optics 15-027 aspheric lenses are used. The lenses are 200 mm in diameter. The first lens is oriented with its flat face towards the chamber. The flat face is located 335 mm from the center of the chamber. The second lens is oriented with its flat face away from the chamber. The flat faces are spaced by 99 mm. A photodiode is located 230 mm away from the flat face of the second lens. This configuration should produce an approximate relay system. The photodiode is located on a three-axis translation stage, and is aligned to maximize the fluorescence signal, so it is not always located a fixed 230 mm away from the flat face of the second lens. The photodiode is sampled using an analog-to-digital converter built into the ARTIQ control system described in Section 4.3.

The transmission efficiency of the lens system is estimated to be 80% using Zemax OpticStudio. When the photodiode is placed 230 mm from the second lens flat face, the magnification of the system is estimated to be 0.78 when considering an object located at the

center of the chamber.

4.1.4 Interferometry chamber

Above the detection chamber is the interferometry chamber. The interferometry chamber is the longest chamber and contains no viewports. Inside the chamber there are three layers of magnetic Mu-metal, an alloy engineered to have high permeability, which shields the atom interferometer from magnetic fields. When the magnetic field in the interferometry chamber is close to zero, small differences in the magnetic field along each interferometer arm lead to m_F states mixing and produce a zero contrast interferometer. To prevent this a solenoid is placed inside the Mu-metal shielding and is run with a constant current to produce a small field, which sets a quantization axis for the atoms in the interferometer.

A series of black cupric oxide baffles 0.11 m in inner radius are placed inside the solenoid and absorb any light from the interferometry beam that has propagated to the chamber walls [24, 64].

4.1.5 Retroreflection chamber

The final chamber component is the retroreflection chamber, which is located above the interferometry chamber. The retroreflection chamber contains a 0.1 m radius retroreflection mirror mounted to the top flange of the experiment chamber through a bellows. The bellows allows the mirror mount to be angled from outside the chamber using a kinematic mount built into the top of the chamber. The mount enables tilts in two directions using Thorlabs DRV517 micrometers with closed-loop feedback piezos.

4.1.6 Compensation coils

Three pairs of Helmholtz compensation coils are used to zero the magnetic field in the 3D MOT region. Two of these pairs are located on the walls of the pit, and are driven using the output of a current-limited DC power supply. These pairs zero the magnetic field component in the North-South and East-West directions. The up-down magnetic field is compensated using a Helmholtz coil pair attached to the vacuum chamber. This pair is controlled using the current driver circuit shown in Fig. B.4. The coil current is adjusted throughout the state preparation sequence depending on the magnetic field configuration required for each state preparation stage.

4.1.7 Vacuum pumping

The vacuum chamber was initially pumped down using a roughing pump and turbopump connected to a gate valve attached to the retroreflection chamber. After the chamber bakeout the gate valve was closed and the mechanical pumps were disconnected from the chamber. This happened in spring 2021. Since then the chamber has been pumped using ion pumps,

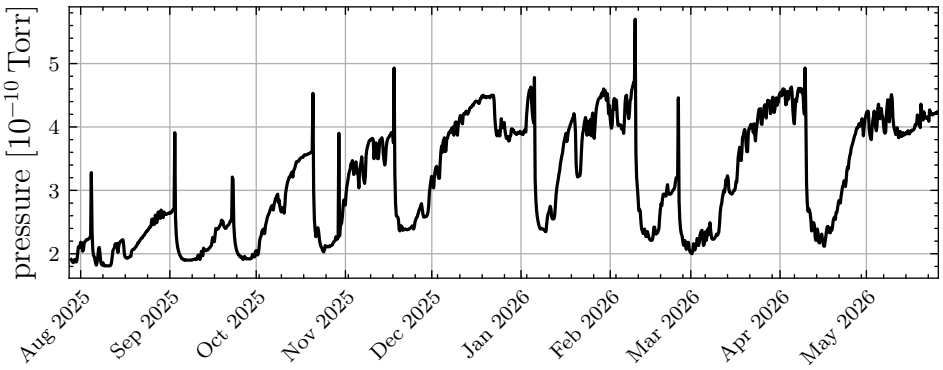


Figure 4.4: Recorded vacuum pressure over time. Spikes correspond to times the titanium sublimation pump was fired.

two titanium sublimation pumps (tisubs), and two non-evaporable getter (NEG) pumps. The chamber also has a residual gas analyzer (RGA), attached to the 3D MOT chamber through an elbow connection, that can be used to determine the composition of gases in the vacuum. The flange the RGA is attached to also has an ion gauge and a tisub. Normally only a single ion pump located on the retro chamber is on. The tisubs are fired as needed and both NEG pumps were activated in spring 2021. Since their initial activation only a single NEG pump has been reactivated.

The ion gauge located on the elbow attached to the 3D MOT region read $\sim 0.7 \times 10^{-11}$ Torr in 2021. The baseline pressure when the experiment has not been running has slowly crept up since then. As of May 2026 the pressure hovers around $\sim 4.3 \times 10^{-10}$ Torr if no actions are taken to reduce the pressure. After firing the tisub the pressure is reduced to $\sim 2.3 \times 10^{-10}$ Torr, and after 4-5 weeks the pressure climbs back up to $\sim 4.3 \times 10^{-10}$ Torr. Several months of ion gauge measurements are shown in Fig. 4.4. Continual measurements only started in late July 2025. Note that since this ion gauge is located close to the 3D MOT region, its recorded pressure might not be representative of other regions of the chamber, i.e. the interferometry region. The 3D MOT lifetime and the lifetime of an atom launched in the atomic fountain are observed to be different, which could be explained by different vacuum pressures in each region.

In March 2025 we decided to reactivate the SAES CapaciTorr D2000 NEG pump in an attempt to lower the vacuum pressure back under 10^{-10} Torr, after determining that repeated tisub firing did not bring the vacuum down to a satisfactory level. NEG pumps consist of a large surface area getter which reacts with gas molecules such that they stick to the getter surface. Eventually the getter saturates, as its surface has become completely covered by the compounds formed by the gas molecules. When this happens the NEG pump can be activated by heating the getter to a temperature above 400° C for approximately an hour. During activation the compounds on the getter surface are pulled inside the getter

material, effectively resetting the getter to an unsaturated state [65]. We were concerned that activating the getter could produce temperature gradients across the rest of the chamber that could fracture chamber viewports or compromise the vacuum in some other way.

Before reactivating the NEG pump, the RGA was used to measure the gas composition in the vacuum chamber, which indicated the presence of water and hydrogen in the vacuum, as well as a few unidentifiable species. The NEG cycle was set to a temperature of 450° C, a rise time of 2 hours, and a hold time of 2.5 hours. The observed heating of the chamber was minimal. After this reactivation cycle there was a small improvement to the observed 3D MOT lifetime. The reactivation process does not appear to have any negative effect on the chamber, and can be easily performed in the future.

4.2 Laser systems

The laser frequencies required for operating the experiment are shown in Fig. 4.6. All frequencies operate somewhere close to the cesium D_2 line and target transitions between states in the $6\,^2S_{1/2}$ and $6\,^2P_{3/2}$ manifolds. In this section I describe the systems used for generating and stabilizing these frequencies. Note that the optical schematics shown in this section omit beam steering mirrors, fiber polarization alignment waveplates, collimation optics, and various other optical elements that are not critical for understanding system functionality. In Fig. 4.5 we show the typical optical setup used for each fibercoupling, which exists before almost every fibercoupling described here.

4.2.1 Repump laser system

The repump laser system is responsible for generating light for transitions from the $6\,^2S_{1/2}\,F = 3$ state to the upper $6\,^2P_{3/2}$ manifold. This light is used as a repumper in the MOT, a pump beam during RSC, and the $F = 3$ state blowaway. The repump laser frequency is stabilized to a cesium cell and serves as the optical reference for all other lasers in the experiment.

The optical layout of the repump laser system is shown in Fig. 4.7. The repump laser is an AOSense AOS-IF-ECDL-852 external cavity diode laser (ECDL), which produces 40 mW of light close to 852 nm. The light is passed through a Thorlabs IOT-5-850-VLP tandem near-IR optical isolator. The tandem isolator is used to minimize the effect of any retroreflections entering the ECDL cavity or the diode itself, which could destabilize the laser frequency. The light after the isolator is then fibercoupled. Light is generally fibercoupled soon after a laser or optical amplifier in the experiment so that if a laser or amplifier fails, the entire downstream optical chain does not need to be realigned.

The fibercoupled light is then outcoupled onto a polarizing beamsplitter cube (PBS). Around 10 mW is reflected by the cube and sent through an optical isolator before being used to seed a free-space tapered amplifier (TA). The TA is a 3 W Coherent DILAS TA-

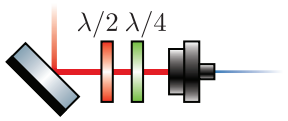


Figure 4.5: Typical optics used before every polarization-maintaining fiber in the experiment. Both the mirror and fiberport should be placed on a kinematic mount, which enables the beam location and angle relative to the fibercoupler to be adjusted properly. Ideally the fibercoupler also has an adjustable position aspheric lens that can be moved to maximize fibercoupling efficiency. If polarized light is being fibercoupled, a $\lambda/2$ and $\lambda/4$ waveplate should be placed before the fiber. These waveplates should be aligned to make the polarization linear on the fiber output and to minimize any polarization instability. The polarization drift can be recorded using a polarimeter such as the Thorlabs PAX1000IR1. If the fiber is undamaged, it should be possible to use the waveplates to keep the output polarization stable to within 2° in both azimuth and ellipticity on the Poincaré sphere. This can be tested by straining or heating the fiber (easily done by compressing the fiber in your hand), which changes the birefringence of the fiber and leads to a polarization shift if the input light is not perfectly coupled to a single fiber axis. It is also good practice to add a PBS cube to the fiber output to further ensure polarization stability.

0850-3000 in a C-mount package[1]. The C-mount is attached to a homebuilt mount which uses two Thorlabs aspheric lenses to couple into the TA waveguide and collimate the TA output vertically. A cylindrical lens after the TA (omitted from Fig. 4.7) collimates the beam horizontally. The TA output is then split by successive PBS cubes and is coupled into fibers going to the cooling laser system described in Section 4.2.2, the Ti:Sa laser phase-locking beatnote described in Section 4.2.5, and a utility beatnote, which can optionally be used to phase lock temporarily needed laser systems. A homebuilt current driver is used to deliver 2 A of current to the TA, and a Team Wavelength PTC5K-CH temperature controller is used to stabilize the TA temperature using a thermoelectric cooler.

The other PBS output is sent through a double-passed 141 MHz acousto-optic modulator (AOM), labeled AOM01 in the diagram. The returning light has a 90 deg polarization shift after double passing a $\lambda/4$ waveplate, and is reflected again by the PBS into a fiber. This fiber is sent to another repump laser subsystem, where it is split again by a PBS. The reflected light is sent through AOM03 and AOM04, which diffract light for the $F = 3$ blowaway and RSC pump, respectively. The zero-order of the double-passed AOM01 is coupled into a fiber

[1]Historically the experiment used several free-space TA systems. However, during my PhD Toptica Eagleyard released the MiniTA tapered amplifier, which provides a fiber connected to the TA waveguide and a built-in free-space collimation package. The MiniTA only provides 2 W compared to the 3 W system from Coherent DILAS. However, it takes far less time to set up (hours compared to the 1−2 weeks of the free-space TA), has a higher fiber coupling efficiency (60−70% compared to 40−50%), and appears to have more stable operation. The remaining TA systems in the experiment have been replaced with the MiniTA.

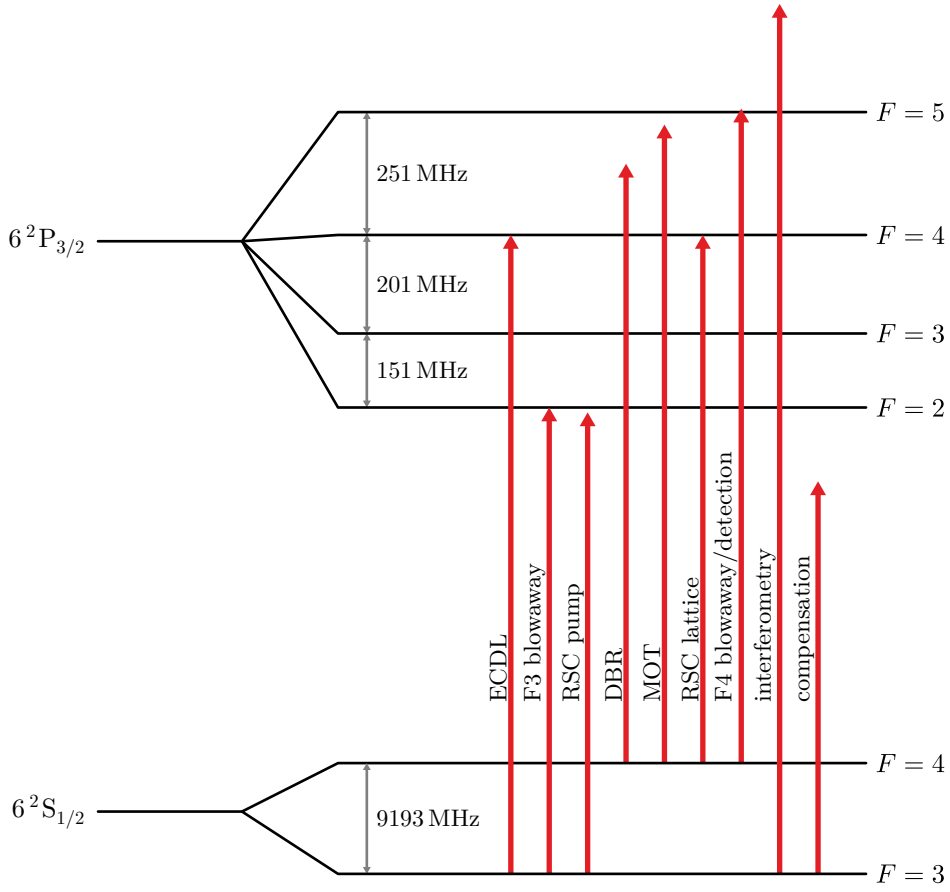


Figure 4.6: Cesium level structure with most laser frequencies used in the experiment shown.

that is optionally sent to the frequency comb for long-term recording.

The transmitted light from the fiber is used to lock the laser frequency using frequency modulation (FM) spectroscopy. In the FM spectroscopy setup used in the experiment, two counterpropagating beams derived from the same laser pass through a cesium cell. The first beam with frequency ω_1 passes through a free-space electro-optic modulator (EOM) driven by a tank circuit at $\omega_{\mathrm{EOM}} = 2\pi \times 7.186\,\mathrm{MHz}$. The second beam is diffracted by AOM02 such that $\omega_2 = \omega_1 + \omega_{\mathrm{AOM2}}$. The intensity of the first beam is recorded on the FM photodiode (PD) in Fig. 4.7. It is then high pass filtered, mixed with the EOM drive frequency ω_{EOM}, and finally low pass filtered. The resulting voltage is proportional to the difference in absorption between the two sidebands $\alpha(\omega_1 - \omega_{\mathrm{EOM}}) - \alpha(\omega_1 + \omega_{\mathrm{EOM}})$. When ω_1 is perfectly centered on resonance, these absorptions are equal and the resulting voltage is zero, providing an error signal that can be used to lock the ECDL. A theoretical treatment of this scheme is provided in [66], and a description of a similar locking configuration is given in [67].

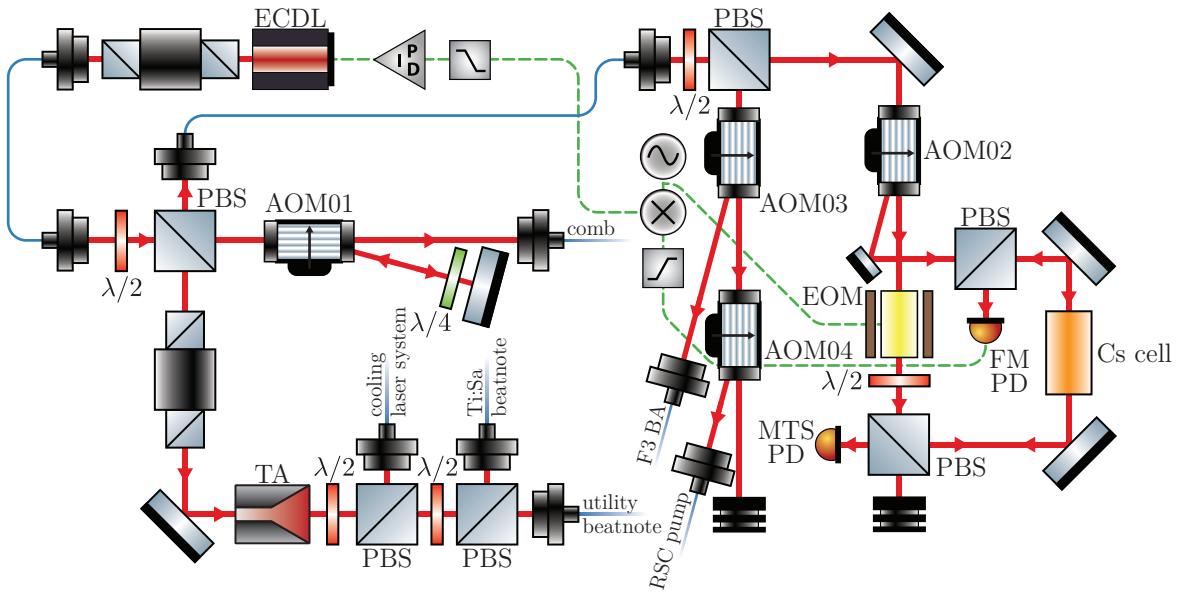


Figure 4.7: Optical schematic of the ECDL repump laser and cesium spectroscopy.

This signal is fed into the AOSense current driver and used as an error signal. This stabilizes the ECDL frequency at whatever frequency keeps the error signal at zero. AOM01 and AOM02 are driven at 141 MHz, and so AOM01 shifts the observed resonance by 141 MHz/2. With these frequency choices the error signal is zeroed $2\times141\,\text{MHz}+141\,\text{MHz}/2 = 352.5\,\text{MHz}$ below the ECDL frequency. There are several resonances that the error signal could lock to, and we choose $F = 3 \rightarrow F' = 2$, which makes the ECDL resonant with the $F = 3 \rightarrow F' = 4$ transition.

AOM03 is driven at 70.45 MHz, making the $F = 3$ blowaway light resonant with the $F = 3 \rightarrow F' = 2$ transition. AOM04 is driven at 65 MHz, making it approximately 5.5 MHz red detuned from the $F = 3 \rightarrow F' = 2$ transition. This is done in order to reduce the scattering rate during RSC. This is effectively the same as further attenuating the beam intensity.

The entire repump laser system is in an acrylic enclosure and only requires minor realignments every 6 months or so.

4.2.2 Cooling laser system

The cooling laser system is responsible for generating light from the $6\,^2\text{S}_{1/2}\,F = 4$ state to the upper $6\,^2\text{P}_{3/2}$ manifold. This light is used to generate the 2D and 3D MOT beams, the RSC lattice beams, the $F = 4$ blowaway, and the fluorescent detection light. The cooling laser system is phase-locked to the repump ECDL described in the previous section.

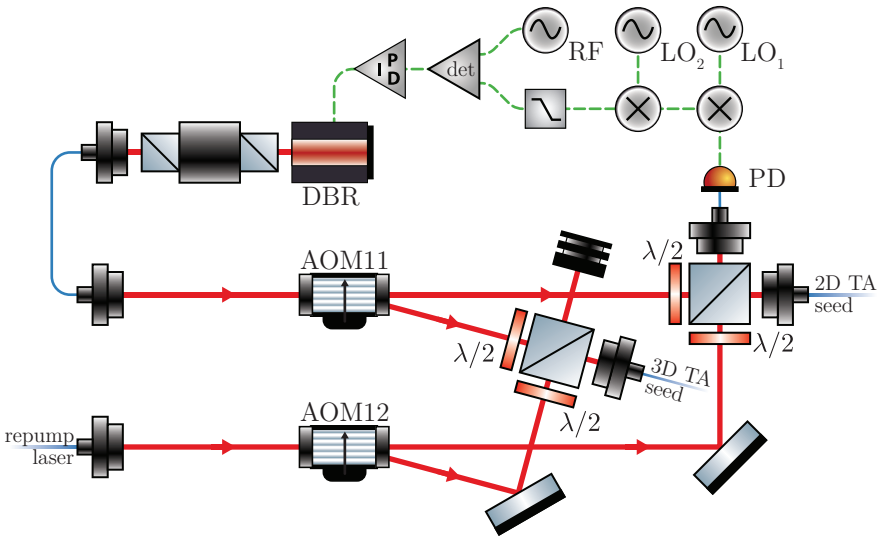


Figure 4.8: Optical schematic of the DBR cooling laser.

The optical layout of the cooling laser system is shown in Fig. 4.8. The cooling laser is a Photodigm 852.347DBRH-L-TO8 distributed Bragg reflector (DBR) laser which outputs 240 mW. Like the ECDL, the DBR is passed through a Thorlabs IOT-5-850-VLP tandem near-IR optical isolator and then fibercoupled. The fibercoupled light is then sent through AOM11, which is driven at 95.5 MHz.

Light from the repump laser system is also outcoupled on the cooling laser system, and is sent through AOM12, which is driven at 80 MHz. The diffracted cooling light from AOM11 and the diffracted repump light from AOM12 are then combined on a PBS cube and coupled into a fiber used to seed the 3D TA system (described in Section 4.2.3). AOM12 can be turned off to remove repump light from the 3D TA seed. Not pictured is an additional PBS cube before the 3D TA seed fiber used to project both cooling and repump light onto the same polarization axis. The waveplates before the first PBS and the RF drive power of AOM11 and AOM12 are tuned to couple ∼13.5 mW of cooling light and ∼1.5 mW of repump light into the fiber.

The zero-order light from AOM11 and AOM12 is combined on another PBS cube. Each output port of the PBS cube is coupled into a fiber. One fiber is used to seed the 2D TA system (described in Section 4.2.3), and the other is used to measure the beatnote between the repump ECDL and cooling DBR light. Not pictured are additional PBS cubes before the 2D TA seed fiber and the beatnote fiber used to project both cooling and repump light onto the same polarization axis. The waveplates before the PBS are tuned to couple ∼13.5 mW of cooling light and ∼1.5 mW of repump light into the 2D TA seed fiber.

The beatnote fiber is tuned to have a roughly equal ratio of cooling and repump powers. A fiberized 15 GHz Optilab SPD-15-A-DC photodiode is used to measure the beatnote, which

is stabilized using a phase lock to 8947 MHz. Care is taken to ensure the DBR is red detuned relative to the ECDL. The phase locking system also provides the ability to rapidly jump the laser frequency, which will be utilized when performing different state preparation stages. The phase lock is described in Section 4.2.6.

Since the ECDL light is resonant with the $F = 3 \rightarrow F' = 4$ transition and the cesium hyperfine splitting is 9193 MHz, the DBR light is 5 MHz red detuned from the $F = 4 \rightarrow F' = 5$ transition. Therefore the 2D TA seed light is composed of two frequencies: repump light resonant with the $F = 3 \rightarrow F' = 4$ transition and cooling light 5 MHz red detuned from the $F = 4 \rightarrow F' = 5$ transition. The 3D TA seed light contains repump light 80 MHz red detuned from the $F = 3 \rightarrow F' = 4$ transition and cooling light 100.5 MHz red detuned from the $F = 4 \rightarrow F' = 5$ transition. The switchyard system described in Section 4.2.4 will reduce these detunings when generating light for the 3D MOT, RSC, $F = 4$ blowaway, and detection.

4.2.3 Tapered amplifier systems

Tapered amplifiers (TAs) consist of a semiconductor gain medium that is single passed to amplify incident seed light. To prevent the TA from lasing on its own it is important to ensure that no cavity is accidentally created around the gain medium, as this can cause the TA to lase at an unwanted frequency and spatial mode. For this reason optical isolators should be placed on at least one side of the TA.

Historically homebuilt free-space TAs were used in the experiment, which had various drawbacks. Now most TAs are Toptica Eagleyard EYP-TPA-850-02000 MiniTAs, with the exception being the TA described in Section 4.2.1. These have fiberized inputs and built-in free-space output collimation packages. They require seed powers of 10−15 mW and output 2 W of optical power. The output beam profile is near circular and can often be directly fibercoupled with efficiencies of 60−80% without any beamshaping. The MiniTA current is supplied by an Arroyo Instruments 7162-04-04 combo laser driver and temperature controller.

Three TA systems are currently used in the experiment, and are shown in Fig. 4.9. The layouts of the 3D and 2D TA systems are identical. The 3D and 2D TA schematics contain a MiniTA used to amplify a seed. The amplified beam is then sent through an optical isolator, then through a PBS plate, and then is fibercoupled. The fibercoupled light is then split into the relevant fibers (the switchyard for the 3D TA system, and the 2D MOT telescopes for the 2D TA system). The initial fiber coupling makes TA replacement more straightforward if the MiniTA breaks, and makes any post-TA fiber-coupling power drift in the 3D and 2D systems common to all beams produced by that TA. The compensation TA system is coupled into a single fiber after AOM41. AOM41 implements the pulse shaping system described in Section 4.2.5.

The PBS plates used here are rated to higher power than the PBS cubes used elsewhere, and have a higher extinction ratio. The PBS plates are custom ordered Laseroptik L-19574 PBS plates, which consist of a circular optic with a 0.5° wedge and a coating optimized for operation at 55.4° at 852 nm. The plate angle is aligned to the angle that minimizes the

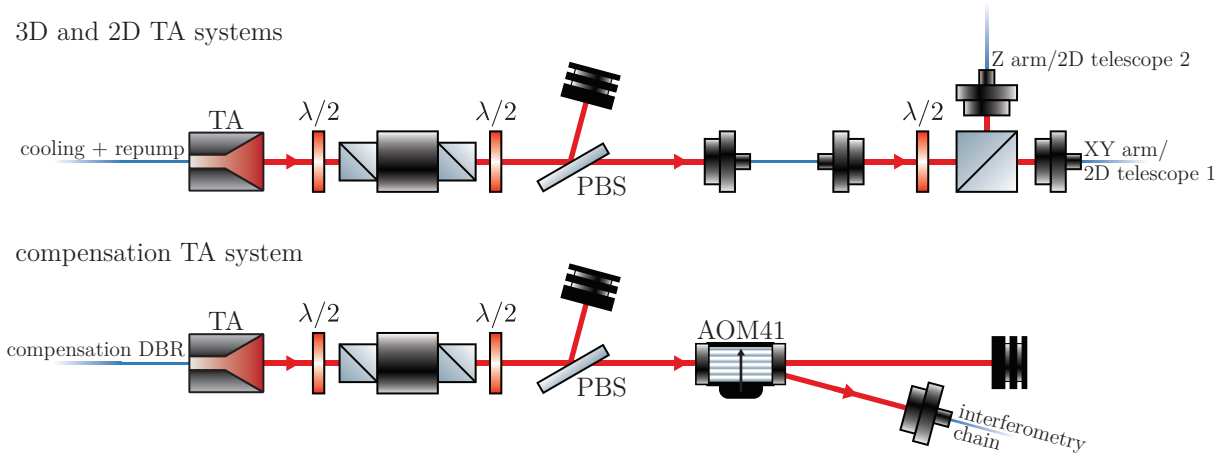


Figure 4.9: Optical schematic of the tapered amplifier systems.

transmission as the input polarization is scanned. When the MiniTA is initially fibercoupled, it is set to output full power and the PBS plate is used to attenuate the optical power directed at the fiber to a safe level for fibercoupling. The PBS is used to attenuate the optical power instead of reducing the MiniTA power because reducing the MiniTA power can change the beam profile, which would mean the fibercoupling would not be optimal at high power. Very little power drift is observed over time. The MiniTA systems routinely operate for weeks without requiring tune-ups of the seed power or adjustments of the post-TA fibercoupling.

The compensation TA is seeded with a single frequency. However, the 2D and 3D TAs are seeded with both cooling and repump light. It has been shown that TAs can act as a frequency mixer, producing sidebands on the input light (see [68] Section 4.4.5). It is not known if this is a problem in this experiment. Importantly, neither the 2D nor 3D TAs are used during the interferometry sequence, so this has no effect on the frequency composition of the interferometry beam.

4.2.4 Switchyard optical system

The switchyard optical system shown in Fig. 4.10 is responsible for routing light generated by the 3D TA to the 3D MOT system, the RSC system, and the blowaway/detection fibers. The 3D TA routes 500 mW to the XY switchyard (the upper optical chain in the figure) and 300 mW to the Z switchyard (the lower optical chain in the figure).

While the 3D MOT is on, AOM21, AOM22, and AOM26 blueshift light from the 3D TA by 80 MHz. Recall that the 3D TA system is seeded with light 80 MHz red detuned from the $F = 3 \rightarrow F' = 4$ transition and light 100.5 MHz red detuned from the $F = 4 \rightarrow F' = 5$ transition. Therefore the MOT X1, X2, Y1, Y2, Z1, and Z2 fibers contain repump light

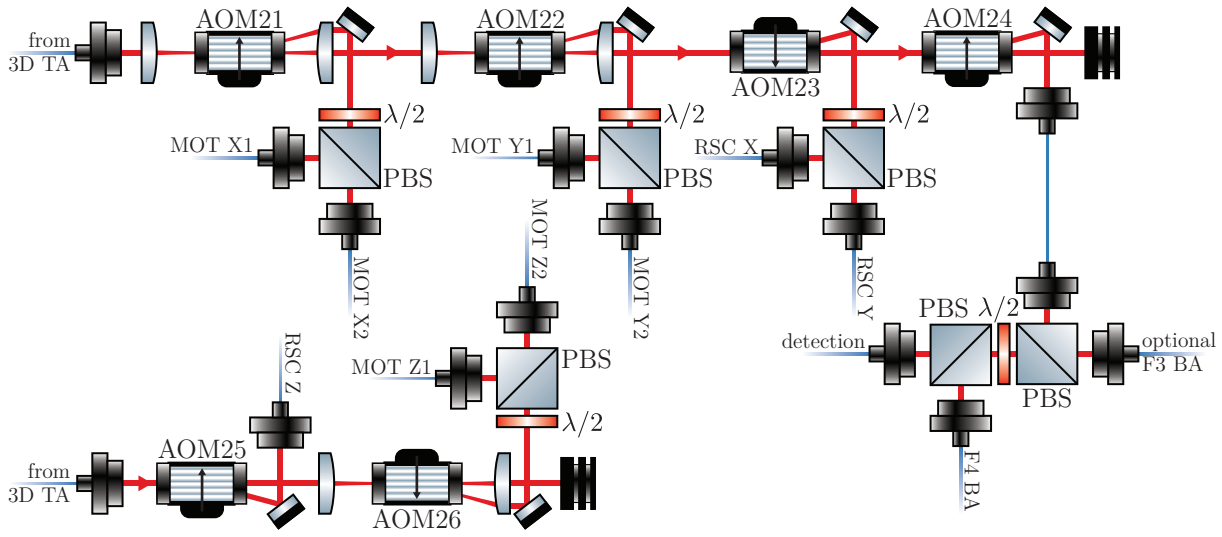


Figure 4.10: Optical schematic of the switchyard used to generate light for the 3D MOT, RSC, blowaway, and detection.

resonant with the $F = 3 \rightarrow F' = 4$ transition and cooling light 20.5 MHz red detuned from the $F = 4 \rightarrow F' = 5$ transition. The detuning of 20.5 MHz during the 3D MOT sequence was chosen as it was found to produce the largest interferometry signal at detection. During the molasses launch the drive frequency of AOM21 and AOM22 is jumped to create a net upward force. To keep the fiber coupling efficiency stable during this jump, AOM21 and AOM22 are placed inside of telescopes. AOM21, AOM22, and AOM26 are all aligned at the same duty cycle used in the experiment cycle, as the thermal load on the AOMs can affect their fibercoupling efficiency.

During the RSC stage AOM21, AOM22, and AOM26 are turned off, and AOM23 and AOM25 are turned on. AOM23 redshifts its diffracted light by 120.4 MHz, making RSC X and RSC Y 30.1 MHz blue detuned from the $F = 4 \rightarrow F' = 4$ transition. The RSC X and RSC Y are brought into resonance with the $F = 4 \rightarrow F' = 4$ transition by shifting the DBR laser frequency as described in Section 4.2.6. The RSC Z beam enters the chamber through the interferometry chain described in Section 4.2.5, which produces a net shift of $-f_{\mathrm{AOM25}} + f_{\mathrm{AOM35}}$. AOM35 is always driven at 80 MHz. The RSC Z beam frequency, after accounting for the Doppler shift, must be the same as RSC X and Y in the frame of the atoms. For the typical 6.2 m/s launch velocity we set AOM25 to 193.3 MHz, which produces light 37.2 MHz blue detuned from the $F = 4 \rightarrow F' = 4$ transition.

At the end of the XY switchyard is AOM24, which is driven by 98.4 MHz to produce repump light 18.4 MHz blue detuned from the $F = 3 \rightarrow F' = 4$ transition and cooling light 2.56 MHz red detuned from the $F = 4 \rightarrow F' = 5$ cycling transition used for detection. This light is coupled into a fiber where it is then split and sent to a detection and $F = 4$

blowaway beam. When the blowaway beam is used AOM12 is switched off to remove any light close to resonance with the $F = 3$ state. When detection is used the repump is left on to enable repumping of any atoms that enter the $F = 3$ state. In the past the $F = 4$ and $F = 3$ blowaways entered the experiment through the same beam, and so light from the $F = 3$ blowaway beam generated by the repump laser system was combined with the $F = 4$ blowaway light here. However, now these two blowaways are separate. Note that the cooling light is detuned by around the natural linewidth, and has a similar scattering rate to what it would have on resonance. While the repump light is detuned, changing the repump light frequency to be closer to resonance is not found to have a large impact on detected fluorescence.

4.2.5 Interferometry laser system

Laser and pulse shaping

The previously discussed laser systems are all used for state preparation, but not for generating atom interferometry pulses. The interferometry laser system prepares all light directed vertically along the vacuum chamber. It is used for generating Bragg beamsplitter pulses, Bloch oscillation pulses, Raman velocity selective (VS) pulses, and the RSC Z beam. The interferometry beam light is directed into the viewport at the bottom of the vacuum chamber described in Section 4.1.2.

All matterwave optics are generated by pulsing light from a Sirah Matisse CS titanium-sapphire CW ring laser (Ti:Sa). The advantage of using a single laser to generate the light required for Bragg, Bloch, and Raman pulses is that it reduces the phase noise between the different frequencies participating in these processes. The Ti:Sa crystal gain medium is pumped using a Spectra Physics Millennia EV 25 W 532 nm laser. Automatic control of internal cavity elements of the Ti:Sa laser is provided by the Matisse controller, which enables automatic relocking as described in Section 4.2.6. We find it takes around 12 hours for the Ti:Sa cavity elements to stabilize after first turning on, as shown in Fig. 4.12. For this reason the Ti:Sa is normally left on continuously. The Ti:Sa is capable of providing 6 W of power at 852 nm.

Directly after the Ti:Sa a plate PBS picks off a small amount of power which is coupled into a fiber. This light is further split into two paths that are sent to a wavemeter and used to create a beatnote against the ECDL laser. This beatnote is used to lock the bare Ti:Sa optical frequency 6825 MHz blue of the ECDL laser, before the Ti:Sa light passes through any of the interferometry-chain AOMs. The locking scheme is described in Section 4.2.6. This Ti:Sa frequency was chosen as it was originally believed to put the light entering the chamber at the wavelength where $\Delta_{\mathrm{AC}} = 0$ (we refer to this wavelength condition as the VS magic wavelength). Recently Yuno Iwasaki showed this was incorrect. The VS magic wavelength mechanism is further discussed in Section 7.1.1. In the future the interferometry laser system will be upgraded so that the Ti:Sa is used for Bragg and Bloch pulses, and an

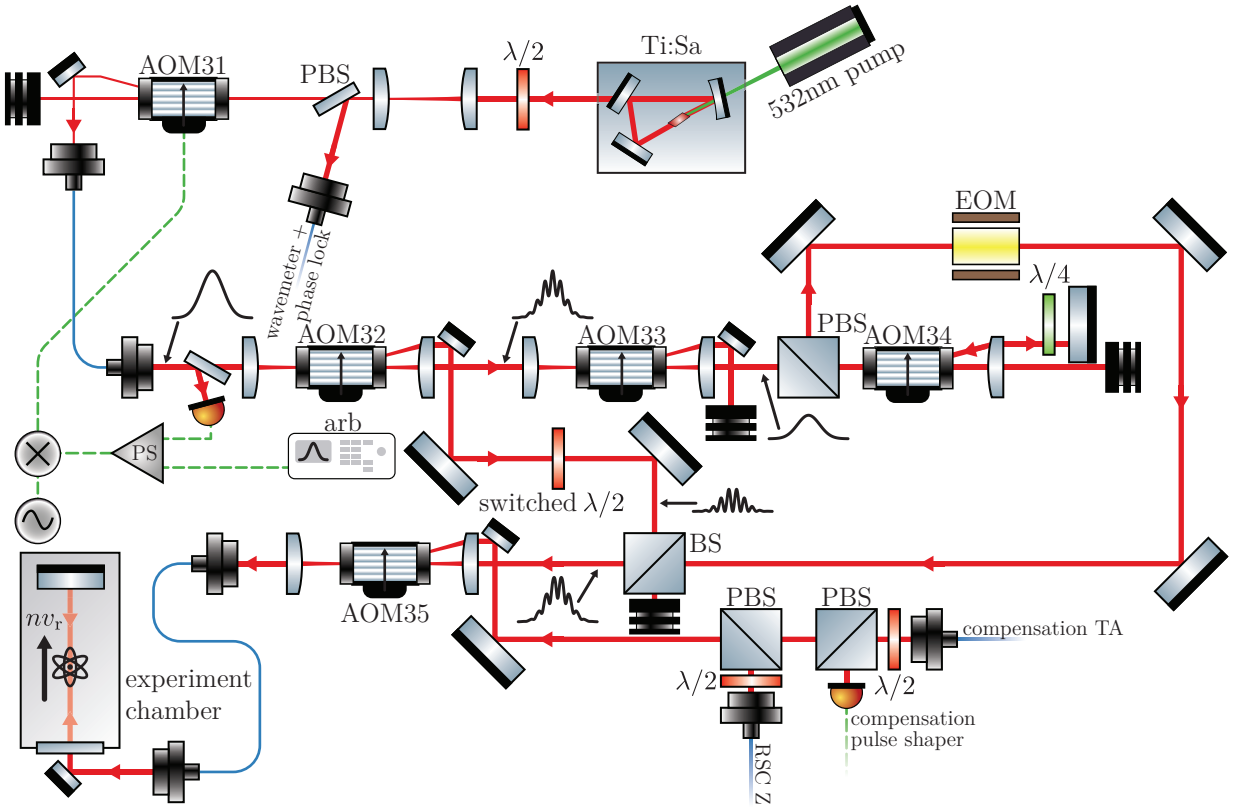


Figure 4.11: Optical schematic of the interferometry laser and the optical chain used to generate light for Raman transitions, Bragg diffraction, and Bloch oscillations. The diagram also displays how RSC Z and Bloch compensation light are injected into the interferometry fiber.

independent laser (operating at the VS magic wavelength) is used for Raman VS pulses. This will allow independent selection of the single photon detuning used for Bragg and Bloch.

After the plate PBS the Ti:Sa light passes through AOM31, whose diffracted order is fibercoupled and sent to the frequency preparation stage of the interferometry chain. The RF drive of AOM31 is pulsed throughout the duration of an experiment shot to produce VS, Bragg, and Bloch pulses. A telescope before AOM31 is used to reduce the beam size by a factor of 4, as otherwise it was found that AOM31 only diffracted the center of the Ti:Sa beam. When AOM31 and the fiberport are optimally aligned, 82% of the power output by the Ti:Sa can be coupled into the fiber after AOM31. AOM31 is driven at 80 MHz.

After being outcoupled onto the frequency preparation stage of the chain, a beam sampler reflects a small fraction of the light onto a photodiode, which feeds back onto the RF drive power of AOM31 to produce the desired pulse shape. The homebuilt circuit shown in Fig. 4.13 uses a summing amplifier to take the difference between the photodiode voltage

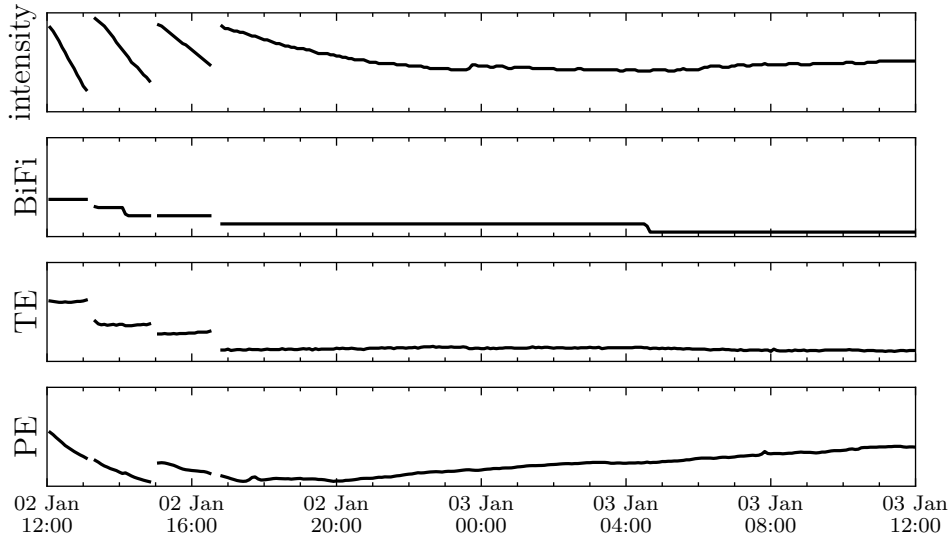


Figure 4.12: Ti:Sa laser intensity and cavity element positions over time after the Ti:Sa was turned on on 2 January 2026. All element positions are set by active feedback loops that stabilize the laser frequency. The intensity and cavity elements take around 12 h to stabilize. BiFi refers to the birefringent filter, TE refers to the thin etalon, and PE refers to the piezo etalon in the cavity.

and the reference voltage. The summing amplifier has a capacitive feedback element which integrates the voltage difference to provide an I gain. The size of this I gain is tuned using a potentiometer after the capacitor. The output of the potentiometer is fed into an inverting amplifier which sources a current signal. A diode at the output of the inverting amplifier limits the current output to a single direction, and a resistor limits the maximum current that can flow. The circuit output is sent to the I port of a MiniCircuits ZFM-3+ mixer, which acts as a current variable attenuator for near-DC I port inputs (with zero current corresponding to full attenuation). The mixer is used to attenuate the RF drive of AOM31, completing the feedback loop. The diode prevents the current direction from flipping. If the current direction was able to flip, the negative feedback loop could turn into an unstable positive feedback loop. An analog switch is used to clear the integrator between interferometry pulses.

Current will only flow into the mixer once the forward voltage of the diode is exceeded. This has the effect of adding a slight delay to the RF drive response at the beginning of the pulse, since the accumulated error must exceed the accumulated forward voltage. In practice this is only noticeable for pulses with a sharp turn on. All interferometry pulses turn on smoothly. Recently, this circuit was upgraded by Yuno Iwasaki to a version that uses an analog absolute value circuit instead of a diode to remove the forward voltage requirement,

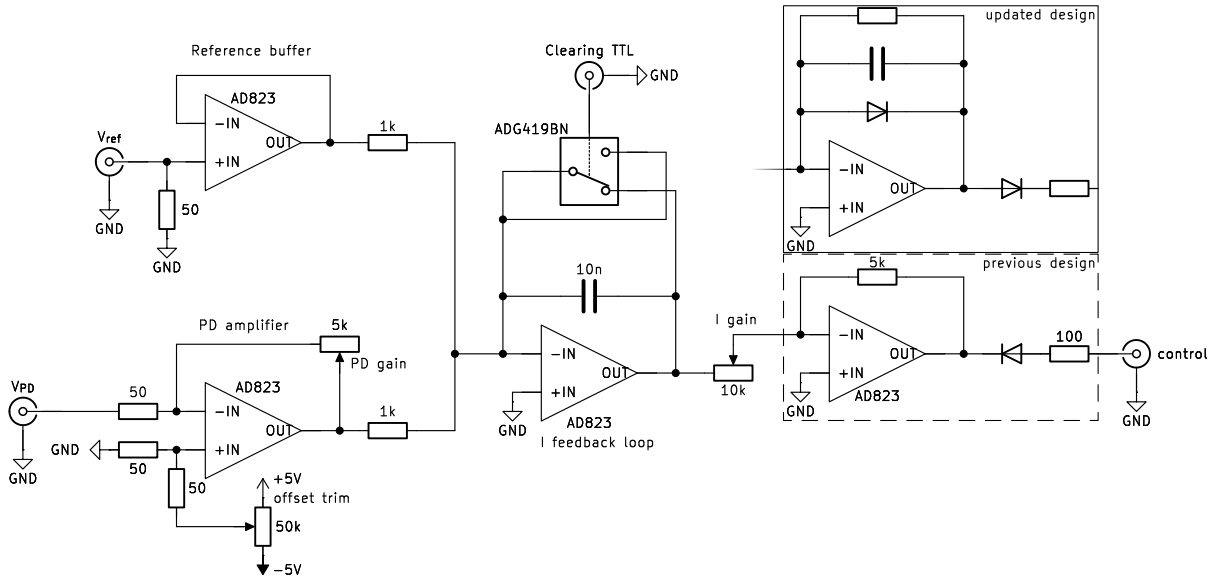


Figure 4.13: The circuit used to stabilize laser intensity to the required pulse shape. The dashed box is the original final op-amp stage. Yuno Iwasaki developed an upgraded output stage shown in the solid boxed region.

which replaces the region of the circuit in the dashed box with the region of the circuit in the solid box. The bandwidth of the circuit-AOM system is limited by the speed of sound in AOM31. For this reason AOM31 is mounted on a translation stage so that the beam enters the AOM as close to the acoustic transducer as possible.

The reference voltage is provided by several arbitrary waveform generators (AWGs) whose outputs are summed together using a basic operational amplifier (op-amp) circuit. The AWGs are programmed with the desired pulse shapes and then triggered whenever a pulse is needed. The Bragg pulse AWGs typically use Gaussian waveforms, while VS pulse AWG typically uses a trapezoidal waveform.

Bragg, Bloch, and Raman pulse preparation

We have now reviewed how Ti:Sa light is generated, phase-locked, and intensity shaped in time. Now we will review what is needed to perform Bragg diffraction, Bloch oscillations, and VS Raman transitions. In Chapter 2 we showed that Bragg diffraction and Bloch oscillations require a far-detuned bichromatic field with a controllable time profile. For VS Raman transitions we needed two far-detuned fields with a frequency difference equal to some ground state energy splitting. We have already described a system that provides most of these ingredients. After the nominal 80 MHz blueshift from AOM31 and 180 MHz blueshift from the downstream frequency-preparation stage, the light entering the chamber has a nominal

single-photon detuning $\Delta = 2\pi \times 7085\,\mathrm{MHz}$ blue of the $6\,^2\mathrm{S}_{1/2}\,F = 3 \rightarrow 6\,^2\mathrm{P}_{3/2}\,F = 4$ transition to which the ECDL is locked. The pulse shaping circuit provides the ability to control the intensity profile in time. Therefore, to implement Bragg and Bloch we must create a frequency splitting on the order of the recoil frequency $\omega_\mathrm{r} = 2\pi \times 2.066\,\mathrm{kHz}$. To implement VS we need a splitting on the order of the cesium hyperfine frequency 9193 MHz.

To produce the two field frequencies required for Bragg diffraction and Bloch oscillations, the beam after the pulse shaper is sent through AOM32, which splits the beam onto two arms. The zero-order arm is referred to as the gravity ramp arm. The first-order arm is referred to as the diffraction arm. The beam on the diffraction arm passes through a Thorlabs LCC1221-B variable waveplate and then is combined with the gravity ramp arm on a non-polarizing beamsplitter cube. It then passes through AOM35 but is not diffracted. Finally it is coupled into a fiber that is sent to the bottom viewport of the vacuum chamber. The gravity ramp arm passes through AOM33 undiffracted, is diffracted through the double-passed AOM34, and then is coupled into the same fiber as the diffraction arm. The variable waveplate is set to make both the diffraction arm and gravity ramp arm beams have the same linear polarization. A cat eye configuration is used to make the double pass beam pointing relatively insensitive to the AOM34 drive frequency. The frequency difference between the diffraction and gravity ramp arms is on the order of kHz to MHz, and is used to create the two field frequencies required for Bragg diffraction. The two-photon detuning is given by $\delta = \omega_d - 2\omega_g(t)$, where ω_d is the frequency of AOM32 and $\omega_g(t)$ is the frequency of AOM34. In the experiment ω_d is tuned to make δ match the resonance condition from Eq. (2.18). AOM32 is built inside of a telescope to increase the robustness of downstream fibercoupling to changes in ω_d, and is nominally driven at $\omega_d = 2\pi \times 180\,\mathrm{MHz}$.

For multifrequency Bragg diffraction and Bloch oscillations, as described in Section 2.2.2, AOM32 is driven with two RF tones at $\omega_d \pm \omega_\mathrm{m}$. The measured intensity of the diffracted beam therefore has an intensity beatnote with frequency $2\omega_\mathrm{m}$, as shown in Fig. 4.11. Because the total intensity passing through AOM32 must be conserved, the zero-order also picks up an intensity beatnote. In multifrequency mode, AOM33 is driven using the same RF source as AOM32, but with the ω_m component 90° out of phase. This means that AOM33 diffracts away the beatnote on the zero-order, leaving a beatnote free profile afterwards. The circuit used to generate the 90° phase shift is a newer implementation of the one described in [21] Fig. (4.13). During Bloch oscillations, ω_m is swept linearly in time to achieve the time-dependent two-photon detuning given in Eq. (2.37).

Bloch oscillations are long enough that the AC Stark shift potential they induce shifts the wavepacket position. For separated wavepackets experiencing slightly different potentials this can prevent wavefunction overlap at the end of the interferometer, reducing interferometer contrast. To prevent this, AOM35 is turned on during Bloch oscillations and light from the compensation TA is coupled into the interferometry beam via AOM35. The intensity and detuning of this light are chosen to produce the reverse AC Stark shift potential, stopping the wavepacket position shift. Note that AOM35 is also used to inject RSC Z light into the interferometry beam fiber, and to reduce the intensity of the pulses used for VS.

To remove the Doppler shift that accumulates as the atom cloud accelerates in the foun-

tain, $\omega_g(t)$ is swept at the rate at which a Doppler shift is accumulated. This allows ω_d to be set without consideration of the velocity of the atom cloud. Atoms accumulate a Doppler shift at rate kg. Accounting for the Doppler shift accumulated by both the upgoing and downgoing beams, as well as the double pass configuration, we ramp $\omega_g(t)$ at $kg \approx 2\pi \times 11.5\,\mathrm{MHz/s}$ to keep the Bragg lattice velocity constant in the frame of the atoms. Note that $\omega_g(t)$ is ramped down in frequency, meaning that the gravity ramp light is continually being redshifted. This means that atoms interact with the lattice formed by the upwardly propagating gravity ramp arm light and the downwardly propagating diffraction arm light.

To perform Raman transitions in the experiment the EOM on the gravity ramp arm is driven with frequency ω_{EOM} and the variable waveplate is set to produce a 90° phase difference between the diffraction arm and gravity ramp arm beams in the fiber. This is required by the Raman transition selection rules, as described in [22] Section 3.4. To ensure only one EOM sideband can drive the Raman process, ω_{EOM} is set to the cesium hyperfine transition frequency minus 1 MHz and AOM32 is shifted by plus 1 MHz to make the transition resonant.

It should now be clear how the interferometry chain provides all of the ingredients necessary for performing Bragg, Bloch, and Raman in the experiment. Finally, note that before the interferometry beam enters the vacuum chamber, it reflects off of a backpolished mirror. A photodiode is placed behind the backpolished mirror in order to record the intensity of each interferometry pulse entering the chamber. The intensity is recorded and saved for each experiment shot.

4.2.6 Phase-locking and automatic relocking

A phase lock is employed to lock the phase of both the DBR and Ti:Sa lasers. The same error signal and control loop electronics are used for both lasers. The control loop is partially implemented in software, enabling different relocking procedures for each laser type. Each laser provides a different set of feedback elements that can be used to tune the laser frequency, with the DBR temperature and diode current determining its frequency, and the Ti:Sa cavity element positions determining its frequency. It is straightforward to provide fast feedback on the DBR laser current, making the DBR laser frequency agile, which is necessary during the state preparation procedure. The Ti:Sa frequency remains fixed throughout the experiment.

The phase lock is performed by generating a beatnote between the ECDL and DBR laser (as described in Section 4.2.2) or the ECDL and Ti:Sa laser (as described in Section 4.2.5). Recall that the ECDL frequency is set to the $F = 3 \rightarrow F' = 4$ transition. Both phase locks record the beatnote signal using a fiberized 15 GHz bandwidth Optilab SPD-15-A-DC photodiode. By stabilizing the frequency and phase of the beatnote sinusoid to a reference sinusoid, a stable phase relationship between the ECDL and phase-locked laser is achieved.

To generate an error signal from the beatnote sinusoid and reference sinusoid the home-built phase-frequency detector shown in Fig. 4.14 is used. The phase-frequency detector outputs a voltage proportional to the phase difference between its two inputs. The two

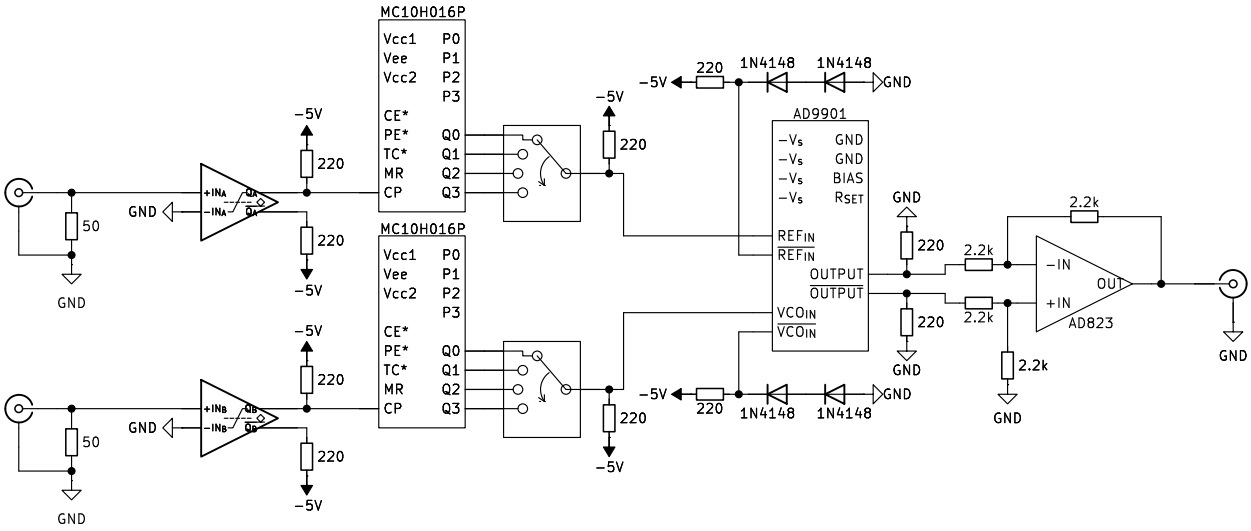


Figure 4.14: The circuit used to produce an error signal from the beatnote through use of the AD9901 phase-frequency detector.

input ports of the phase-frequency detector are terminated with 50 Ω and are each fed into a MAX9693 ECL comparator, which digitizes the signal. The output of each comparator is fed into a MC10H016P binary counter, which acts as a frequency divider. By selecting the desired output of the MC10H016P, the frequency can be divided by 2, 4, 8, or 16. The DBR uses a total division factor of 4 on each input signal, and the Ti:Sa uses a total division factor of 8 on each signal. The divided signals are then sent to an AD9901 phase-frequency detector, which outputs a voltage proportional to the phase difference between its two inputs. If the two inputs have different frequencies the output will rail either high or low, depending on which input frequency is higher. To convert the differential ECL signal into a single ended signal an AD823 op-amp configured as a differential amplifier is used. Two diodes and a 220 Ω resistor are used to fix the AD9901 ECL port voltages at approximately −2 V.

The ECL components used in the phase-frequency detector have a bandwidth of 200 MHz, and so the order-GHz beatnote detected by the fiberized photodiode cannot be sent directly into the circuit. Instead the signals are demodulated. The Ti:Sa beatnote is demodulated using a Vaunix Lab Brick typically set to 6675 MHz. The beatnote is then sent into the phase-frequency detector circuit using a 150 MHz signal as a reference. The phase-lock for the Ti:Sa therefore seeks to stabilize the beatnote frequency 6825 MHz blue of $F = 3 \rightarrow F' = 4$.

The DBR beatnote is doubly demodulated, once with microwave frequency of 8597 MHz, and then with RF frequency 200 MHz. The microwave frequency is provided by a Windfreak SynthHD clocked to a GPS stabilized 10 MHz source, and the 200 MHz is provided by an ARTIQ DDS (see Section 4.3). The double demodulation is employed because the ARTIQ

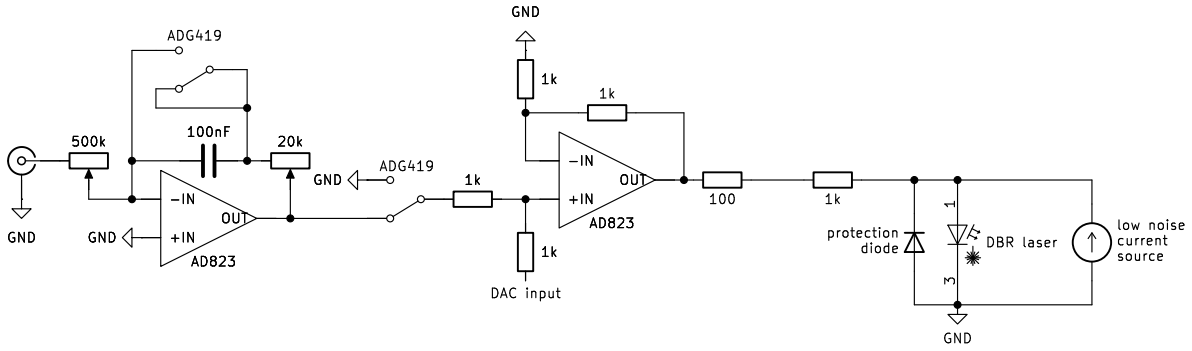


Figure 4.15: The circuit used to lock the DBR laser using the error signal produced by the circuit in Fig. 4.14.

DDS provides the ability to switch the RF frequency with precise timing, which can be used to rapidly shift the DBR laser frequency. The phase-frequency detector is fed the doubly demodulated signal and a 150 MHz reference signal. The phase-lock for the DBR therefore seeks to stabilize the beatnote frequency to 8947 MHz red of $F = 3 \rightarrow F' = 4$. The double demodulation RF chain is displayed in Fig. 4.8, where LO_1 is the microwave source, LO_2 is the RF source, and RF is the 150 MHz reference frequency.

The error signals generated by the DBR and Ti:Sa phase-frequency detectors are fed into the homebuilt PI loop shown in Fig. 4.15. The PI controller stage consists of a single inverting AD823 op-amp with a resistor and capacitor in series along the returning path. The 500 kΩ and 20 kΩ potentiometers control the PI and P gains, respectively. A digitally controlled switch parallel to the capacitor can be used to disable the I gain, leaving only the P gain. Following the PI controller stage is a second digitally controlled switch used to enable/disable the PI controller from entering the subsequent summing amplifier. The summing amplifier sums the control signal with the signal from a DAC. The output of the summing amplifier is used to modulate the current of the DBR laser by injecting the PI control signal at the DBR laser diode anode through a 1.1 kΩ resistor. The current driver for the DBR is a homebuilt low-noise current source designed by Brian Estey based on a BJT current mirror. The current noise on the current driver must not broaden the laser linewidth beyond the bandwidth of the feedback loop. In the case of the Ti:Sa, the output of the PI loop is fed to the intracavity fast piezo driver.

In order to produce a phase lock the bandwidth of the feedback loop must be higher than the natural linewidth of both lasers being phase locked together. Otherwise one of the two laser phases will drift away faster than the other can follow or than the feedback loop can counteract. In general, to narrow the linewidth of some laser, you must be able to perform feedback on it at a bandwidth higher than its natural linewidth.

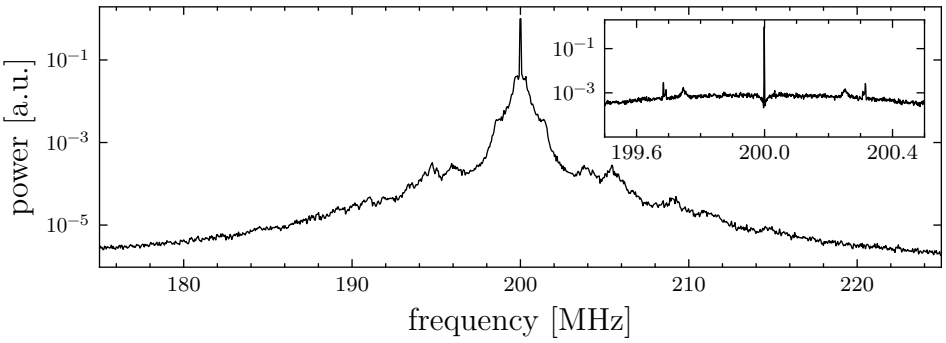


Figure 4.16: Measured spectrum of the beatnote between the ECDL and DBR lasers when locked. Inset displays the same signal measured with a lower resonant bandwidth.

DBR phase-lock performance

The characterization of the DBR phase-lock resulted in the paper [25]. Three approaches were used to characterize the phase-lock: a measurement of the ECDL-DBR beatnote spectrum, a measurement of the phase slippage of the phase-lock, and a measurement of the DBR linewidth with and without the phase lock. To measure the beatnote on a spectrum analyzer, the beatnote was demodulated to an RF signal at 200 MHz. The recorded spectrum is pictured in Fig. 4.16. Both the main figure and the inset show a peak at exactly 200 MHz. The width of the peak is limited by the 300 Hz resolution bandwidth of the spectrum analyzer, which is consistent with the delta function that would be expected from a phase-lock. The peak sits on top of a broad pedestal, indicating that there is some residual phase noise between the DBR and ECDL laser frequencies. This residual phase noise might affect the capability of the ECDL-DBR laser pair to drive coherent processes, as observed in [69, 70].

The phase slippage of the phase-lock is measured by recording the error signal produced by the phase-frequency detector on an oscilloscope and using the experimentally measured conversion factor between the error signal voltage and phase. The RMS phase difference is 1.4 radians. Note that this is measured after the frequency divider in Fig. 4.14. A histogram of the measured phase difference over an 8 ms period recorded at a sample rate of 0.5 GS/s is shown in Fig. 4.17.

The DBR linewidth is measured using the Sirah EagleEye optical spectrum analyzer, and found to be 93(21) kHz when phase locked. The natural linewidth was found to be 404(46) kHz, and the ECDL linewidth was measured to be 54(12) kHz. Therefore the phase-lock narrows the DBR linewidth. This can only occur if the DBR feedback loop bandwidth is higher than the DBR natural linewidth.

The response of the DBR phase lock to discrete jumps in the demodulation frequency

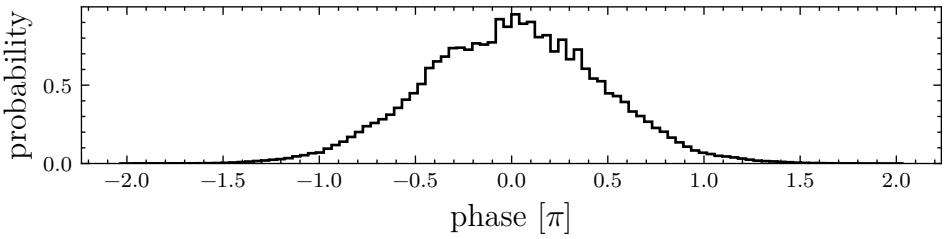


Figure 4.17: Measured distribution of the phase between the ECDL and DBR frequencies when locked. The phase difference is measured after the frequency divider. Distribution is recorded over an 8 ms period at a 0.5 GS/s sample rate. The distribution is normalized and reported as a probability distribution.

was also investigated, as this is relevant to the frequency jumps applied during the PGC and RSC stages of the experiment. The beatnote between the DBR laser and ECDL was recorded as the ARTIQ-controlled RF demodulation source was jumped by Δf. Fig. 4.18 shows short-time Fourier transforms generated using the recorded beatnote, along with a dashed blue line that shows the ideal instantaneous beatnote response. In Fig. 4.18a, Δf was set to 20 MHz at $t = 0$ µs. At the selected integral gain it took approximately 30 µs to reacquire the phase-lock. Therefore the PI loop swept the DBR laser frequency at around 0.66 MHz/µs. In Fig. 4.18c the integral gain was decreased and Δf was set to 50 MHz at $t = 0$ µs. In this condition it took approximately 110 µs to reacquire the phase-lock, indicating a reduced sweep rate of 0.45 MHz/µs.

In Fig. 4.18b, Δf was jumped by 20 MHz every 30 µs for 10 cycles starting at $t = 0$ µs, producing a linear optical frequency sweep from 0 to 200 MHz in 300 µs. In Fig. 4.18d Δf was jumped by 20 MHz every 30 µs for 5 cycles (again starting at $t = 0$ µs), followed by a delay of 60 µs. Then, Δf was jumped by -20 MHz every 30 µs for 5 cycles. This sequence produced a triangle-wave like optical frequency sweep. Using a direct digital synthesizer (DDS) capable of frequency sweeps instead of discrete frequency jumps would allow for smoother transitions in frequency.

Relocking procedures

Automatic relocking is made possible using the DAC and electronically controlled analog switch in Fig. 4.15. The figure omits ADCs located at the PI loop error signal input and control signal output that are also used for automatic relocking. The DAC, ADC, and analog switch are all controlled by a Raspberry Pi over an I^2C connection. This allows the relocking logic to be easily implemented in python, while the feedback loop is kept all-analog (and therefore fast). The Raspberry Pi is also used to log the recorded error and control signal values to a centralized server, enabling the state of each lock to be easily monitored by

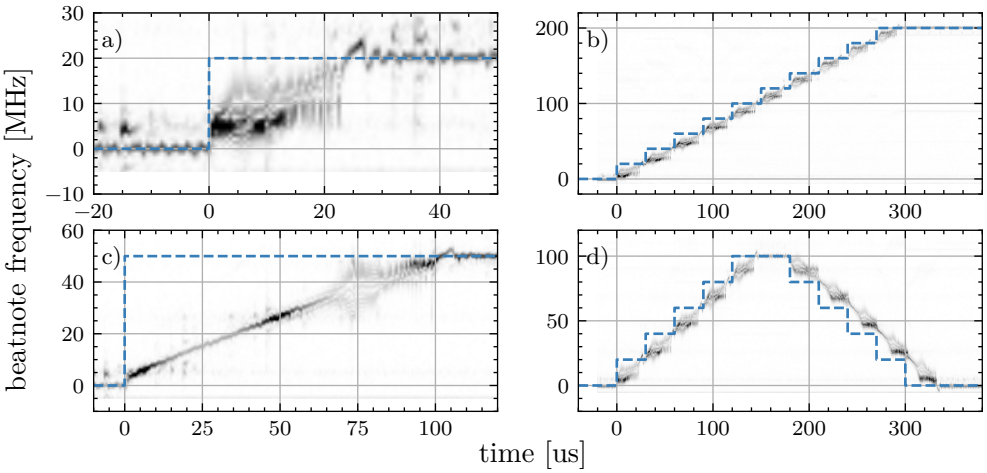


Figure 4.18: Short-time Fourier transforms of the beatnote measured between the ECDL and DBR while the lock is activated and the reference beatnote frequency is jumped. The reference beatnote frequency is given by the dashed blue line.

the experimenter. The relocking code used for both the DBR and Ti:Sa was developed by Andrew Christensen.

The DBR is automatically locked using the following procedure: first the PI loop is switched off using the analog switch before the summing amplifier. The DAC then performs a voltage scan until the error signal ADC detects that a zero-crossing has been passed. The DAC then moves the summing amplifier to this zero crossing and enables the PI loop using the analog switch. The PI loop then zeros the error signal, locking the DBR. Unlock events are detected by measuring error signals above some threshold or large jumps in the control signal. If either of these events occurs the automatic locking procedure is restarted.

To lock the Ti:Sa the following procedure is used: the Ti:Sa frequency is constantly recorded on a wavemeter. Locking works by positioning each cavity element (the birefringent filter, the thin etalon, and the piezo etalon) at the location that brings the Ti:Sa frequency closest to the target frequency. This is performed using a python program on the Raspberry Pi connected to the Ti:Sa via a custom python-to-LabVIEW interface. The piezo etalon has the smallest effect on the cavity mode spacing, and is the last element to be positioned. It is positioned at the location where the Ti:Sa goes from being red detuned relative to the target frequency to blue detuned. Once the birefringent filter, thin etalon, and piezo etalon are positioned, the Raspberry Pi enables the previously disabled PI loop. The PI loop scans the fast piezo such that the Ti:Sa becomes more red detuned and eventually the piezo is positioned to zero the error signal, at which point the PI loop stops scanning and enters a negative feedback loop. Simultaneously, the Raspberry Pi DAC is used to perform a digital PI loop on the slow piezo position, using the fast piezo control signal as an error signal. This

keeps the fast piezo location centered within its range of motion. The Raspberry Pi detects unlocks by querying the wavemeter-measured Ti:Sa frequency and comparing it to the target frequency. Note that the absolute frequency of the Ti:Sa is guaranteed by the phase lock; the wavemeter is only used as a monitor.

4.2.7 Frequency comb and wavemeter

A Bristol Instruments 871A wavemeter with a fiber-optic switch is used to continually monitor the Ti:Sa laser frequency. It is also frequently used to tune the temperature at which various laser diodes are set, and for debugging. A Menlo Systems SmartComb is used to make the precise measurement of the Ti:Sa laser frequency needed for a measurement of the fine-structure constant.

4.2.8 Interferometry beam alignment

The interferometry beam and atom cloud trajectory must both be colinear with gravity to keep the atom cloud and interferometry beam centered on one another throughout the fountain trajectory. To align the interferometry beam, a mirror mounted to a two-axis tilt sensor is placed in a kinematic mount below the bottom chamber viewport. The mount is adjusted until the tilt sensor is aligned to gravity. The mirror is then rotated in the kinematic mount while observing the tilt sensor reading to measure and remove any bias. The interferometry beam is retroreflected off of the mirror and backfibercoupled. This ensures the beam is orthogonal to the mirror, which in turn is aligned to gravity. The tilt-sensor-mounted mirror is then removed, and the interferometry beam propagates up the vacuum chamber and reflects off of the in-vacuum retro mirror, which is then aligned to achieve backfibercoupling. After this procedure both the interferometry beam and retro mirror are aligned with gravity. The tilt sensor is a Model 755 from Applied Geometrics/Jewell Instruments.

A pair of linear stages attached to mirrors on the bottom of the vacuum chamber are used to align the interferometry beam with the atom cloud. After this alignment the backfibercoupling is checked to ensure the interferometry beam is still aligned to gravity.

4.3 Control system

Control of the experiment is provided by the open-source Advanced Real-Time Infrastructure for Quantum physics (ARTIQ) system. The core of the ARTIQ system is an FPGA that executes a queue of timestamped instructions provided by the experimenter. The ARTIQ system includes peripherals providing TTL inputs/outputs, RF DDS outputs, DAC outputs, and ADC inputs. Timestamped instructions are of the form “turn on TTL10 at time 50 us” or “set DDS5 frequency to 200 MHz at time 1320 us”. The FPGA ensures that each of these instructions is executed exactly at its specified timestamp.

The FPGA is provided with instructions by an Ubuntu computer running ARTIQ control software. The ARTIQ software compiles instructions written in python to instructions that can be interpreted by the ARTIQ FPGA bitstream. This means that not all python features and packages can be used when writing these python functions, which can lead to some confusion for novice ARTIQ users. To achieve the desired timing when writing ARTIQ sequences, the user must also be aware of the time it takes for the FPGA to execute each instruction. If this is ignored the user can encounter so-called "slack" errors, which occur when an instruction is given a timestamp in the past.

To enable the straightforward development of new experiment sequences, we use a class system for ARTIQ that defines different experiment "blocks". This allows different experiment subsystems, for example the cooling laser or the interferometry chain, to be abstracted into a single object instead of a collection of TTL and DDS channels. This enables users to write new experiment sequences without needing to remember every aspect of the experiment hardware. It also allows for existing experiment sequences to be read in a much more intuitive manner.

The ARTIQ DDSs provide the RF drive used for most of the AOMs in the experiment. The ARTIQ system is provided with a 125 MHz clock signal generated from a Vaunix Lab Brick clocked to a stable 10 MHz signal. The 10 MHz signal is generated from a Stanford Research Systems FS752 GPS-disciplined oscillator that utilizes an antenna mounted on the roof of Campbell Hall. The ARTIQ system also provides control over several third-party devices, such as arbitrary waveform generators, linear stages/flipper mounts, and power supplies, used elsewhere in the experiment. These devices are controlled over serial or LAN, and are not precisely timed. Any third-party device that must be triggered with precise timing is triggered via an ARTIQ TTL.

4.4 Atom cooling, trapping, and state preparation

Cesium atoms are introduced into the experiment from a room-temperature cesium ampule located in the 2D MOT chamber. To use these atoms for atom interferometry, their wavefunctions must be narrow enough in momentum space to undergo the Bragg diffraction process described in Section 2.2. Specifically, the wavefunction must be smaller than the momentum state spacing $2\hbar k$. The thermal de Broglie wavelength [71] is given by

$$\lambda_{\text{thermal}} = \frac{h}{\sqrt{2\pi k_{\text{B}} m T}} \tag{4.1}$$

To achieve $h/\lambda_{\text{thermal}} \ll 2\hbar k$ for cesium using 852 nm light, we find

$$T \ll \frac{2\hbar^2 k^2}{\pi k_{\text{B}} m} \approx 126\,\text{nK} \tag{4.2}$$

To create an atom cloud with temperature less than 100 nK we employ a variety of laser cooling and trapping techniques developed from the late 1980s onwards. The general process is

1. Cool room-temperature atoms in 2D using the 2D MOT, and direct them towards the 3D MOT.
2. Trap and further cool atoms in the 3D MOT to 125 µK.
3. Use PGC to cool the atoms to ∼1 µK during the atom launch. Atoms are in $6\,^2\mathrm{S}_{1/2}\,F = 4$ and spread out across many m_F states after PGC.
4. Use RSC to cool the atoms to ∼450 nK. This is the temperature observed when the system is very well tuned, and corresponds to a velocity spread of $1.5v_\mathrm{r}$. Atoms are in $6\,^2\mathrm{S}_{1/2}\,F = 3, m_F = 3$ after RSC.
5. Use adiabatic passage to transfer the atoms to $6\,^2\mathrm{S}_{1/2}\,F = 3, m_F = 0$.
6. Use two velocity selective Raman transitions to select a vertical momentum distribution roughly $0.1\hbar k$ in width.
7. Start interferometry!

We show the detected time-of-flight signals after the molasses launch, RSC, and VS in Fig. 4.19. The atom numbers in the figure are computed using the procedure described in Section 5.3.2. An illustration of the sequence from state preparation to detection is shown in Fig. 4.20. Note that this sequence only creates a momentum distribution much less than $2\hbar k$ in the vertical direction (i.e. along the chamber axis). Since this is the dimension along which Bragg diffraction/Bloch oscillations are performed, it is not necessary to cool the transverse momentum distributions below $2\hbar k$ (although doing so would reduce systematic effects).

In the remainder of this section we will discuss how each of these state preparation stages is implemented experimentally.

4.4.1 Magneto-optical traps

The experiment utilizes a 2D MOT to load the 3D MOT. The 2D MOT optical configuration is discussed in Section 4.1.1. There are three pairs of 3D MOT beam telescopes, referred to as the X, Y, and Z pairs. The X1 beam is colinear and counterpropagating w.r.t. Y2, and X2 is colinear and counterpropagating w.r.t. Y1. The Z telescopes create colinear counterpropagating beams. The X1-Y2 beam pair is orthogonal to the X2-Y1 beam pair. Both pairs are oriented at approximately 45° to gravity. The Z beams are orthogonal to both the X1-Y2 and X2-Y1 pairs. Note that both X beams propagate downwards and both Y beams propagate upwards. This allows the beams to be used for the molasses launch described in the next section.

The location of the 3D MOT is set by the location where the magnetic field is zeroed [72]. Ideally the atom beam formed by the 2D MOT and the 3D MOT beams would intersect at this location. The direction of the atom beam is defined by the differential pumping

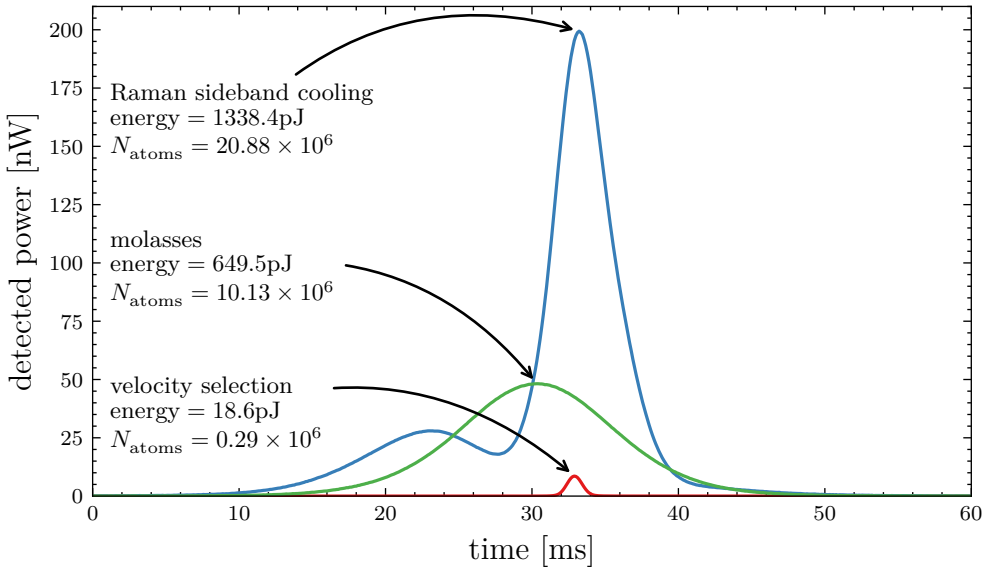


Figure 4.19: Falling cloud time-of-flight traces observed after a molasses launch, RSC, and VS, reported in units of optical power. A basic model of the detection optical system and the number of photons emitted per atom allows the total detected energy to be converted into an atom number. This figure is further discussed in Section 5.3.2.

tube, fixing its propagation in the 3D MOT chamber. The 3D MOT coils, 3D MOT beams, and compensation coils must be aligned/configured to form a MOT somewhere along the atom beam. The difficulty with performing this alignment comes from the uncertainty surrounding the alignment of each of these components: How can we detect if the 3D MOT beams intersect with the atom beam? How do we know where the magnetic field is zeroed?

To perform these alignments, we use the atom cloud formed by the 3D MOT and the atom beam formed by the 2D MOT as a tool. For this procedure to work, there must already be some detectable, although possibly poorly aligned, 3D MOT. First, the switchyard optical system is tuned to make all six 3D MOT beams have equal power entering the chamber. The waveplates on the 3D MOT telescopes are also aligned to produce circularly polarized light. The 3D MOT location is recorded on three cameras located at orthogonal viewports. Two of the three pairs of 3D MOT beams are then blocked. The single remaining counterpropagating beam pair causes the atom beam to fluoresce within its volume. This pair is then translated until the fluorescing volume is overlapped with the 3D MOT location that was measured on each of the three cameras. An aperture can be used to reduce the volume if higher precision is desired. This process is then repeated for the remaining beam pairs. Each beam pair is made colinear by overlapping the apertured beams on a beam card outside of the chamber.

An example image of the fluorescing volume recorded by the X1-Y2 beam pair is shown in Fig. 4.21. The image was taken by a camera located next to one of the Z beams. The red

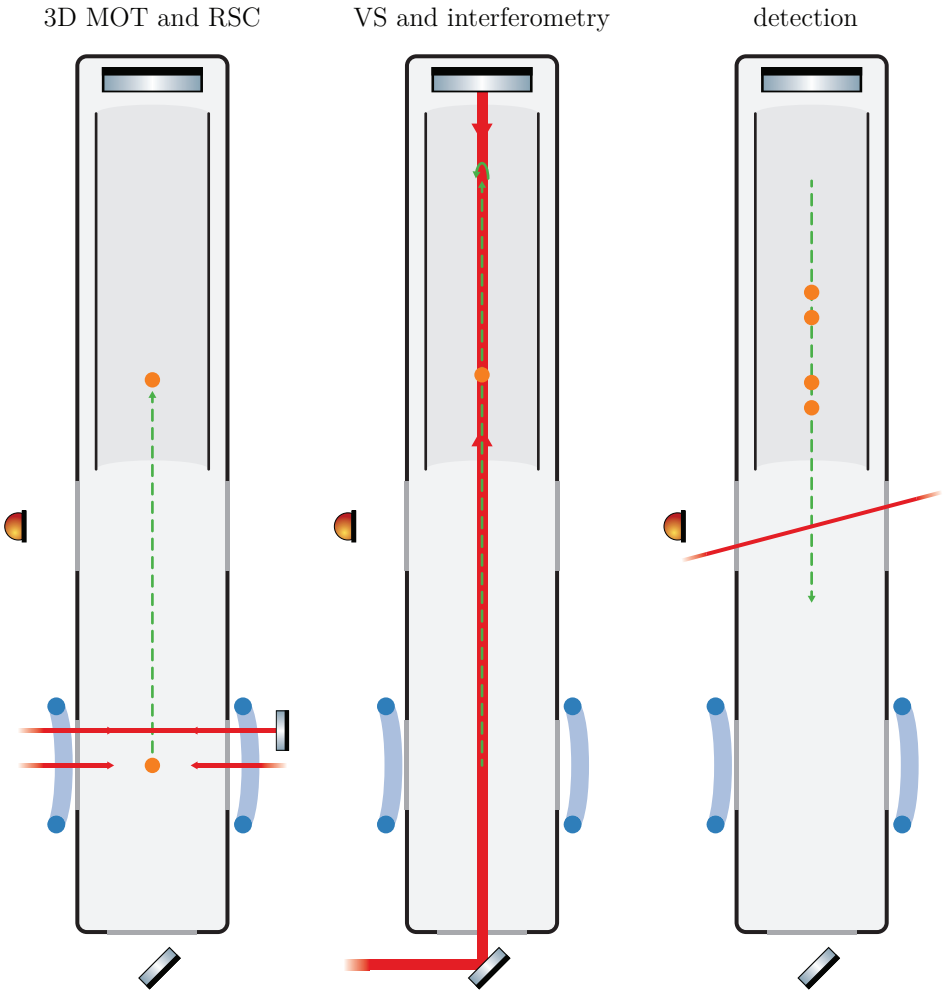


Figure 4.20: Schematic of the experiment chamber during different stages of the experiment. The 3D MOT coils are shown in blue, lasers in red, and viewports in gray. Atom clouds are orange, and their trajectory is shown in green. The three layers of Mu-metal shields 2.34 m in length [24], solenoid, and baffles are represented by the cylinder in the interferometry chamber. The position degree of freedom in Eq. (2.1) and the position axis in Fig. 3.3 correspond to the vertical dimension in the chamber. *Left:* After the 3D MOT is formed the molasses launch and PGC occur. Shortly afterwards the atom cloud passes through the RSC region (the retroreflected RSC X beam is shown) and is further cooled. *Middle:* After the VS pulses the atoms enter the magnetically shielded region of the chamber and the interferometry beam is pulsed to form Bragg beamsplitters and Bloch oscillations. *Right:* After the interferometer has completed, the spatially separated output ports of the interferometer fall through resonant light and fluoresce. The fluorescence is detected by a photodiode.

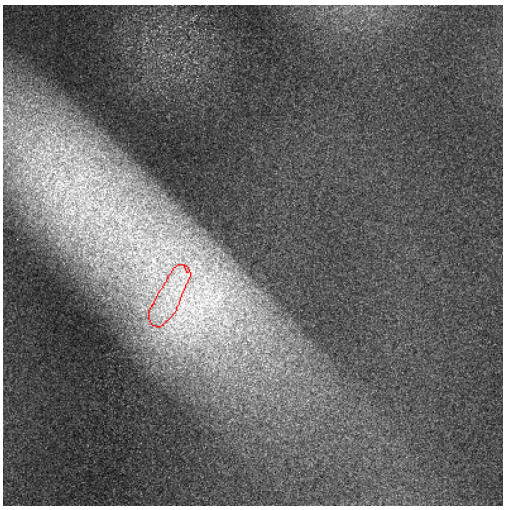


Figure 4.21: Fluorescence observed in the 3D MOT region by a camera when all but one pair of 3D MOT beams is turned off. The fluorescent region reveals the volume that the 3D MOT beam pair illuminates, and can be used to center the beam on the 3D MOT location.

circle is a manually drawn outline of the recorded MOT location.

The 3D MOT coils are positioned and the compensation coils are adjusted to zero the magnetic field at the location where the atom beams and 3D MOT beams have been aligned. However, likely due to the fluctuating magnetic field that exists during the 3D MOT coil turn off, the location of the atom cloud shifts as the coil is turned off. It is not clear if the magnetic field zero should be aligned to the atom cloud location when the coils are on or are in the process of turning off.

Each of these alignment procedures can change the optimal alignments determined during previous procedures. For example, aligning the 3D MOT beams to each other can affect overlap with the atom beam. For this reason it is optimal to repeat each alignment procedure several times until the alignments have converged. In practice different parts of this procedure are performed when the system needs a tune-up, but usually the full procedure is not performed.

Note that the X1-Y2 beam pair is close to being colinear with the 2D MOT atom beam. It seems possible that the Y2 beam is partially acting as a push beam to slow the atom beam down [62]. This is partially backed up by the observation that intense fluorescence is observed when all 3D MOT beams but the Y2 beam are blocked.

4.4.2 Molasses launch

A MOT provides both a restoring force and a friction-like force proportional to velocity. When the 3D MOT coils are turned off the restoring force turns off as well, leaving only a

friction-like force. When the magnetic field is zeroed PGC occurs–a Sisyphus effect that can cool atoms to the recoil limit [73, 74]. An intuitive description of PGC is provided in [75].

PGC occurs in the experiment sequence automatically once the 3D MOT coils have fully turned off and $B = 0$ at the 3D MOT location (it takes eddy fields produced by the coil turn off around 30 ms to die off). In normal operation the 3D MOT is loaded for 1.15 s before launch. 42 ms before the intended launch time the 3D MOT coils are turned off, giving PGC 8 ms to cool the atom cloud before the launch.

A moving molasses is used to launch the atoms. This works by blue detuning the bottom two Y 3D MOT beams and red detuning the top two X 3D MOT beams. This produces a net upward scattering force that launches the atoms. The exact velocity that the atoms will be accelerated to is determined by computing the detunings that yield equal frequencies in the frame of the launched atoms. In the frame of the moving atoms the detuned X and Y laser frequencies ω'_x and ω'_y become

$$\omega'_x + \frac{kv_{\text{launch}}}{\sqrt{2}} = \omega'_y - \frac{kv_{\text{launch}}}{\sqrt{2}} \rightarrow \sqrt{2}kv_{\text{launch}} = \omega'_y - \omega'_x = 2\delta_{\text{launch}} \tag{4.3}$$

The factor of $\sqrt{2}$ comes from the X and Y beams being at 45° to gravity. For a launch velocity of 6.2 m/s this yields a frequency shift of $\delta_{\text{launch}} = 2\pi \times 5.14\,\text{MHz}$ on each beam. This shift is applied at the time of launch. At the same time a shutter blocks the 2D MOT light, deactivating the atom beam.

3 ms after the launch the cooling laser frequency is blueshifted by jumping the second demodulation frequency in the phase lock by 150 MHz (described in Section 4.2.6). This lowers the size of the Sisyphus potential created by the polarization gradients and enables further cooling. This all occurs as the atoms are moving upwards in the 3D MOT beams, meaning that the atoms are moving into regions of lower optical intensity. This further lowers the PGC potential. 4.3 ms after the launch the 3D MOT beams are turned off, at which point the atoms should be ∼27 mm above the 3D MOT location and outside the extent of the 3D MOT beams. The DBR phase lock is then moved back to its nominal location.

The MOT alignment procedure overlaps the 3D MOT beams and B-field zero with the atom beam. However, it does not ensure that the X and Y beams are at 45° to gravity. To perform this alignment the detection photodiode shown in Fig. 4.3 is translated to maximize the rising fountain signal. The rising fountain signal is the fluorescence signal produced by the atom cloud as it rises upwards through the detection region. The falling fountain signal is then measured, and the X, Y, and Z beams are translated and angled to maximize the signal. The photodiode is then again aligned to maximize the rising fountain signal, and the X, Y, and Z beam translating/angling is repeated. This is done in a loop until the maximum of the rising and falling fountain signals occur at the same location. At the end of this alignment procedure, the 3D MOT beam alignment procedure described in Section 4.4.1 might need to be repeated.

The main degrees of freedom useful for optimizing PGC are the compensation coil currents, as these tune the B-field at the atom cloud location. The currents are adjusted to

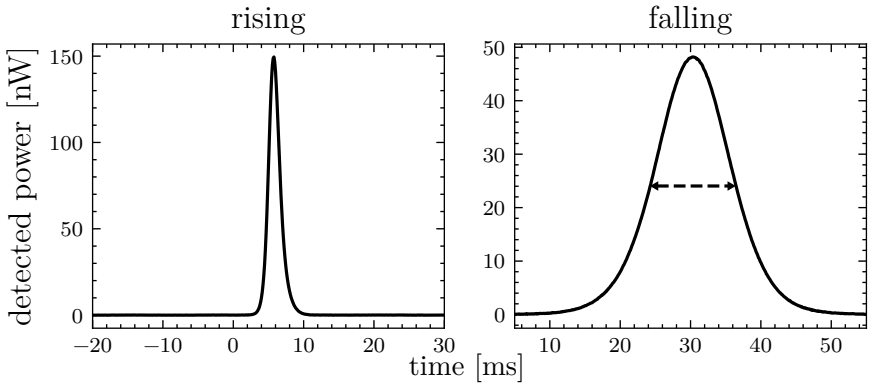


Figure 4.22: Time-of-flight traces of the rising and falling molasses clouds measured in the detection region. The dashed line on the falling cloud represents the point at which the full width at half maximum is measured.

minimize the width of the falling fountain signal. Smaller width signals correspond to atom clouds with smaller velocity spreads. The vertical compensation coil current is set using a current driver, and its value during the atom launch can be set independently from its value during MOT loading.

An example rising and falling fountain signal is shown in Fig. 4.22, with the FWHM of the falling cloud shown as a dashed line. Note that the time axis for each plot has an arbitrary offset. The falling atom cloud has velocity $-4\,\mathrm{m/s}$ and a measured FWHM of 12.2 ms. This means the cloud is 49 mm in length. The time from launch to detection is roughly 1 s, so for a cloud initially located at a single point[2] the expansion rate is approximately 25 mm/s in each direction, or roughly 7 recoil velocities. This corresponds to a temperature of $\sim$10 µK. The exact velocity width/temperature of the cloud varies depending on how well tuned the system is. After the molasses launch the atom cloud is in the $F = 4$ hyperfine level and spread out over all m_F states.

Selecting the launch velocity requires balancing fountain duration, atom loss, and cloud expansion. The faster the atom launch, the more time is available to run an interferometer. But faster launches also produce lossy molasses launches, which degrades the final interferometry signal. We also must operate the interferometer within the magnetically shielded region of the vacuum chamber. The interferometer must also be timed to avoid operating Bragg diffraction or Bloch oscillations when unwanted standing wave lattices are formed at velocities close to the atom velocity. Fig. 4.23 shows how the choice of launch velocity affects the interferometer duration that can be run inside the magnetically shielded region. We see that 6.2 m/s gives us around 400 ms in the shielded region.

All lattices formed in the chamber during Bloch oscillations are plotted in Fig. 4.24 for

[2]This is a fair approximation; in Section 5.3 it is shown that the atom cloud in the MOT is only a few millimeters in extent.

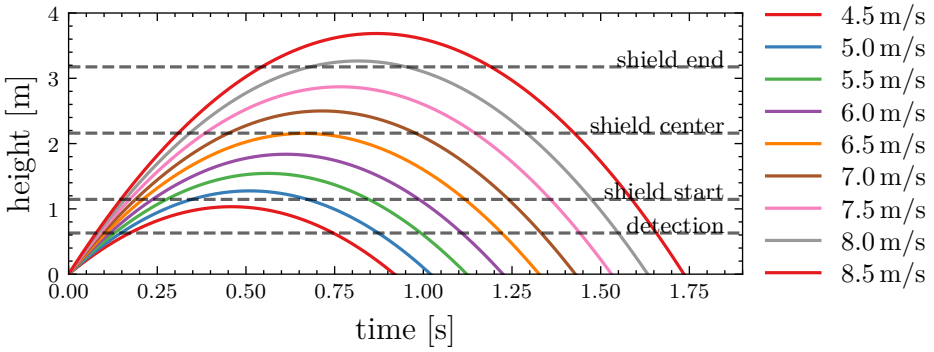


Figure 4.23: Height of the atom cloud as a function of time for several launch velocities. Horizontal dashed lines denote various relevant heights in the interferometry chamber. A launch velocity of 6.2 m/s is chosen as it balances sufficient time in the shielded region and interferometer duration with signal size, which goes down with larger launch velocities.

$v_{\text{launch}} = 6.2\,\text{m/s}$ and 4th-order Bragg, 100th-order Bloch, a Bloch ramp rate of $14g$, and a first Bragg pulse triggered 225 ms after launch (just inside the magnetically shielded region). The Bloch lattices are the lattices that actually perform Bloch oscillations. The junk lattices are formed by unwanted pairs of upgoing and downgoing laser frequencies. Care must be taken to ensure that the junk lattice velocities do not intersect the Bloch lattice velocities, as otherwise the Bloch process will be disrupted. These conditions must also be met during Bragg diffraction.

4.4.3 Raman sideband cooling

RSC works by pumping atoms into the ground state of an optical lattice [76–80]. The RSC X, Y, and Z beams form a 3D optical lattice. The Z beam is detuned so that all beams are resonant with $F = 4 \rightarrow F' = 4$ in a frame moving at the velocity of the atoms (see Section 4.2.4). The resonant light quickly pumps atoms into the $F = 3$ state, where they see a lattice detuned by around 9 GHz. A non-zero magnetic field is tuned to make the splitting between m_F states equal to the vibrational splitting of the lattice. This enables the RSC X standing wave to drive Raman transitions between different m_F states. Finally, the RSC pump beam is used to pump atoms into progressively higher m_F states. The whole process has the effect of pumping atoms into $F = 3, m_F = 3$ and the vibrational ground state of the lattice.

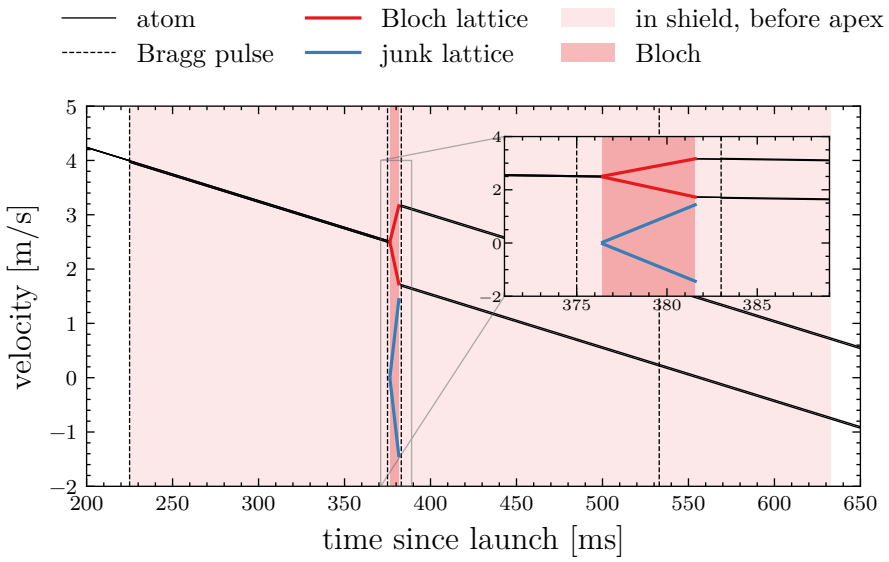


Figure 4.24: Velocity of atoms in the experiment chamber as a function of time during an SCI with Bloch oscillations. The interferometer starts 225 ms after launch, has a launch velocity of 6.2 m/s, 4th-order Bragg diffraction, 100th-order Bloch oscillations, a Bloch oscillation lattice acceleration of $7kg$, $T = 150$ ms, and $T' = 8$ ms. The velocities of other lattices formed in the interferometer during Bloch oscillations are shown. Care must be taken to ensure the velocity of these lattices does not intersect the velocity of the atoms, as otherwise the atoms will be ejected from their Bloch lattice as the two lattices accelerate past each other. We see that the Bloch lattice velocities do not intersect while Bloch is on, although they become relatively close (this is clear from the inset). After Bloch oscillations the third Bragg pulse occurs and the atom velocity lines bifurcate.

Lattice structure

We will compute the lattice potential for the case in our experiment, as it provides insight into some of the curiosity surrounding the RSC process. For our three-beam RSC configuration the total field is

$$\begin{aligned}\mathbf{E}_{\text{RSC}} = &\left[2\sqrt{a_x}E_0 \sin(\theta)\cos(kct)\cos(kx)\right.\\ &\left.+\sqrt{a_z}E_0\left(\cos(kct)\cos(kz)+\sin(kct)\sin(kz)\right)\right]\hat{y}\\ &+\left[-2\sqrt{a_x}E_0\cos(\theta)\cos(kct)\cos(kx)\right.\\ &\left.-\sqrt{a_y}E_0\left(\cos(kct)\cos(ky)+\sin(kct)\sin(ky)\right)\right]\hat{z}\end{aligned} \tag{4.4}$$

where a_i denotes the relative intensity of beam i and θ sets the angle between the RSC X polarization and the $\hat{z}$ direction. Note that there is no component in $\hat{x}$. The resulting optical lattice potential is proportional to $\mathbf{E}_{\text{RSC}} \cdot \mathbf{E}_{\text{RSC}}$. This will produce a potential with

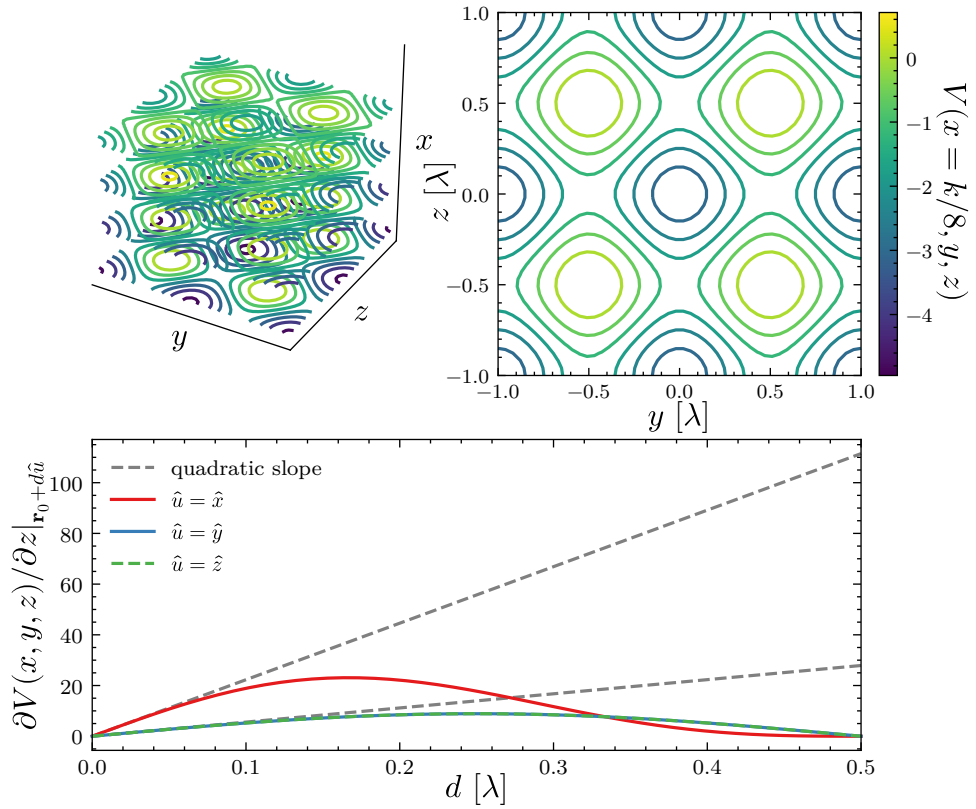


Figure 4.25: The potential formed by the RSC lattice, and its derivative as a function of displacement from the potential minimum for different displacement directions.

several terms rotating at optical frequencies, which can be dropped as these dynamics are fast relative to the dynamics of the atoms. Alternatively the potential can be time averaged over one optical period. This yields

$$V_{\mathrm{RSC}}(x,y,z;\theta) \propto -\frac{a_x+a_y+a_z}{2} - a_x\left(\frac{1}{2}+\cos(2kx)\right) \\ -2\sqrt{a_x a_y}\cos(kx)\cos(ky)\cos(\theta) - 2\sqrt{a_x a_z}\cos(kx)\cos(kz)\sin(\theta) \tag{4.5}$$

where a negative sign was introduced to account for the fact that RSC beams are red detuned, meaning that the potential minimum is located in high intensity regions. This produces the potential shown in Fig. 4.25 for $a_x = a_y = a_z$ and $\theta = 45\,\mathrm{deg}$. On the upper left a contour plot of the potential at $x = 0, \lambda/16, 2\lambda/16, 3\lambda/16, 4\lambda/16$ is shown. The upper right shows a contour plot at $x = \lambda/8$, where it is apparent that the potential forms an eggshell-like pattern in these slices. For a more in-depth discussion of optical lattice potentials, see [81].

We find the potential minimum numerically and plot the derivative of the potential evaluated at $\mathbf{r}_0 + \hat{u}d$, where $\mathbf{r}_0$ is the potential minimum and $\hat{u}$ is either the x, y, or z

direction. For a quadratic potential the derivative will be linearly increasing with d. The plot shows the numerically computed derivatives as a function of d for each direction, as well as the linearly increasing derivative that we would expect for a perfectly quadratic potential. We see the numerically computed slope is close to that of a perfectly quadratic potential for $d \lesssim 0.1\lambda$, meaning that the optical lattice can be approximated as a harmonic oscillator potential in this region. We see that the derivative slope is different for x and y/z directions, meaning that the harmonic potential is anisotropic.

Because the harmonic potential is anisotropic, the vibrational splitting will be unequal in different directions. Given this, it is unclear how the RSC Raman transitions and pumping process cools atoms into the vibrational ground state in all three dimensions. Clearly there is some additional process coupling dimensions together or some other form of energy exchange between dimensions, such as collisions between atoms. This is discussed in [76, 80].

Alignment procedure and timing

The anisotropic RSC potential, along with other strange observations that have occurred when attempting to align the RSC lattice beams (such as the observation of multiple "cooled" atom clouds at different velocities away from the RSC lattice velocity), has led us to the conclusion that the RSC process in our experiment acts differently from the idealized process described in the literature. The alignment procedure and timing sequence used in the experiment is the product of many confusing attempts to align RSC. While this alignment procedure does often produce RSC signals higher than the molasses signal, this is not guaranteed. Sometimes the RSC signal is lower than the molasses signal, which is a sign that the RSC process is selecting rather than cooling atoms into a narrower velocity distribution. When this occurs the process is often restarted and other state preparation stages are checked for errors in alignment, timing, or more.

To start the RSC alignment procedure, the RSC Z beam is aligned to be parallel to gravity using the procedure described in Section 4.2.8. Then the RSC X beam is aligned to be orthogonal to gravity, and the RSC X retroreflection is aligned to overlap with the incident RSC X beam. To make the RSC X beam orthogonal to gravity, one of two methods can be used. The first method retroreflects the beam off of a mirror coplanar with gravity[3]. The second method aligns the beam so that it passes through the same location in the incident and exiting viewport. This assumes the chamber tolerances are precise enough and that the chamber is vertical with respect to gravity (this can be accomplished using a tilt sensor). The beam is then apertured and blowaway light is temporarily used instead of RSC light. The beam is then translated (not tilted) so that the blowaway minimizes the falling molasses signal, centering the beam on the atom cloud at the time the blowaway light turns on (which is set to be the same time that RSC happens). The same alignment procedure is used on the RSC Y beam.

[3] A quick and dirty method for doing this is to use the tilt sensor of a smartphone to align its screen, and then to use the screen as a mirror.

Next, the RSC Y beam is blocked so that a 2D RSC lattice is formed instead of a 3D one. The idea here is to do the initial alignment on a system with fewer degrees of freedom than the full 3D RSC lattice. The RSC X and RSC Z beams are translated until some RSC signal is visible on the falling fountain. The RSC pump beam is optimized as well. The experimenter then tunes the RSC degrees of freedom until the RSC signal is larger than the molasses launch signal. Ideally only small tilt adjustments are needed. Then it is checked that blocking the RSC Y beam lowers the RSC signal. If blocking RSC Y raises the RSC signal or leaves it unchanged, it is either hurting or not contributing to the RSC process. At some point it might be necessary to tilt the RSC X, Y, and Z beams to finely tune the direction of the atom cloud after the RSC process. The goal is to direct the cloud such that the detected signal height is maximized at the rising molasses photodiode location.

Once this has been done we switch from optimizing the RSC signal height to optimizing the double VS signal height using the RSC degrees of freedom. Past RSC alignments attempted to maximize the falling RSC signal size instead of the VS signal. However, it was found that the alignment configuration that maximized the RSC signal did not always maximize the VS signal. The reason for this is unknown[4]. Once a VS signal exists, the RSC beam alignments should be walked to optimize the falling VS signal with the detection photodiode located at the rising cloud location.

Once it seems RSC X, Y, and Z are all contributing positively to the VS signal, it seems best to tune all RSC degrees of freedom in a cycle (including all RSC coil currents and RSC waveplates), as well as the adiabatic passage microwave horn. Doing this in a somewhat random way seems to yield reasonable results, as there are so many degrees of freedom and the RSC process is so complicated that there is no clear argument for optimizing in a particular order. There are several RSC parameters controlled by ARTIQ that should be optimized as well, such as the vertical compensation field current, the RSC Z lattice velocity, and the RSC timing parameters. The vertical compensation field current and the RSC Z lattice velocity can have a huge effect when tuning RSC.

The RSC X, Y, Z, and pump beams are turned on 9.75 ms after the molasses launch. 4 ms before the lattice is turned on the vertical compensation coil current is set and the RSC coil is switched on. 1 ms before the lattice is turned on the cooling laser frequency is jumped to make the RSC beams resonant with $F = 4 \rightarrow F' = 4$ and repump light is removed from the RSC beams by turning AOM12 off.

The RSC lattice beams remain on for 3 ms before their intensity is ramped down by gradually reducing the AOM23 and AOM25 RF drive powers in ARTIQ in an attempt to release the atoms adiabatically from the lattice. The adiabatic release does not seem to have a huge effect, which could imply that the atoms are already at the edge of the RSC beams by the time the rampdown occurs. 0.1 ms before the lattice beams are ramped down the RSC pump beam is turned off. The RSC falling cloud signal looks like the molasses signal

[4]In fact, it can even be beneficial to optimize RSC and molasses launch knobs using a delayed Bragg diffraction inversion signal, as described in Section 4.5. One general lesson of operating this experiment has been to optimize on the furthest downstream signal possible.

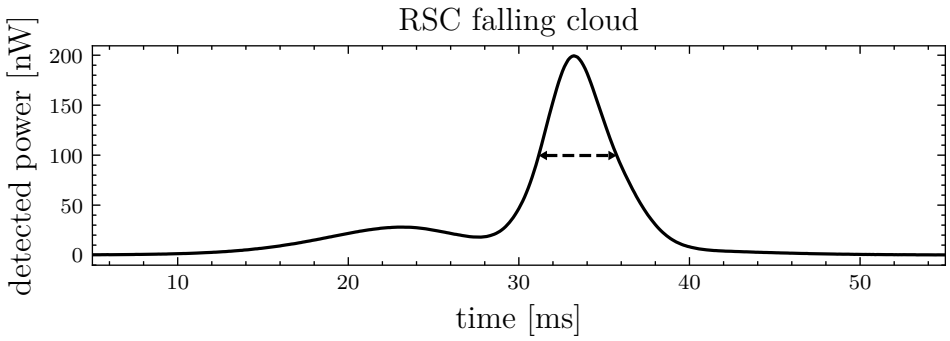


Figure 4.26: Time-of-flight trace of the falling RSC clouds measured in the detection region. The dashed line on the falling cloud represents the point at which the full width at half maximum is measured.

with a narrower peak superimposed on top, as shown in Fig. 4.26. The measured FWHM of the RSC peak is 4.6 ms. The same back-of-the-envelope calculation used to estimate the velocity spread after PGC yields an expansion rate of 9 mm/s in either direction, or roughly 2.5 recoil velocities. This corresponds to a temperature of 1.2 µK, and corresponds to an adequately, but not optimally, tuned RSC signal. In Fig. 4.19 we plot the same RSC trace as in Fig. 4.26 along with the falling molasses launch shown in Fig. 4.22. We see that the RSC signal is substantially higher than the molasses signal, which must be due to its lower transverse temperature.

4.4.4 Adiabatic passage

After RSC, atoms are in $F = 3, m_F = 3$. We wish to perform interferometry in a magnetically insensitive state, and so we must transfer the atoms to the $F = 3, m_F = 0$ state. Each m_F level is coupled by a microwave transition to m_F states in the other hyperfine state, as shown in Fig. 4.27. The population in state $F = 3, m_F = 3$ can be transferred to $F = 3, m_F = 0$ by adiabatically sweeping from slightly above the $F = 3, m_F = 3 \rightarrow F = 4, m_F = 3$ frequency to slightly below the $F = 3, m_F = 0 \rightarrow F = 4, m_F = 1$ transition. This will transfer atoms along the path with the blue dashed line in the figure.

If desired, the frequency can be swept further to transfer the atoms to $F = 4, m_F = 0$. An alternate transition pathway from $F = 3, m_F = 3$ to $F = 3, m_F = 0$ exists by sweeping the microwave frequency down from slightly above the $F = 3, m_F = 3 \rightarrow F = 4, m_F = 2$ frequency to slightly below the $F = 3, m_F = 0 \rightarrow F = 4, m_F = 0$ transition.

In the experiment adiabatic passage is performed using a microwave horn. The horn is attached to the bottom of the vacuum chamber and directed through the bottom viewport. The horn orientation and position are aligned on the double VS signal. There is effectively

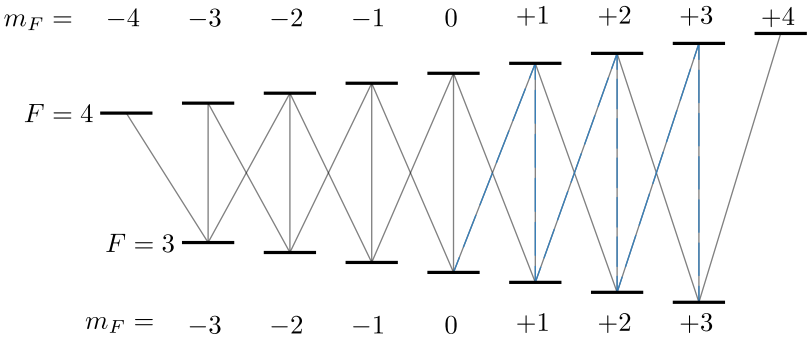


Figure 4.27: m_F states in the $F = 3$ and $F = 4$ hyperfine ground states of cesium, along with their allowed transitions. During the adiabatic passage stage, population is transferred from $F = 3, m_F = +3$ to $F = 3, m_F = 0$ along the dashed blue line.

no strategy for aligning the horn–it is wiggled around in a somewhat random way until sufficient VS signal is observed. Rotating the horn can have a large effect on the transfer efficiency. The double VS signal is the best signal to optimize the horn alignment with, as it only measures atoms transferred to $F = 3, m_F = 0$. The microwave frequency is sourced using a circuit identical to the one described in Section 3.5.2 of [82].

1 ms after the RSC lattice is turned off the vertical compensation coil current is increased in order to create a magnetic field splitting of approximately 150 kHz between adjacent m_F states. 38 ms later the microwave source is swept from 1568 kHz above to 336 kHz below the $F = 3, m_F = 0 \rightarrow F = 4, m_F = 0$ transition over 40 ms. An absorptive switch is used to switch the microwave horn on from 1118 kHz above to 78 kHz above the $F = 3, m_F = 0 \rightarrow F = 4, m_F = 0$ transition. This corresponds to a sweep from slightly above the $F = 3, m_F = 3 \rightarrow F = 4, m_F = 3$ transition to slightly below the $F = 3, m_F = 0 \rightarrow F = 4, m_F = 1$ transition, and leaves atoms in the $F = 3, m_F = 0$ state. By toggling the switch on and off at different times, the atoms are exposed to different microwave scans.

The signal observed after adiabatic passage (after an $F = 4$ blowaway but before VS) looks similar to the observed signal after RSC, but is approximately a factor of 2 smaller in height. By performing microwave spectroscopy of the m_F states after RSC it appears that RSC leaves atoms spread across several m_F states, which explains most of the factor of 2 loss. This spread is unwanted, and is due to an imperfect RSC process. Microwave spectroscopy is performed using the same microwave source used for adiabatic passage, but where the source is fixed in frequency instead of swept.

4.4.5 Blowaway light

Blowaway light is light resonant with a cycling transition used to remove atoms in a particular hyperfine ground state from the detected time-of-flight signal by transferring enough momentum to the atoms that they are "blown" outside of the detection region. To blow away atoms in the $F = 4$ state, light resonant with the $F = 4 \rightarrow F' = 5$ transition is used. The $F' = 5$ state can only decay to the $F = 4$ state, and so the transition is closed. Occasionally a photon will off-resonantly transition the $F = 4$ atom off of the $F = 4 \rightarrow F' = 4$ transition, from which the atom can decay to $F = 3$. For this reason the $F = 4$ blowaway is not 100% efficient. This is also the reason why repump light is needed in addition to cooling light in the MOT.

To blow away atoms in $F = 3$, the $F = 3 \rightarrow F' = 2$ transition is used. However, because $F' = 2$ has fewer m_F states than $F = 3$, the $F = 3$ blowaway has the effect of pumping atoms into a superposition of m_F states that is dark to the blowaway light, as described in [82]. This, along with off-resonant transitions, prevents the $F = 3$ blowaway from being 100% efficient.

The $F = 4$ blowaway beam enters the experiment at the bottom of the chamber. It is aligned at a slight angle to gravity in order to intersect with the atom cloud from after RSC to after VS. It is first fired after the adiabatic passage sweep for a duration of 10 ms. The $F = 3$ blowaway beam enters the experiment next to the detection beam. It is angled downwards to intersect the rising atom cloud before it passes through the detection region. It is fired for a duration of 20 ms after the first VS pulse to blow away atoms remaining in $F = 3$. The final $F = 4$ blowaway fires for a duration of 40 ms after the second VS pulse to blow away atoms remaining in $F = 4$. Both blowaway beams are aligned by minimizing the background of the double VS signal.

The combined efficiency of the $F = 4$, $F = 3$, and $F = 4$ blowaway process is 97% when the system is well tuned. The $F = 3$ blowaway is close to 100% efficient as it is orthogonal to the launch, which means it requires fewer scatters to remove $F = 3$ atoms from the time-of-flight signal. Since the $F = 4$ is almost colinear with the launch (this is necessary for it to intersect the atom cloud at multiple times), it requires many more scatters to remove $F = 4$ atoms from the time-of-flight signal. If we assume that $F = 3$ is indeed 100% efficient, then $F = 4$ is 98.5% efficient.

4.4.6 Velocity selection

After adiabatic passage and the first $F = 4$ blowaway, atoms are prepared in the $F = 3, m_F = 0$ state. The role of VS is to remove all atoms outside a narrow region of the velocity distribution of the atom cloud. The VS process operates by utilizing two subsequent velocity selective Raman transitions of the kind described in Section 2.4. The laser frequencies used to drive the velocity selective Raman transition are described in Section 4.2.5.

To understand the behavior of the Raman process given by Eq. (2.42), we integrate the momentum space wavefunction of the two ground states in time for a resonant two-photon

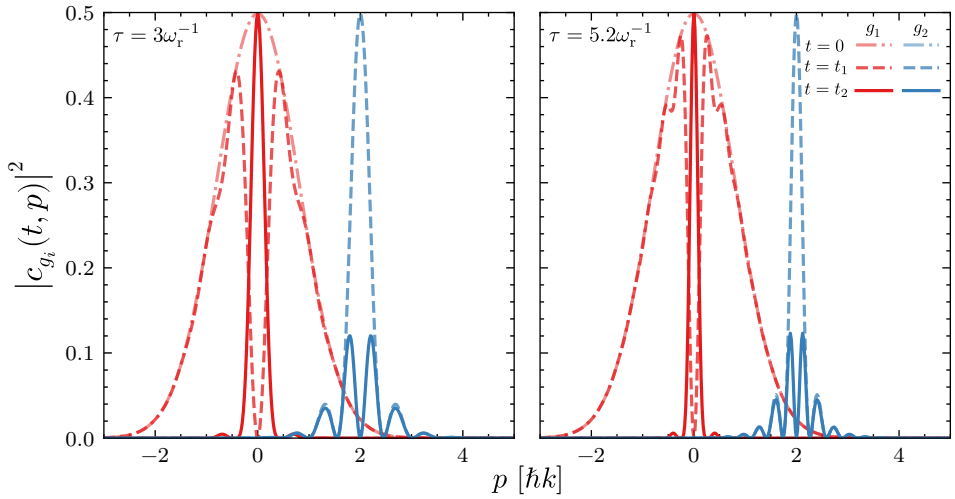


Figure 4.28: Wavefunction in each of the two ground states computed at various points of the VS process. Each panel is computed with a different Raman pulse duration τ (both panels are square pulses in time). In the experiment, $\tau = 5.2\omega_r^{-1}$ (the panel on the right) is used.

detuning and a constant effective Rabi frequency Ω_{eff} tuned to produce a π pulse (where $\Omega_{\text{eff}} = \Omega_{\text{eff}}^{(1)} = \Omega_{\text{eff}}^{(2)}$). This is shown in Fig. 4.28. At $t = 0$, the wavefunction is described by a Gaussian wavepacket in the g_1 state with a FWHM of approximately $2\hbar k$. At time $t = t_1$ the first VS pulse has transferred the center of the wavefunction to the g_2 state. The remaining amplitude in the g_1 state is then set to zero (corresponding to the blowaway), and a second VS pulse is performed to transfer population back into g_1. At $t = t_2$ we see that the population in g_1 has a FWHM of around $0.4\hbar k$.

From Fig. 4.28 we see why double VS is used, as after the first velocity selective pulse (at $t = t_1$) there are momentum sidebands on the g_2 wavefunction. Switching from a pulse with a constant Ω_{eff} of duration τ to a pulse with a Gaussian $\Omega_{\text{eff}}(t)$ smooths the frequency composition of the pulse, which suppresses these sidebands. The wavefunctions resulting from such a Gaussian pulse are shown in Fig. 4.29.

The longer the Raman pulse duration is, the narrower its frequency composition becomes. This means that its bandwidth in momentum space is reduced and it transfers a narrower band of velocities. This is apparent from Fig. 4.30, where we see that the computed population extent after double VS decreases as the pulse duration increases. In the experiment we use a pulse duration of 400 µs, for which the calculation predicts a selected population with a momentum width of approximately $\sigma = 0.073\hbar k$. This momentum width is sufficient for use in Bragg diffraction in the quasi-Bragg regime, as the bandwidth of the Bragg process is broader than this width. Making VS more velocity selective increases interferometer contrast

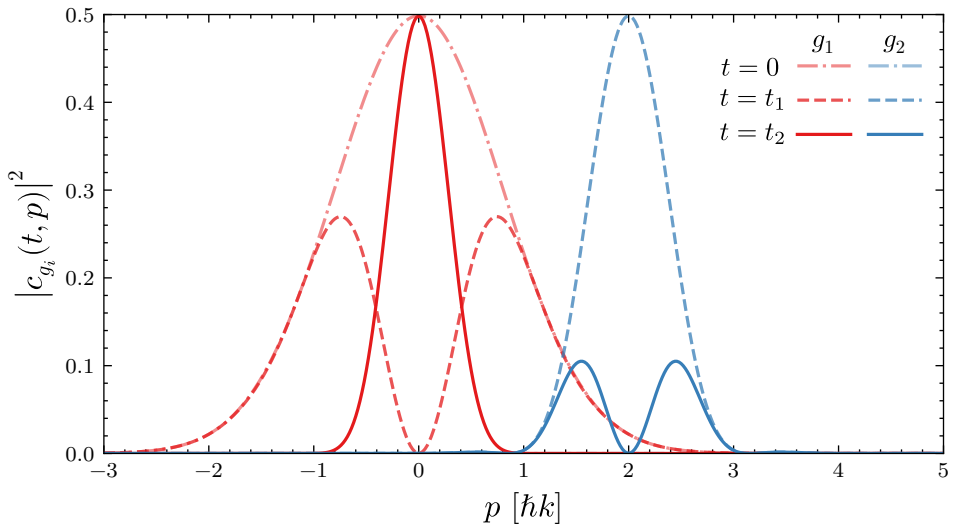


Figure 4.29: Population in each of the two ground states computed at various points of the VS process when a Raman pulse with a Gaussian profile in time is used. The Gaussian has $\sigma = 0.5\omega_r^{-1}$.

at the cost of signal. From Fig. 4.30, we see that making the pulse duration longer than 400 µs only marginally reduces the selected momentum width.

Note that while the VS integration above is for a single atom wavefunction and not an ensemble of wavefunctions, the result is the same. The part of each atom's momentum space wavefunction resonant with the Raman process will be transferred, regardless of where the center of the atom's wavefunction is. Therefore, VS makes the width of the momentum space wavefunction for all atoms in the interferometer equal, assuming all initial momentum widths are much less than the velocity selected momentum width. This also sets the width of the atom cloud velocity distribution.

The VS pulse should already have been aligned when optimizing the RSC alignments. If not, the only available degrees of freedom are the translation stages that move the location of the interferometry beam at the bottom of the chamber. The drive frequency of AOM32 is scanned to determine the Raman two-photon detuning δ that transfers the largest number of atoms. This scan selects what part of the RSC velocity distribution is transferred by VS.

The first VS pulse is performed 10.2 ms after the adiabatic passage frequency sweep ends, bringing the atoms from $F = 3, m_F = 0 \rightarrow F = 4, m_F = 0$. 20.2 ms later, only 0.2 ms after the $F = 3$ blowaway has turned off, the second VS pulse is performed to bring the atoms back to $F = 3, m_F = 0$ from $F = 4, m_F = 0$. After this a final $F = 4$ blowaway is fired, and then the system is ready for interferometry.

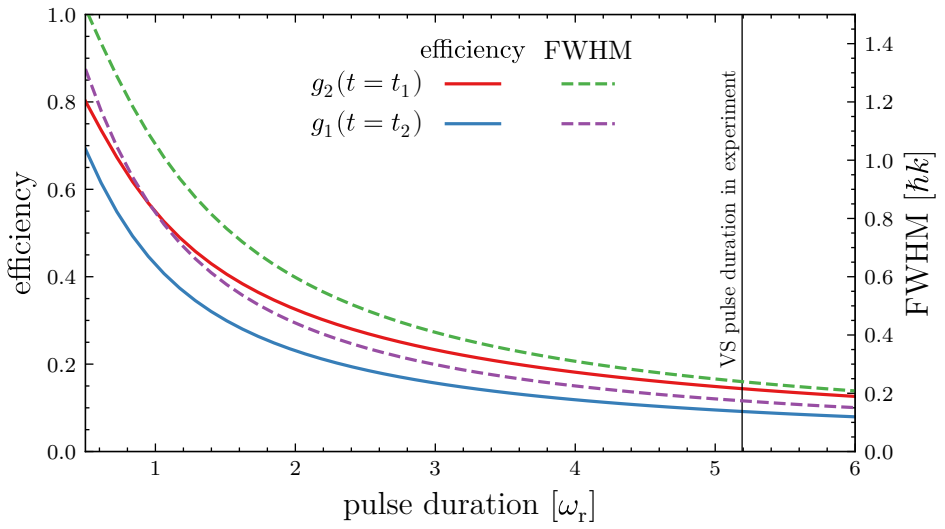


Figure 4.30: The efficiency and selected population width of the VS process computed as a function of pulse duration. A square pulse is used in time, with Ω_{eff} adjusted to maintain a π pulse. The efficiency and selected population width are shown both after a single VS pulse (g_2), and after two VS pulses (g_1).

4.5 Interferometry sequences

In this section we discuss how the matterwave optics used to form an interferometer are tuned. Typically the first Bragg diffraction pulse is fired 225 ms after launch, by which point the atoms have entered the magnetically shielded region of the interferometry chamber. Typically a simultaneous conjugate Ramsey-Bordé interferometer with $T = 100\,\text{ms}$, $T' = 8\,\text{ms}$, $n = 4$, and $N = 100$ is used, as this balances accumulated phase with interferometer contrast and atom number. These were the parameters used to record the detection photodiode trace shown in Fig. 4.31. The interferometer peak heights are extracted to compute the population of atoms in each output port of the simultaneous interferometers.

Vibrations on the top mirror lead to a phase shift of the retroreflected beams. Because vibrations are inconsistent between experiment shots, this phase shift varies shot-to-shot. This adds a random common-mode phase to each interferometer. This is discussed in further detail in Section 5.1.

4.5.1 Bragg diffraction

To select the Bragg pulse intensity, a single Bragg pulse is operated and the inversion between the initial and final momentum states is measured as a function of the peak effective

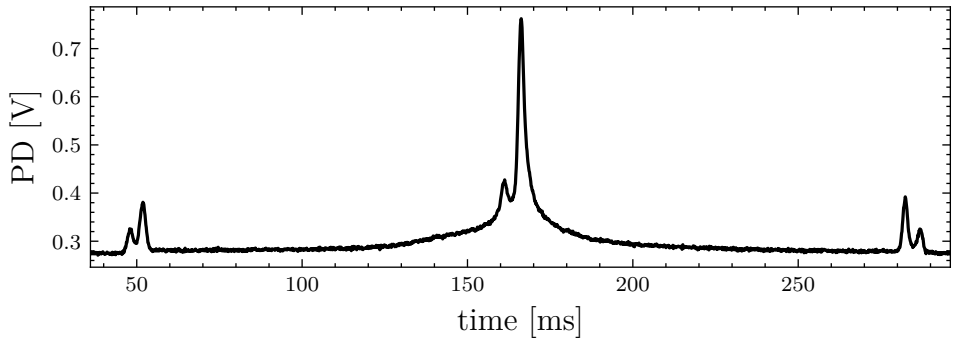


Figure 4.31: Measured time-of-flight trace of the SCI signal. For this interferometer, $T = 100\,\text{ms}$, $n = 4$, and $N = 100$. The pair of peaks around 50 ms are the outputs of the lower interferometer, and the pair of peaks around 235 ms are the outputs of the upper interferometer. The first pair of peaks have FWHM $\sim$2 ms, which is limited by the cloud size rather than the cloud expansion rate.

Rabi frequency $\Omega_{\text{eff},0}$. One of these measurements is shown in Fig. 4.32. Experimentally this is accomplished by modifying the amplitude of the waveform generated by the pulse shaping function generator described in Section 4.2.5, and so the x-axis is labeled with an amplitude in millivolts, not $\Omega_{\text{eff},0}$. This figure is referred to as a Bragg flop. The $\pi/2$ pulse condition is taken to be the inversion midpoint, which occurs at around 315 mV in the figure. The conversion factor between the interferometry beam power and the function generator amplitude is roughly 0.25 mW/mV, and so the $\pi/2$ pulse occurs at a power of $\sim$80 mW. Recall from Section 4.1.2 that the interferometry beam has a $1/e^2$ radius of 7.7 mm.

Multifrequency Bragg $\pi/2$ pulses are typically operated at twice this intensity, as expected from the factor of two relating the multifrequency and single-frequency effective Rabi frequencies derived in Section 2.2.2. Since the atom cloud expands as it moves upwards in the vacuum chamber, the average intensity experienced by the atoms decreases. This pushes out the $\pi/2$ pulse condition to higher peak effective Rabi frequency values. To counteract the cloud expansion, the peak effective Rabi frequency could be set for each Bragg pulse individually. However, in practice we leave it fixed.

Optimizing the Bragg pulse parameters also requires scanning the two-photon detuning δ. The measured inversion as δ is scanned across a downward 4th-order Bragg resonance is shown in Fig. 4.32, and δ is chosen by selecting the value that produces the maximum inversion. Both $\Omega_{\text{eff},0}$ and δ can be tuned more precisely to maximize the interferometer contrast. However, this only produces a small contrast improvement over the contrast achieved by selecting $\Omega_{\text{eff},0}$ and δ from inversion plots.

The Bragg pulse waveform is described by Eq. (2.20). The pulse width σ can also be tuned. σ is inversely correlated with the beamsplitter momentum bandwidth as well as

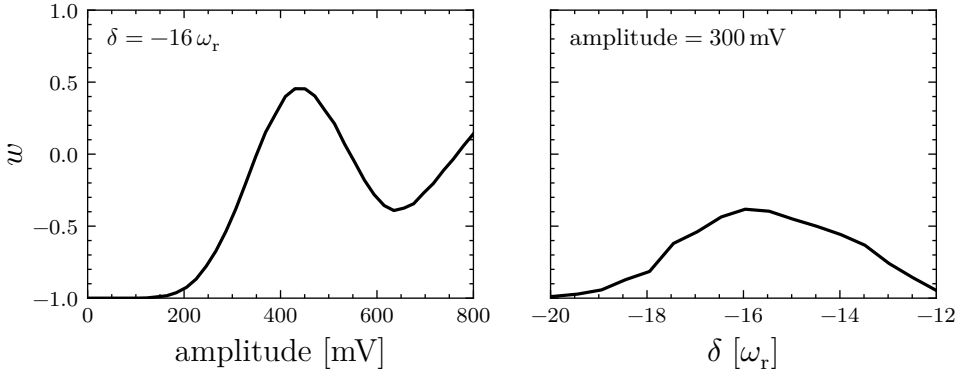


Figure 4.32: The measured population inversion as a function of the Bragg pulse amplitude and the Bragg two-photon detuning δ. Note the inversion here is defined as $w = (p_n - p_0)/(p_0 + p_n)$, where $p_0 = |\langle 0\hbar k|\phi\rangle|^2$ and $p_n = |\langle 2n\hbar k|\phi\rangle|^2$ in whatever frame puts the initial state at $n = 0$. The Bragg pulse amplitude is used to set the reference waveform amplitude for the Gaussian Bragg diffraction pulse, as described in Section 4.2.5. This is proportional to the effective Rabi frequency. Note that a downward 4th-order Bragg pulse is used, for which $\delta = -16\omega_r$.

the required $\Omega_{\text{eff},0}$ to drive a $\pi/2$ pulse. For values of σ above a certain threshold the interferometer contrast will start to decrease, as the momentum bandwidth of the Bragg pulse will drop below the momentum distribution of the atom cloud. Tuning the VS pulse duration can change where this threshold is. Generally, longer σ values correspond to smaller diffraction phases, as was seen in Fig. 2.6 and Fig. 2.7. Typically $\sigma = 20\,\mu s$ or $\sigma = 40\,\mu s$ is used in the experiment.

Optimal transverse alignment of the interferometry beam can be found by delaying the first Bragg diffraction time to $\sim$400 ms after launch, aperturing the detection beam, and maximizing the inversion by translating the interferometry beam. The larger transverse extent of the atom cloud at this delay makes the inversion more sensitive to translations. In certain cases, it can even be beneficial to optimize 3D MOT and RSC degrees of freedom on the delayed Bragg inversion signal, as the inversion is sensitive to the transverse temperature of the atom cloud.

4.5.2 Bloch oscillations

Several degrees of freedom must be tuned to implement Bloch oscillations, including the Bloch order N (yielding momentum transfer $2N\hbar k$), the value of Ω_{eff} during the Bloch oscillation pulse, and the Bloch duration τ (see Section 2.3). Additionally, the compensation light intensity or frequency must be tuned to produce interferometers with maximal contrast. The Bloch timing must also account for the presence of unwanted Bloch lattices, as discussed

in Section 4.4.2.

The values used for Ω_{eff} and τ are coupled. Higher Ω_{eff} creates a deeper lattice that can drive Bloch oscillations at shorter τ values. However, higher Ω_{eff} also results in more single photon scattering, which decoheres atom wavefunctions and lowers the contrast of the interferometer. For this reason it is useful to tune Bloch parameters on the full interferometer signal. Initially this should be done at short T (a few milliseconds), as at short T Bloch compensation light is unneeded. Once initial values for Ω_{eff} and τ have been selected, T can be increased and the compensation light intensity can be tuned to maximize the contrast. In practice $N = 100$ is used and the ramp duration is $\tau = 5.1\,\text{ms}$ (this corresponds to a ramp rate of $14kg$). For $N = 100$, Bloch oscillations require roughly 400 mW peak power. Higher Bloch orders are often not possible to achieve due to the junk lattice intersection described in Section 4.4.2.

The choice of Bloch parameters is constrained. To get higher Bloch efficiency, the lattice depth can be increased, or the ramp rate can be reduced. Both of these lead to higher single-photon scattering, which reduces interferometer contrast, canceling the benefits of higher Bloch efficiencies. Single-photon scattering can be mitigated by increasing the single photon detuning of the Ti:Sa, which can only happen with increased laser power. Currently, Bloch oscillations require more power than the Bragg and Raman pulses, and limit the size of the interferometry beam that can be used.

The Bloch oscillation pulses have a trapezoidal profile in time, as the linear ramp on and off is intended to adiabatically load and unload the atoms in the lattice. If the ramp is too short, then atoms are heated by the lattice. If the ramp is too long, then it will begin contributing to single photon scattering and reduce the contrast of the interferometer. In practice a ramp duration of 200 µs is used. The Bloch frequency sweep occurs during the flat region of the trapezoidal profile.

4.6 Developing a robust experiment

Several years of work went into making the experimental apparatus described in this chapter robust and capable of continuous operation. In order to achieve this robustness we attempted to run overnight experiments for several weeks, which provided immediate information about which experiment subsystems failed over long operation times. It also revealed that certain systems, such as the TAs and the Ti:Sa, became more stable after a day or two of operation.

The original experiment used an injection locked bare laser diode instead of the DBR laser to seed the TA cooling light. This lock was incredibly finicky and did not stabilize after long periods of continuous operation. This was likely due to its more complex optical setup and lack of active feedback. To simplify the cooling laser the phase-locking electronics were developed to produce the existing DBR cooling laser. The optical simplicity of this locking method makes it much more robust. The cooling laser system often runs for months at a time without needing to be tuned. The automatic relocking provided by the Raspberry Pi controlled PI loop was also necessary for continuous operation.

At this stage, the most unreliable part of the experiment became the ARTIQ control system, which would occasionally crash and could only be recovered from a power cycle. This was due to a combination of errors in how ARTIQ sequences had been programmed, as well as fundamental problems with the ARTIQ hardware overheating and the FPGA gateware. After fixing the overheating issue by modifying the ARTIQ board to include an additional fan, the entire control system codebase was rewritten to avoid the previously encountered errors and to make an ARTIQ version upgrade with gateware fixes possible.

At this stage the experiment has been successfully run for weeks at a time with minimal intervention, which enables the collection of large datasets for characterizing various systematic effects. The only intervention required is occasional tune ups of the 2D and 3D TA fiber couplings, which only takes a few minutes once a week. Several other practices help provide this robust performance:

- *Monitor the experiment where possible.* Record laser powers, laser frequencies, etc. to a centralized location, such as an InfluxDB database, when possible. This is often trivial to accomplish, as many experiment systems (i.e. digital laser diode drivers) are already operated by a computer.

- *Build modular optical systems.* Try to confine optical systems to a specific function. Build these systems on separate optical breadboards so that they can be moved to different locations in the lab without requiring full reassembly.

- *Add fibercoupling in strategic locations.* Some optical systems (such as lasers or amplifiers) have a much higher probability of failure than others (such as a mirror or a lens). For optical systems with a potentially high failure rate, add fibercouplers soon after to make replacement easier.

- *Only reflect s and p polarizations off of dielectric mirrors.* Dielectric mirrors are designed to conserve s and p polarization. If light is not purely s or p polarized, the reflected light will no longer have a purely linear polarization.

- *Fibercouple properly.* Ensure that fibercoupling of polarized light is performed with the optics described in Fig. 4.5 for optimal polarization stability. When outcoupling polarized light in a fiber, a PBS cube should be used to minimize downstream polarization drifts and to ensure that only s or p polarization reflects off of downstream dielectric mirrors.

Chapter 5

Experiment Characterization

In this chapter we will discuss how the uncertainty in our measurement of the cesium recoil frequency is determined. We will first determine what the maximum possible precision of the measurement could be. In the atom interferometer assuming that all interferometry processes are noiseless–i.e. there is no noise contributed by fluctuations in the atom cloud position/velocity distribution, the Bragg/Bloch pulse intensities, detection light intensity, etc.–the maximum precision is given by the amount of information provided by a single atom exiting the interferometer averaged over the total number of atoms. We compute this idealized precision, quantify the number of atoms in the experiment, and discuss several sources of noise that limit experiment precision.

5.1 Extracting ω_{r} from interferometer output

In Eq. (3.34) we computed the interference fringes w_+ and w_- in the upper and lower interferometers of an SCI accounting for contrast loss. Plugging the upper and lower interferometer phases from Eq. (3.19) into w_+ and w_- we find

$$w_- = b_- + A_- \cos(\phi_c + \phi_d) \qquad w_+ = b_+ + A_+ \cos(\phi_c - \phi_d) \tag{5.1}$$

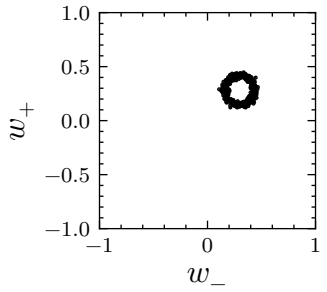


Figure 5.1: An experimentally recorded ellipse.

where A_-, A_+, b_-, and b_+ are determined by the contrast loss of the interferometer and

$$\phi_c = 2nT(T+T')kg \qquad \phi_d = 8n(N+n)T\omega_\mathrm{r} - nT\omega_\mathrm{m} \tag{5.2}$$

We are interested in measuring ϕ_d, as that will let us determine ω_r. We could scan ω_m, observe the $w_\pm$ fringes, and then fit each fringe to determine ϕ_d. Unfortunately, vibrations on the retro mirror of the experiment prevent us from doing this. These vibrations shift the mirror position between Bragg beamsplitter pulses, which shifts the laser phase imprinted on the diffracted wavepackets after each Bragg beamsplitter. If this phase is included in our phase calculation, we will see that it becomes part of the common mode phase ϕ_c, and not part of the differential phase ϕ_d. Because vibrations are different shot-to-shot, this means each experiment shot has an effectively random ϕ_c. Further effects stemming from vibrations are discussed in Section 7.1.2.

For this reason, we want a way of plotting $w_\pm$ so that we can extract ϕ_d without needing to worry about ϕ_c. You might have noticed that when w_- and w_+ are plotted parametrically as a function of ϕ_c, they describe an ellipse. Specifically, if we define

$$w'_- = (w_- - b_-)/A_- \qquad w'_+ = (w_+ - b_+)/A_+ \tag{5.3}$$

then the equation for an ellipse

$$1 = \left(\frac{w'_- + w'_+}{2\cos(\phi_d)}\right)^2 + \left(\frac{w'_- - w'_+}{2\sin(\phi_d)}\right)^2 \tag{5.4}$$

is satisfied. In Fig. 5.1 these quantities are plotted parametrically for experimental data, showing an ellipse. Plugging our expressions for $w'_\pm$ into this equation yields

$$0 = a_1 w_-^2 + a_2 w_- w_+ + a_3 w_+^2 + a_4 w_- + a_5 w_+ + a_6 \tag{5.5}$$

where

$$\begin{aligned}
a_1 &= \frac{\csc^2(2\phi_d)}{A_-^2}c \\
a_2 &= -\frac{2\cot(2\phi_d)\csc(2\phi_d)}{A_- A_+}c \\
a_3 &= \frac{\csc^2(2\phi_d)}{A_+^2}c \\
a_4 &= \frac{\csc^2(\phi_d)\sec^2(\phi_d)(A_- b_+ \cos(2\phi_d) - A_+ b_-)}{2A_-^2 A_+}c \\
a_5 &= \frac{\csc^2(\phi_d)\sec^2(\phi_d)(A_+ b_- \cos(2\phi_d) - A_- b_+)}{2A_- A_+^2}c \\
a_6 &= \frac{\csc^2(\phi_d)(A_+^2 b_-^2 + A_-^2 b_+^2 - 2A_- A_+ b_- b_+ \cos(2\phi_d))}{A_-^2 A_+^2}c - c
\end{aligned} \tag{5.6}$$

where c is an arbitrary constant and $c > 0$. This shows that when w_- and w_+ are plotted parametrically, including the effect of junk ports and contrast loss, the result is an ellipse described by the more general equation for a conic. Computing the distance between a point and an ellipse requires numerical root finding, and so we determine the parameters a_1 through a_6 through a fitting routine that relies on minimizing the algebraic distance between a set of points and the conic equation [83].

Solving this system of equations for the ellipse parameters yields

$$\begin{aligned}
\phi_d &= \frac{1}{2}\arccos\left(-\frac{a_2}{2\sqrt{a_1 a_3}}\right) \\
A_+ &= 2\sqrt{a_1}\sqrt{\frac{-a_2 a_4 a_5 + a_1 a_5^2 + a_2^2 a_6 + a_3(a_4^2 - 4a_1 a_6)}{(a_2^2 - 4a_1 a_3)^2}} \\
A_- &= 2\sqrt{a_3}\sqrt{\frac{-a_2 a_4 a_5 + a_1 a_5^2 + a_2^2 a_6 + a_3(a_4^2 - 4a_1 a_6)}{(a_2^2 - 4a_1 a_3)^2}} \\
b_+ &= -\frac{a_2 a_4 - 2a_1 a_5}{a_2^2 - 4a_1 a_3} \\
b_- &= -\frac{a_2 a_5 - 2a_3 a_4}{a_2^2 - 4a_1 a_3}
\end{aligned} \tag{5.7}$$

The first expression is used to convert a fit of the ellipse recorded by the interferometer to a measurement of the differential phase. Recently Andrew Christensen developed an efficient method for extracting the ellipse parameters using a Bayesian fit that accounts for various noise sources that can shift the measured populations. To develop this method, a probability distribution describing the probability of measuring a particular w_- and w_+ value was derived by considering points on an ellipse with Gaussian noise. The fit is performed by maximizing the likelihood of the observed w_- and w_+ points over the parameters of the distribution. Several numerical tricks make this fit method fast. This is now used as the primary ellipse fitting routine in the experiment, as it is subject to fewer biases than the standard algebraic minimization ellipse fitting routine.

In order to determine ω_r from the fitted value of ϕ_d we must be confident that ω_m is zeroing the phase to within 2π. As discussed in Section 3.1, this can be done by scanning T and ensuring that the ellipse phase remains constant. Once this has been done and ϕ_d has been determined via a fit, we solve for ω_r via

$$\omega_\mathrm{r} = \frac{\phi_d/nT + \omega_\mathrm{m}}{8(N+n)} \tag{5.8}$$

5.2 Projection noise limit

We want to compute the error on the measured fringes $\hat{w}_\pm$ in an SCI. Note that

$$\hat{p}_0 = \frac{N_0}{N_0 + N_{-n}} \tag{5.9}$$

where N_0 is the number of atoms detected in momentum state $0\hbar k$ at the end of the interferometer and N_{-n} is the number of atoms detected in momentum state $-2n\hbar k$ at the end of the interferometer. Each detected atom has collapsed into a single output momentum state and is an independent measurement. Provided a fixed total detected atom number $N_- = N_0 + N_{-n}$, N_0 is drawn from a binomial distribution with $N_0 + N_{-n}$ trials and a probability of success given by p_0 from Eq. (3.17). Therefore the variance of $\hat{p}_0$ is given by

$$\mathrm{var}(\hat{p}_0) = \frac{\mathrm{var}(N_0)}{(N_0 + N_{-n})^2} = \frac{p_0(1-p_0)}{N_0 + N_{-n}} \tag{5.10}$$

The observed fringe variance is given by

$$\mathrm{var}(\hat{w}_-) = \mathrm{var}\left(\frac{N_0 - N_{-n}}{N_0 + N_{-n}}\right) = \mathrm{var}\left(2\hat{p}_0 - 1\right) = \frac{4p_0(1-p_0)}{N_0 + N_{-n}} \tag{5.11}$$

This gives the observed fringe variance in terms of the probability of an atom exiting the lower interferometer in momentum state $0\hbar k$. We can rewrite this in terms of Eq. (3.18), our equation for the inversion

$$\mathrm{var}(\hat{w}_-) = \frac{1 - w_-^2}{N_0 + N_{-n}} \tag{5.12}$$

The equivalent result is found for the upper interferometer, and so we find the standard error on each fringe is

$$\sigma_{w_\pm} = \sqrt{\frac{1 - w_\pm^2}{N_\pm}} \tag{5.13}$$

where $N_\pm$ is the number of atoms in the lower or upper interferometer, as appropriate. Note that in deriving this result we assumed perfect contrast and no offset. Defining $\phi_\pm = \phi_c \mp \phi_d$, we compute

$$\frac{\partial w_\pm}{\partial \phi_\pm} = -\sin(\phi_\pm) \tag{5.14}$$

Therefore the square of the standard error on $\phi_\pm$ is

$$\sigma^2_{\phi_\pm} = \left|\frac{\partial w_\pm}{\partial \phi_\pm}\right|^{-2} \sigma^2_{w_\pm} = \frac{1 - w_\pm^2}{N_\pm \sin^2(\phi_\pm)} = \frac{1}{N_\pm} \tag{5.15}$$

From this result we can compute the error on the differential phase ϕ_d by noting $\phi_d = (\phi_- - \phi_+)/2$

$$\sigma_{\phi_d} = \sqrt{\frac{\sigma_{\phi_-}^2 + \sigma_{\phi_+}^2}{4}} = \sqrt{\frac{1}{4N_-} + \frac{1}{4N_+}} \tag{5.16}$$

For a well-tuned $\pi/2$ beamsplitter an equal number of atoms exit the upper and lower interferometers, so $N_- = N_+ = N_{\text{atoms}}/2$, leading to

$$\sigma_{\phi_d} = \frac{1}{\sqrt{N_{\text{atoms}}}} \tag{5.17}$$

This is the shot-noise limit, or quantum projection noise limit, for the atom interferometer. Since $\phi_d = 8n(N+n)T\omega_{\text{r}} - nT\omega_{\text{m}}$, the error in ω_{r} is

$$\sigma_{\omega_{\text{r}}} = \frac{1}{8n(n+N)T\sqrt{N_{\text{atoms}}}} \tag{5.18}$$

For a single interferometer shot with N_{atoms} total atoms, this is the lowest possible error we can obtain on ω_{r}. It assumes unity detection efficiency, no contrast loss, and no other noise sources. Detection inefficiencies, contrast loss, and other sources of noise (such as detection noise, pulse instability, changes in the atom cloud distribution shot-to-shot, etc.) mean that the actual error obtained per-shot in the interferometer is higher.

To characterize sources of experiment noise, the experimenter can determine how the error scales with various parameters, such as the number of atoms in the experiment or the number of shots over which ω_{r} is averaged. If the experimenter sees, for example, that the error does not scale with $1/\sqrt{N_{\text{atoms}}}$, this is a sign that another noise source is dominating over the shot-noise.

5.3 Atom number measurements

To determine how closely the experiment approached the theoretical precision limit in Eq. (5.17), the total number of atoms propagating through the interferometer had to be determined. Both absorption- and fluorescence-based methods were used to measure the atom number.

5.3.1 Absorption measurements

To measure the atom number in the MOT, an absorption measurement is performed similarly to the one described in [84]. The 3D MOT is loaded for the standard 1.15 s, after which the repump light is turned off using AOM12, causing the atom cloud to accumulate in the $F = 3$ state. 800 µs later the cooling light is turned off. 400 µs later, a single "probe" beam

tuned to the $F = 3 \rightarrow F' = 4$ transition is turned on to illuminate the atoms for 100 µs. From the $F' = 4$ state atoms can decay to either $F = 3$ or $F = 4$ with a known per-atom probability. By measuring the light absorbed by the probe beam, an accurate atom number can be determined.

We sum over the transition strengths of each possible decay path to compute its relative weight

$$W_{F'=4,m'_F\rightarrow F} = \sum_{q=-1,0,1} \sum_{m_F=-F}^{F} \left|\langle F, m_F|er_q|F', m'_F\rangle\right|^2 \tag{5.19}$$

The decay probability of each path is then given by

$$P_{F'=4,m'_F\rightarrow F} = \frac{W_{F'=4,m'_F\rightarrow F}}{W_{F'=4,m'_F\rightarrow F=3} + W_{F'=4,m'_F\rightarrow F=4}} \tag{5.20}$$

The matrix element can be reduced using the identity given by [50]

$$\langle F, m_F|er_q|F', m'_F\rangle = \langle J\|e\mathbf{r}\|J'\rangle\,(-1)^{m_F+J+I}\sqrt{2F+1}\begin{pmatrix} F' & 1 & F \\ m'_F & q & -m_F\end{pmatrix} \tag{5.21}$$

$$\times\sqrt{(2F'+1)(2J+1)}\begin{Bmatrix} J & J' & 1 \\ F' & F & I\end{Bmatrix} \tag{5.22}$$

where the round-bracketed expression is the Wigner 3j symbol and the curly-bracketed expression is the Wigner 6j symbol. Evaluating the probabilities for the cesium D_2 line, where $I = 7/2$, $J = 1/2$, and $J' = 3/2$, yields

$$P_{F'=4,m'_F\rightarrow F=3} = \frac{5}{12} \qquad P_{F'=4,m'_F\rightarrow F=4} = \frac{7}{12} \tag{5.23}$$

Note that these probabilities are independent of the excited m'_F state. The probe beam is tuned to the $F = 3 \rightarrow F' = 4$ transition, so we are interested in computing how many photons an atom will scatter on average before falling into the $F = 4$ dark state (ignoring off-resonant excitation, etc.).

$$P\,(n \text{ scatters until dark}) = P^{n-1}_{F'=4\rightarrow F=3}P_{F'=4\rightarrow F=4} \tag{5.24}$$

The expected number of photon scatters per atom is therefore

$$\text{photon scatters per atom} = \sum_{n=1}^{\infty} nP\,(n \text{ scatters until dark}) = \frac{1}{P_{F'=4\rightarrow F=4}} = \frac{12}{7} \tag{5.25}$$

If the number of photons absorbed from the probe beam is known, then this factor can be used to determine the number of atoms present.

We begin by positioning a 7.8 mm $1/e^2$ diameter probe beam to intersect the MOT, and aligning a Thorlabs PDA100A2 photodiode on the beam exit. A 25 mm lens is used to

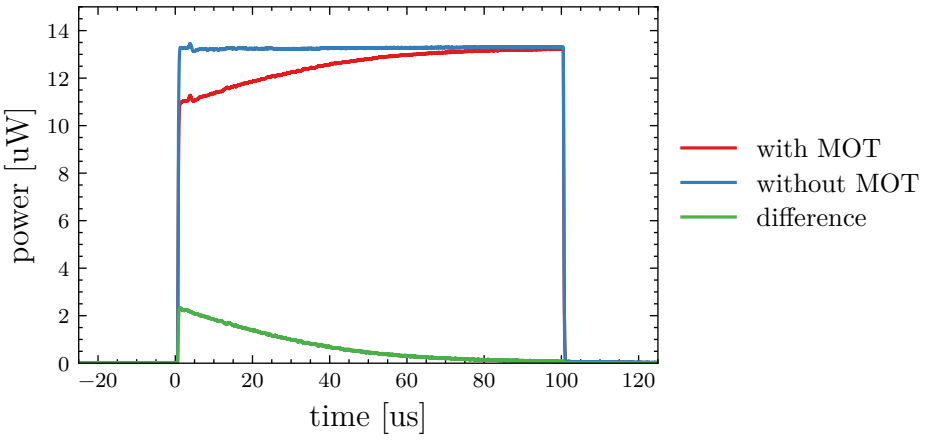


Figure 5.2: Optical power of the probe beam detected on a Thorlabs PDA100A2 photodiode as a function of time. When a MOT has been loaded, there is a dip in the detected power due to absorption. By taking the difference of the cases without and with the MOT and integrating, the total optical energy absorbed by atoms in the MOT is measured.

focus the beam onto the active area of the photodiode. The photodiode is then calibrated using laser light of a known power, and then the absorption sequence is run. The resulting photodiode trace is shown in Fig. 5.2, where we can see that without the MOT the detected power in the probe beam is constant, and with the MOT the detected power increases and then asymptotes to the power without the MOT. The difference between the two shows the power absorbed by the atoms in the MOT. This is integrated to find the total absorbed energy, divided by the energy of a photon at 852 nm to determine the number of absorbed photons, and then divided by the expected number of scatters to determine the atom number.

A second measurement is performed by using a camera to measure the probe beam absorption rather than a photodiode. The same sequence is performed, and the camera is exposed for the duration of the 100 µs probe beam pulse. The camera images with and without the MOT, as well as the difference between the two images, are shown in Fig. 5.3. The camera detection efficiency is calibrated using a laser beam of known power, and the image can be integrated to determine the absorbed energy and therefore the atom number. This yields a value around 31% higher than the photodiode measurement for an unknown reason. Neglecting the expansion of the probe beam as it propagates after passing through the MOT, this reveals that the atom cloud in the MOT is around 1 mm by 2 mm in the plane orthogonal to the probe beam. Measurements of the atom number after the molasses launch were performed, finding the launch to be 83% efficient.

When the system is well tuned, the camera method measures an atom number of ∼440 million atoms. This is within the expected range for a 3D MOT. This approach makes an

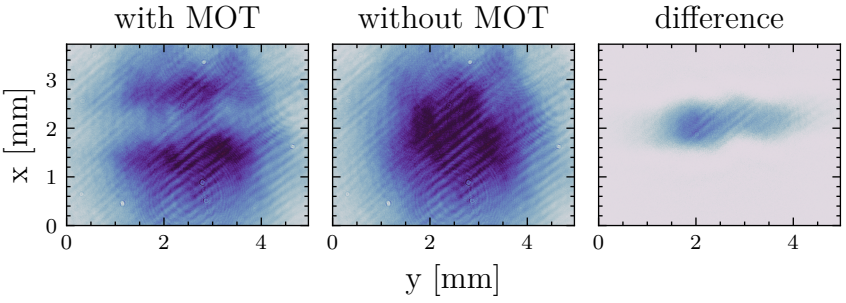


Figure 5.3: Images recorded by a Thorlabs Zelux CMOS camera by exposing the probe beam with and without a 3D MOT. The total optical power absorbed by the 3D MOT is computed by integrating over their difference. The difference also reveals the size of the 3D MOT projected onto a plane perpendicular to the propagation of the probe beam. The position units specify pixel positions on the camera.

accurate measurement of the atom number provided that the photodiode/camera measurements are accurate.

5.3.2 Fluorescence measurements

It is significantly harder to produce an accurate fluorescence-based measurement of the atom number than an absorption-based one, as a fluorescence-based measurement requires modeling the detection apparatus and fluorescence rate accurately, as well as calibrating whatever detector is used. However, it is difficult to perform absorption measurements at all stages of state preparation as 1) few atoms remain after certain stages, such as VS, and 2) the atoms are moving. Here we discuss how fluorescence measurements are made in the detection region.

The detection system uses a Thorlabs FDS100 photodiode, which has a sensitivity of 0.55 A/W at 852 nm. The photodiode current is fed into a transimpedance amplifier and a 1 GΩ feedback resistor, resulting in a power sensitivity of 1.8 nW/V. A line filter ensures that the photodiode only detects 852 nm photons, meaning that the number of photons detected per volt-second is 7.7×10^9 photons/Vs. Integrating the photodiode trace and converting from volt-seconds to photons using this conversion factor produces the number of detected photons.

Computing an atom number from the number of detected photons requires an estimate of the photon collection efficiency η and the total number of photons scattered by a single atom. η is estimated to be $\eta = 0.022$ by considering the detection optics described in Section 4.1.3. The detection beam is roughly 3 mm in height, and atoms travel through it at 4 m/s, meaning

that atoms fluoresce for 0.75 ms. The saturated on-resonance scattering rate is $\Gamma/2$, which for cesium is 16.4×10^6 photons/s. Thus, each atom should scatter approximately ∼12300 photons, ∼271 of which should be collected by the detection optics and detected on the photodiode.

Falling molasses launch, RSC, and VS traces taken from the same experiment run are displayed in Fig. 4.19, showing the relative atom number in each trace. Here we do not assume saturation, but compute the expected scattering rate from the detection laser detuning. Note that the observed atom number after falling molasses launch is only around 10 million, around a factor of 30 less than what was measured by absorption after the molasses launch. The reduction is partially caused by transverse atom cloud expansion, which causes detected atoms to fall outside of the active area of the detection photodiode. It is likely that the remaining loss is due to collisions with background gas particles in the vacuum.

Because RSC provides additional cooling in the transverse direction, the detected atom number on falling RSC is higher than the detected atom number on falling molasses launch. After double VS, the atom number is reduced by a factor of approximately 70, leaving 290,000 atoms (half of which are lost to junk ports). Plugging this into Eq. (5.18) we estimate a sensitivity of $\sigma_{\omega_r}/\omega_r = 610$ ppt per shot for an interferometer with $n = 4$, $N = 100$, $T = 100$ ms. In practice, Bloch significantly reduces the atom number in each output port of the interferometer, junk ports reduce the number of detected atoms actively contributing to interference, and various mechanisms leading to contrast loss reduce the actual sensitivity further. This leads to a substantially lower per-shot sensitivity.

5.4 Time-averaging

In the ideal case, averaging together multiple experiment shots should improve the measurement precision by $\sqrt{N_{\text{shots}}}$. This is because, in the ideal case, running two experiment shots is equivalent to just doubling the atom number in Eq. (5.18) (ignoring any subtleties with the ellipse fit). However, in practice shot-to-shot drifts in the experiment apparatus can reduce the size of this improvement. To characterize the measurement precision and the size of these drifts, an Allan deviation can be plotted.

There are many drifts that could conceivably cause suboptimal averaging-down rates, such as drifts in the atom cloud (position, velocity distribution, atom number, etc.) and drifts in the interferometry pulses (fiber coupling drifts leading to power fluctuations in detection or beamsplitter intensity, etaloning of the interferometry beam in the vacuum chamber, etc.). One method of characterizing the ellipse noise is to operate the interferometer in an open configuration by delaying the final Bragg pulse to prevent the interferometer from closing. This is shown in Fig. 5.4, where interferometers with several values of T are operated in closed and open configurations. This enables the noise stemming from processes unrelated to interference to be isolated. The standard deviation of the open ellipse points in each direction is $\sigma_{w_-} = 1.5 \times 10^{-2}$ and $\sigma_{w_+} = 1.1 \times 10^{-2}$.

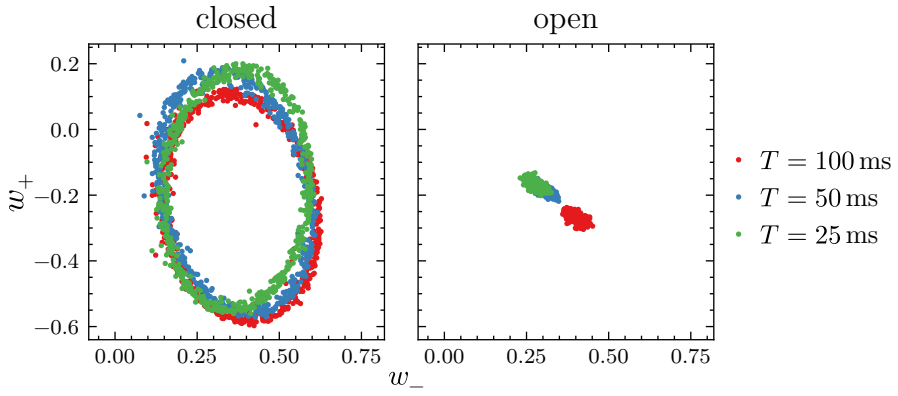


Figure 5.4: *Left:* Ellipses recorded by an SCI for several values of T. *Right:* An experiment run in an identical configuration to the left panel, but with the final beamsplitter pulse delayed so that interference does not occur. Taken together, these datasets help determine which noise sources are due to phase-sensitive processes and which are phase-insensitive.

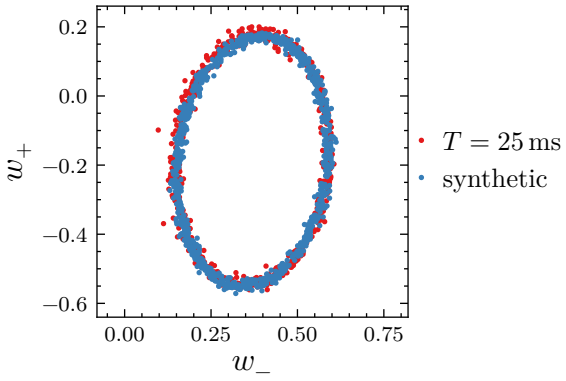


Figure 5.5: A synthetically generated ellipse overlaid on top of the $T = 25\,\mathrm{ms}$ ellipse in Fig. 5.4. Gaussian noise is injected into the synthetic ellipse to mimic the shot-to-shot variation in the single Bragg pulse inversion.

This noise is largely explained by the shot-to-shot variation of the single Bragg pulse inversion, which is measured to be $\sigma_w = 0.9 \times 10^{-2}$. A synthetic ellipse generated with this noise level on both w_- and w_+ is shown in Fig. 5.5. The synthetic ellipse is found to have half the noise of an actual ellipse (determined by characterizing uncertainty in the ellipse fit, not by eyeballing the figure), implying that the shot-to-shot Bragg beamsplitter variation is responsible for a substantial portion of the total ellipse noise.

The shot-to-shot Bragg beamsplitter variation is higher than the shot-noise limit. Several attempts were made to determine the source of this noise, including variations in the Bragg

beamsplitter pulse peak intensity, the atom number, and the background on the recorded fluorescence signal. None of these were found to explain the observed variation.

5.5 Current experiment precision

Actual measurements of ω_r are performed at a variety of interferometer durations T as certain systematic effects are independent of T. When the differential phases of these ellipses are plotted as a function of T, they should form a line with slope $n(8(n+N)\omega_r - \omega_m)$ and intercept $\phi_{\text{systematic}}$. This is easily seen by restating the expression for the differential phase with the systematic shift included.

$$\phi_d = n(8(n+N)\omega_r - \omega_m)T + \phi_{\text{systematic}} \tag{5.26}$$

In this way a linear fit can be used to determine the value of ω_r, which is the only unknown in the slope. Of course, not all systematics are independent of T, and so $\phi_{\text{systematic}}$ must be Taylor expanded in T to distinguish the T-dependent and T-independent parts.

Currently the experiment achieves a 24-hour precision of 400 ppt in ω_r, or 200 ppt in the fine-structure constant, when averaging across different values of T. In these data runs, the longest duration interferometer uses $T = 100\,\text{ms}$, $n = 4$, and $N = 100$ to accumulate 8.6 Mrad of SCI phase (see Eq. (3.16)). The 400 ppt in ω_r corresponds to a phase error of 3.5 mrad in this condition. Therefore, systematic effects must be understood to this level to report a measurement. Characterizing these systematics is largely the responsibility of simulations matched to experimental data, as discussed in Chapter 6. We use these simulations to accurately predict experiment output and to compute systematic shifts in Chapter 7.

We can optimistically estimate possible performance of the interferometer by assuming that the shot-to-shot Bragg deviation is fixed (reducing the fringe standard deviation by a factor of two), improvements to Bloch result in three times as many atoms in the interferometer (a factor of $1/\sqrt{3}$), T can be increased from 100 ms to 120 ms (a factor of 5/6), and the time required per shot decreases from 2.7 s to 2.5 s (a factor of $\sqrt{2.5\,\text{s}/2.7\,\text{s}}$). We also consider the idealized performance given by *Eq.* (5.18) assuming that the per shot sensitivity of the system is shot noise limited (i.e. no detector noise, etc.), and that there is no time-dependent drift in the measured recoil. For this case we take $N_{\text{atoms}} = 150000$, $n = 4$, $N = 100$, $T = 120\,\text{ms}$, and an experiment shot duration of 2.5 s. This estimate ignores contrast loss effects. These estimates, along with our experimentally observed precision, are given in Table 5.1.

Table 5.1: Comparison of current experiment performance with possible (but optimistic) and idealized performance.

Performance metric	**Current**	**Possible**	**Idealized**
Per-shot sensitivity	72 ppb	16 ppb	500 ppt
24-hour precision in ω_r	400 ppt	90 ppt	3 ppt
24-hour precision in α	200 ppt	45 ppt	1.5 ppt

Chapter 6

Simulations

To understand various systematic effects present in the experiment, a simulation is needed. In the 2018 Berkeley fine-structure constant measurement, systematic effects were considered in independent calculations and simulations. In Section 6.2.1, I develop a tree-based simulation method capable of tracking most known systematic effects in the experiment simultaneously. This is valuable because certain systematic effects are not independent of each other (this is described in Section 7.1), and because it allows all experiment data to be compared against a single unified simulation. The tree-based simulation generalizes the simulation method developed for the 2018 Berkeley fine-structure constant measurement and the method developed by Zachary Pagel [24], making it straightforward to include novel systematic effects and to simulate different interferometer geometries. For example, the tree-based method makes it trivial to simulate parasitic trajectories, like those shown in Fig. 6.1. In Section 7.1.5, Section 7.1.6, and Section 7.1.7 we compare predictions made with the tree-based simulation to measured experiment output, showing agreement. The inclusion of all systematic effects in a single simulation and the extensive verification of the simulation against the experiment significantly improve on the simulation capabilities of the previous generation fine-structure constant measurement.

In Section 6.1 we will discuss methods for simulating single matterwave optics. We will then use these methods to simulate entire interferometers in Section 6.2. The development of one of these methods resulted in a paper [26]. We will end with a discussion of how spatial separation is encoded in the momentum-space wavefunction in Section 6.3.

To begin, we will provide an example of how a single matterwave optic can be simulated. Consider the Bragg diffraction equations of motion in Eq. (2.19), which describe couplings between momentum states spaced by $2\hbar k$. Let's say we want to simulate 4th-order Bragg diffraction from the state $|0\hbar k\rangle$ to $|8\hbar k\rangle$. This is straightforward to do: we define our initial state as

$$|\phi(t=0)\rangle = |0\hbar k\rangle \tag{6.1}$$

and integrate Eq. (2.19) using a standard Runge-Kutta type solver. It is helpful to perform the simulation with dimensionless units by taking $t \to \tau\omega_{\mathrm{r}}^{-1}$, which allows us to rewrite

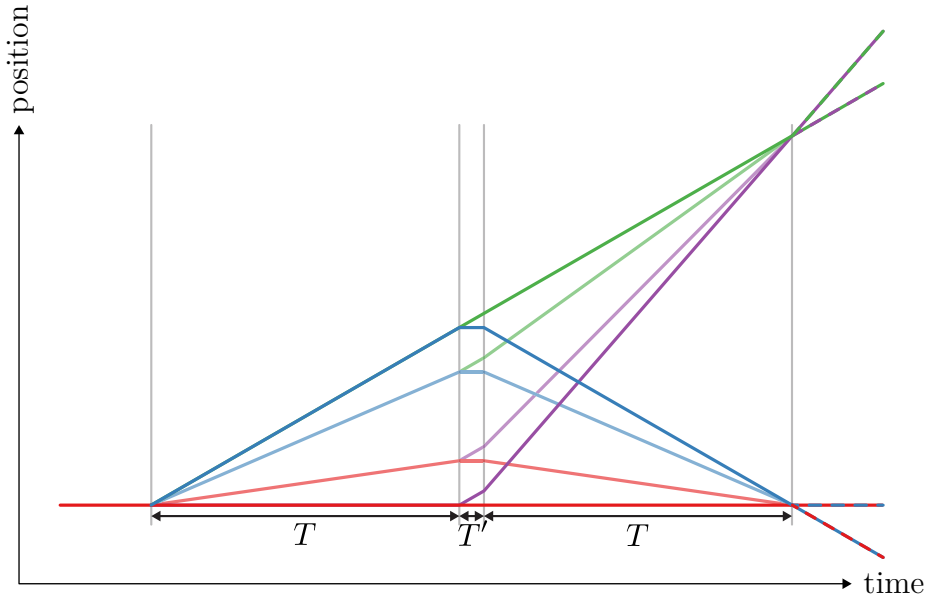


Figure 6.1: An SCI composed of 4th-order Bragg pulses. Parasitic trajectory paths propagate through the interferometer and interfere at the time of the final beamsplitter pulse. These are shown as partially transparent lines. These parasitic momentum states can be produced during Bragg diffraction pulses.

Eq. (2.19) as

$$i\hbar\dot{\tilde{c}}_j(\tau) = -\hbar\frac{\Omega_{\text{eff}}(\tau)\omega_{\text{r}}^{-1}}{2}\left[\exp(i(\delta\omega_{\text{r}}^{-1}-4-8j)\tau)\tilde{c}_{j+1}(\tau)+\exp(i(-\delta\omega_{\text{r}}^{-1}-4+8j)\tau)\tilde{c}_{j-1}(\tau)\right] \tag{6.2}$$

Now instead of defining Ω_{eff} and δ in usual time units, we define them in terms of ω_r. We integrate Eq. (6.2) with $\Omega_{\text{eff}}(t)$ given by Eq. (2.20), $\sigma = 0.259\omega_{\text{r}}^{-1}$, $\Omega_{\text{eff},0} = 19.406\omega_{\text{r}}$, and $\delta = 16\omega_{\text{r}}$ from $t = 0$ to $t = 6\sigma$. The result of this integration is displayed in the second from the top panel of Fig. 2.3, where we see a $\pi/2$ pulse transfer population from the $|0\hbar k\rangle$ state to the $|8\hbar k\rangle$ state. All of the differential phase calculations in Chapter 2 use this same integration approach. When we integrate a process with a time-dependent phase (i.e. Bloch oscillations), we take $\delta\tau \to \delta\tau + \theta(t)$.

The states described by Eq. (6.2) are plane waves in position space. That is, they are described by wavefunctions of infinite extent in position space and delta functions in momentum space. This can be seen by computing the position space wavefunction of some

state $|2j\hbar k\rangle$ [85]

$$\psi(x) = \langle x|2j\hbar k\rangle = \frac{1}{\sqrt{2\pi\hbar}}\exp(i2kjx) \tag{6.3}$$

Because the states are plane waves, they never spatially separate. As discussed in Section 3.2, the actual states in the interferometer are described by wavepackets with finite spatial extent. Clearly, while integrating Eq. (6.2) allows us to simulate single matterwave optics, it does not allow us to simulate full interferometers as it does not capture spatial separation. The third panel in Fig. 2.3 simulates spatial separation by simulating a complete wavepacket instead of plane wave momentum states (see Section 6.2.2). Our simulations must capture the effect of spatial separation to properly describe the interferometer.

6.1 ODE integration methods

There are many standard algorithms for integrating Eq. (6.2). In the past a Crank-Nicolson-like time-differencing scheme was used to integrate Eq. (6.2) in time [24]. The method was chosen as it preserved unitarity. This was a poor choice of integration method, as it did not use an adaptive step size and required very small steps for the phase of the output wavefunction to converge. A discussion of the phase error resulting from this integration method can be found in [86]. A better choice is a Runge-Kutta-type method. Currently we use the Dormand-Prince eighth-order method [87] using the `solve_ivp` function from [88]. The system of ODEs is defined using a Numba-compiled function for faster evaluation [89]. The Dormand-Prince eighth-order method was found to produce higher accuracy than the Crank-Nicolson method in far less time. This is consistent with the approaches used for integrating the Bragg equations of motion in the literature [46, 47]. Here we discuss two methods for integrating Eq. (6.2) quickly for an ensemble of atoms.

6.1.1 Lookup tables

In certain cases, such as integrating a cloud of atoms propagating through the interferometer, it is possible to attain a further speedup through the use of a lookup table. For a cloud of atoms propagating through a single interferometer shot, each atom will have a distinct $\Omega_{\text{eff}}(\tau)$ based on its position at each Bragg beamsplitter and a distinct δ based on its vertical velocity. By interpolating over a precomputed grid of possible $\Omega_{\text{eff}}(\tau)$ and δ values, we can greatly speed up the simulation. In this case, $\Omega_{\text{eff}}(\tau)$ will take the form

$$\Omega_{\text{eff}}(\tau; x, y) = \begin{cases} \Omega_{\text{eff},0}(x,y)\exp\left(-\frac{(\tau-\tau_0)^2}{2\sigma^2}\right), & |\tau-\tau_0| < 3\sigma \\ 0, & \text{otherwise} \end{cases} \tag{6.4}$$

This is similar to Eq. (2.20), except that the atom location is used to compute the peak effective Rabi frequency $\Omega_{\text{eff},0}(x,y)$. In a laser beam $\Omega_{\text{eff},0}(x,y)$ would be proportional to

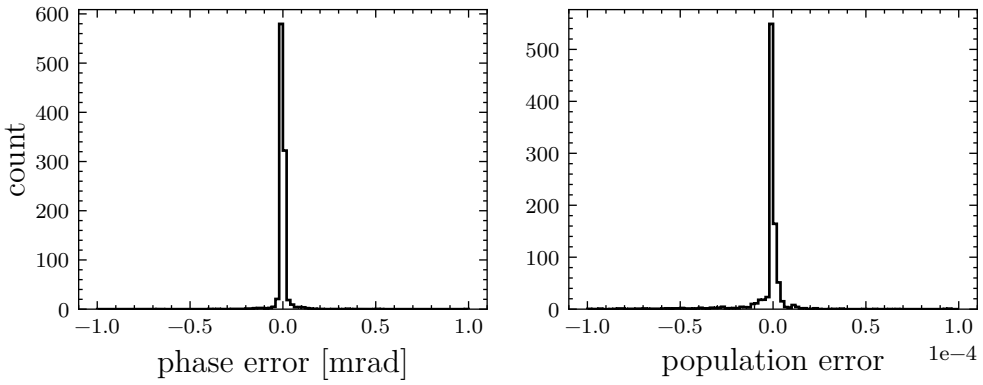


Figure 6.2: Histogram of the error on interpolated Bragg beamsplitter integrations with respect to direct Bragg beamsplitter integrations. Phase error is computed as $\arg(\phi_{\text{integrated}}/\phi_{\text{interpolated}})$, and population error is computed as $|\phi_{\text{interpolated}}|^2 - |\phi_{\text{integrated}}|^2$.

the transverse beam intensity profile. This means that $\Omega_{\text{eff},0}(x, y)$ is bounded between 0 and some peak beam intensity. For an atom moving with velocity v relative to the Bragg beamsplitter resonance, the two-photon detuning δ will take the form

$$\delta = 4\,(2n_0 + n)\,\omega_\text{r} + 4\omega_\text{r} v/v_\text{r} \tag{6.5}$$

where $v_\text{r} = \hbar k/m$ is the recoil velocity. We integrate Eq. (6.2) over a 2D grid of $\Omega_{\text{eff},0}$ and δ values, and use the SciPy `RegularGridInterpolator` with `method="cubic"` to perform the interpolation. In practice this 2D grid ends up being around 10 GB in size. Separate lookup tables must exist for single and multifrequency Bragg beamsplitters. A computer with sufficient RAM must be used so that both lookup tables fit in computer memory. The grid spacing must be selected to produce a sufficiently low error on the interpolated wavefunction.

The interpolation accuracy is tested by selecting random points in the 2D parameter space and comparing the interpolated values of those points with direct computation via the Dormand-Prince method. In Fig. 6.2 we show the differences between these values for a Bragg beamsplitter with Gaussian width $0.259\omega_\text{r}^{-1}$. The lookup table grid contains 800 points from $\Omega_{\text{eff},0} = 0\omega_\text{r}$ to $45\omega_\text{r}$ and 712 points from $\delta = 4\omega_\text{r}$ to $28\omega_\text{r}$. The spacing of the δ points is nonlinear, with more points close to the 4th-order Bragg pulse resonance condition and fewer points away. We sample 1000 points in the 2D parameter space and compute the phase and population errors. The phase error plot reveals that the lookup table has well under a milliradian of error for most points. The population error reveals that most population errors are less than a part in 10000. Since an SCI is composed of four Bragg beamsplitter pulses, it is reasonable to estimate a phase error of under a milliradian when using the lookup table.

The disadvantage of the lookup table approach is that scanning certain parameters, for example the beamsplitter pulse width σ, requires computing an entirely new lookup table.

For multifrequency beamsplitters, the modulation frequency ω_{m} must also be fixed, as must the phase of the modulation beatnote (this is the phase of the cosine term in Eq. (2.36)). It is possible to fix the timings and modulation frequency experimentally, so it is possible to compute lookup tables that are properly matched to the experiment. Lookup tables are used whenever possible. Once the Bragg order n, Bloch order N, and Bragg pulse time profile has been selected, several T values are chosen to be simulated (selecting several values of ω_{m}). Precomputed tables are generated for each T value and can then be used to simulate various effects, such as changing atom cloud parameters or different spatial beam profiles. Tables can also be used to fit these parameters to observed experiment data.

6.1.2 Parallelized ODE method

When the lookup table method cannot be used and an ensemble of atoms needs to be integrated through the interferometer, we can use a parallelized Runge-Kutta type method. For a fixed step size Runge-Kutta method, this is straightforward. The wavefunction of each atom in the ensemble is stacked into a single array, and every atom is stepped forwards in time on a shared time grid. Because the evolution of each atom is independent of every other atom, the per-atom update for each step can be vectorized (meaning it is evaluated in parallel across available CPU cores).

To make this method work with an adaptive step size, we set the step size for all atoms using the adaptive step size computed for the worst-case atom. The Dormand-Prince 5(4) method is used for this integration [90]. It advances the solution with the 5th-order estimate and uses the embedded 4th-order estimate to provide an estimate of the truncation error. This estimate is usually used to control the adaptive step-size of the integrator. To use the adaptive step size for the entire ensemble of atoms at once, the worst-case truncation error is computed across all atoms

$$\mathrm{err} = \max_{i,j} \left| \tilde{c}_{i,j}^{(5)} - \tilde{c}_{i,j}^{(4)} \right| \tag{6.6}$$

where $\tilde{c}_{i,j}^{(5)}$ and $\tilde{c}_{i,j}^{(4)}$ are the 5th and 4th-order solutions for $\tilde{c}_j(t)$ for atom i in momentum state $2j\hbar k$. This worst-case error is used to set the step size using the controller

$$\Delta t_{\mathrm{new}} = \Delta t \cdot \mathrm{clamp}\left(0.9 \left(\frac{\mathrm{tol}_{\mathrm{step}}}{\mathrm{err}} \right)^{1/5}, 0.1, 10 \right) \tag{6.7}$$

where the exponent $1/5$ follows from the 4th-order error estimate, and the clamping prevents the step size from changing by more than a factor of ten in either direction. When the error estimate exceeds the specified tolerance, smaller step sizes are attempted until the error estimate is sufficiently low. The integration loop is implemented using Numba, minimizing python overhead.

I found this integration scheme using a custom LLM-in-a-loop environment. The loop I built was inspired by the AlphaEvolve algorithm discovery and optimization coding agent

Table 6.1: Per-atom integration times for the DOP853, parallelized ODE, and lookup table methods.

Atoms	DOP853	Parallelized	Lookup table
1	25.0 ms	44.0 ms	28.1 µs
10	15.5 ms	44.9 ms	3.7 µs
100	15.6 ms	4.9 ms	0.8 µs
1000	15.6 ms	0.8 ms	0.5 µs
10000	15.5 ms	0.4 ms	0.5 µs

[91]. I built an environment in which an LLM could propose new algorithms for integrating an ensemble of atoms through an SCI. The proposed algorithm was used to propagate a randomly sampled atom cloud through an SCI. The result was verified against the known solution computed using the Dormand-Prince solver. Proposed algorithms were promoted if they 1) had numerical error below the tolerance of the Dormand-Prince solver; and 2) had runtimes no more than 50% slower than the leading proposed algorithm. Otherwise the algorithm was rejected. After the algorithm was either promoted or rejected, a new LLM session either iterated on a promoted algorithm to produce further speedups, or proposed a new algorithm. This loop was run for roughly 24 hours.

A wide variety of integration methods were implemented during this process, including several other promising methods. However, the method described above was simple, and its implementation was straightforward to check by hand. The method seems to be equivalent to the `EnsembleGPUKernel` method in the Julia `DiffEqGPU.jl` library, which was proposed in [92] in 2024. It is likely the LLM did not invent this method, but rather ported it from [92].

A timing comparison between the eighth-order Dormand-Prince method (the SciPy `DOP853` integrator), the parallelized method, and the lookup table method is provided in Table 6.1. The table displays the effective per-atom integration time. We see that for an ensemble of 10,000 atoms the parallelized ODE method is roughly 40 times faster than the `DOP853` integrator, and the lookup table is roughly 800 times faster than the parallelized ODE method. Clearly the lookup table method is best to use when possible. The parallelized ODE method is useful to use when the previously mentioned lookup table limitations apply. To compute this table, an Apple M3 Pro chip was used. The reported times will vary depending on the particular system used and the exact implementation in code.

6.2 Full interferometer simulations

6.2.1 Tree approach

As mentioned earlier, the states described by Eq. (6.2) are plane waves that extend throughout all space. Since atom interferometers are described by wavepackets with finite spatial

extent, we must adapt the plane wave calculation so that it is able to represent spatial separation. Recall from Section 3.2.2 that between beamsplitter pulses, wavepackets travel along the trajectories that would be expected from classical mechanics. This provides a straightforward prescription for simulating complete atom interferometers: integrate Eq. (6.2) to compute the population transfer and phase shift introduced by each beamsplitter; track the position of each state classically between beamsplitters; at the end of the interferometer construct the final wavefunction by considering which states end in the same position and momentum state.

In this picture, the interferometer is naturally described using a directed tree structure. Nodes of the tree describe different states (containing the state position, momentum, amplitude, and phase), and tree edges describe different operations applied to the interferometer. Nodes contain a position z_j, a momentum p_j, and a time t_j, where j indexes the node. Beamsplitter operations create child nodes with the same position and time as the parent node, but different momenta. Free evolution creates a single child node with the same momentum as the parent node, but with a different position and time.

The initial state of the wavefunction is represented by the root node. We consider the case where the initial state is described by some wavepacket $\left|\psi_{z_0,p_0}(t=t_0)\right\rangle$, which describes a wavepacket with position and momentum expectation values z_0 and p_0 at time t_0. Applying a Bragg beamsplitter to node j yields

$$U_{\mathrm{BS}}\left|\psi_{z_j,p_j}(t_j)\right\rangle = \sum_{m=-\infty}^{\infty} a_{\mathrm{BS},m}(z_j,p_j,t_j)\left|\psi_{z_j,\,p_j+2m\hbar k}(t_j)\right\rangle \tag{6.8}$$

where $a_{\mathrm{BS},m}(z,p,t)$ is the matrix element describing the amplitude and phase shift after the beamsplitter operation in the momentum state indexed by m. It is computed by integrating Eq. (6.2) using one of the methods described in Section 6.1. In practice, we truncate the sum over some reasonable range of momentum values (typically around 100).

Applying free propagation of duration T to node j yields

$$U_{\mathrm{free}}(T)\left|\psi_{z_j,p_j}(t_j)\right\rangle = a_{\mathrm{free}}(p_j,T)\left|\psi_{z_j+p_jT/M,\,p_j}(t_j+T)\right\rangle \tag{6.9}$$

where M is the particle mass. The coefficient $a_{\mathrm{free}}(p_j,T)$ is given by

$$a_{\mathrm{free}}(p_j,T) = \exp(-ip_j^2T/2M\hbar) \tag{6.10}$$

Note that nodes can optionally contain an additional property δp_j, which tracks the deviation of the wavepacket momentum from $p_0 + 2m\hbar k$ for some integer m. This is used to track slight changes in particle momentum that result from Gouy shifts in the beam. When free evolution is computed, the user has the option to adjust p_j as appropriate to account for this shift. The tree built by these operations, in the case of an SCI, can be seen in Fig. 6.3. The x-axis displays the level of the tree, and the y-axis displays the index of each node at that level. Black lines denote paths that exit the interferometer in one of the expected ports. Gray lines denote junk or parasitic paths. In the figure we only include paths that

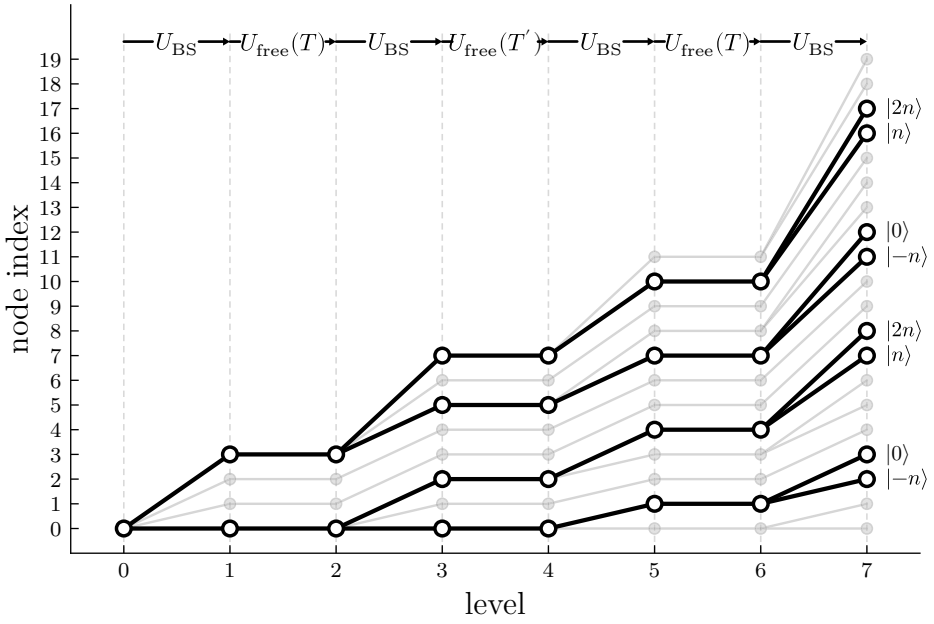


Figure 6.3: Diagram of the nodes used to represent an SCI when using the tree-based simulation method. Levels denote the interferometer state after different operations. For example, a beamsplitter is applied to go from level 0 to 1. At the end of the interferometer, nodes with equal positions and momenta interfere. These are represented by the nodes connected with black lines. Other nodes do not interfere, and are so-called "junk" outputs. These are represented in gray. Not all junk outputs are shown here, as we only include paths that end the interferometer with a population greater than 0.001 for the chosen Bragg parameters (parameters were chosen to minimize parasitic output ports). If parasitic interferometers are included in the tree, additional paths will appear in the tree that will ultimately interfere.

end the interferometer with a population greater than 0.001 for the chosen Bragg parameters (parameters were chosen to minimize parasitic output ports). We see that the interferometer starts with a single node, and subsequent operations form a tree from this node.

We have seen that applying U_{BS} and $U_{\mathrm{free}}(T)$ to our initial node produces a series of child nodes. After the application of ℓ total operations to the initial node, we will be at depth ℓ of the tree. At any given depth, the wavefunction describing the interferometer is given by the superposition of all wavepackets represented by the nodes at that level

$$\left|\Psi^{(\ell)}\right\rangle = \sum_{j \in \text{level } \ell} a_j \left|\psi_{z_j,p_j}\right\rangle \tag{6.11}$$

where a_j is the product of the $a_{\mathrm{BS},m}(z,p,t)$ and $a_{\mathrm{free}}(p,T)$ edge coefficients along the path from the root node to node j. When the interferometer is over, we compute the wavefunc-

tion at the bottommost level and group nodes sharing the same momentum and position expectation values. This produces a wavefunction of the form

$$\left|\Psi^{(\text{bottom})}\right\rangle = \sum_{z,p} \left[\sum_{j \in \Gamma(z,p)} a_j \right] \left|\psi_{z,p}\right\rangle \tag{6.12}$$

where a_j is the product of edge coefficients as before, and $\Gamma(z,p)$ is the set of all paths in the tree that result in a node at position z and momentum p. This expression might be a bit confusing. The idea is that the inner sum is a sum over the wavefunction amplitude of all nodes that end up at position z and momentum p. When this is a sum over more than one node, interference occurs. This sum sets the amplitude of the state $\left|\psi_{z,p}\right\rangle$. The outer sum is a sum over all the different states $\left|\psi_{z,p}\right\rangle$ that exit the interferometer. From this wavefunction it is trivial to compute the population in each output position and momentum state. Since the interferometer detects output momenta but is mostly insensitive to the output position of each state, the simulation computes the population exiting the interferometer with momentum p via

$$P(p) = \sum_{z} \left| \sum_{j \in \Gamma(z,p)} a_j \right|^2 \tag{6.13}$$

where the inner sum coherently adds the amplitudes of all nodes ending at position z and momentum p, and the outer sum incoherently adds the resulting populations over all output positions z.

The actual simulation code works in two stages. In the first stage the order of beamsplitter and free propagation operations is specified and the tree is constructed. The edge coefficients $a_{\text{BS},m}(z,p,t)$ and $a_{\text{free}}(p,T)$ are not computed at this stage. The user then supplies functions for computing $a_{\text{BS},m}(z,p,t)$ and $a_{\text{free}}(p,T)$ numerically. In the second stage, the user requests the population in particular output momentum states. The code then determines which nodes are required for computing the population in that state. The ancestry of those nodes is then traversed, and the edge coefficients are computed. The final population is then computed and returned to the user. To efficiently simulate a cloud of atoms propagating through the interferometer, the vectorized ODE integrator or lookup table methods described in Section 6.1 can be used.

The code also provides a separate implementation that symbolically computes the interferometer phase for an idealized beamsplitter using the SymPy symbolic math library. This implementation was used to compute the phases stated in Section 3.2.

The tree-based approach can be used to propagate parasitic interferometers. In Fig. 6.1 we see an SCI with parasitic interferometer paths (these are slightly transparent in the figure). These are naturally included by including a child node for every possible momentum state output by each beamsplitter operation. This would increase the number of paths shown in the SCI tree diagram in Fig. 6.3. In addition to propagating parasitic interferometers like the ones shown in the figure, the tree-based approach also allows junk ports of the

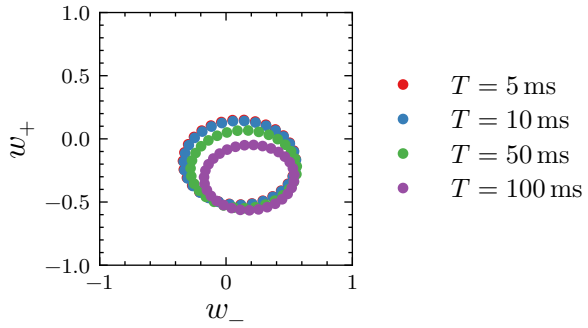


Figure 6.4: Simulated ellipses generated using a tree-based simulation method of an SCI. Several values of T are simulated, and, as expected, a decay of contrast is observed as T increases.

interferometer to be easily included in the simulation. Furthermore, Gouy phase effects can be easily included as the tree-based method records the wavevector at the position of each simulated atom for each Bragg beamsplitter pulse. This wavevector is then used to determine the exact momentum of each atom when free evolution phase is computed. The modular approach of the tree-based simulation method makes it relatively trivial to incorporate new physical effects as they become apparent.

To produce a simulated simultaneous conjugate Ramsey-Bordé ellipse using the tree-based method, we take $\delta\tau \to \delta\tau + \theta$ in Eq. (6.2) and scan θ to simulate the motion of the retro mirror. This phase imparts a laser phase onto ϕ_c, which results in w_- and w_+ forming an ellipse when plotted parametrically. We display several simulated ellipses generated from a cloud of atoms using the tree-based method in Fig. 6.4.

6.2.2 Wavepacket integration approach

In Appendix A.6 we derived the position-space equation of motion, reproduced below

$$i\hbar\dot{\psi}_g(t,z) = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\psi_g(t,z) - \hbar\left[\Omega_1^{\mathrm{AC}}(t) + \Omega_2^{\mathrm{AC}}(t) + \Omega_{\mathrm{eff}}(t)\cos(2kz - \delta t)\right]\psi_g(t,z) \tag{6.14}$$

This looks like the usual one-dimensional Schrödinger equation with a sinusoidal potential. Therefore, standard approaches for solving the 1D Schrödinger equation can be used, such as the split-step method [93, 94] or the Crank-Nicolson method [95]. Other methods are described in [86]. We will briefly describe these two methods, as well as another method that solves the Schrödinger equation position-space PDE by stacking together multiple solutions to the Schrödinger equation momentum-space ODE.

Split-step method

The split-step method relies on the following approximation for the time-evolution operator

$$\exp\left(-\frac{i}{\hbar}\hat{H}\Delta t\right) \approx \exp\left(-\frac{i\hat{H}_x\Delta t}{2\hbar}\right)\exp\left(-\frac{i\hat{H}_p\Delta t}{\hbar}\right)\exp\left(-\frac{i\hat{H}_x\Delta t}{2\hbar}\right) \tag{6.15}$$

where $\hat{H}_x$ is the part of the Hamiltonian dependent on position and $\hat{H}_p$ is the part of the Hamiltonian dependent on momentum. The approximation comes from applying Glauber's theorem twice and approximating away the term of order Δt^3 [93, 94]. Alternatively the Baker-Campbell-Hausdorff theorem can be used. We make an additional approximation by treating our time-dependent Hamiltonian as constant over each time step.

This approximation is used to propagate the wavefunction by using a fast Fourier transform (FFT) to convert between the position- and momentum-space wavefunctions. The rightmost term is applied to the position-space wavefunction. Then an FFT is performed, the middle term is applied, an inverse FFT is performed, and then the leftmost term is applied. This process is repeated over multiple time steps to evolve the wavefunction in time [93, 94].

Crank-Nicolson method

The Crank-Nicolson method for the one-dimensional time-dependent Schrödinger equation can be derived by approximating the unitary evolution operator $U(\Delta t) = \exp(-i\widehat{H}\Delta t/\hbar)$ via the $(1,1)$-order Padé approximant [86, 96, 97]. As with the split-step method, we approximate our time-dependent Hamiltonian as time-independent over the duration of a single time step. Applying this to some initial wavefunction $\psi(x,t)$ we find

$$\psi(x,t+\Delta t) = U(\Delta t)\psi(x,t) \approx \frac{1-iH\Delta t/2\hbar}{1+iH\Delta t/2\hbar}\psi(x,t) = B^{-1}A\psi_j^\ell$$

where in the last step we discretized $\psi(x,t)$ via $\psi_j^\ell = \psi(j\Delta x, \ell\Delta t)$. Using a second-order central finite difference to treat the $\partial^2/\partial x^2$ in the momentum term in the Hamiltonian results in tridiagonal A and B

$$A_{j,j} = \left[1 - iH^\ell\Delta t/2\hbar\right]_{j,j} = 1 - i\frac{\Delta t}{2}\left(V_j^\ell + \frac{1}{\Delta x^2}\right) \tag{6.16}$$

$$B_{j,j} = \left[1 + iH^\ell\Delta t/2\hbar\right]_{j,j} = 1 + i\frac{\Delta t}{2}\left(V_j^\ell + \frac{1}{\Delta x^2}\right) \tag{6.17}$$

$$A_{j,j-1} = A_{j-1,j} = \left[1 - iH^\ell\Delta t/2\hbar\right]_{j,j-1} = i\frac{\Delta t}{4\Delta x^2} \tag{6.18}$$

$$B_{j,j-1} = B_{j-1,j} = \left[1 + iH^\ell\Delta t/2\hbar\right]_{j,j-1} = -i\frac{\Delta t}{4\Delta x^2} \tag{6.19}$$

where $H^\ell = H(\ell\Delta t)$ and $V^\ell = V(\ell\Delta t)$. For a discussion of higher-order Crank-Nicolson type routines see [95, 98]. A single time step can be executed in linear time through the use of the Thomas tridiagonal matrix division algorithm. Note that the $(1,1)$-order Padé approximant $U(\Delta t) \approx (1 - iH\Delta t/2\hbar)/(1 + iH\Delta t/2\hbar)$ is unitary, and so the Crank-Nicolson method preserves wavefunction normalization [95–98].

Stacked ODE method

The stacked ODE method [26, 94, 99] takes advantage of the property that a sinusoidal potential has the form of a delta function in momentum space, which produces the discrete ladder of states in Eq. (6.2). Deriving this method is straightforward. When deriving Eq. (2.14) we defined $c_j(t)$ as $c_j(t) = c_g(t, 2j\hbar k)$. If we instead use the discretization scheme $c_{j,s}(t) = c_g(t, 2\hbar kj + \hbar ks/S)$, where S sets the number of discretized states to include between intervals of $2\hbar k$ in momentum, we find

$$\begin{aligned} i\hbar\dot{c}_{j,s}(t) =& \hbar\omega_{\mathrm{r}}(2j + s/S)^2 c_{j,s}(t) - \hbar\frac{\Omega_{\mathrm{eff}}(t)}{2}\left[\exp(i\delta t)c_{j+1,s}(t) + \exp(-i\delta t)c_{j-1,s}(t)\right] \\ &- \hbar\Omega_1^{\mathrm{AC}}(t)c_{j,s}(t) - \hbar\Omega_2^{\mathrm{AC}}(t)c_{j,s}(t) \end{aligned} \tag{6.20}$$

This is nearly identical to Eq. (2.14), except that the kinetic energy term has an s/S term. Here s sets how far in momentum space a particular state is from a multiple of $2\hbar k$, and j sets roughly the number of $2\hbar k$ a particular state is from zero momentum. We enforce $j \in \mathbb{Z}$, $S \in \mathbb{Z}, S > 0$, and $s \in \{-S, -S+1, \ldots, S-2, S-1\}$. In the limit that $S \to \infty$ this definition for $c_{j,s}(t)$ becomes continuous and covers all momenta. As in the single ODE case, we truncate the number of j values to J.

To solve for the evolution of a wavepacket, we set S high enough to describe the wavepacket's features and solve the $2S$ systems of J ODEs and stack together their solutions to reconstruct the wavepacket.

To select reasonable values for J and S, note that in momentum space the grid spacing is $\hbar k/S$, and the total extent of the simulation is $2\hbar kJ$. Therefore, in position space the grid spacing is $1/Jk$, and the total extent of the simulation is S/k. For a wavefunction with features at the resolution of δ_x in position space and δ_p in momentum space propagating through an atom interferometer with spatial extent Δ_x and momentum extent Δ_p, the following inequalities must be satisfied:

$$\begin{aligned} \delta_p &\gtrsim \hbar k/S \\ \delta_x &\gtrsim 1/Jk \\ \Delta_p &\lesssim 2\hbar kJ \\ \Delta_x &\lesssim S/k \end{aligned} \tag{6.21}$$

Therefore reasonable choices for J and S are $J \gtrsim \max(1/k\delta_x, \Delta_p/2\hbar k)$ and $S \gtrsim \max(\hbar k/\delta_p, \Delta_x k)$.

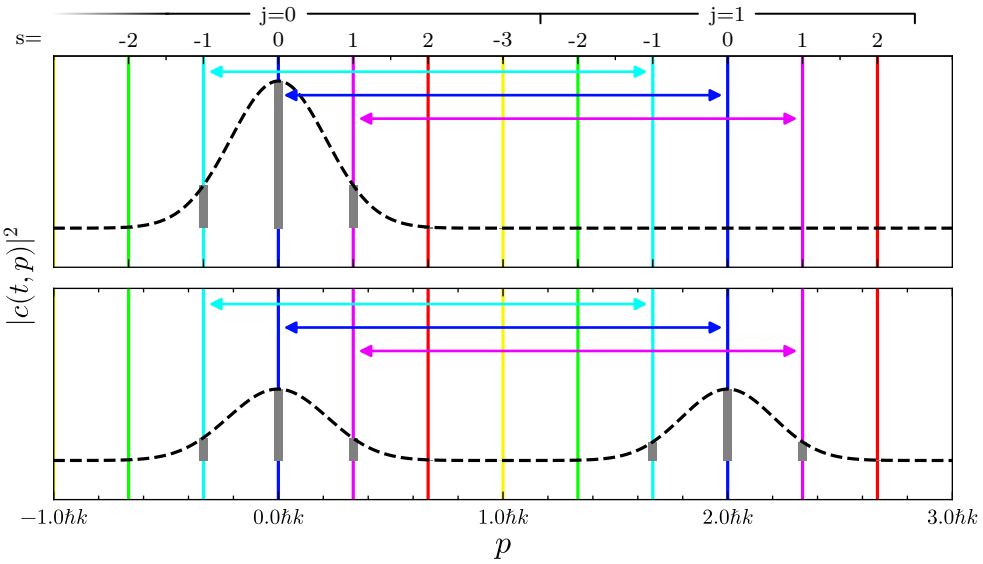


Figure 6.5: The discretization scheme used when setting $\phi_{j,s} = \phi(2\hbar kj + \hbar ks/S)$ for the case where $S = 3$.

The discretization scheme is shown graphically in Fig. 6.5 for the case where $S = 3$. The upper plot shows an initial momentum-space wavefunction, and the lower plot shows the same wavefunction after undergoing a 1st-order Bragg pulse. Each discretized point is indicated by a vertical line. For each value of j there is a bin between $2j\hbar k - \hbar k$ and $2j\hbar k + \hbar k$ which contains 6 discretized points. The value of s at each of these points is shown on the top axis. The value of j for each bin is also shown on the top axis. The arrows are each $2\hbar k$ long and indicate which discretized points are coupled to each other in momentum space (coupled discretized points share the same color). Only three arrows are shown here for simplicity, but in reality every discretized point has two arrows pointing out of it in opposite directions as it is coupled to two other momentum states. Note that coupled states share the same value of s, but have different values of j. The gray bars indicate the population attributed to each particular discretized state. A similar picture is provided in [94] Fig. (10b).

In the figure, we see that there are values of s that have little or no amplitude. These values of s can be removed from the calculation, truncating the number of ODEs that must be solved. For initial wavefunctions with a momentum spread much smaller than $2\hbar k$, this can provide a computational speedup in exchange for a loss in accuracy. The effect of this truncation on accuracy is tested in the next section.

In Section 6.1.1 we discussed using a lookup table to store ODE integration results. By

applying the frame transformation $\tilde{c}_{j,s}(t) = \exp(i(2j + s/S)^2\omega_r t)c_{j,s}(t)$ to Eq. (6.20) we see

$$\begin{aligned} i\hbar\dot{\tilde{c}}_{j,s} = &-\hbar\frac{\Omega_{\text{eff}}(t)}{2}\left[\exp(-i(4-8j)\omega_r t)\exp(-i(\delta - 4\omega_r s/S)t)\tilde{c}_{j-1,s}(t)\right. \\ &\left. + \exp(-i(4+8j)\omega_r t)\exp(i(\delta - 4\omega_r s/S)t)\tilde{c}_{j+1,s}(t)\right] \\ &-\hbar\left(\Omega_1^{\text{AC}} + \Omega_2^{\text{AC}}\right)\tilde{c}_{j,s}(t) \end{aligned} \tag{6.22}$$

In this frame we see that the effect of changing s is equivalent to a change in the two-photon detuning. For a fixed pulse envelope $\Omega_{\text{eff}}(t)$, this allows us to produce a lookup table for the initial state $\tilde{c}_{j=0,s=0}(t=0)$ over δ and $\Omega_{\text{eff},0}$. By adjusting δ and applying further frame transformations, we can recover the evolution of $\tilde{c}_{j,s}(t)$ for any j and any s from the result for $\tilde{c}_{j=0,s=0}(t)$. This enables the full wavefunction propagation to be computed using a single lookup table.

Method comparison

The accuracy of the ODE stacking method was tested by simulating a Ramsey-Bordé interferometer constructed from Bragg beamsplitters using the split-step, Crank-Nicolson, and ODE stacking methods. Note that the split-step and Crank-Nicolson methods are capable of integrating an arbitrary potential instead of a purely sinusoidal one. The simulation is performed for 50 values of ω_m and the resulting wavefunction is integrated across output ports to yield populations

$$p_j = \int_{2\hbar k(j-1/2)}^{2\hbar k(j+1/2)} c_g(t,p')^* c_g(t,p')\,\mathrm{d}p' \tag{6.23}$$

These populations are used to compute interference fringes as a function of ω_m, which are fit to extract the fringe phase. The deviation of this phase between different integration methods is compared.

The interferometer uses $T = 10\omega_r^{-1}$, $T' = 2\omega_r^{-1}$, the $\Omega_{\text{eff}}(t)$ defined in Eq. (2.20), and the initial wavefunction

$$c_g(t,p) = \frac{1}{\sqrt{\sqrt{2\pi}\sigma_p}}\exp\left(-\frac{p^2}{4\sigma_p^2}\right) \tag{6.24}$$

where $\sigma_p = 0.14\hbar k$. We simulate the full interferometer for three Bragg orders: $n = 3$ with $\sigma = 0.4966\omega_r^{-1}$ and $\Omega_{\text{eff},0} = 9\omega_r$; $n = 4$ with $\sigma = 0.224\omega_r^{-1}$ and $\Omega_{\text{eff},0} = 20\omega_r$; $n = 5$ with $\sigma = 0.247\omega_r^{-1}$ and $\Omega_{\text{eff},0} = 30\omega_r$. These parameters produce a $\pi/2$ pulse for their respective beamsplitters. The two-photon detuning δ is chosen using the Bragg resonance condition (see Eq. (2.18)).

Note that all integration methods use the same momentum-space discretization set by choosing $S = 60$ and $J = 60$ (for the Crank-Nicolson method this is converted into a position-space discretization). The DOP853 adaptive-step-size integrator is used for the ODE method

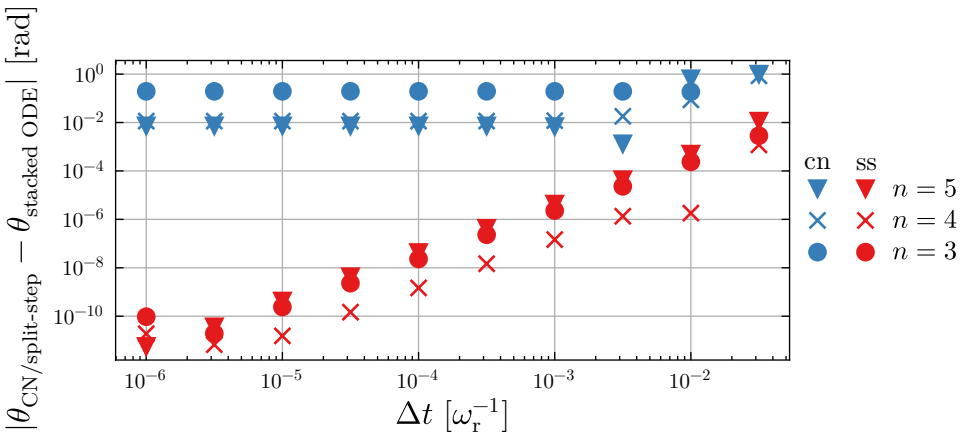


Figure 6.6: Difference between computed interferometer phases for the stacked ODE method, the split-step method (ss), and the Crank-Nicolson method (cn).

with double precision numbers and relative and absolute tolerances of 10^{-10}. Fig. 6.6 shows the absolute value of the difference between the phase simulated by the stacked ODE method and the split-step and Crank-Nicolson methods. The split-step method converges to the stacked ODE method with a difference of around 10^{-9} rad at a step size of approximately $\Delta t = 10^{-5}\omega_{\mathrm{r}}^{-1}$. The Crank-Nicolson method does not converge, as its accuracy is already limited by the grid spacing and not the time step.

We can also determine how truncating the space of considered s values affects the ODE stacking result. In Fig. 6.7 we remove values of s where $|c_{j,s}(t=0)|^2 < \varepsilon$ for any j, where ε is the chosen truncation threshold. We then compare the simulated interferometer phase with and without truncation. For the case considered here of an initially Gaussian wavefunction, it is clear that depending on the precision requirement of the user, there is potentially a substantial degree of truncation that can be applied.

6.3 Spatial separation of wavepackets

It is informative to reconcile what differences between the plane wave and stacked ODE integrations give rise to the physics of spatial separation. To recap: atom interferometers operate in the quasi-Bragg regime, where the momentum width of the wavepacket describing an atom is much less than the bandwidth of the Bragg diffraction process. During this process, the diffracted and undiffracted states do not separate spatially. This means that the plane wave basis accurately describes Bragg beamsplitter dynamics. However, after the Bragg beamsplitter, the wavepackets separate spatially and no longer overlap, meaning the plane wave basis is no longer sufficient.

If we only consider $s = 0$ in Eq. (6.20), then the simulation describes a plane wave

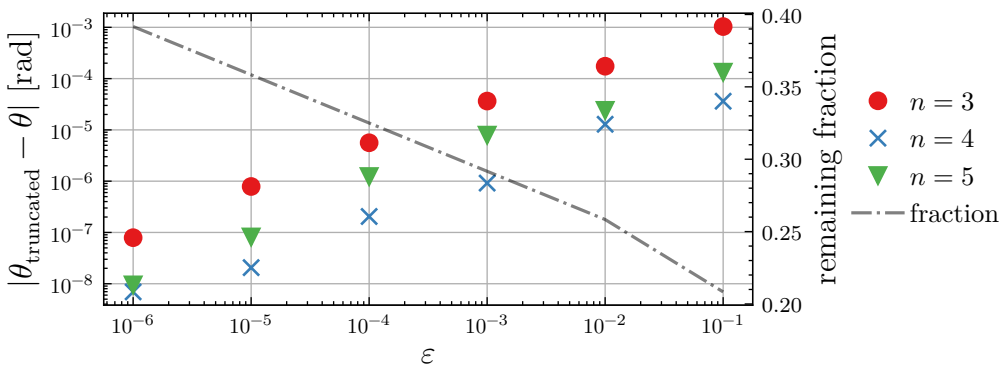


Figure 6.7: Error in the computed interferometer phases as states are removed from the stacked ODE method. The left-hand-side axis displays the error between the truncated and untruncated simulations. The right-hand-side axis displays the remaining fraction of s states for a particular truncation level.

with infinite spatial extent and no spatial separation. By including additional values of s, wavepacket separation appears because wavepacket position is encoded in the slope of the momentum-space wavefunction phase. The grid over s provides the resolution in momentum space required to represent this slope.

To see this explicitly, consider two spatially separated Gaussian wavepackets. A single Gaussian wavepacket with momentum and position expectation values $2n_0\hbar k$ and x_0 is described by

$$\left\langle x \middle| \psi_{n_0,x_0} \right\rangle = A \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2}\right) \exp(i2n_0kx) \tag{6.25}$$

where A is the relevant normalization factor. If we consider a state composed of two Gaussian wavepackets at different positions and momenta $|\psi\rangle = |\psi_{0,0}\rangle + |\psi_{m,d}\rangle$ and transform to momentum space, we find

$$\begin{aligned}
\langle p|\psi\rangle =& \langle p|\psi_{0,0}\rangle + \langle p|\psi_{m,d}\rangle \\
=& \int_{-\infty}^{\infty} \left[\langle p|x'\rangle\langle x'|\psi_{0,0}\rangle + \langle p|x'\rangle\langle x'|\psi_{m,d}\rangle\right] \mathrm{d}x' \\
=& \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} \left[\exp(-ipx'/\hbar)\langle x'|\psi_{0,0}\rangle + \exp(-ipx'/\hbar)\langle x'|\psi_{m,d}\rangle\right] \mathrm{d}x' \\
=& \frac{A\sigma_x}{\sqrt{\hbar}} \left[\exp\left(-\frac{p^2\sigma_x^2}{2\hbar^2}\right) + \exp\left(-\frac{(p-2m\hbar k)^2\sigma_x^2}{2\hbar^2}\right) \exp\left(i(2mk - p/\hbar)d\right)\right]
\end{aligned} \tag{6.26}$$

We immediately see that if the two wavepackets have the same momentum expectation value but have separation d, the wavefunction phase has the form $\arg(1+\exp(-ipd/\hbar))$–i.e. a line

with slope $-d/\hbar$. In the case where we only consider $s = 0$, we undersample this linearly ramped phase, as it can wrap around 2π between neighboring momentum states.

To prevent this undersampling, we discretize p onto a momentum grid with spacing $\hbar k/S$. The phase between neighboring grid points is then kd/S, and by requiring $2\pi > kd/S \rightarrow S > kd/2\pi$ we ensure the linearly ramped phase is sampled accurately. Note that this condition is consistent with the one given in Eq. (6.21). If this condition is not met, we will continue to undersample the phase in momentum space and fail to capture the physics of spatial separation.

Interestingly, the insight that including additional s values allows spatial separation to be represented has implications for the tree-based method described in Section 6.2.1. In the tree-based method spatial separation was enforced by tracking the classical motion of each wavepacket between Bragg beamsplitters and then simulating separate Bragg beamsplitter processes for each set of wavepackets that spatially overlapped. To simulate an ensemble of atoms, this could be repeated over a distribution of initial momentum states. This is equivalent to incorporating multiple values of s in the ODE stacking approach. This leads to the perhaps surprising conclusion that simulations of an ensemble of atoms with a distribution of initial momentum states do not need to manually track spatial separation, as the effect of spatial separation is automatically incorporated by averaging over the momenta of the ensemble. This will only work if a significant number of initial momentum states are sampled, and so in practice it is computationally cheaper to track the separations using the tree-based method.

Summary

In this chapter we have reviewed the simulation methods developed for simulating atom interferometers. In the next chapter we will fit tree-based simulations to experimentally measured Bragg pulses (Section 7.1.5), Bloch oscillations (Section 7.1.6), and the full SCI interferometer (Section 7.1.7). This will show that the simulation accurately describes experiment output. We will also use the tree-based simulation to perform a theoretical study of the systematics in the SCI and a novel interferometer geometry (Section 7.2.1). We will also use the stacked ODE method to simulate systematic effects associated with the width of the momentum wavefunction in Section 7.1.3.

Chapter 7

Investigations

This chapter is broadly split into a systematics section and a methods section. In the systematics section we review the systematic effects present in the experiment, and simulate several of these. In Section 7.1.5, Section 7.1.6, and Section 7.1.7 we perform an extensive verification of the simulation described in Section 6.2.1 against the experiment. These extensive simulation-experiment comparisons were not performed in the 2018 Berkeley fine-structure constant measurement, and provide increased confidence in absolute systematic shifts reported by the simulation.

The methods section reviews several novel interferometer geometries and investigates their performance (both through experiment and simulation), and introduces a new beamsplitter technique. In Section 7.2.1 we show, using simulation, a modification to the usual SCI geometry that suppresses systematic effects related to the diffraction phase. A novel beamsplitter technique based on serrodyne-modulation is described in Section 7.2.2.

7.1 Systematics

The 2018 Berkeley fine-structure measurement identified several systematic effects listed in Table 7.1. In this section we will discuss progress towards understanding these systematic effects in the new apparatus as well as several newly discovered systematic effects. The largest shifts to the measurement are produced by the Gouy phase, the acceleration gradient, and the laser frequency. The largest uncertainties stem from the thermal motion of atoms and the speckle phase shift.

We will quickly review the mechanism behind several of these systematic effects. The thermal motion of atoms systematic captures the diffraction phase shifts produced by the transverse motion of atoms between Bragg beamsplitter pulses. This transverse motion results in atoms sampling different beam intensities for each Bragg beamsplitter, which we see can produce a variation in the diffraction phase from Fig. 2.7. This effect can be suppressed by increasing the Bragg beamsplitter pulse duration, using a larger interferometry beam to reduce intensity variations, or further cooling the cloud to reduce the transverse motion [47].

Table 7.1: Error budget of the 2018 Berkeley fine-structure constant measurement [3]. The thermal motion of atoms and speckle phase shift systematics produce systematic shifts via the Gouy phase and diffraction phase mechanisms. For this reason, the shifts they produce are included in the Gouy phase shift (and their shifts are listed as zero), but their errors are listed separately. The thermal motion of atoms error, for example, describes the uncertainty in the Gouy phase shift (and therefore α) resulting from the uncertainty surrounding atom positions in the cloud over the interferometer duration. This is true for the speckle phase shift and parasitic interferometer systematics as well.

Effect	$\delta\alpha/\alpha$ **(ppb)**
Thermal motion of atoms	0 ± 0.08
Speckle phase shift	0 ± 0.04
Parasitic interferometers	0 ± 0.03
Gouy phase	-2.60 ± 0.03
Non-Gaussian waveform	0 ± 0.03
Bloch oscillation light shift	0 ± 0.002
Modulation frequency wave number	0 ± 0.001
Laser frequency	-0.24 ± 0.03
Index of refraction	0 ± 0.03
Acceleration gradient	-1.79 ± 0.02
Beam alignment	0.05 ± 0.03
Sagnac effect	0 ± 0.001
Density shift	0 ± 0.003
Total systematic error	-4.58 ± 0.12

In the 2018 measurement this effect was bounded via a simulation [3].

The "speckle phase shift" is produced by spatial noise on the interferometry beam, as well as the edges of the interferometry beam reflecting off of the sides of the vacuum chamber back towards the atom cloud. This spatial noise produces phase shifts through diffraction phase and through Gouy shifts (discussed later on). The 2018 measurement suppressed this effect by using a spatial filter to clean the beam mode, and bounded the effect through empirical observation [3].

We have already discussed the mechanism behind parasitic interferometers in Section 6.2.1. In the 2018 measurement these were suppressed by selecting a Bragg beamsplitter pulse duration that suppressed the generation of population in unwanted momentum states [3, 100].

The Gouy phase shift results from the variation of the laser wavevector $\mathbf{k}$ over the laser wavefront. For an idealized Gaussian beam, this variation is known analytically. For a realistic beam with spatial noise, small scale variations in $\mathbf{k}$ occur that are hard to model.

The 2018 measurement primarily characterized this systematic by imaging the interferometry beam and simulating its effect [3].

The non-Gaussian waveform systematic captures any shift resulting from imperfectly shaped Bragg pulse shapes in time. Ideally, waveforms should have an effective Rabi frequency given by Eq. (2.20) in time. However, the pulse shaping circuit cannot have infinitely high bandwidth, and so the actual pulse shape will be slightly different. This was bounded in the 2018 measurement by using the result of Section 5 in [100].

Each of these five effects are coupled. For example, the population transferred into unwanted momentum states after a Bragg beamsplitter is highly dependent on the local pulse intensity. This means that the occurrence of parasitic interferometers is coupled to the thermal motion of atoms, and to a lesser extent the intensity speckle on the interferometry beam. Part of the Gouy phase shift stems from speckle on the interferometry beam. The thermal motion of atoms causes atoms to experience different speckle and Gouy phase shifts. For all of these reasons, it is logical to build a simulation framework capable of capturing all of these effects simultaneously. This is the goal of the tree-based simulation method described in Section 6.2.1. In the 2018 measurement these effects were considered separately despite these couplings (see the caption of Table 7.1). The remaining systematic effects in the table are more straightforward to disentangle.

To suppress the thermal motion of atoms and the speckle phase shift, the vacuum chamber was made much larger and baffles were added to support a larger beam and reduce speckle. Furthermore, the longer chamber gives small features on the interferometry beam distance to expand away before interacting with the atoms. This is the motivation behind much of the vacuum chamber design discussed in Section 4.1. These changes directly target the two largest uncertainties in the error table.

The new chamber design will also help suppress the thermal motion of atoms systematic, although it is not clear by how much. Currently the experiment utilizes a 7.7 mm $1/e^2$ waist radius beam, while the 2018 measurement used a 5.7 mm $1/e^2$ waist radius beam (further reduced in size to a $1/e^2$ radius of 3.4 mm by an apodizing filter, although the resulting beam did not have a Gaussian profile). The beam size is limited by intensity, and the experiment is operated at a smaller single-photon detuning than the 2018 measurement. Furthermore, the taller vacuum chamber gives the atom cloud a longer time to expand between release from the 3D MOT and the first interferometry pulse, producing a larger cloud. The detection photodiode size and detection beam width determine which atoms in the cloud are ultimately detected. As the photodiode is 1 cm in width, this bounds the range of beam intensities that can be sampled by the atoms.

In addition to the effects listed in Table 7.1, several new systematic effects have been identified, which are referred to as the VS magic wavelength (the systematic effect occurs when operating *away from* this magic wavelength), the Bragg vibration systematic, and the wavefunction width systematic. All of these systematic effects stem from mechanisms that shift atoms propagating through the interferometer slightly away from resonance with a Bragg diffraction pulse. The VS magic wavelength systematic, in particular, possibly increased the error bar on the thermal motion of atoms systematic in the 2018 measurement. These

systematics can be precisely simulated, which will enable further suppression of systematic uncertainties. We will begin by discussing the VS magic wavelength systematic.

Finally, recall from Section 5.4 that the full interferometer phase $\Phi = 2\phi_d$ is measured to a precision of 3.44 mrad, and so systematic effects must also be understood at this level if the experiment precision is to be bounded by statistics and not systematics.

7.1.1 Velocity selection magic wavelength

In Eq. (2.46) we derived the Raman VS resonance condition, restated below

$$\frac{2kp}{m} = \omega_1 - \omega_2 - 4\omega_\mathrm{r} - \omega_{g_2} - \Delta_\mathrm{AC} \tag{7.1}$$

Only one value of p is resonant with the VS process. As we saw in Fig. 4.30, the VS process has some bandwidth, and so some sliver of the wavefunction around this resonance is transferred.

Consider what it means for $\Delta_\mathrm{AC} \neq 0$. Recall that $\Delta_\mathrm{AC} = \Omega_\mathrm{AC}^{g_1} - \Omega_\mathrm{AC}^{g_2}$. The two AC Stark shift terms are proportional to the laser intensity, and so Δ_AC must also be proportional to laser intensity, which varies over the laser beam. Therefore $\Delta_\mathrm{AC} \to \Delta_\mathrm{AC}(x, y)$. This means our resonance condition becomes

$$\frac{2kp}{m} = \omega_1 - \omega_2 - 4\omega_\mathrm{r} - \omega_{g_2} - \Delta_\mathrm{AC}(x, y) \tag{7.2}$$

This tells us that if the differential AC Stark shift is non-zero, the momentum p selected by VS becomes spatially dependent! It turns out (as we will calculate), that this can produce shifts in the selected atom momentum that result in noticeably different diffraction phases. Thankfully, it is possible to operate the interferometry Ti:Sa laser at a single photon detuning where $\Delta_\mathrm{AC} = 0$, zeroing any systematic effect stemming from this mechanism. To allow the Bragg diffraction single photon detuning to be set independently of the VS single photon detuning, a new laser will be installed on the experiment.

The shift in the Bragg two-photon detuning produced by a momentum offset of p is given by Eq. (6.5). Plugging in our expression for p above, we find

$$\delta = \underbrace{4\,(2n_0 + n)\,\omega_\mathrm{r}}_{\text{usual Bragg two-photon detuning}} + \underbrace{\omega_1 - \omega_2 - 4\omega_\mathrm{r} - \omega_{g_2}}_{\text{tuned to sum to zero}} - \Delta_\mathrm{AC}(x, y) \tag{7.3}$$

where δ is the Bragg two-photon detuning and ω_1 and ω_2 denote the laser frequencies used during the VS process, not the Bragg process. So δ is modulated by $-\Delta_\mathrm{AC}(x, y)$. We will see that $\Delta_\mathrm{AC}(x, y)$ can vary on the order of a kilohertz. From Fig. 2.8 we note that, depending on the Bragg pulse duration, this can produce experimentally relevant shifts on the interferometer phase by shifting the diffraction phase on the order of 10 mrad to 100 mrad.

To compute Δ_AC we must know all frequencies in the VS beam, as well as their relative intensities. This calculation was performed in detail in [51]. Recall from Section 4.2.5 that the

VS laser is formed from two paths on the interferometry chain with orthogonal polarizations. One of the paths passes through an EOM, and produces an infinite series of frequencies. The other contains a single frequency. We must also account for the multiple levels that exist in $6\,^2\mathrm{P}_{3/2}$ manifold for the case of cesium. This corresponds to having multiple excited states in Eq. (2.41). Since the excited states are not coupled together and are all far from resonance, each excited state can be adiabatically eliminated and is described by a term analogous to $c_e(t,p)$ in Eq. (2.41). The $\dot{c}_{g_1}(t,p)$ and $\dot{c}_{g_2}(t,p)$ terms in Eq. (2.41) gain a sum over different excited states, as well as laser fields. After eliminating the excited states, this yields

$$\begin{aligned}
i\hbar\dot{c}_{g_1}(t,p) =& \frac{p^2}{2m}c_{g_1}(t,p) - \frac{\hbar}{2}\sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_1,e_{F'},\ell}(t)\Omega_{g_1,e_{F'},\ell'}(t)}{2(\omega_{e_{F'}}-\omega_{\ell'})}e^{i(\omega_\ell-\omega_{\ell'})t}c_{g_1}(t,p+\hbar k_\ell-\hbar k_{\ell'}) \\
&- \frac{\hbar}{2}\sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_1,e_{F'},\ell}(t)\Omega_{g_2,e_{F'},\ell'}(t)}{2(\omega_{e_{F'}}-\omega_{g_2}-\omega_{\ell'})}e^{i(\omega_\ell-\omega_{g_2}-\omega_{\ell'})t}c_{g_2}(t,p+\hbar k_\ell-\hbar k_{\ell'}) \\
i\hbar\dot{c}_{g_2}(t,p) =& \frac{p^2}{2m}c_{g_2}(t,p) - \frac{\hbar}{2}\sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_2,e_{F'},\ell}(t)\Omega_{g_2,e_{F'},\ell'}(t)}{2(\omega_{e_{F'}}-\omega_{g_2}-\omega_{\ell'})}e^{i(\omega_\ell-\omega_{\ell'})t}c_{g_2}(t,p+\hbar k_\ell-\hbar k_{\ell'}) \\
&- \frac{\hbar}{2}\sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_1,e_{F'},\ell'}(t)\Omega_{g_2,e_{F'},\ell}(t)}{2(\omega_{e_{F'}}-\omega_{\ell'})}e^{i(\omega_\ell+\omega_{g_2}-\omega_{\ell'})t}c_{g_1}(t,p+\hbar k_\ell-\hbar k_{\ell'})
\end{aligned} \tag{7.4}$$

This elimination is valid in the case where the coupling terms in the two equations have the same denominators [15]. Here F' indexes the excited hyperfine level. Assuming that no pair of fields is resonant with a Bragg process, each expression is given by a sum over an AC Stark shift type term and a Raman coupling type term. The AC Stark shift terms are given by

$$\begin{aligned}
\Omega_{\mathrm{AC}}^{g_1} =& \sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_1,e_{F'},\ell}(t)\Omega_{g_1,e_{F'},\ell'}(t)}{4(\omega_{e_{F'}}-\omega_{\ell'})}e^{i(\omega_\ell-\omega_{\ell'})t} \\
\Omega_{\mathrm{AC}}^{g_2} =& \sum_{F'}\sum_{\ell,\ell'}\frac{\Omega_{g_2,e_{F'},\ell}(t)\Omega_{g_2,e_{F'},\ell'}(t)}{4(\omega_{e_{F'}}-\omega_{g_2}-\omega_{\ell'})}e^{i(\omega_\ell-\omega_{\ell'})t}
\end{aligned} \tag{7.5}$$

As usual, the difference between the AC Stark shift type terms in each equation defines Δ_{AC}. Depending on the required accuracy of the calculation, the rotating wave approximation can be made again so that only terms without the rotating exponential remain. However, depending on the precision at which Δ_{AC} must be computed, this approximation may not be valid.

The Rabi frequency terms are given by [50, 51]

$$\Omega_{g,e,\ell}(t) = \sqrt{2F+1}(-1)^{m+J+I}\begin{pmatrix} F' & 1 & F \\ m' & q & -m\end{pmatrix}\sqrt{(2F'+1)(2J+1)}\begin{Bmatrix} J & J' & 1 \\ F' & F & I\end{Bmatrix}\Gamma\sqrt{\frac{I(t)}{2I_{\mathrm{sat}}}} \tag{7.6}$$

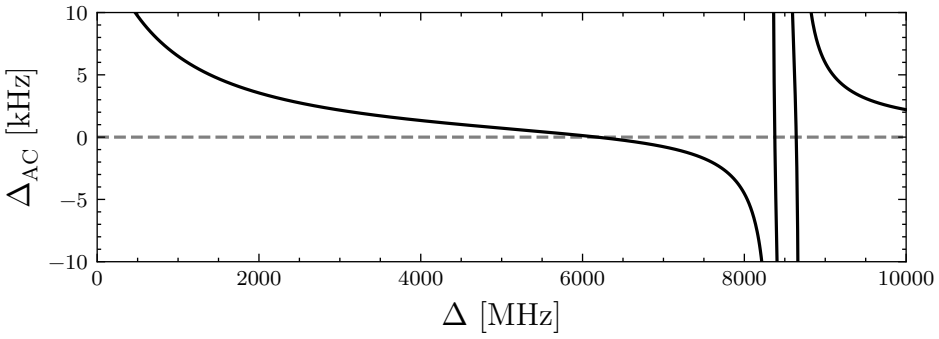


Figure 7.1: Calculated dependence of the Δ_{AC} on the single photon detuning Δ for the cesium $6\,^2S_{1/2}$ to $6\,^2P_{3/2}$ transition. We see that Δ_{AC} goes to zero around $\Delta = 6000\,\text{MHz}$. Δ_{AC} is computed via Eq. (7.5) for the configuration used in the experiment. Note that the calculation result depends strongly on the relative beam intensities and EOM modulation depth, and so this result is not expected to describe the actual experiment. A similar calculation is performed in [51].

where the unprimed quantum numbers are those of the ground state g and the primed quantum numbers are those of the excited state e. The quantum number q is that of the field. As discussed in Section 4.4.6, atoms enter VS in the $F = 3, m_F = 0$ state. Since the beam is colinear with the magnetic field during VS (and is composed of orthogonal linear polarizations), this means that an equal superposition of $\sigma_\pm$ couplings drive the transition.

The relative intensity of each beam participating in the VS transition is accounted for by setting $I(t)$ appropriately. It is assumed that all beams share the same time profile, but have different peak intensities. The intensity of the fields exiting the EOM is given by the Bessel function of the first kind $I_{\text{EOM},j} = I_{\text{EOM}} J_j(\beta)^2$, where j indexes the sideband order and β sets the modulation depth of the EOM. The relative intensity of the light entering the EOM, I_{EOM}, and the light exiting AOM32 in Fig. 4.11, I_{AOM}, must also be defined.

Taking all of this together, we can compute the differential AC Stark shift as a function of the single photon detuning Δ (measured from the $F = 3 \to F' = 4$ transition, with positive Δ denoting blue detuning). The calculation is performed for $I = 7/2$, $J = 1/2$, $J' = 3/2$, $\Gamma = 2\pi \times 5.2\,\text{MHz}$, $I_{\text{sat}} = 1.1\,\text{mW/cm}^2$, $I_{\text{EOM}} = I_{\text{AOM}} = 11\,\text{mW/cm}^2$, and $\beta = 1.1$ in Fig. 7.1. In this calculation the rotating wave approximation was made to remove all terms with rotating exponentials in Eq. (7.4). We see from the figure that $\Delta_{AC} = 0$ when $\Delta \approx 6000\,\text{MHz}$. The value of Δ where this crossing occurs is highly dependent on the parameters I_{EOM}, I_{AOM}, and β. Note from the form of Eq. (7.5) that scaling the intensity from the Ti:Sa linearly scales the y-axis in the plot.

Yuno Iwasaki investigated this effect experimentally and determined that $\Delta_{AC} = 0$ when the Ti:Sa beatnote is set to 5825 MHz, which is referred to as the "magic wavelength".

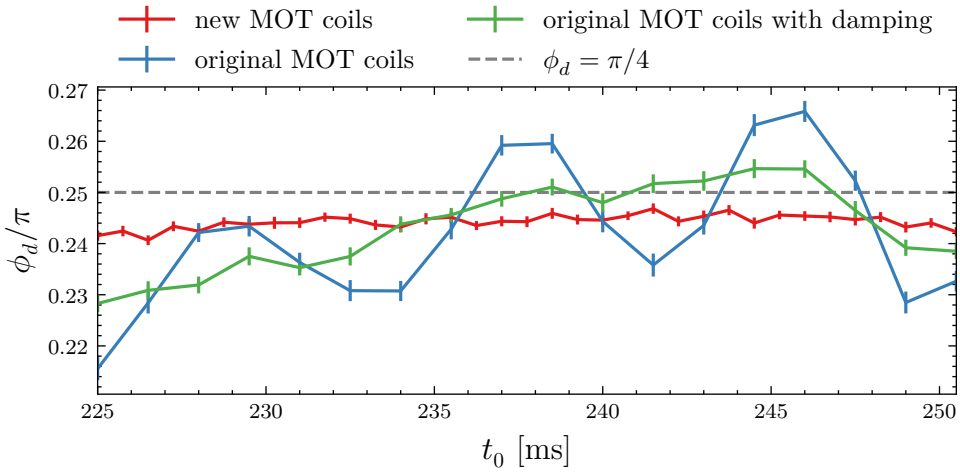


Figure 7.2: The differential phase ϕ_d measured as a function of the time t_0 of the first Bragg beamsplitter pulse after the molasses launch. The differential phase is shown in three different conditions: with the new 3D MOT coils installed on the experiment, with the original 3D MOT coils installed on the experiment, and with damping sandbags placed on top of the retro mirror chamber while the original 3D MOT coils were installed on the experiment. The $\Phi = 2\phi_d = \pi/2$ condition is also shown. It is clear that the new 3D MOT coils remove much of the observed oscillation.

She also measured the systematic phase shift produced by this effect. Note that if the experiment is operated away from the magic wavelength, it is possible to simulate this effect by appropriately initializing the atom cloud used in the tree-based simulation method.

7.1.2 Bragg vibration systematic

Before the original 3D MOT coils were swapped for the new coils (as described in Section 4.1.2), it was found that when the timing of the first Bragg beamsplitter pulse time t_0 was scanned, an oscillation in the extracted differential phase ϕ_d was observed on the order of a hundred milliradians. A similar oscillation was found in ϕ_d as T was scanned. These observations implied the presence of a time or position dependent phase shift in the experiment. Several attempts were made to determine the source of this oscillation, which was eventually attributed to vibrations in the retro mirror. Sandbags were placed on top of the retro mirror chamber in an attempt to damp the vibrations, which changed the recorded oscillation but did not remove it. After switching from the original 3D MOT coils to the new 3D MOT coils, as described in Section 4.1.2, the observed oscillation in the SCI differential phase was suppressed. Both the sandbag and new 3D MOT coil conditions are shown in Fig. 7.2, alongside the original 3D MOT coil condition.

The mechanism by which the mirror vibration produces this shift is as follows: a mirror

moving at velocity v produces a frequency shift on the reflected beam given by $2kv$. From Section 2.2.2, we note that this changes the two-photon detuning via $\delta \to \delta + 2kv = \delta + 4\omega_r v/v_r$. Since the upper interferometer diffracts up, and the lower interferometer diffracts down, this offset does not affect the two beamsplitters equivalently, and produces differential diffraction phase in the interferometer. The vibrations in the mirror were consistent enough shot-to-shot to produce a phase shift that did not average away. Andrew Christensen reproduced this effect by modulating the retro mirror velocity using a piezo actuator during a Bragg diffraction pulse, confirming that this mechanism produced the observed phase dependence.

To characterize these vibrations we built an optical Michelson interferometer, where one arm was directed vertically in the experiment chamber and reflected off of the retro mirror (slightly more than 8 m round trip) and the other arm had a round trip of a few centimeters. A spare DBR laser was used to provide interferometry light, and a Red Pitaya-based PI loop was used to keep the interference fringe at a constant intensity by modulating the DBR wavelength. The PI loop provided current feedback, which shifted the laser wavelength enough to keep the phase difference between the two arms constant.

The phase difference between the two arms is given by the difference in the number of waves in each arm. Because the arm reflecting off the retro mirror is much longer than the other arm, this difference is approximately

$$\Delta n_\lambda = L/\lambda = 8\,\mathrm{m}/852\,\mathrm{nm} \tag{7.7}$$

Because the mirror vibrates, we must take L to be time dependent. To keep Δn_λ constant we must modulate λ in time.

$$\Delta\lambda = \lambda(t) - \lambda_0 = \frac{c}{f_0 + \Delta f(t)} - \lambda_0 = \frac{c/f_0}{1 + \Delta f(t)/f_0} - \lambda_0 \approx -\lambda_0 \Delta f(t)/f_0 \tag{7.8}$$

Therefore, because the PI loop keeps Δn_λ constant, we find

$$\Delta L(t) = \Delta n_\lambda \Delta\lambda(t) = -\Delta n_\lambda \lambda_0 \Delta f(t)/f_0 \tag{7.9}$$

By measuring the frequency shift of the DBR, the time-dependent path length was read out. This is equal to twice the displacement of the retro mirror. The recorded displacement as a function of time averaged over 50 shots is shown in Fig. 7.3. Note that the mirror velocity spans a range on the order of 0.1 mm/s, which, from Eq. (6.5), we see produces frequency shifts on the order of $0.1\omega_r$. From Fig. 2.8 we see this is substantial enough to introduce phase shifts between 10 mrad and 100 mrad, consistent with what was observed experimentally.

By driving the retro mirror using the piezo actuators described in Section 4.1.5 over a range of drive frequencies, it was determined that the mirror or mirror mount had a mechanical resonance that was not observable using an accelerometer placed on the top of the experiment chamber. This resonance was excited by the turn off of the original 3D MOT coils described in Section 4.1.2, and produced a frequency shift on the retroreflected beam as described previously. It is likely that these vibrations reduced the experiment averaging

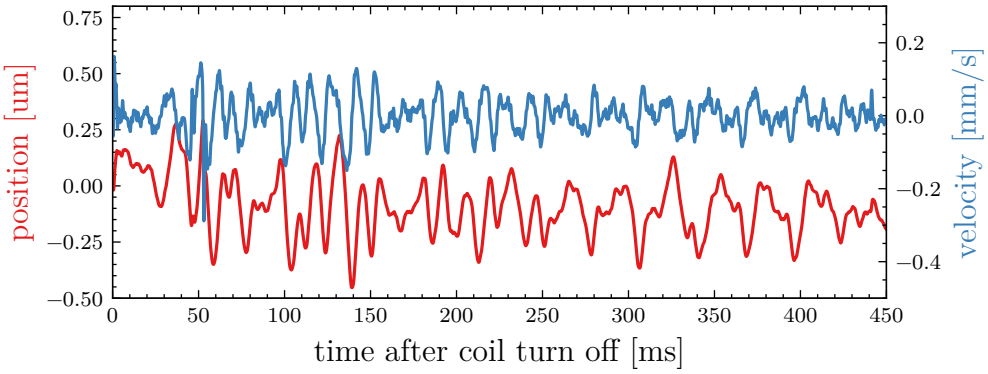


Figure 7.3: The position and velocity of the retro mirror as a function of time after the 3D MOT coil turn off. The position is recorded by an optical Michelson interferometer placed underneath the 3D MOT chamber, and the velocity is determined by numerically taking the derivative of the position.

down rate. It is also possible that they were responsible for some of the shot-to-shot variation of the Bragg beamsplitter pulses described earlier. However, it was difficult to make a clear comparison of this shot-to-shot variation before and after the 3D MOT coil replacement.

7.1.3 Wavefunction width

Another systematic effect was found while developing the stacked ODE simulation method described in Section 6.2.2. By scanning the width of the momentum space wavefunction integrated using the stacked ODE method, we found the interferometer phase depended on the wavefunction momentum width. The resonance condition for the Bragg beamsplitter is only met for a single momentum. Since the wavefunction has width in momentum space, it cannot be that the entire wavefunction is resonant with the Bragg beamsplitter. Instead, two photon detuning varies across the wavefunction, introducing a systematic shift through the same mechanism described in the previous two sections and shown in Fig. 2.8.

The stacked ODE method was used to investigate the dependence of the interferometer phase on the wavefunction width. A Ramsey-Bordé interferometer composed of 4th-order Bragg pulses with $\sigma = 0.26\omega_r^{-1}$, $T = 64.9\omega_r^{-1}$, $T' = 64.9\omega_r^{-1}$, $\Omega_{\text{eff},0} = 20\omega_r$, $\sigma_{\text{cloud}} = 1.5\,\text{mm}$, and $\sigma_{v_T,\text{cloud}} = 3.5\,\text{mm/s}$ was simulated. For cesium, these numbers correspond to a $\sigma = 20\,\mu\text{s}$, $T = 5\,\text{ms}$, and $T' = 5\,\text{ms}$ interferometer. The cloud parameters are used to sample 1000 atoms that are then propagated using the stacked ODE method with $S = 1102$ and $J = 9$. Each atom in the cloud is initialized with a Gaussian wavefunction centered at

$n_0 = 0$. This leads to the following discretized momentum space wavefunction:

$$\phi_{j,s} = \frac{1}{\sqrt{\sigma_{\text{wf}}\sqrt{\pi}}} \exp\left(-\frac{(2j + s/S)^2}{2\sigma_{\text{wf}}^2}\right) \tag{7.10}$$

where σ_{wf} is the Gaussian amplitude-width in units of $\hbar k$; the corresponding momentum probability distribution has standard deviation $\sigma_{\text{wf}}\hbar k/\sqrt{2}$. Each atom's position and velocity is drawn independently from a normal distribution centered at the origin. Atom positions are sampled via $x, y \sim \mathcal{N}(0, \sigma_{\text{cloud}}^2 = (1.5\,\text{mm})^2)$. Transverse atom velocities are sampled via $v_x, v_y \sim \mathcal{N}(0, \sigma_{\text{transverse}}^2 = (3.5\,\text{mm/s})^2)$, and longitudinal velocities are sampled via $v_z \sim \mathcal{N}(0, \sigma_{\text{v}_z}^2 = (0.35\,\text{mm/s})^2)$. The laser beam has a $1/e^2$ beam radius of $w_0 = 6.2\,\text{mm}$. These are realistic parameters for the atomic fountain described in this thesis. The output wavefunctions are converted to momentum state populations using Eq. (6.23).

The results of this simulation are shown in Fig. 7.4. To isolate the wavefunction width effect from the effect of parasitic momentum states, a "mask" is applied in momentum space that removes unwanted momentum states from the simulation after each Bragg pulse. Comparing the "masked" and "unmasked" interferometer simulations gives insight into how much of the interferometer phase is due to parasitic interferometers. The wavefunction width effect is responsible for the phase deviations measured in the "masked" case. The simulation is performed for both a single atom and an atom cloud propagating through the interferometer. From the figure we see that in all cases there is a dependence on the momentum space wavefunction width on the order of 10 mrad for the wavefunction width selected by the VS process used in the experiment (see Section 4.4.6).

Note that the crossing of the single atom and cloud cases occurs at the location where the wavefunction width is equivalent to the cloud width. That is, because each atom is represented by a wavefunction with an independent center in momentum space, the ensemble of atoms has a velocity width set by this distribution of centers. When the distribution of wavefunction centers is equal to the wavefunction width, the single atom and cloud cases cross. Note that this is related to the discussion in Section 6.3.

7.1.4 Parasitic interferometers

The 2018 Berkeley fine-structure measurement suppressed parasitic interferometers by determining a "magic duration" for Bragg diffraction, where parasitic ports are not heavily populated. This was reported in Fig. (12) of [100], where it is shown that for a particular σ a 5th-order Bragg pulse does not populate the $p = 2\hbar k$ momentum state neighboring the initial $p = 0\hbar k$ momentum state. This result is replicated for a 5th-order Bragg pulse in Fig. 7.5, where we see that when $\sigma = 0.22\omega_{\text{r}}^{-1}$, we suppress the population in the $p = 2\hbar k$ state. The loss (defined as $1 - |\langle 0\hbar k|\phi\rangle|^2 - |\langle 10\hbar k|\phi\rangle|^2$ for 5th-order Bragg diffraction) is also displayed. Note that for each σ in this plot, $\Omega_{\text{eff},0}$ is optimized to produce a $\pi/2$ pulse.

Unfortunately, this "magic duration" does not suppress the population of parasitic momentum states for all intensities. This is shown in Fig. 7.6, where $\Omega_{\text{eff},0}$ is scanned for the

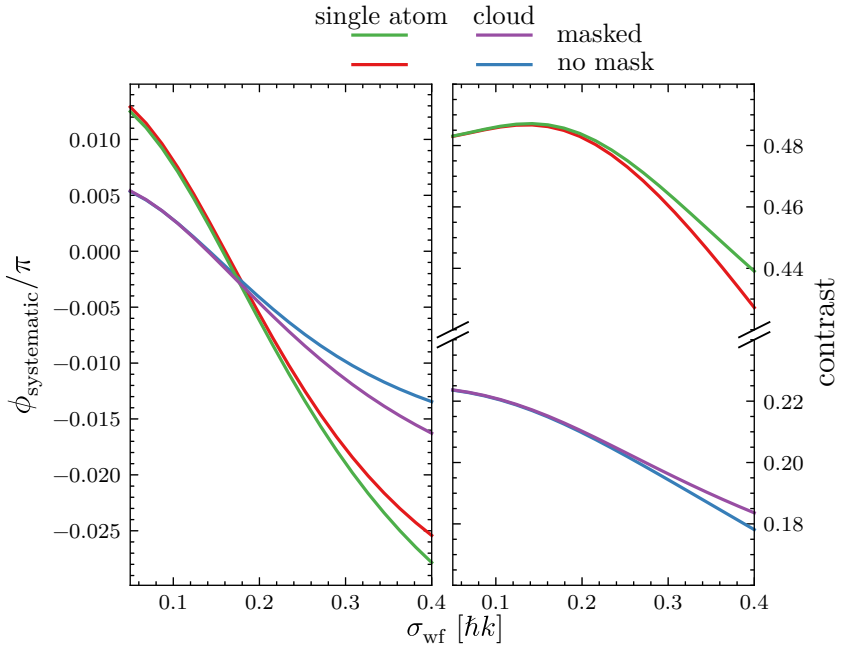


Figure 7.4: Simulated interferometer phase and contrast as a function of wavefunction width for a single atom and cloud of atoms. "Masked" denotes an interferometer where parasitic momentum states are removed from the simulation after each beamsplitter pulse. The phase is reported as $\phi_{\text{systematic}}$, which denotes the deviation of the interferometer phase from the expected phase for a Ramsey-Bordé interferometer.

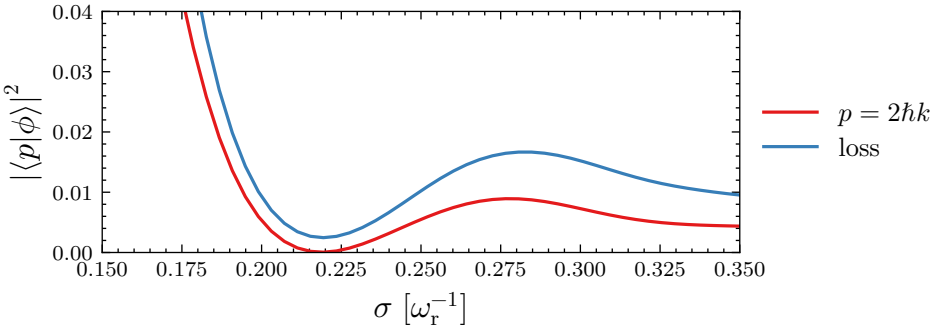


Figure 7.5: Simulated population in the $2\hbar k$ momentum state for a 5th-order Bragg diffraction pulse as a function of the pulse width σ. The loss is any population not in $0\hbar k$ or $10\hbar k$. The figure is intended to replicate Fig. (12) of [100].

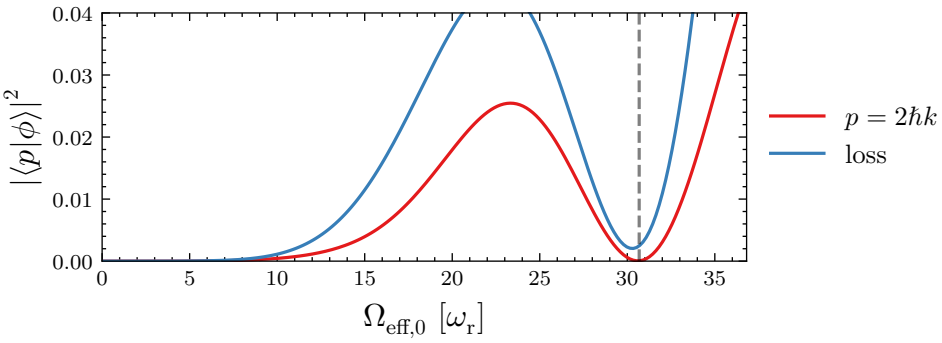


Figure 7.6: Simulated population in the $2\hbar k$ momentum state for a 5th-order Bragg diffraction pulse with $\sigma = 0.22\omega_r^{-1}$ (which suppressed parasitic populations in Fig. 7.5) as a function of the peak effective Rabi frequency. The loss is any population not in $0\hbar k$ or $10\hbar k$.

magic pulse duration of $\sigma = 0.22\omega_r^{-1}$ found in Fig. 7.5. The vertical dashed line in the plot corresponds to the value of $\Omega_{\text{eff},0}$ that produces a $\pi/2$ pulse when $\sigma = 0.22\omega_r^{-1}$. The plot reveals that for intensities lower than the optimal, substantial population of the $p = 2\hbar k$ state can occur. For this reason, parasitic interferometers cannot be fully suppressed through choice of σ, and should be characterized in simulation. As discussed in Section 6.2.1, the tree-based simulation is able to capture the effect of parasitic interferometers.

7.1.5 Bragg pulse matching

Before developing the tree-based simulation framework described in Section 6.2.1, I compared simulated and experimental single Bragg pulses. This was a simpler approach than immediately attempting to simulate a full interferometer, as the simulation was much less complex. Furthermore, the simulation methods used to simulate a full interferometer must also replicate a single Bragg pulse accurately, and so this comparison served as a necessary intermediate test.

One of the most basic comparisons that can be made is of the populations output by a single Bragg pulse as a function of pulse intensity. In order to make this comparison, the time-of-flight traces recorded by the detection system must be converted into momentum state populations. The first step in extracting the populations is assigning an arrival time to each detected momentum. As an example consider Fig. 7.7. First a peakfinder is used to determine the approximate peak locations. In the figure three peaks would be detected–one at $p = 0\hbar k$, one at $8\hbar k$, and one at $6\hbar k$. The peakfinder finds the sample index of the peak. To allow for the possibility that the peak maximum was located between ADC samples, a Gaussian is fit to the data points surrounding the detected peak, and the fitted Gaussian center is stored as the peak time for that particular peak. This process is performed for each

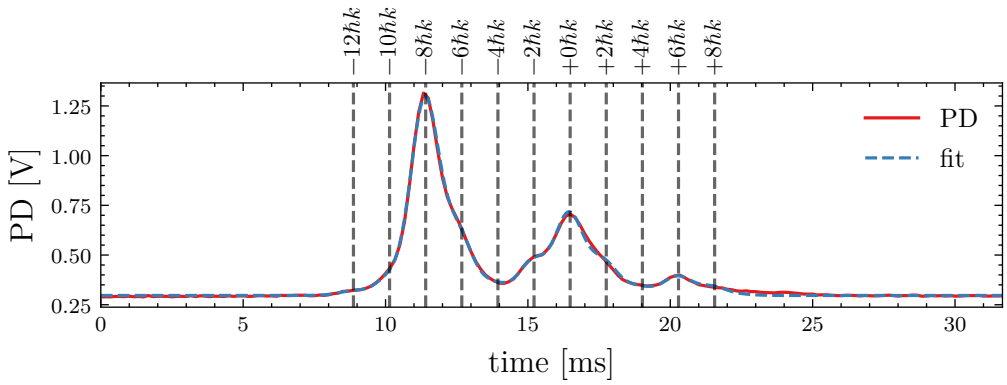


Figure 7.7: Method by which momentum state populations after a Bragg pulse are measured. The PD trace is fit to a sum of Gaussians centered at each of the vertical dashed lines. Each line corresponds to a particular momentum state (labeled at the top of the figure). The extracted Gaussian peak height is taken to be the population in that state.

time-of-flight trace independently.

From the launch velocity, Bragg pulse timing, and location of the detection system, the expected timing between detected momenta can be estimated. These estimates are used to cluster the extracted peak timings into bins associated with each momentum. Each bin can then be averaged to determine the expected arrival time $\hat{t}_j$ for its corresponding momentum.

Several variants of this process exist. For example, the peakfinder and Gaussian fits can be run on an average of many time-of-flight traces. Or the peakfinder can be run on an average of many time-of-flight traces, while the Gaussian fit is still run on a per-trace basis. The averaged arrival times in each bin can also be constrained by kinematics. This is done by minimizing[1]

$$\sum_j |\hat{t}_j - t_j| \tag{7.11}$$

where $\hat{t}_j$ are the arrival times computed by averaging across each bin, and t_j are the arrival times computed from kinematics. Here t_j is given from kinematics as

$$t_j = \frac{v_j + \sqrt{2g(h_{\mathrm{BS}} - h_{\mathrm{det}}) + v_j^2}}{g} \qquad v_j = v_{\mathrm{BS}} + \frac{2j\hbar k}{m} \tag{7.12}$$

The minimization is carried out by varying $h_{\mathrm{BS}} - h_{\mathrm{det}}$, v_{BS}, and (optionally) g. The exact combinations of these options used to determine a mapping of the momentum $2j\hbar k$ to an

[1] Alternatively, the sum of squared errors could be used as a loss function. This would penalize large residuals produced by outliers more strongly. In practice, this choice does not greatly affect the fit of the simulation to experiment data.

arrival time depend on the dataset under consideration. For example, in a full interferometer it is easier to use $\hat{t}_j$ than t_j, as the exact form for t_j is nontrivial to derive.

Once a set of timings has been determined, each time-of-flight trace is fit to a sum of Gaussians with centers located at each peak arrival time. This is seen in Fig. 7.7, where a Gaussian is fit at each of the 11 arrival times (denoted by dashed vertical lines). A curve fitting algorithm is then used to fit the peak height of each Gaussian. All Gaussians share a σ value, which is also fit. The peak height of each Gaussian sets the relative population in its corresponding momentum state. This process produces relatively good fits to experiment data, as seen in Fig. 7.7, where the fit line follows the detected PD trace quite well.

There are several approximations made in this method. First, not all peaks are described by Gaussians. In reality the peaks should be asymmetric due to their acceleration under gravity [101–103]. This can be partially accounted for by fitting the peak after VS to a sum of some set of basis functions, i.e. the Hermite polynomials, and then fitting the resulting curve to the data. Andrew Christensen proposed this approach and it seems to work quite well. Second, in certain regimes the width of the peak selected by VS will be different from the Bragg diffracted peaks, as the momentum bandwidth of each process is different.

Once the peaks have been extracted, the measured momentum state populations can be plotted. For example, in Fig. 7.8 we show the measured momentum state populations as a function of the peak effective Rabi frequency $\Omega_{\text{eff},0}$ for an $n = 1$ Bragg pulse. The figure omits many of the momentum states extracted by the fitting algorithm for clarity[2]. Our first goal is to simulate the experimentally observed populations using the simulation methods developed in Chapter 6.

We integrate Eq. (6.2) to produce the basic simulation predictions in Fig. 7.8. From the figure, it is clear that this integration does not capture the experimentally measured contrast decay. Contrast decay can come from a variety of sources not included in this single integration, such as vertical momentum spread of the atom cloud/wavefunction or transverse position spread. The decay observed in the figure can be explained by the transverse spread of the atoms. We now describe how to model a cloud of atoms with transverse spread.

Atom positions are described by a normal distribution with width σ_{cloud}. Because $I \propto \Omega_{\text{eff}}$, the transverse intensity profile of the beam is used to map positions to Rabi frequencies. The intensity profile is taken to be that of an ideal Gaussian beam, so

$$\Omega_{\text{eff}}(t, x, y) = \Omega_{\text{eff},0}(t) \exp\left(-2\frac{x^2 + y^2}{w_0^2}\right) \tag{7.13}$$

The intensity as a function of time is given by a model of the Bragg pulse shape. Recall from Section 4.2.5 that the intensity of the interferometry beam is set using an AWG programmed with a particular pulse shape. To scan $\Omega_{\text{eff},0}$ as done in Fig. 7.8, the AWG amplitude is scanned. Pulse shapes measured before the beam enters the chamber for two values of $\Omega_{\text{eff},0}$ (denoted by their corresponding AWG amplitudes) are displayed in Fig. 7.9. From the figure we see that the Gaussian profile in time sits on top of a slight pedestal. Therefore our model

[2]This is apparent as the populations do not sum to 1 for all $\Omega_{\text{eff},0}$.

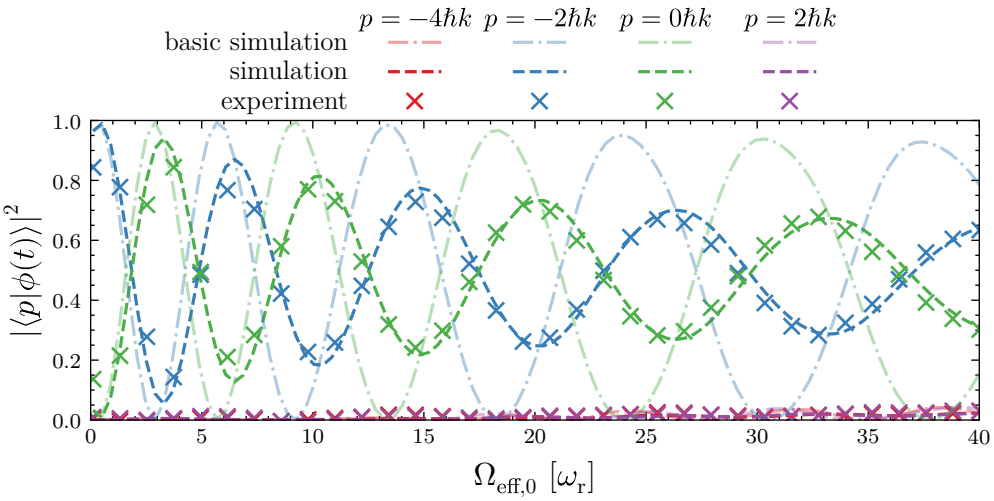


Figure 7.8: Experimentally measured momentum state populations as a function of $\Omega_{\text{eff},0}$, along with a simulation of an atomic cloud undergoing Bragg diffraction (solid dashed lines). A basic simulation of a single atom undergoing Bragg diffraction in a uniform beam is also shown (transparent dashed lines).

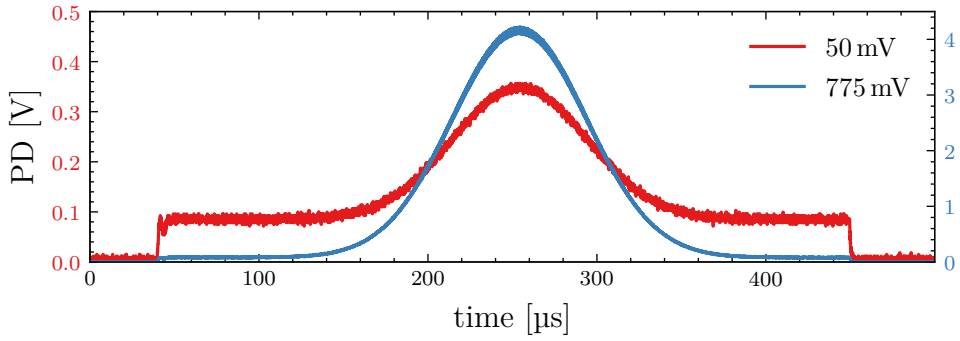


Figure 7.9: Measured Bragg pulse waveforms for two different AWG pulse shape amplitudes. The Gaussian waveform sits on a pedestal, which must be considered for the simulated Bragg dynamics to match the experimentally measured dynamics. The left y-axis shows the measured photodiode voltage for the 50 mV amplitude AWG pulse, and the right y-axis shows the measured photodiode voltage for the 775 mV amplitude AWG pulse.

for the Bragg pulse shape is given by a Gaussian pulse with a constant offset. We also need to know the two photon detuning δ to integrate Eq. (6.2), which is assumed to be on resonance for an $n = 1$ Bragg beamsplitter.

To integrate Eq. (6.2) over the atom cloud we could either sample atom positions, or solve over a grid of positions and weight the output populations according to the position distribution. In this case we use the latter approach. This produces the simulated population curves displayed in Fig. 7.8, which we see agree well with the experimentally measured populations. The remaining deviations between the simulation and experiment are either due to measurement error or effects not yet included in the simulation (i.e. vertical momentum spread, etc.). Note that the x-axis scale is determined from a fit to the simulation. For higher order Bragg pulses modeling the previously mentioned pulse pedestal is not as important, as the higher order of these processes requires a much larger $\Omega_{\text{eff},0}$ to start transferring momentum.

It is important to ensure that the simulation is able to replicate a variety of experimental parameters, beyond just our choice of $\Omega_{\text{eff},0}$. To test this, a three-dimensional scan over $\Omega_{\text{eff},0}$, δ, and the timing of the first Bragg pulse is performed. As atoms in the cloud have some transverse temperature, scanning the timing is effectively a scan over cloud size (and potentially center position relative to the laser beam). In Fig. 7.10 we show the measured population inversion as a function of $\Omega_{\text{eff},0}$ and this delay for 1st through 4th-order Bragg diffraction from this dataset. The arbitrary units of $\Omega_{\text{eff},0}$ are the pulse amplitudes programmed into the AWG in millivolts. The unit is listed as arbitrary as the constant of proportionality between the programmed Bragg pulse amplitude and $\Omega_{\text{eff},0}$ is unknown.

We fit a simulation to the entire three dimensional dataset by minimizing

$$\sum_C \sum_j \left| \widehat{P}_{j,C} - P_{j,C}(\theta) \right| \tag{7.14}$$

where $\widehat{P}_{j,C}$ and $P_{j,C}$ denote the measured and simulated populations in condition C, respectively. The sum is minimized by varying the following free parameters in the simulation:

1. The constant of proportionality between the Bragg pulse amplitude programmed into the AWG and $\Omega_{\text{eff},0}$.
2. The position offset of the cloud from the beam center at time zero (taken to be in direction $\hat{x}$).
3. The net transverse velocity of the cloud.
4. The angle of the transverse velocity of the cloud relative to $\hat{x}$.
5. The transverse velocity spread of the cloud.
6. The spatial extent of the normal distribution describing the cloud at time zero.
7. The detuning of the cloud center velocity from Bragg resonance.

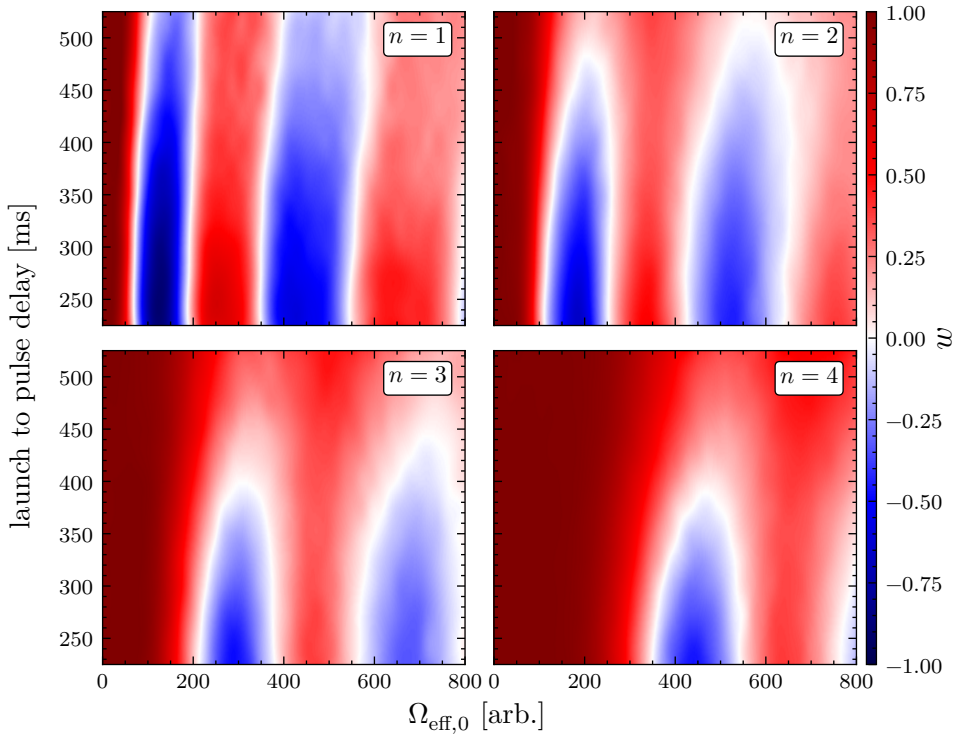


Figure 7.10: Measured population inversion resulting from 4th-order Bragg diffraction as a function of $\Omega_{\text{eff},0}$ and the Bragg pulse delay.

8. The momentum width of the cloud in $\hat{z}$ ($\hat{z}$ is the direction of the interferometer).

This particular dataset produces a sum over ∼20000 different conditions C. To reduce the fitting time, 400 conditions were randomly sampled and used in the minimization. The difference between the experimentally measured inversion $w_{\text{experiment}}$ and the simulated inversion $w_{\text{simulation}}$ is computed. The resulting differences are shown in Fig. 7.11. In this simulation the Bragg pulse pedestal was not included, which we see causes the simulation to underestimate the inversion for low $\Omega_{\text{eff},0}$ for the $n = 1$ Bragg pulse. In Fig. 7.12 we show the same experiment-simulation comparison as a function of δ and $\Omega_{\text{eff},0}$ for a single Bragg pulse delay.

The simulation clearly does not perfectly match the experiment, with max discrepancies of approximately 0.75 occurring between the experimentally measured and simulated inversions for $n = 1$. For $n \neq 1$ and $\Omega_{\text{eff},0} < 400$ (these are the experimentally relevant cases), the maximum observed discrepancy is closer to 0.15. In this simulation, the Bragg pulse pedestal shown in Fig. 7.9 was ignored. Ignoring the pedestal is particularly detrimental for the $n = 1$ simulations, since the pedestal is high enough to partially drive 1st-order

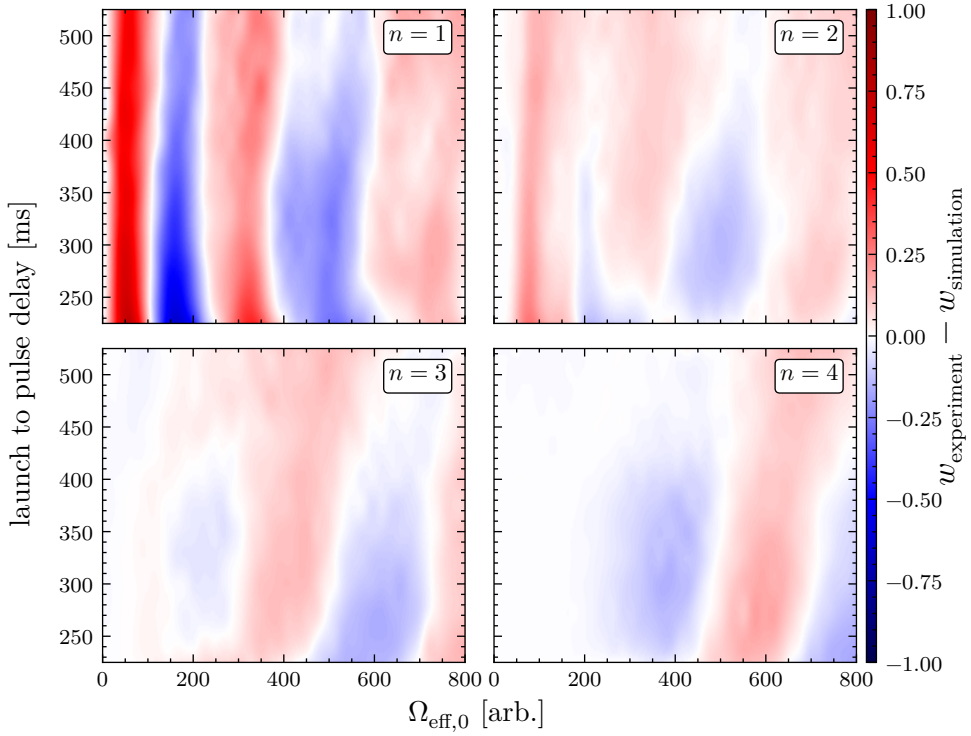


Figure 7.11: Residuals between the measured and simulated Bragg diffraction inversion as a function of $\Omega_{\text{eff},0}$ and the Bragg pulse delay. The measured inversion is displayed in Fig. 7.10.

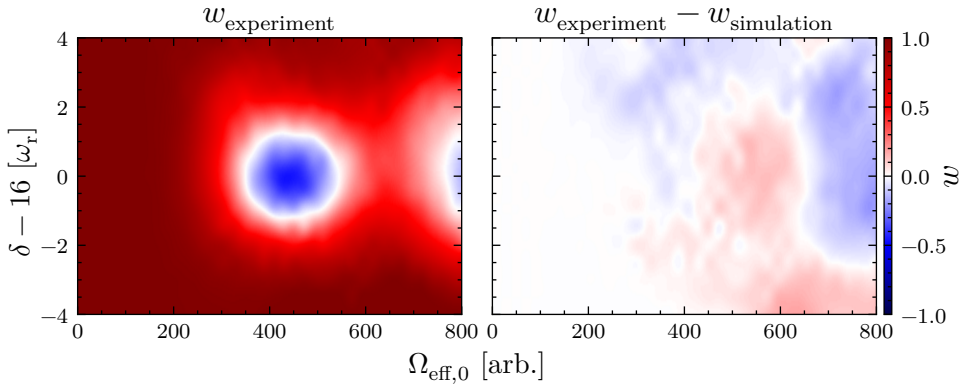


Figure 7.12: *Left:* Measured population inversion from 4th-order Bragg diffraction as a function of $\Omega_{\text{eff},0}$ and δ. *Right:* Difference between the experimentally measured and simulated population inversions.

Bragg diffraction. This disagreement is particularly apparent in Fig. 7.11. Furthermore, the simulation assumed unity detection efficiency and the dataset was not recorded at the magic wavelength, meaning that the transverse position might be slightly correlated with detuning from the Bragg resonance across the cloud. All of these effects will reduce the quality of the experiment-simulation comparison. All have since been incorporated in the simulation codebase.

One goal of the experiment-simulation comparisons should be to minimize the number of free parameters in the simulation, either through direct measurement in the experiment or by using simpler experiment-simulation comparisons to extract parameters that can be used in more complex experiment-simulation comparisons. For example, in the above simulation the interferometry beam size was fixed. We will come back to this idea in the next section, when we will discuss using single Bragg pulse experiment-simulation comparisons to produce parameters for full ellipse experiment-simulation comparisons.

To further test our ability to simulate Bragg pulses, we can simulate sequences of Bragg pulses that produce trajectories that do not close. For example, we can perform the first two or three Bragg pulses in an SCI and check if the simulated populations match the experimentally measured ones. We can even simulate four Bragg pulses and slightly delay the final Bragg pulse so that the interferometer does not close. A dataset was recorded with two, three, and four 4th-order Bragg pulses. In each case, a 2D scan over the peak effective Rabi frequencies $\Omega_{\text{eff},0}^{(1,2)}$ and $\Omega_{\text{eff},0}^{(3,4)}$ was performed, where $\Omega_{\text{eff},0}^{(1,2)}$ and $\Omega_{\text{eff},0}^{(3,4)}$ denote the pulse amplitudes used for the single and multi-frequency Bragg pulses, respectively. Note that in the two Bragg pulse condition $\Omega_{\text{eff},0}^{(3,4)}$ has no effect.

The tree-based interferometer simulation method described in Section 6.2.1 is used to reproduce the experiment output. The simulation is fit to the results of the two Bragg pulse case. Those fitted parameters are then used to predict the experiment output for the four Bragg pulse case and are shown in Fig. 7.13. As in the single Bragg pulse plots, the arbitrary units of $\Omega_{\text{eff},0}$ correspond to the pulse amplitudes programmed into the AWG in millivolts. We see that, similarly to the single Bragg pulse comparison, the simulation agrees relatively well with the experiment, although it is not perfect. Again, this is likely due to the assumption of unity detection efficiency and operation away from the magic wavelength.

These experiment-simulation comparisons show that the tree-based interferometer method is capable of simulating certain aspects of a full ellipse. In particular, we expect the method to faithfully reproduce the center of the ellipse in w_- and w_+. We also should expect the simulation to predict any contrast loss stemming from Bragg pulse dynamics, as well as any phase shifts resulting from the Bragg pulse dynamics. We expect even better agreement when operating at the magic wavelength, pedestals are included, and the position dependent detection efficiency is accounted for. Note that a similar experiment-simulation comparison was performed in [46].

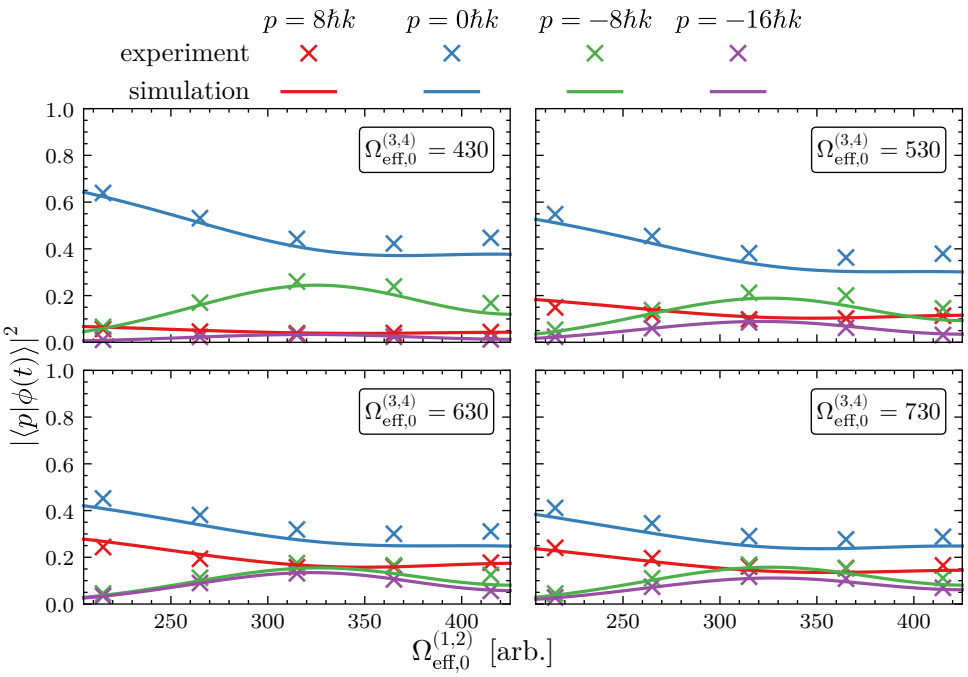


Figure 7.13: Measured populations after an SCI beamsplitter sequence, where the last beamsplitter pulse is delayed in time to prevent interference. Populations are shown as a function of the peak Bragg pulse intensity for the first two ($\Omega^{(1,2)}_{\text{eff},0}$) and last two ($\Omega^{(3,4)}_{\text{eff},0}$) Bragg pulses. The simulation parameters are determined by fitting a dataset with the first two Bragg beamsplitters to a simulation, and then using those parameters to simulate the more complex four beamsplitter case shown here.

7.1.6 Bloch matching

In a similar vein to the Bragg pulse matching, we perform a single Bloch pulse, extract the resulting momentum state populations (using the same method as in Bragg), and compare them with simulation. This comparison is shown in Fig. 7.14. Because Bloch oscillations occur over a much longer duration than Bragg diffraction, single photon scattering must be accounted for, and so the density matrix is integrated using the von Neumann equation

$$i\hbar\frac{\partial\rho}{\partial t} = \left[\hat{H}_{\text{eff}}, \rho\right] - i\hbar\Gamma\mathbf{X}\rho/2 \tag{7.15}$$

where $\hat{H}_{\text{eff}}$ is given by Eq. (2.15) and $\mathbf{X}_{i,j} = \delta_{i,j+1} + \delta_{i,j-1}$. The resulting ODEs are integrated using the DOP853 solver for a cloud of atoms, as described in Chapter 6, using a scattering rate Γ computed using Eq. (1.83) from [50]. The simulation cloud parameters are then fit to the experiment results. In the case considered here, $N = 10$ Bloch oscillations are used with $\tau = 7.2\,\text{ms}$. This simulation approach does a relatively good job of replicating the

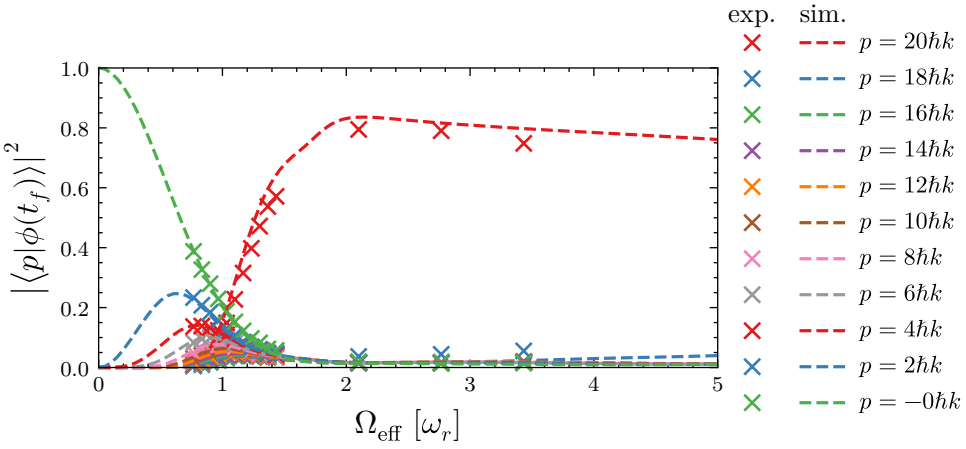


Figure 7.14: Measured and simulated populations after a single $N = 10$ Bloch oscillation pulse with $\tau = 7.2\,\text{ms}$.

experimentally measured results, correctly predicting the population of intermediate states between $0\hbar k$ and $20\hbar k$. The populations measured in the $20\hbar k$ state are overestimated by the simulation. This might be due to the fact that Γ is directly computed instead of measured.

7.1.7 Ellipse matching

To simulate ellipses we use the tree-based simulation described in Section 6.2.1. We record a dataset similar to the multiple Bragg pulse dataset described in Section 7.1.5, where a 2D scan over the peak effective Rabi frequencies $\Omega_{\text{eff},0}^{(1,2)}$ and $\Omega_{\text{eff},0}^{(3,4)}$ was performed. These frequencies are reported in arbitrary units corresponding to the amplitudes programmed into the AWG used to generate the pulse profiles, as discussed in the previous section. In each condition the experimentally recorded ellipse is fit and the ellipse parameters b_-, b_+, A_-, A_+, and ϕ_d (defined in Section 5.1) are extracted. An ellipse is then simulated for each experimental condition. The simulated ellipse is fit and the same ellipse parameters are extracted. The simulation is fit to the experiment by minimizing the difference between each of these five parameters across all conditions. Note that neither the experiment nor the simulation uses Bloch oscillations in this comparison. The fit reports very small position offsets of the cloud from the beam center (under 1 mm).

The same parameters are allowed to vary in this simulation as were allowed to vary in the simulation described in Section 7.1.5, the only difference being that the initial extent of the interferometry cloud is allowed to vary in each direction as opposed to being the same in each direction. The experimentally measured and simulated ellipses are shown in Fig. 7.15. We see that the simulation does a good job of capturing the location and extent of each

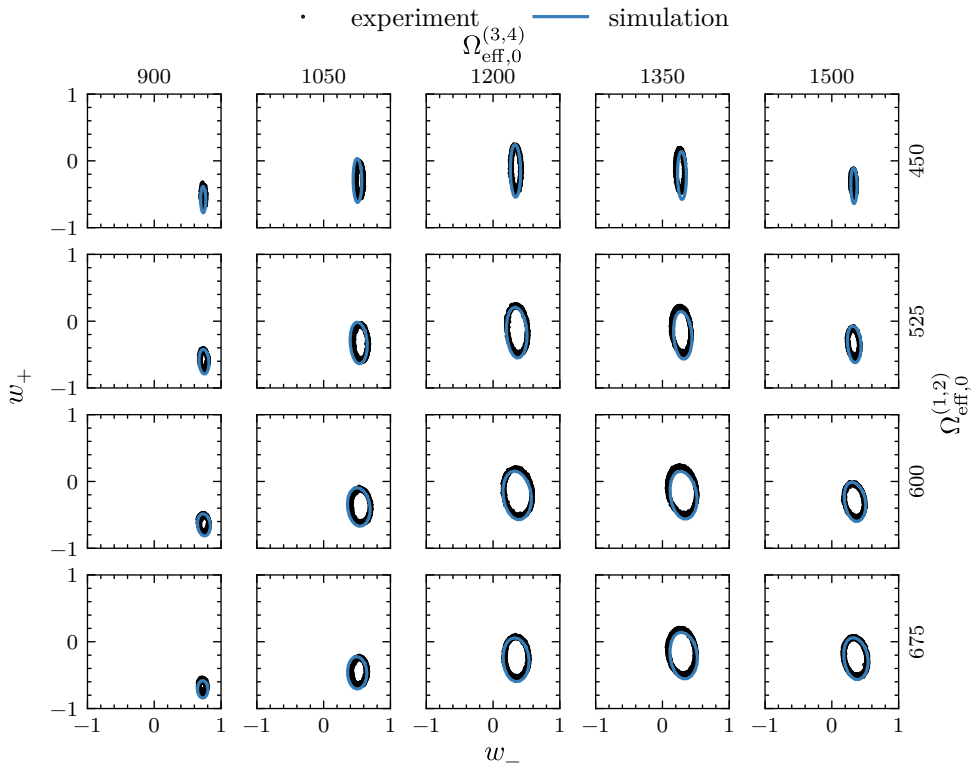


Figure 7.15: Experimentally measured and simulated SCI ellipses.

ellipse in w_- and w_+. This is a sign that the simulation properly accounts for the leading sources of contrast loss, as well as effects that can shift the ellipse center. The ellipse is primarily shifted through the existence of junk ports that produce non-interfering baselines of population that the SCI output ports are detected on top of. This shift is partially mitigated by Bloch oscillations, which shift the output momenta away from the junk port momenta.

While the ellipse location and extent are well captured by the simulation, the ellipse angle does not agree as closely. This is apparent from Fig. 7.16, where the fitted parameters are plotted for each condition. We see that b_-, b_+, and A_- agree relatively well between simulation and experiment. However, the predicted contrast in the upper output ports given by A_+ shows relatively large disagreement between simulation and experiment (although it qualitatively seems to follow the same trend), and, more worryingly, large disagreements persist in the ellipse angle/interferometer phase ϕ_d. I do not report a quantitative measure of fit here because the error in the interferometer phase ϕ_d is at least $0.1\pi \approx 300\,\text{mrad}$, which is much larger than the desired ϕ_d accuracy of 1.72 mrad (equivalent to 3.44 mrad in the full interferometer phase Φ, as discussed in Section 5.4). It is clear that either the simulation does not capture all effects at play or systematic shifts were introduced during the analysis

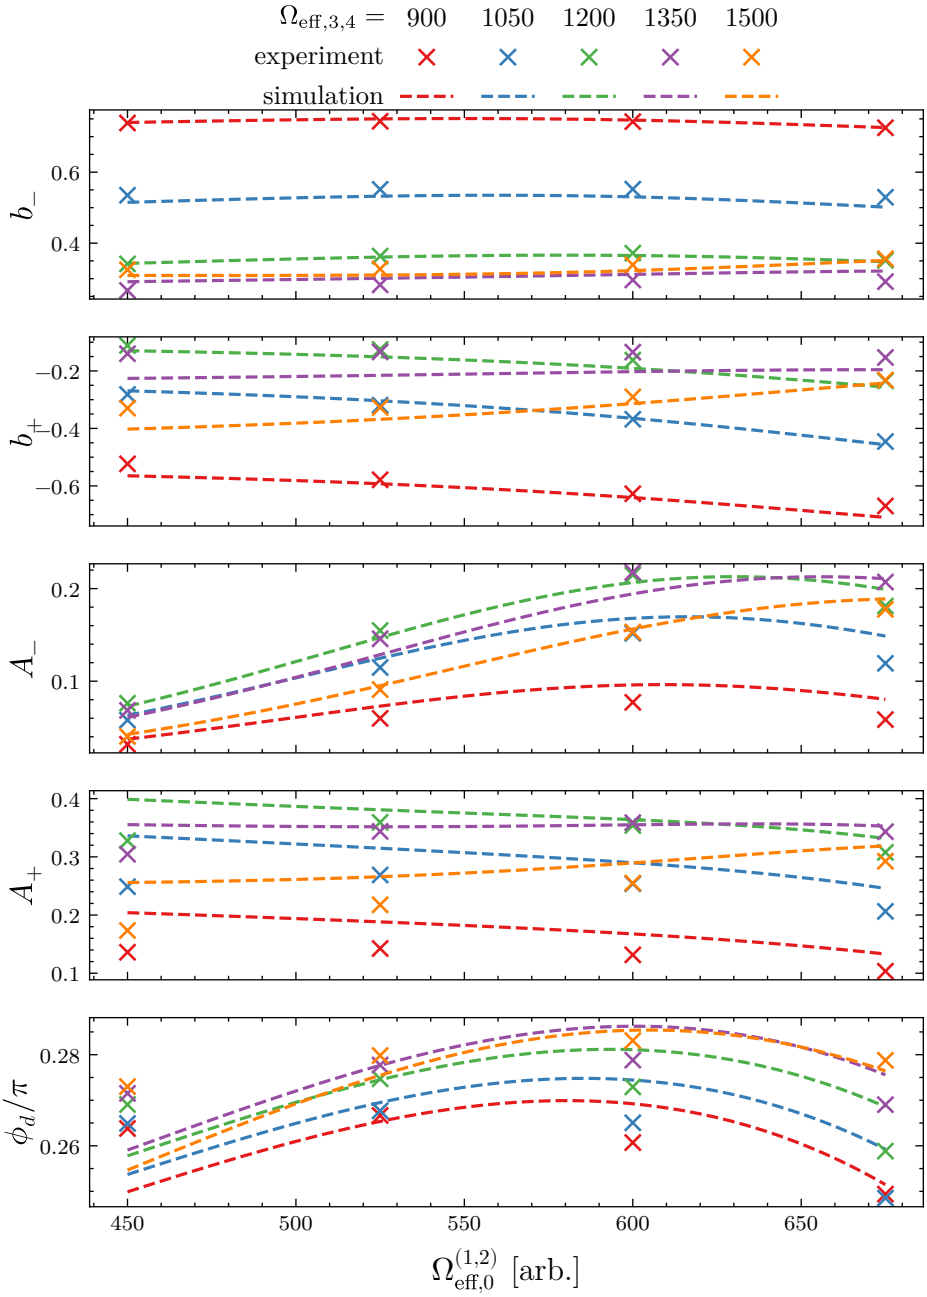


Figure 7.16: Experimentally measured and simulated SCI ellipse parameters extracted via a fit.

of the experiment data.

The simulation used here is the same as the one used in the previous section and therefore assumed unity detection efficiency. The data were recorded away from the magic wavelength, and the simulation did not include free evolution phase (i.e. gravity gradients, etc.), Gouy phase, separation phase, or parasitic interferometers. The tree-based simulation codebase has since been updated to handle all of these effects in unison, and so future comparisons should account for all of these properly.

Furthermore, the minimization routine used in the fit is a naive choice, as it treats b_-, b_+, A_-, A_+, and ϕ_d on equal footing. To measure the recoil frequency, it is much more important to predict ϕ_d accurately than the other parameters. Also, as mentioned in the previous section, ideally we don't fit the simulation to any phase-sensitive measurement process, as that raises the slight possibility that we are just fitting the phase by having enough free parameters instead of really predicting it. To avoid fitting the simulation, we must extract the free parameters used in the simulation by performing other experiment-simulation comparisons. The cloud size and expansion rate, for example, can be captured by running a single Bragg pulse at multiple delays (effectively scanning the cloud size), and fitting a simulation of a single Bragg pulse to it. This approach ensures parameters are consistent across different experimental tests, reduces the number of free parameters, and increases our confidence that the simulation captures all effects present in the interferometer.

7.2 Methods

7.2.1 Novel interferometer geometries

In addition to the SCI used in the experiment, several other interferometer geometries sensitive to ω_r can be investigated. We discuss several geometries and then review several simulations of systematic effects present in each.

Simultaneous Mach-Zehnder interferometer

The simultaneous conjugate Mach-Zehnder interferometer (SCMZI) is similar to the simultaneous conjugate Ramsey-Bordé interferometer (SCRBI), but where the direction of final two Bragg pulses has been flipped. This is seen in Fig. 7.17, where we see the SCMZI (not transparent) and SCRBI (partially transparent). The phase accumulated in the SCMZI is

$$\Phi = 16nNT\omega_\mathrm{r} - 2nT\omega_\mathrm{m} \tag{7.16}$$

Later on we will compare the performance of the SCRBI and the SCMZI through simulation, which shows that the SCMZI has some favorable properties over the SCRBI. Note that the SCMZI only accumulates phase when Bloch oscillations are implemented.

The SCMZI was implemented in the experiment, and the ellipses it generated were compared to those of the SCRBI. Both ellipses were found to have the same contrast, which was predicted from simulation (shown later on).

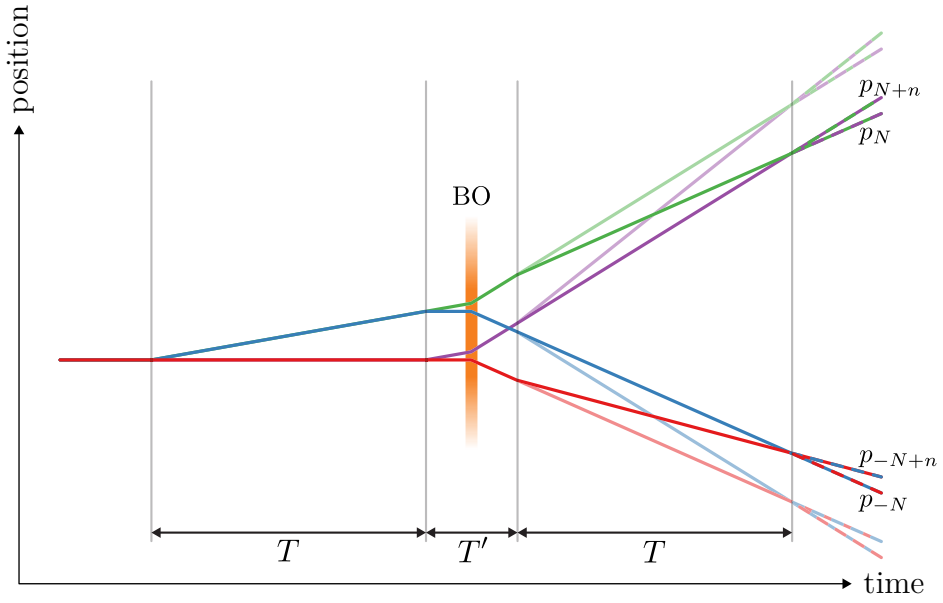


Figure 7.17: The simultaneous conjugate Mach-Zehnder interferometer geometry overlaid on a partially transparent simultaneous conjugate Ramsey-Bordé interferometer geometry. Bloch oscillations (BO) are shown in orange. The trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters are applied.

Multiple $n = 1$ interferometer

Instead of implementing an SCI using a single nth-order Bragg pulse, we could consider implementing an SCI using n 1st-order Bragg pulses. We refer to the target Bragg order as n_{target}. This results in the geometry shown in Fig. 7.18, where a $\pi/2$ pulse is initially used to split the interferometer along two arms, and then successive Bragg pulses further pull these arms apart. The potential advantage of this geometry is that, because 1st-order Bragg pulses require a smaller $\Omega_{\text{eff},0}$ than higher order pulses, this could be implemented with a much larger interferometry laser beam (assuming enough power remains for Bloch oscillations).

The following pulse sequence is used to form the interferometer geometry:

1. Split the initial state into two arms and mirror one of the outputs up to Bragg order n_{target}.
2. Mirror the state at Bragg order n_{target} back down to 1, apply a $\pi/2$ pulse to bring it to the correct state. This also diffracts the initially undiffracted arm to Bragg order 1. Then mirror 1 up to n_{target}.

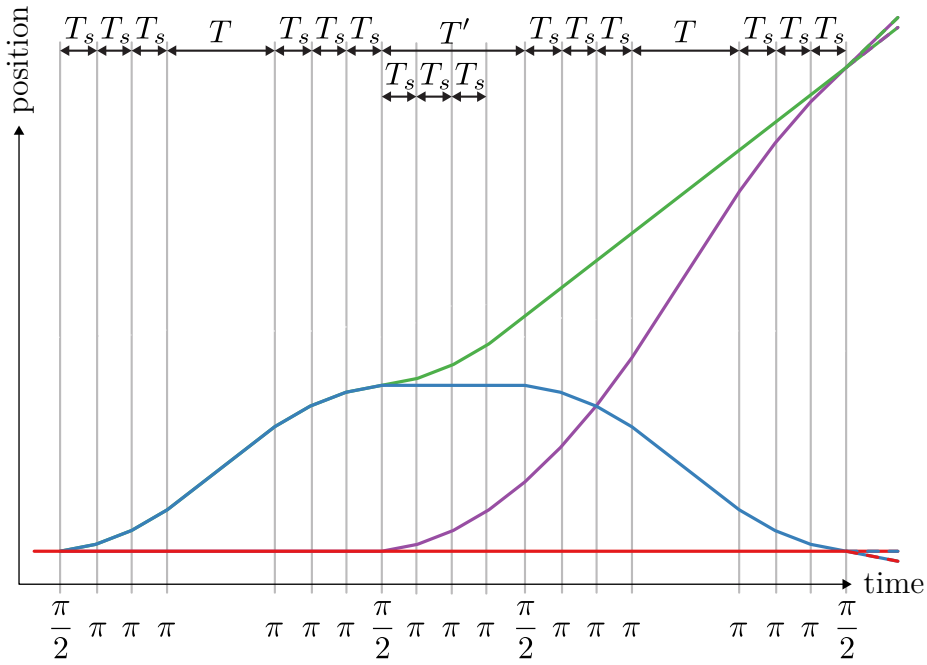


Figure 7.18: An SCI implemented with multiple 1st-order Bragg diffraction pulses. Trajectories are shown without a gravitational potential. Light vertical lines denote the times at which beamsplitters or mirrors are applied. The pulse type ($\pi/2$ beamsplitter or π mirror) is shown underneath each vertical line.

3. Apply a $\pi/2$ pulse to diffract the lower interferometer's upper arm down (the $\pi/2$ is necessary to avoid diffracting away the lower arm), followed by a π pulse to $-n_{\text{target}}$. Simultaneously do the same for the upper interferometer's lower arm.

4. Apply a π pulse to diffract the lower interferometer's upper arm up to 0, then a $\pi/2$ pulse to interfere the 0 and -1 states. Simultaneously do the same for the upper interferometer's lower arm.

This pulse sequence is shown in Table 7.2. In the case where the separation between successive Bragg pulses, T_s, is small relative to T, this geometry accumulates the same phase as a normal SCI. The 1st-order Bragg pulses have no intermediate momentum states, and therefore produce smaller diffraction phases (although, as we will see, this is counteracted by the need to perform more than just 4 Bragg pulses to form the SCI). The ability to use a larger interferometry beam reduces the variation in diffraction phase across the cloud and suppresses Gouy phase.

Note that at the end of the interferometer the output ports of each interferometer are only spaced by $2\hbar k$ in momentum, which means that their arrival times are much closer in

Table 7.2: Pulse sequence for the stacked first-order Bragg SCI geometry.

Pulse	**Type**	**Pulse sequence**						
1	Index	1	2	$\cdots$	n_{target}			
	Pulse type	$\pi/2$	π	$\cdots$	π			
2	Index	1	$\cdots$	$n_{\text{target}}-1$	n_{target}	$n_{\text{target}}+1$	$\cdots$	$2n_{\text{target}}-1$
	Pulse type	π	$\cdots$	π	$\pi/2$	π	$\cdots$	π
3	Index	1	2	$\cdots$	n_{target}			
	Pulse type	$\pi/2$	π	$\cdots$	π			
4	Index	1	2	$\cdots$	n_{target}			
	Pulse type	π	π	$\cdots$	$\pi/2$			

the time-of-flight trace. This can make it difficult to extract the relative populations in each interferometer.

We perform a 2D scan of the intensity used for the $\pi/2$ pulses, $\Omega_{\text{eff},0}^{(\pi/2)}$, and the π pulses, $\Omega_{\text{eff},0}^{(\pi)}$. In this scan we set the value of $\Omega_{\text{eff},0}$ used for the multifrequency Bragg pulses to exactly twice what was used for the single frequency Bragg pulses. We use 1st-order Bragg pulses and set $n_{\text{target}} = 4$. The results of this scan are shown in Fig. 7.19, where we see that $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ form a 2D parabola in the contrast. We fit the resulting data to a 2D parabola and find its maximum in $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ to determine which pulse intensities to operate the interferometer at.

We perform several scans using the experiment apparatus to characterize the performance of this interferometer geometry. Bloch oscillations are not performed. In Fig. 7.20 we display the measured interferometer contrast as a function of n_{target}. We see a linear relationship between the interferometer contrast and the target order. Unsurprisingly higher values of n_{target} produce lower contrasts. For reference, a typical 4th-order SCRBI contrast without Bloch oscillations is 0.25, which is similar to what is observed in Fig. 7.20. Therefore the two geometries achieve similar contrasts for the same total accumulated phase.

We additionally perform a scan of T for different Bragg orders, n_{target}, and interferometry beam sizes. To change the interferometry beam size, we switched between the Schäfter Kirchhoff 60FC-T-4-M75-37 $f = 75\,\text{mm}$, 60FC-T-4-M100-02 $f = 100\,\text{mm}$, and 60FC-T-4-M125-54 $f = 125\,\text{mm}$ collimators. These collimators produce $1/e^2$ beam radii of 7.7 mm, 10.2 mm, and 12.8 mm, respectively. Recall from Section 4.1.2 that the 60FC-T-4-M75-37 $f = 75\,\text{mm}$ is the standard collimator used on the experiment. For each Bragg order and interferometry beam size we perform a scan equivalent to that performed in Fig. 7.19 in order to determine the pulse intensities to use for the $\pi/2$ and π pulses.

The recorded contrasts from this scan are shown in Fig. 7.21. Generally, it appears that larger beams produce higher contrasts, as expected. However, this is still counteracted by

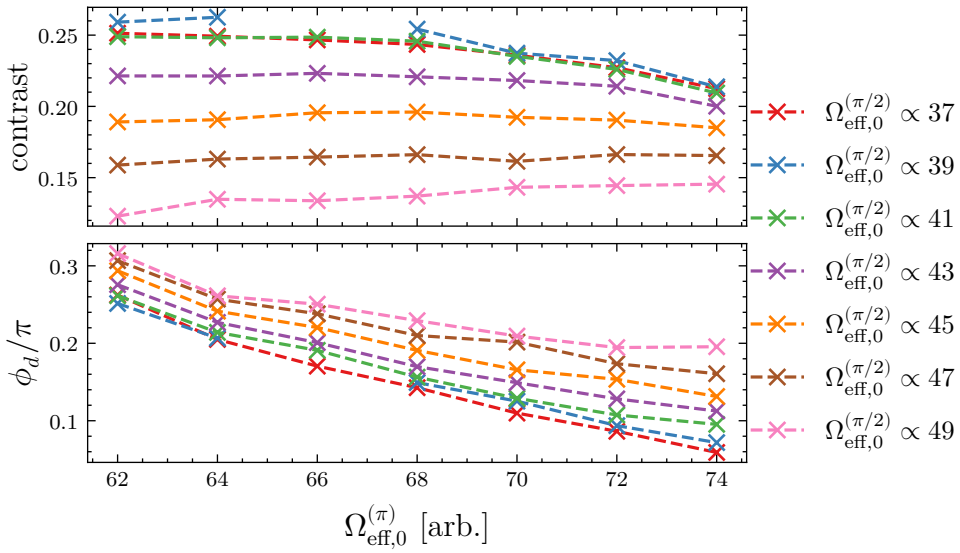


Figure 7.19: Experimentally measured contrast and differential phase ϕ_d for several $\pi/2$ and π pulse intensities used for the multiple 1st-order Bragg pulse SCI geometry for $n_{\mathrm{target}} = 4$. The interferometer uses $T = 10\,\mathrm{ms}$, $T' = 2\,\mathrm{ms}$, and $T_s = 250\,\mu\mathrm{s}$. All pulses are described by Eq. (2.20) with $\sigma = 30\,\mu\mathrm{s}$.

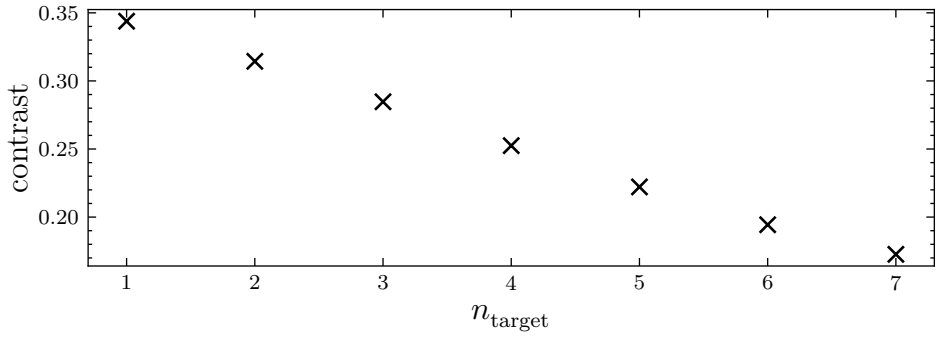


Figure 7.20: Experimentally measured ellipse contrast as a function of n_{target}. Here $\sigma = 30\,\mu\mathrm{s}$, $T = 10\,\mathrm{ms}$, $T' = 2\,\mathrm{ms}$, $T_s = 250\,\mu\mathrm{s}$, $\Omega^{(\pi/2)}_{\mathrm{eff},0}$ uses a 39 mV AWG peak pulse amplitude and $\Omega^{(\pi)}_{\mathrm{eff},0}$ uses a 66 mV AWG peak pulse amplitude.

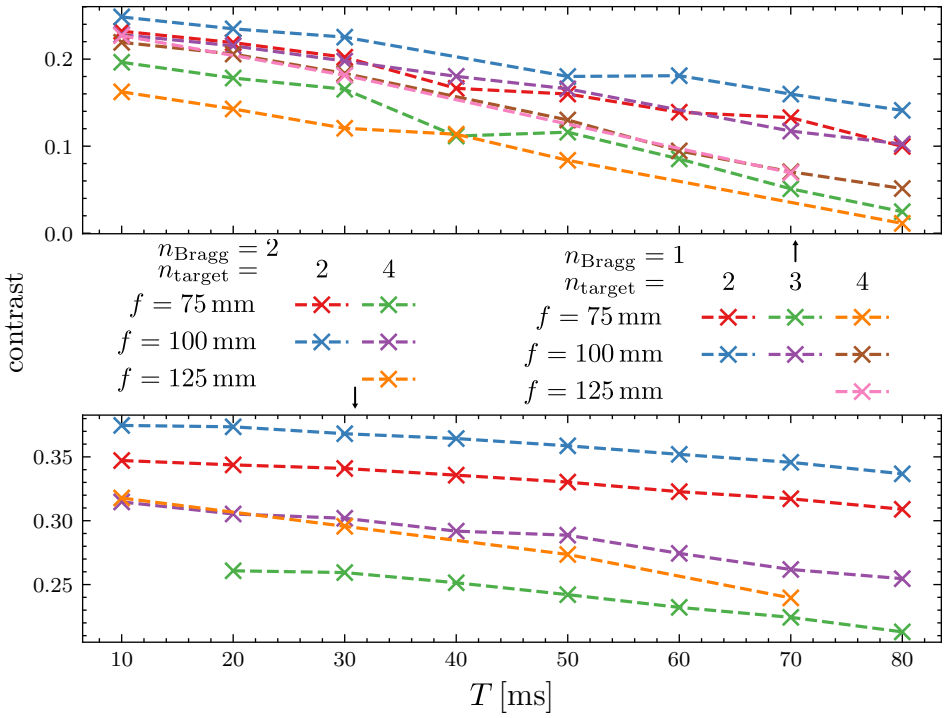


Figure 7.21: Experimentally measured ellipse contrast as a function of T for the SCI composed of multiple 1st-order Bragg pulses for several different n_{target} and interferometry beam fiber outcoupler focal lengths. Here $\sigma = 30\,\mu\text{s}$, $T' = 2\,\text{ms}$, and $T_s = 250\,\mu\text{s}$. $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ are set by performing the scan in Fig. 7.19 and choosing the values that maximize the contrast.

the choice of n_{target}, as should be expected from Fig. 7.20. In Fig. 7.22 we show the same scan, but plotted as a function of the per-shot sensitivity to the interferometer phase (note that the total interferometer phase depends on both the choice of Bragg order and n_{target}). In this picture we see that the best choice of beam size is consistently the largest one (i.e. even if the contrast is lower for some of these conditions, the measurement sensitivity is higher). A similar geometry was investigated in [104–106], which characterized systematic effects stemming from a stacked Bragg pulse geometry.

Performance comparison

Using the tree-based simulation method developed in Section 6.2.1 we simulate the performance of an SCRBI, SCMZI, and an SCRBI composed of multiple 1st-order Bragg pulses. The SCRBI and SCMZI use 4th-order Bragg pulses, and 10th-order Bloch oscillations. The SCRBI composed of multiple 1st-order Bragg pulses applies four pulses in succession to pro-

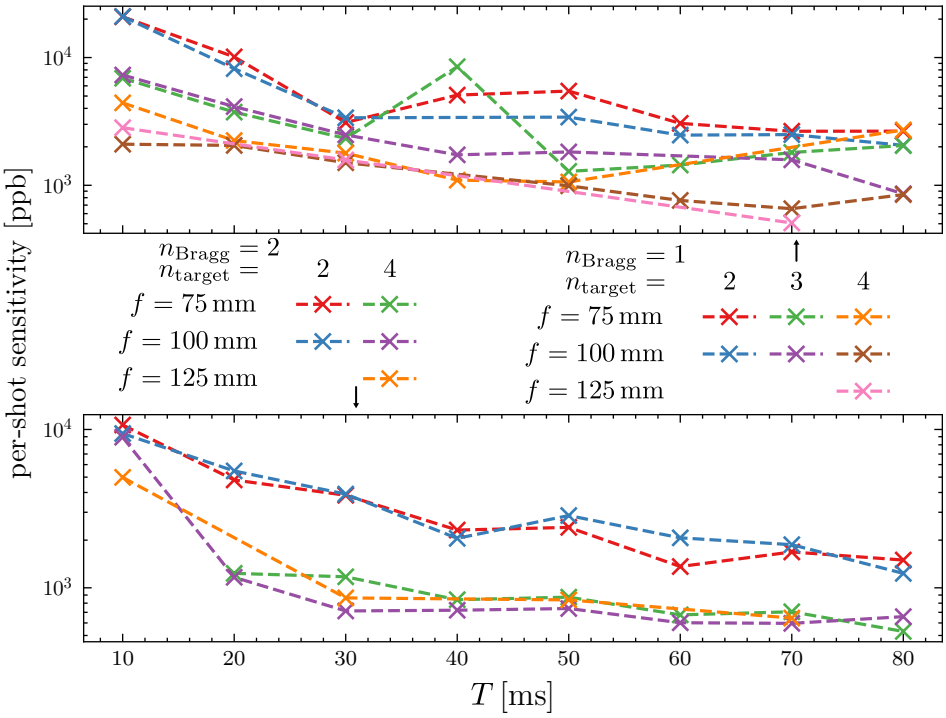


Figure 7.22: Experimentally measured interferometer per-shot sensitivity to the recoil frequency as a function of T for the SCI composed of multiple 1st-order Bragg pulses for several different n_{target} and interferometry beam fiber outcoupler focal lengths. Here $\sigma = 30\,\mu\text{s}$, $T' = 2\,\text{ms}$, and $T_s = 250\,\mu\text{s}$. $\Omega_{\text{eff},0}^{(\pi/2)}$ and $\Omega_{\text{eff},0}^{(\pi)}$ are set by performing the scan in Fig. 7.19 and choosing the values that maximize the contrast. Note that no Bloch oscillations are used here.

duce the equivalent momentum splitting of a 4th-order Bragg pulse. No Bloch oscillations are used for this interferometer.

Unless otherwise stated, all interferometer simulations use $T = 50\,\text{ms}$, $T' = 10\,\text{ms}$, $t_0 = 225\,\text{ms}$ (the time from the cloud initialization to the first Bragg pulse), a peak effective Rabi frequency of $\Omega_{\text{eff},0} = 19.39\omega_{\text{r}}$, $\sigma = 20\,\mu\text{s}$, a Gaussian pulse truncation of 5σ, and a target phase of $\phi_d = \pi/4$. A cloud of 1000 atoms is simulated, with positions and velocities drawn from normal distributions. Cloud positions are drawn from a three dimensional normal distribution with standard deviation 1 mm. Cloud velocities are drawn from a normal distribution with vertical velocity standard deviation $0.1v_{\text{r}}$, and transverse velocity standard deviation $2v_{\text{r}}$. The laser is modeled as an ideal Gaussian beam with $w_0 = 7.7\,\text{mm}$ centered on the atom cloud. The simulation integrates the Bragg equations of motion to simulate beamsplitters and accounts for Gouy phase shifts produced by k-vector shifts away from the waist of the Gaussian beam.

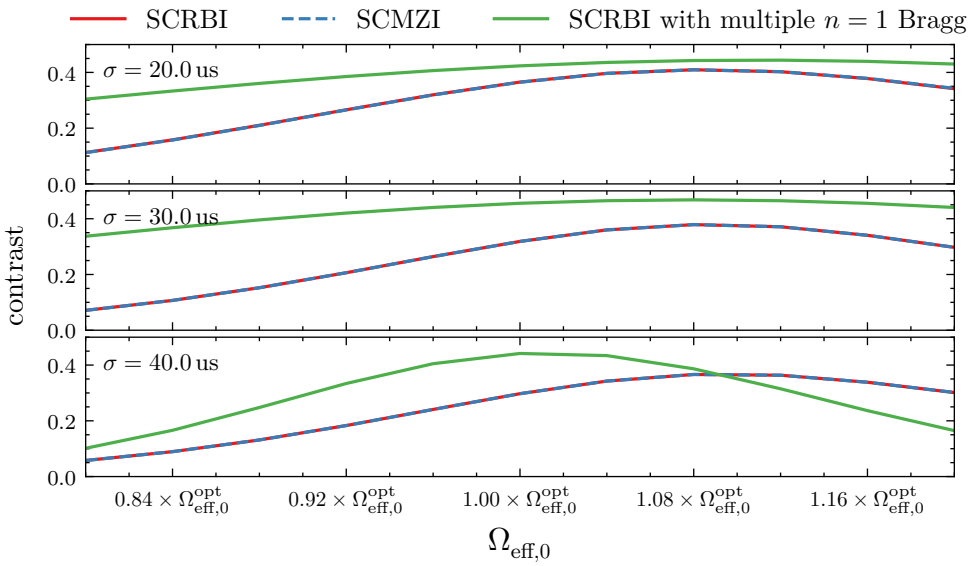


Figure 7.23: Simulated interferometer contrast as a function of the peak effective Rabi frequency $\Omega_{\text{eff},0}$ for several different Gaussian pulse widths σ. As σ is varied, different values of $\Omega_{\text{eff},0}$ are needed to produce a $\pi/2$ pulse. The x-axis is shown relative to the optimal $\Omega_{\text{eff},0}^{\text{opt}}$. Several types of interferometers are shown: SCRBI, SCMZI, and an SCRBI composed of multiple $n = 1$ Bragg pulses.

We begin by displaying the contrast of each interferometer as a function of $\Omega_{\text{eff},0}$ for several values of σ in Fig. 7.23. We see that the SCRBI and SCMZI have identical contrasts, while the SCRBI composed of multiple 1st-order Bragg pulses has a somewhat higher maximum contrast. This difference can be significant or insignificant, depending on the Bragg pulse width σ. Fig. 7.24 shows the simulated value of ϕ_d from the same set of simulations. We see that the SCRBI and SCMZI are relatively close to the expected interferometer phase $\phi_d = \pi/4$, while the SCRBI with multiple 1st-order Bragg pulses can have a large systematic shift. This is introduced by the mirror pulses. We ignore the SCRBI with multiple 1st-order Bragg pulses in the following simulations, as the systematic shift it introduces makes it ill-suited for measuring the recoil to high precision.

We show the simulated phase value of ϕ_d for the SCRBI and SCMZI as a function of the detuning from the Bragg resonance in Fig. 7.25, the delay from launch to the first Bragg beamsplitter in Fig. 7.26, and the interferometer duration T in Fig. 7.27. In all of these, as well as in Fig. 7.24, we see that the SCMZI phase is less sensitive to the changing parameter than the SCRBI phase is for $N \gtrsim 20$. Furthermore, we see that the absolute phase shift from the expected phase of $\phi_d = \pi/4$ is smaller for the SCMZI than the SCRBI. Clearly, as long as we are operating an interferometer with sufficient N, it is advantageous to use the SCMZI geometry instead of the SCRBI geometry.

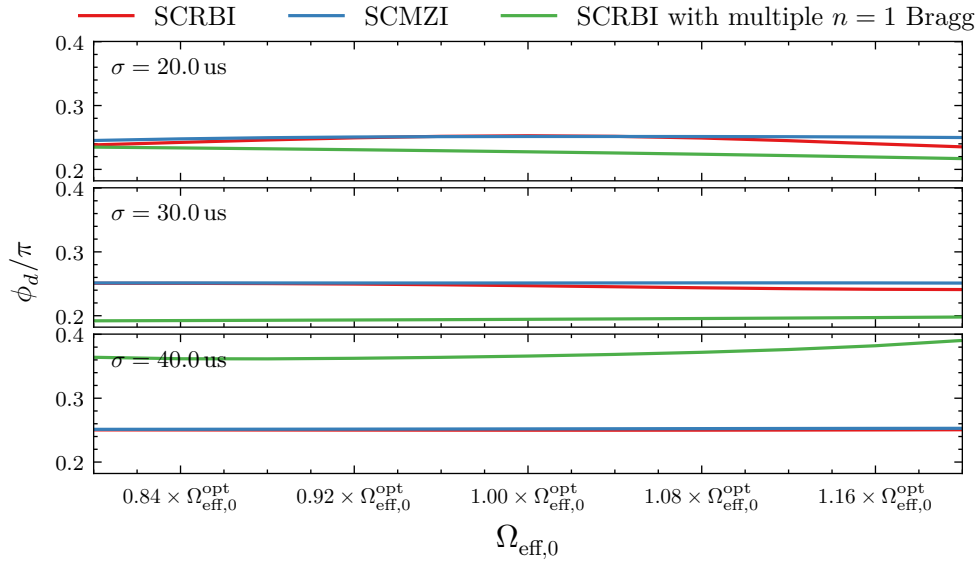


Figure 7.24: Simulated interferometer phase as a function of the peak effective Rabi frequency $\Omega_{\text{eff},0}$ for several different Gaussian pulse widths σ. As σ is varied, different values of $\Omega_{\text{eff},0}$ are needed to produce a $\pi/2$ pulse. The x-axis is shown relative to the optimal $\Omega_{\text{eff},0}^{\text{opt}}$. Several types of interferometers are shown: SCRBI, SCMZI, and an SCRBI composed of multiple $n = 1$ Bragg pulses.

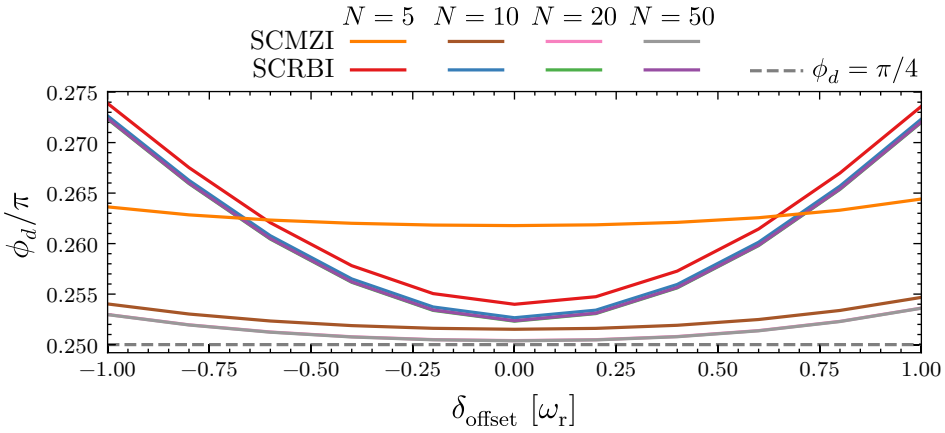


Figure 7.25: Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of δ_{offset}, which is the offset of the Bragg two-photon detuning, δ, from resonance. Several Bloch orders, N, are shown. Note that the $N = 20$ and $N = 50$ cases overlap for both the SCMZI and SCRBI.

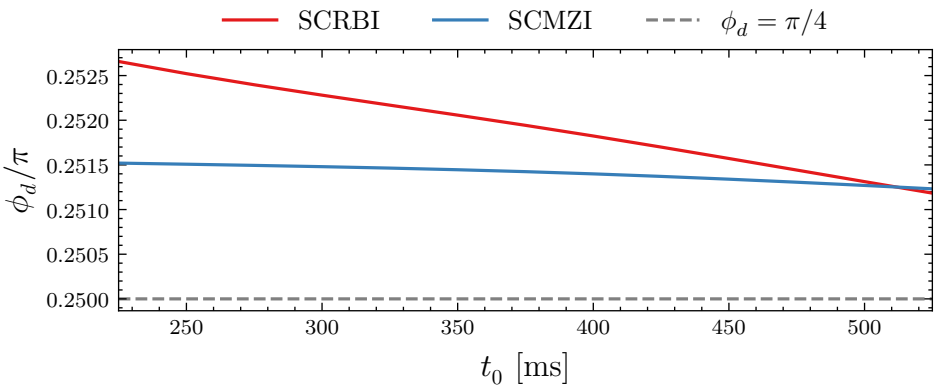


Figure 7.26: Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of t_0, the time of the first Bragg pulse, for $N = 10$.

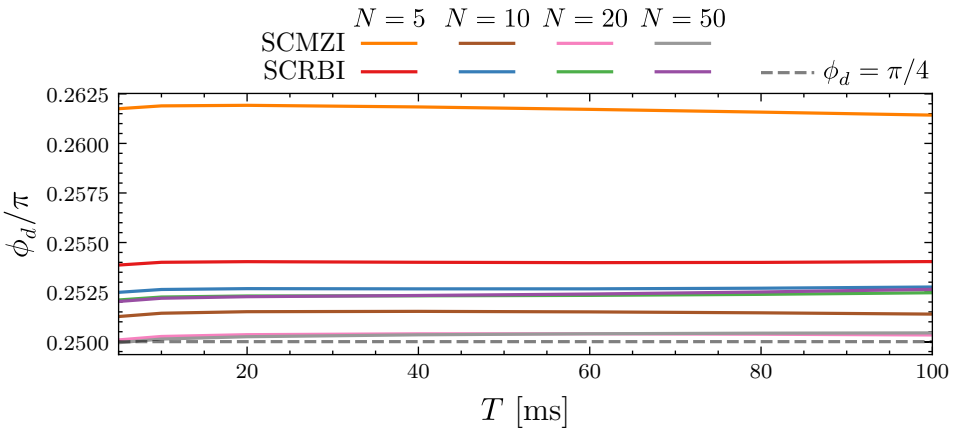


Figure 7.27: Simulated interferometer phase for SCRBI and SCMZI interferometers as a function of T. Several Bloch orders, N, are shown.

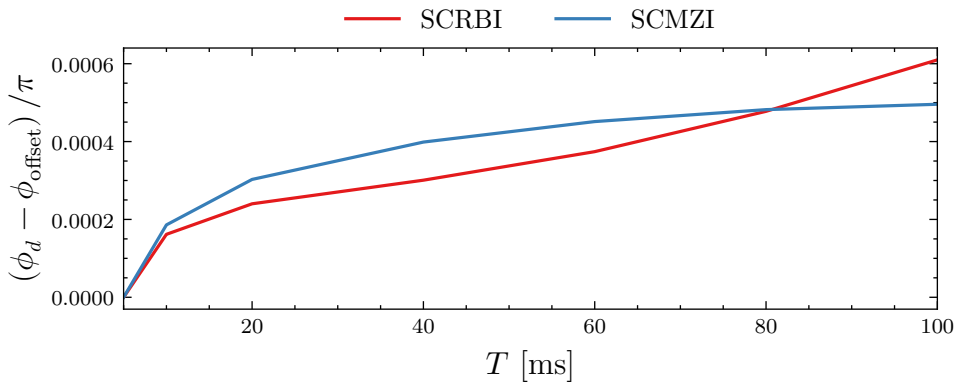


Figure 7.28: Simulated dependence of the interferometer phase on T for SCRBI and SCMZI interferometers with $N = 50$ Bloch oscillations. The phase is the same as that reported in Fig. 7.27, but ϕ_{offset} has been subtracted, where ϕ_{offset} is the interferometer phase for either the SCRBI or SCMZI at $T = 5\,\text{ms}$.

Note that the 2018 Berkeley measurement recorded SCRBI ellipses at multiple T values to measure the T-independent part of the diffraction phase systematic, leaving the T-dependent part to be characterized by simulation [21]. In Fig. 7.28 we show that the SCMZI phase has a smaller dependence on T than the SCRBI, reducing the contribution of the T-dependent part of the systematic. This will make it possible to characterize the systematic shift with lower uncertainty.

We can understand the lower phase dependence by performing the same calculation as done in Section 2.5.2 of Brian Estey's thesis [21], but for an SCMZI instead of an SCRBI. Following along with this calculation, we see that the SCMZI has more substantial diffraction phase cancellations than the SCRBI does for large N. In the limit where $N \to \infty$ and where atoms experience identical, spatially uniform $\Omega_{\text{eff},0}$ for the second and third Bragg pulses, the cancellation is perfect and no absolute phase shift is produced on the interferometer phase by the diffraction phase effect. This is confirmed by simulation.

7.2.2 Serrodyne modulated matterwave optics

In this section we discuss the concept of serrodyne modulated matterwave optics [27, 107–109]. Typically the frequency difference δ between the counterpropagating laser fields used for Bragg diffraction is generated using an acousto-optic modulator (AOM). However, we can consider a retroreflection scheme where it is produced by modulating the retro mirror position in time to imprint a time-dependent phase of δt on the retroreflected beam. The mirror moves at a constant velocity to produce a phase $\theta(t)$ on the retroreflected beam. Once a phase of 2π has been reached, the mirror can be quickly moved back to its initial position. This cycle is repeated indefinitely to produce a frequency difference of δ between the two

beams. The mirror must be reset to its initial position fast enough to avoid disturbing the Bragg diffraction dynamics. In the ideal case, the mirror position is described by

$$z(t) = (n\omega_{\mathrm{r}}\lambda t/\pi) \bmod (\lambda/2) \tag{7.17}$$

which produces phase $\theta(t) = 2kz(t)$ on the retroreflected beam. To see that this results in the usual Bragg diffraction equations of motion, we take $\delta t \to \theta(t)$ in Eq. (2.14) to find

$$i\hbar\dot{c}_j(t) = 4\hbar j^2\omega_{\mathrm{r}}c_j(t) - \hbar\frac{\Omega_{\mathrm{eff}}(t)}{2}\left[\exp(i\theta(t))c_{j+1}(t) + \exp(-i\theta(t))c_{j-1}(t)\right] \tag{7.18}$$

where we dropped the AC Stark shift terms. Plugging in our expressions for phase and position, and dropping the modulus as the exponential terms are periodic in 2π, we recover the usual equations of motion for Bragg diffraction. In practice the mirror moves with some non-zero fall duration from $z = \lambda/2$ to $z = 0$. During this non-zero fall the Bragg drive effectively pauses as the lattice is no longer resonant with the Bragg process. Note that the fall duration must be short relative to the timescale of the effective Rabi frequency. The drive resumes coherently when the ramp rises again. The mirror position $z(t)$ produces phase $\theta(t) = 2kz(t)$ on the retroreflected beam.

To investigate serrodyne modulated Bragg beamsplitters using the experimental apparatus described here, we phase modulate the RF drive of AOM32 with a sawtooth waveform. We show that this is equivalent to modulating the retro mirror position in Appendix A.7. We first investigate the quality of single Bragg serrodyne pulses, and then an MZI.

Single Bragg beamsplitter

Serrodyne-modulated single Bragg pulses are implemented in the experiment by driving $\theta(t)$ with a sawtooth waveform with a constant rising ramp rate of $4n\omega_{\mathrm{r}}$ and a variable duration falling ramp. We investigate the robustness of Bragg diffraction to variations in the falling ramp duration and peak amplitude. We also vary the ramp amplitude, but do not show these experimental results as the ramp amplitude is relatively straightforward to tune experimentally, while the falling ramp duration is limited by experiment hardware.

A simulation of a single serrodyne pulse is shown in Fig. 7.29. The simulation shows a 4th-order Bragg pulse with $\sigma = 0.259\omega_{\mathrm{r}}^{-1}$ and $\Omega_{\mathrm{eff},0} = 20.2\omega_{\mathrm{r}}$. The phase drive $\theta(t)$ in this particular simulation is given by the Fourier series expansion of a sawtooth wave. Here the first 30 terms in the series are included. We can view the population in $0\hbar k$ and $8\hbar k$ as a function of the number of terms in the series to determine the required bandwidth of our phase modulator. Note that the mth term in the series has angular frequency $2\pi m/T$, where $T = 2\pi/4n\omega_{\mathrm{r}}$. We show the result of this simulation in Fig. 7.30, seeing that a $\pi/2$ pulse is recovered at around $m = 15$. This gives us an estimate of $\sim 60n\omega_{\mathrm{r}}$ for the required bandwidth of our phase modulator. We will see that lower bandwidths can be used by controlling the phase of the sawtooth waveform.

We also simulate the population transfer as a function of the ramp amplitude, A_{ramp}, in Fig. 7.31. The results of this plot are not so surprising: when the ramp amplitude is an

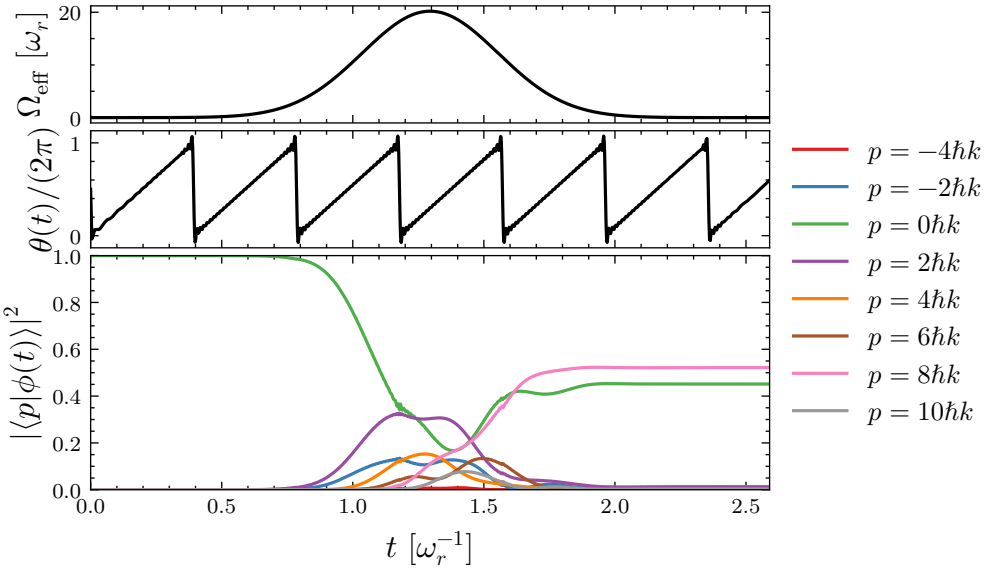


Figure 7.29: Simulation of a 4th-order Bragg diffraction process where the two-photon detuning is generated through serrodyne modulation. The simulation is made by integrating Eq. (A.51). *Top:* The effective Rabi frequency of the Bragg pulse as a function of time. *Middle:* The phase modulation applied to the Bragg process as a function of time. *Bottom:* The population in momentum space.

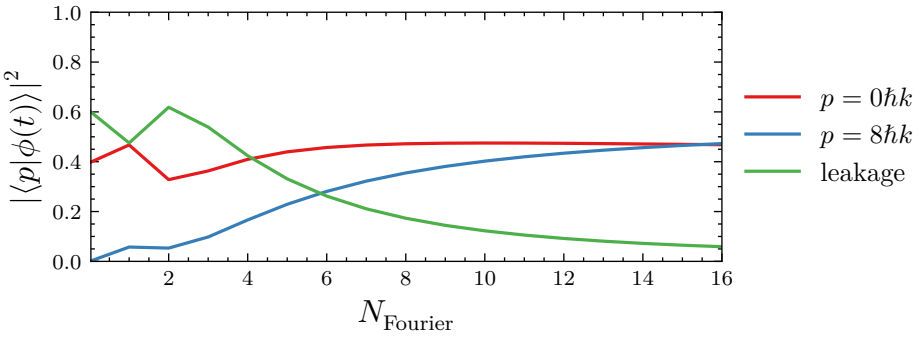


Figure 7.30: Simulated population in the initial $0\hbar k$ state and final $8\hbar k$ state for a 4th-order Bragg pulse as a function of the number of Fourier modes included in the phase modulation waveform. Leakage to unwanted states is also plotted.

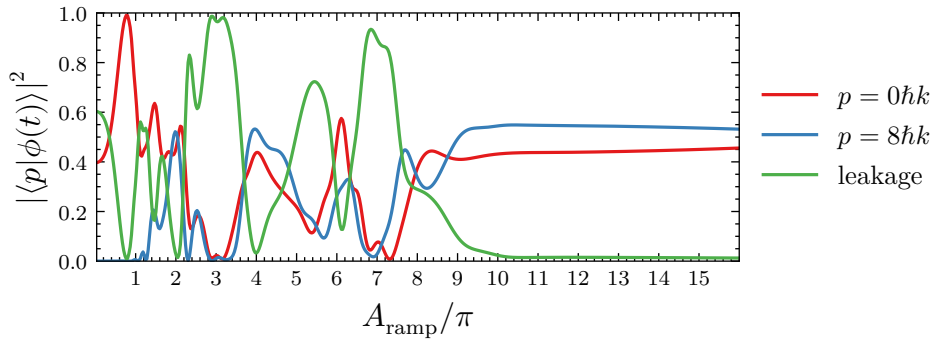


Figure 7.31: Simulated population in the initial $0\hbar k$ state and final $8\hbar k$ state for a 4th-order Bragg pulse as a function of the amplitude of the phase modulation. Leakage to unwanted states is also plotted.

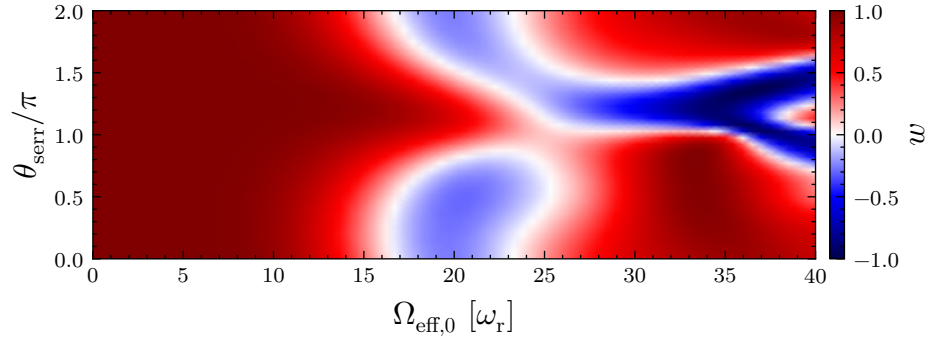


Figure 7.32: Simulated population inversion between the $0\hbar k$ and $8\hbar k$ momentum states as a function of $\Omega_{\text{eff},0}$ and the phase of the sawtooth waveform θ_{serr}.

exact multiple of 2π we see the leakage goes to zero and the $0\hbar k$ and $8\hbar k$ states have roughly equal population. Once the ramp amplitude is many times 2π, the sawtooth period becomes long enough that the entire beamsplitter occurs before the falling edge.

Finally, we can simulate the population inversion between $0\hbar k$ and $8\hbar k$ as a function of $\Omega_{\text{eff},0}$ and θ_{serr}, which is the phase of the sawtooth waveform at the start of the serrodyne pulse. The resulting inversion is shown in Fig. 7.32. We see that this phase can have a significant effect on the inversion, moving the pulse from an effective $\pi/2$ pulse at $\Omega_{\text{eff},0} = 16.5\,\text{r}$ and $\theta_{\text{serr}} = 0$ to a completely ineffective pulse at the same $\Omega_{\text{eff},0}$ and $\theta_{\text{serr}} = \pi$. In our experiments we scan θ_{serr} and select the one that produces the most negative inversion.

In Fig. 7.33 we perform experiments for $n = 2, 3, 4$ with a ramp amplitude of 2π. In this

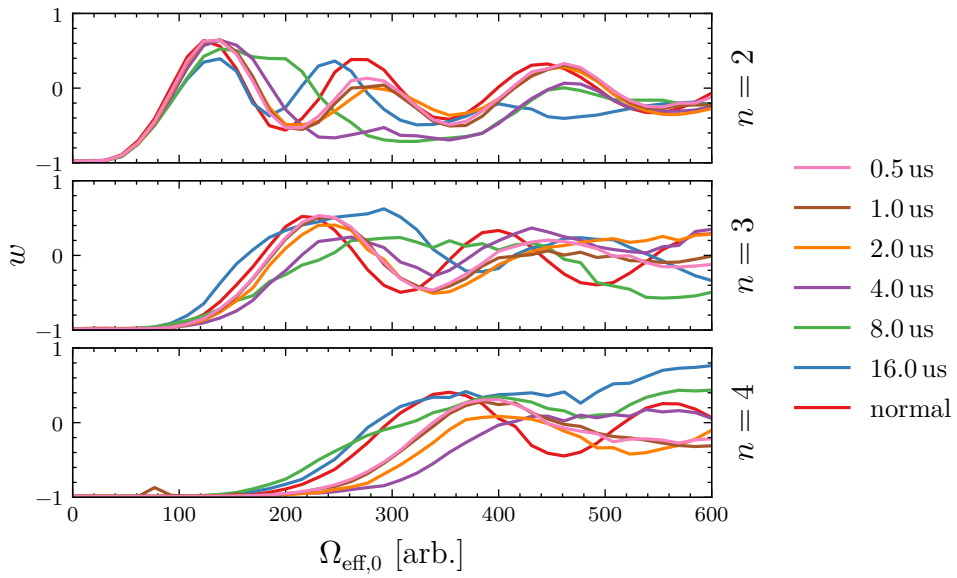


Figure 7.33: Experimentally measured Bragg inversion as a function of $\Omega_{\text{eff},0}$ for several falling ramp durations and Bragg orders. The normal line denotes a Bragg diffraction pulse performed without serrodyne-modulation.

condition, the rising ramp duration of the sawtooth waveform is 60.5 us, 40.3 us, and 30.3 us for $n = 2$, 3, and 4, respectively. The Bragg pulse intensity is given by a Gaussian profile in time with $\sigma = 40$ us. Each Bragg pulse is therefore subject to multiple falling ramps of the sawtooth waveform. For $n = 2$ it seems all fall durations are approximately equivalent between 0 and the π pulse condition. For $n = 3$ we start to see substantial deviation from the normal π pulse condition when the fall duration is greater than 2us, although all fall durations still reach a similar and usable maximum inversion. Similar conclusions are drawn for the $n = 4$ case.

The population leakage to states other than $0\hbar k$ and $2n\hbar k$ is given in Fig. 7.34. We see that there is minimal leakage for the case of the normal Bragg pulse across all Bragg orders. Serrodyne-modulated pulses with fall durations greater than 2 us start to deviate from the normal Bragg pulse leakage before the π pulse condition. If a longer fall duration than 2 us is required, it will come at the cost of producing additional, and likely unwanted, momentum states at the interferometer output. As we will show, however, these fall durations can still produce coherent interferometers.

Note a special case of serrodyne modulated Bragg diffraction occurs when the phase is modulated according to a symmetric triangle wave and the wave peak or trough occurs simultaneously with the peak of $\Omega_{\text{eff}}(t)$. In this condition it is possible to create double-Bragg diffraction [110–113], where an atom initially in state $|0\hbar k\rangle$ is diffracted into $|2n\hbar k\rangle$ and $|-2n\hbar k\rangle$ in equal populations. This beamsplitter configuration is of interest for other

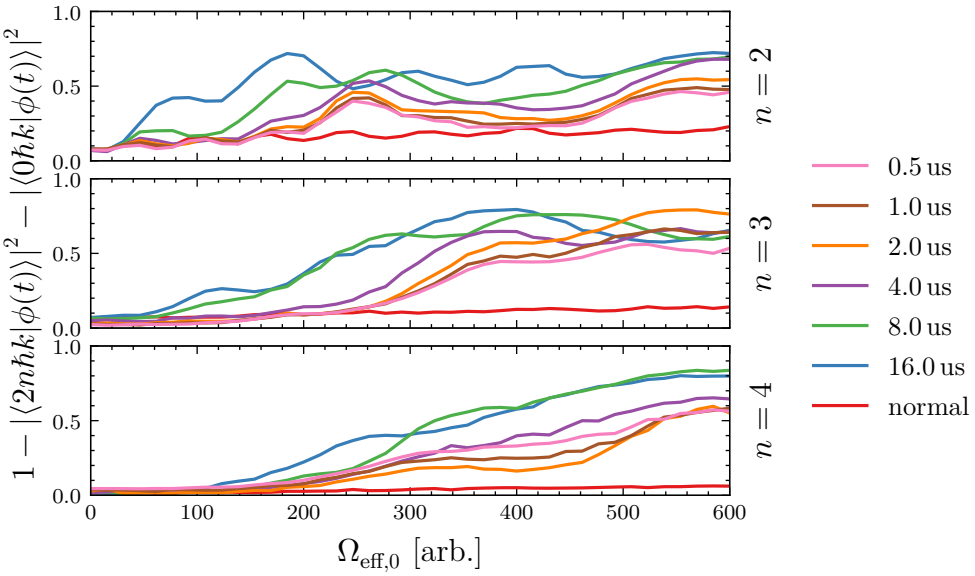


Figure 7.34: Experimentally measured population leakage to unwanted momentum states as a function of $\Omega_{\text{eff},0}$ for several falling ramp durations and Bragg orders. The normal line denotes a Bragg diffraction pulse performed without serrodyne-modulation.

interferometer geometries. This process is simulated and shown in Fig. 7.35.

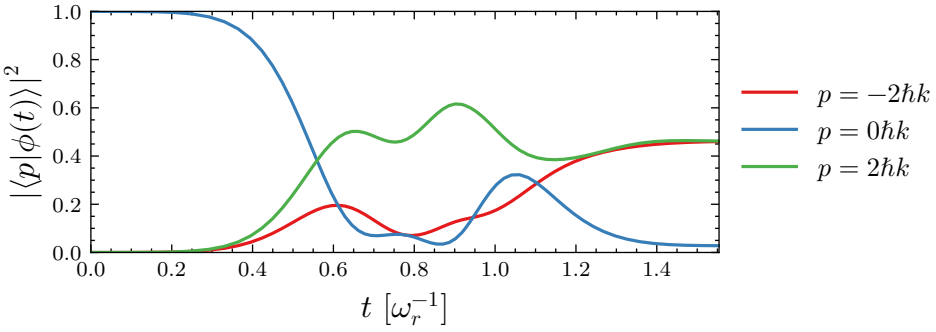


Figure 7.35: Simulated population in the initial $0\hbar k$ state and final $\pm 8\hbar k$ states for 4th-order double Bragg diffraction.

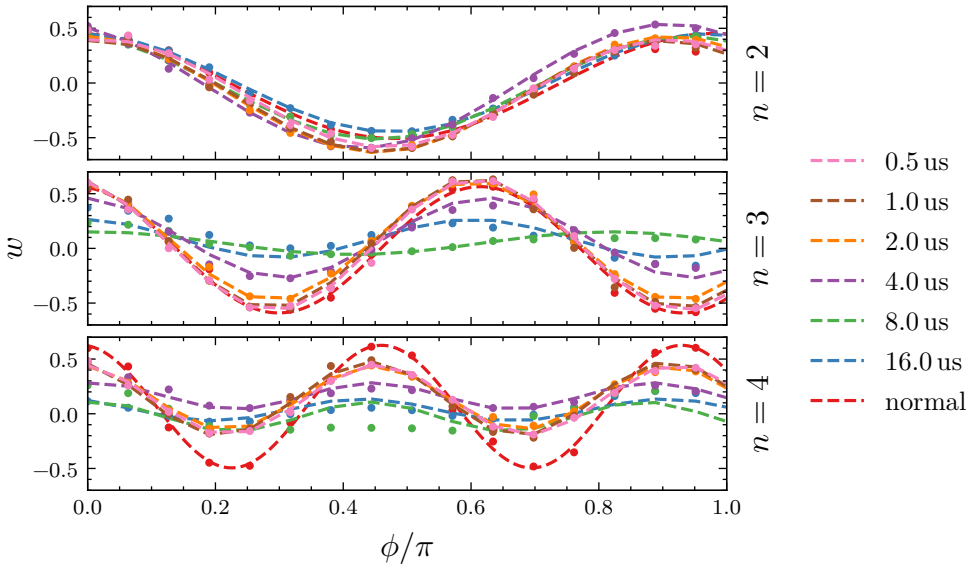


Figure 7.36: Experimentally measured MZI fringes composed of normal and serrodyne-modulated Bragg pulses for different Bragg orders. Dots are experimentally recorded data points. Dashed lines are fits to those points.

Serrodyne Mach-Zehnder interferometer

We form an MZI using a serrodyne-modulated beamsplitter, mirror, and then beamsplitter pulse. The pulse spacing T is set to an exact integer multiple of the sawtooth period to ensure each pulse experiences the same serrodyne waveform phase θ_{serr}. The integer multiple is selected for each n to make T as close to 300us as possible. The interferometer phase is scanned by introducing a phase shift ϕ to the final serrodyne-modulated Bragg pulse (this phase shift is on the RF drive, not the serrodyne waveform). Vibrations on the retro mirror cause the interferometer phase to vary shot-to-shot, which we average away by running over multiple shots for each ϕ value.

The MZI has outputs in states $0\hbar k$ and $2n\hbar k$, with populations given by

$$p_0 = \cos^2((nT^2(2kg + \alpha) - n\phi)/2) \tag{7.19}$$

$$p_n = \sin^2((nT^2(2kg + \alpha) - n\phi)/2) \tag{7.20}$$

As discussed in Section 3.2.4, the frequency ramp rate α is set to zero the $2kg$ term, and so our measured populations are $\cos^2(n\phi/2)$ and $\sin^2(n\phi/2)$. Therefore, the fringe frequency will vary as a function of n as ϕ is scanned. The experimentally observed fringes are shown in Fig. 7.36. These fringes reveal that non-zero fall durations can lead to contrast reductions and interferometer phase shifts.

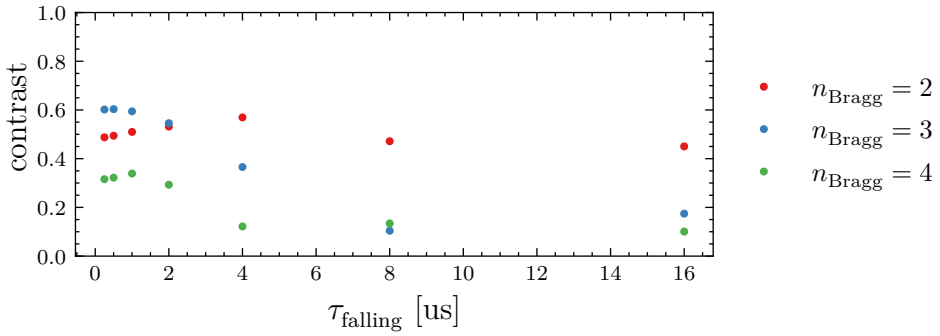


Figure 7.37: Experimentally measured MZI fringe contrast as a function of the falling ramp duration for different Bragg orders. Contrasts are measured by fitting a sinusoid to MZI fringes as in Fig. 7.36 and extracting the amplitude. The theoretical maximum contrast of the interferometer is 1, and in practice is limited to a maximum of approximately 0.6 due to the variation of Ω_{eff} across the atom cloud.

In Fig. 7.37 we show the dependence of the MZI fringe contrast on the falling ramp duration and ramp amplitude. We observe that lower Bragg orders generally achieve a higher contrast at long fall durations, which can be partially attributed to the smaller fraction of the sawtooth period that these fall durations account for.

Significance

From Fig. 7.37 we clearly see that usable interferometers can be formed with significant falling durations. This opens the door to creating Bragg interferometers using single frequency lasers and a position modulated retro mirror. These designs also break the degeneracy produced by directional symmetry in retroreflected atom interferometers [114–118].

The observation of viable single Bragg pulses and full interferometers for poor-quality serrodyne modulation implies that Bragg diffraction does not require a well-behaved, single-frequency phase drive to operate successfully.

Note that this investigation was not made in the pursuit of a fine-structure constant measurement.

7.3 Measuring wavefunction spatial extent

As mentioned in Section 7.1.3, the spatial extent of the wavefunction determines the scale of certain systematic effects. In Section 3.4 we discussed how we must design interferometer geometries so that the wavepackets propagating along different arms of the interferometer spatially overlap at the time of the final beamsplitter pulse. We mentioned that this overlap

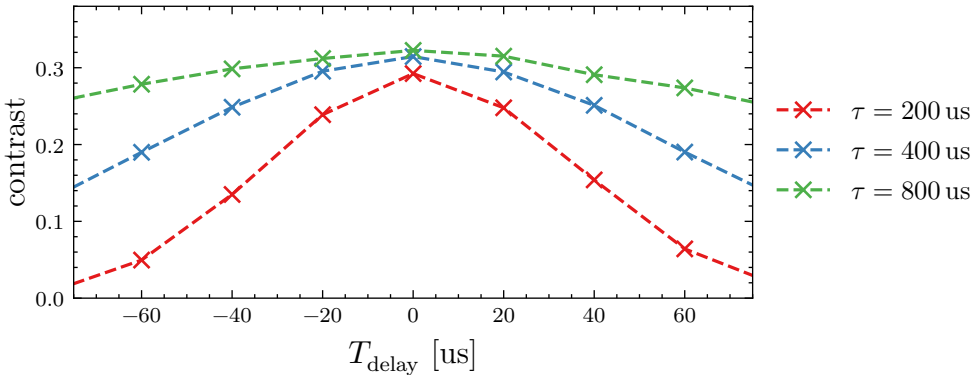


Figure 7.38: Experimentally measured ellipse contrast as a function of the delay, T_{delay}, of the fourth beamsplitter pulse. As the Raman pulse duration, τ, increases and the double VS pulse process becomes more velocity selective, the position space wavefunction becomes wider and the contrast decays more slowly as a function of T_{delay}.

needed to be smaller than the wavepacket spatial extent, and drew an analogy to the idea of coherence length in an optical interferometer. By delaying the time of the final Bragg pulse, we can reduce the overlap between the interfering wavepackets to extract information about their spatial extent.

In Fig. 7.38 we measure the interferometer contrast as a function of the delay, T_{delay}, of the fourth Bragg pulse in an SCI from its nominal timing. Delaying this pulse reduces wavefunction overlap at the end of the interferometer, resulting in a lower contrast. We do this for four VS pulse durations, τ, and see that the shorter-duration pulses (which select a broader momentum width and a narrower spatial width) show a faster contrast decay with T_{delay}, while the longer-duration pulses (which select a narrower momentum width and broader spatial width) show a slower contrast decay with T_{delay}. Note that the effective Rabi frequency used for each VS pulse is tuned to 80% of the value that would produce a π-pulse.

In Eq. (3.32), when computing the separation phase, we saw that two Gaussian wavepackets with wavefunction standard deviation σ_x have a spatial decay in population given by

$$\exp\left(-\frac{\Delta x^2}{4\sigma_x^2}\right) = \exp\left(-\left[\frac{n v_{\text{r}} T_{\text{delay}}}{\sigma_x}\right]^2\right) \tag{7.21}$$

where Δx is the separation between the two wavepackets and n is the Bragg order. Note that the wavefunction width σ_x is related to the population width σ_x' via $\sigma_x' = \sigma_x/\sqrt{2}$. Performing a fit to the data in Fig. 7.38 and noting that the experiment used 4th-order Bragg pulses, we find $\sigma_x' = 0.47\,\mu\text{m}$ for $\tau = 200\,\mu\text{s}$, $\sigma_x' = 0.84\,\mu\text{m}$ for $\tau = 400\,\mu\text{s}$, and $\sigma_x' = 1.68\,\mu\text{m}$ for $\tau = 800\,\mu\text{s}$. In Section 4.4.6 we determined that a 400 µs double VS pulse should select a population with momentum width $\sigma_p' = 0.073\hbar k$.

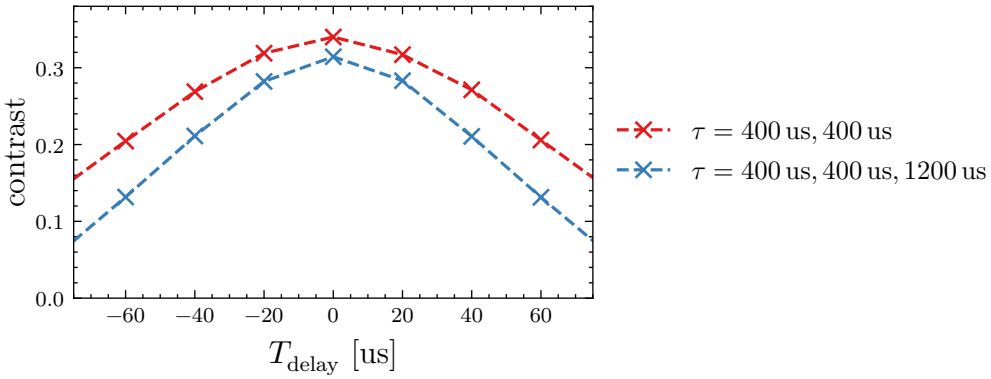


Figure 7.39: Experimentally measured ellipse contrast as a function of the delay, T_{delay}, of the fourth beamsplitter pulse. As the Raman pulse duration, τ, increases and the VS pulse becomes more velocity selective, the position space wavefunction becomes wider and the contrast decays more slowly as a function of T_{delay}. Two VS sequences are shown, one with the standard two $\tau = 400$µs pulses, and another with the standard two pulses followed by a single 1200µs square Raman pulse.

Note that these amount to two independent determinations of the position space wavefunction width and the momentum space wavefunction width. We find $\sigma_x'\sigma_p' = 0.45\hbar$, which slightly violates the uncertainty principle. This is likely due to experimental error. This result is consistent with an experiment performed in [119] by the Mueller group, in which a similar measurement was made. In that comparison σ_x' was found to be $\sigma_x' = 0.81\,\mu\text{m}$ and σ_p' was found to be $\sigma_p' = 0.17\hbar k$.

We can run a second similar experiment by performing a third VS pulse after the standard two VS pulses. The third pulse has a longer duration than the first two, and will remove the center of the wavepacket in momentum space. This will lower the overall contrast and increase the rate at which contrast decays with T_{delay}. We show the result of this experiment in Fig. 7.39, where in one sequence two 400µs square Raman pulses are performed and in the other sequence two 400µs square Raman pulses followed by a single 1200µs square Raman pulse are performed. The result is as expected. This is a basic attempt to engineer the interferometer wavepacket, although in this case it does not produce any improved interferometer properties.

Chapter 8

Summary and outlook

8.1 Summary

In this thesis we discussed the building blocks of atom interferometers in Chapter 2 and Chapter 3. We then reviewed the new experiment apparatus purpose built for a second generation measurement of the fine-structure constant at Berkeley in Chapter 4, discussing the hardware systems developed for the measurement and the alignment procedures used to tune the experiment. The vacuum system described in this chapter was previously described in [24]; however, other systems described in this chapter have not been described in literature or in a thesis. This was followed by a characterization of interferometer performance in Chapter 5, in which measurement techniques and analysis new to the second generation fine-structure constant measurement were applied.

Chapter 6 reviewed the simulation techniques available for studying systematics in an atom interferometer. Several new simulation techniques were developed, including a tree-based simulation capable of simulating different systematic effects in a unified simulation framework. A novel integration method was also introduced. In Section 7.1 this simulation was verified against experiment data, finding agreement. The simulation framework developed in this thesis improves on the simulation capabilities of the second generation fine-structure constant measurement over the first generation measurement. In Chapter 7 we reviewed the systematic effects present in the measurement, and discussed new interferometer geometries and beamsplitters.

The work performed in this thesis led to several publications. These described a novel phase locking technique [25], a novel method for integrating a full interferometer wavefunction [26], and a novel method for producing atom beamsplitters [27].

8.2 Outlook

In this thesis, we discussed the development and operation of an atom interferometer designed to precisely measure the cesium recoil frequency in service of determining the fine-structure

constant. There are several hardware changes that might improve the experiment, such as reducing the diameter of the detection and 3D MOT vacuum chambers, and using a separate laser for Raman velocity selection.

The vacuum chamber inner diameter was chosen to be 0.4 m in order to maximize the size of the interferometry beam. However, the bottom viewport limits the maximum size of the interferometry beam to a maximum diameter of 0.2 m. This means that the 3D MOT and detection chambers could have been made with a diameter of around 0.25 m without reducing the maximum interferometry beam size. This would have provided the following benefits:

- **Less cramped pit environment:** Currently two graduate students or postdocs must be in the pit simultaneously to align the 3D MOT, RSC, and detection beams, as a single person cannot reach all of these optics without moving to the other side of the pit (which takes around 10 seconds). These beams often need to be aligned in pairs, with one person adjusting the beam alignment on either side of the pit. With a smaller diameter chamber, a single person could have made both adjustments simultaneously. This would have led to much faster alignment procedures. Furthermore, in its current state, working in the pit often requires contorting into uncomfortable positions for long periods of time.

- **More space for the 2D MOT:** A smaller diameter 3D MOT and detection chamber would have made more space for the 2D MOT. This would enable the use of a single 2D MOT region, instead of three 2D MOT regions as discussed in Section 4.1.1, which likely would have resulted in a higher 2D MOT atom flux.

- **Smaller 3D MOT coils:** Without such a large chamber diameter, much smaller 3D MOT coils could have been used. This would have likely prevented the vibration problem that occurred with the original 3D MOT coils, and would have enabled higher field gradients to create a more compact 3D MOT. It is also possible that the large eddy currents produced by the current 3D MOT turn off impede the PGC and RSC state preparation stages.

- **Better detection efficiency:** Reducing the detection chamber diameter would increase the detection system solid angle, improving detection efficiency.

I do not see any disadvantages of having a smaller 3D MOT and detection chamber diameter. Future grad students might consider reducing the chamber diameter to the limit set by the bottom viewport, which I expect would lead to faster experiment tuning, more reliable operation, and possibly improved state preparation and atom number.

Additionally, it is worthwhile to consider the effect of moving the 3D MOT chamber above the detection chamber. This would result in a lower experiment shot duration, allowing for a higher averaging rate. The shorter free fall duration would also lead to fewer losses to the vacuum, raising the number of detected atoms. Putting the detection region below the 3D

MOT region would also provide more time for momentum states exiting the interferometer to spatially separate. This would allow for taller detection beams, increasing the time that atoms could be exposed to fluorescent light and increasing the number of photons detected per atom.

Interferometer performance is primarily limited by systematic uncertainties. Reduction of the statistical error will lead to faster averaging and characterization of systematic effects. Shrinking systematic errors will reduce the time needed to characterize the measurement at an interesting level of precision. Improving either or both of these is beneficial to the experiment.

Shrinking systematic errors further requires either a colder atomic cloud or a larger interferometry beam. All of the beam-related systematic effects discussed in Section 7.1 would be suppressed with a larger beam-size-to-atom-cloud ratio. Even if the diffraction phase can be suppressed through novel beamsplitter techniques or similar, the process of transferring momentum unavoidably results in Gouy shifts entering the interferometer phase. These shifts can only be suppressed by reducing the range of $\mathbf{k}$ vectors that atoms sample throughout the interferometer. One (possibly far-fetched) approach to dealing with these Gouy shifts could be to design Bragg beamsplitters that introduce a phase that cancels the Gouy shift phase. Perfect cancellation would result in a revival of interferometer contrast, which in theory could be experimentally detected. This phase cancellation might be accomplished by operating the velocity selection laser away from the magic wavelength.

Lower atom cloud temperatures might be obtained by reducing the 3D MOT vacuum chamber diameter, as this might lead to better aligned and tuned state preparation stages. To achieve an order of magnitude reduction in temperature, it would be necessary to apply more aggressive cooling techniques, such as those used to achieve Bose-Einstein condensation.

Achieving a larger interferometry beam size requires reducing the laser intensity required for Bragg beamsplitters and Bloch oscillations, or increasing the available laser power. We investigated the possibility of performing multiple $n = 1$ pulses in sequence as a way to reduce laser power requirements in Section 7.2.1. Even if the systematic effects introduced by this geometry were understood, the experiment would still be limited by the power required for Bloch oscillations. Therefore, the main path to increasing the interferometry beam size is to increase the available laser power.

Previously an attempt was made in the Mueller group to generate high-power laser pulses at 852 nm for millisecond pulse durations through optical parametric amplification of doubled Nd:YAG light (this used a similar system to the one outlined in [120]). However, while the Nd:YAG system produced several kilowatts of 532 nm light, it did not provide sufficient power to achieve efficient optical parametric amplification. One idea could be to use the same Nd:YAG system to pump a Ti:Sa crystal, similar to what was done in [121]. Instead of placing the Ti:Sa in a cavity, it could be used as a single-pass amplifier for frequency-stabilized seed light. However, this system would need to produce pulses on the order of 10 ms in length to be used for Bloch oscillations. Another option for producing high-power laser light is through coherent beam combining of multiple light sources [122–125].

The other route to higher laser power could come from targeting different transitions

in cesium where higher laser power is available, or by switching to another atomic species where higher laser power is available at some transition from the ground state or a metastable state. Note that while a cycling transition is necessary for the state preparation stages used in this experiment, Bragg diffraction *does not* require a cycling transition to work. All Bragg diffraction needs is some transition which can be used to create an off-resonant potential. Rubidium is perhaps the obvious choice, as doubled telecom lasers can be used [126]. Another consideration might be sodium, where the D line transitions at 589.16 nm and 589.76 nm could be addressed by a doubled Raman fiber laser (Toptica and IPG Photonics produce 1178 nm lasers that output tens of watts). The $2\,^3\mathrm{S}_1 \rightarrow 2\,^3\mathrm{P}_{1,2}$ transition in metastable helium could be addressed directly by a ytterbium fiber laser at 1083 nm. Of course, each of these species comes with additional challenges that we do not face with cesium.

No matter which approach is used, it will be necessary to accurately simulate the beam-related systematic effects in the experiment. For this reason, the comparison between simulated and actual interferometers performed here is necessary for performing an improved measurement of the cesium recoil frequency, and ultimately the fine-structure constant.

Appendix A

Reference math

This appendix contains several identities and derivations that support the derivations performed in the thesis. They are included here and not in the thesis for brevity. I doubt that any of these derivations are novel. This appendix also states some of the conventions used in the thesis. These conventions should be consistent with Robert Littlejohn's notes for Quantum Mechanics course at UC Berkeley [127].

A.1 Unitary transformations

Two representations of the same state $|\psi\rangle$ and $\left|\widetilde{\psi}\right\rangle$ are related through unitary transformation U

$$|\psi\rangle = U^\dagger \left|\widetilde{\psi}\right\rangle \tag{A.1}$$

Theorem A.1. *For a system with two representations of the same state $|\psi\rangle$ and $\left|\widetilde{\psi}\right\rangle$ related through unitary transformation U, the Hamiltonian in each representation is related via*

$$\widetilde{\mathcal{H}} = U\mathcal{H}U^\dagger - i\hbar U\frac{dU^\dagger}{dt} = U\mathcal{H}U^\dagger + i\hbar\frac{dU}{dt}U^\dagger \tag{A.2}$$

Proof. The time-dependent Schrödinger equation for both states is

$$i\hbar\frac{d|\psi\rangle}{dt} = \mathcal{H}|\psi\rangle \qquad\qquad i\hbar\frac{d\left|\widetilde{\psi}\right\rangle}{dt} = \widetilde{\mathcal{H}}\left|\widetilde{\psi}\right\rangle \tag{A.3}$$

The time derivative of $|\psi\rangle$ can be expanded in terms of $\left|\widetilde{\psi}\right\rangle$

$$i\hbar\frac{d|\psi\rangle}{dt} = i\hbar\frac{d\left(U^\dagger\left|\widetilde{\psi}\right\rangle\right)}{dt} = i\hbar\left(U^\dagger\frac{d\left|\widetilde{\psi}\right\rangle}{dt} + \frac{dU^\dagger}{dt}\left|\widetilde{\psi}\right\rangle\right) \tag{A.4}$$

So the Schrödinger equation for $|\psi\rangle$ becomes

$$i\hbar\left(U^\dagger\frac{d\left|\widetilde{\psi}\right\rangle}{dt}+\frac{dU^\dagger}{dt}\left|\widetilde{\psi}\right\rangle\right)=\mathcal{H}|\psi\rangle=\mathcal{H}U^\dagger\left|\widetilde{\psi}\right\rangle \tag{A.5}$$

Multiplying both sides by U

$$i\hbar\frac{d\left|\widetilde{\psi}\right\rangle}{dt}+i\hbar U\frac{dU^\dagger}{dt}\left|\widetilde{\psi}\right\rangle=U\mathcal{H}U^\dagger\left|\widetilde{\psi}\right\rangle\rightarrow i\hbar\frac{d\left|\widetilde{\psi}\right\rangle}{dt}=U\mathcal{H}U^\dagger\left|\widetilde{\psi}\right\rangle-i\hbar U\frac{dU^\dagger}{dt}\left|\widetilde{\psi}\right\rangle \tag{A.6}$$

$$=\left(U\mathcal{H}U^\dagger-i\hbar U\frac{dU^\dagger}{dt}\right)\left|\widetilde{\psi}\right\rangle \tag{A.7}$$

$$=\widetilde{\mathcal{H}}\left|\widetilde{\psi}\right\rangle \tag{A.8}$$

Recognizing this as the Schrödinger equation for $\left|\widetilde{\psi}\right\rangle$ we find

$$\widetilde{\mathcal{H}}=U\mathcal{H}U^\dagger-i\hbar U\frac{dU^\dagger}{dt}=U\mathcal{H}U^\dagger+i\hbar\frac{dU}{dt}U^\dagger \tag{A.9}$$

where we used

$$0=\frac{dI}{dt}=\frac{d\left(UU^\dagger\right)}{dt}=\frac{dU}{dt}U^\dagger+U\frac{dU^\dagger}{dt}\rightarrow U\frac{dU^\dagger}{dt}=-\frac{dU}{dt}U^\dagger \tag{A.10}$$

□

A.2 Time-evolution operator

The time-evolution operator is given by [128]

$$U(t,t_0)=\exp\left(-i\int_{t_0}^{t}\mathrm{d}t'H(t')/\hbar\right) \tag{A.11}$$

A.3 Atom-field interaction

The form of the atom-field interaction is given by [72]

$$H_{\text{atom-field interaction}}=-e\mathbf{r}\cdot\boldsymbol{\epsilon}E(\mathbf{r},t) \tag{A.12}$$

where $\boldsymbol{\epsilon}$ is the polarization vector. The matrix elements of $H_{\text{atom-field interaction}}$ are

$$\left[H_{\text{atom-field interaction}}\right]_{ij}=\langle i|-e\mathbf{r}\cdot\boldsymbol{\epsilon}E(\mathbf{r},t)|j\rangle=-\mu_{i,j}E(\mathbf{r},t) \tag{A.13}$$

where $\mu_{i,j} = e\langle i|\boldsymbol{\epsilon}\cdot\mathbf{r}|j\rangle$ is the transition dipole moment between states i and j. The transition dipole moment can be further reduced using angular momentum algebra to show that different light polarizations $\boldsymbol{\epsilon}$ lead to different transition strengths (see [72, Section 4.5], [50, Section 7.3.7]). We define the Rabi frequency between two levels $\Omega_{i,j}$ as

$$\Omega_{i,j} = -\frac{E_0}{\hbar}\mu_{i,j} \tag{A.14}$$

where E_0 is the field amplitude. If multiple fields are present there is a separate Rabi frequency for each.

A.4 Position and momentum space identities

Theorem A.2. *For some state* $|\psi\rangle$

$$i\hbar\langle p|\frac{d}{dt}|\psi\rangle = i\hbar\dot{\phi}(p) \tag{A.15}$$

$$\langle p|\exp(\pm i2k\hat{x})|\psi\rangle = \phi(p\mp 2\hbar k) \tag{A.16}$$

Proof. We take the following relations as given

$$\langle x|\psi\rangle = \psi(x) \qquad \langle p|\psi\rangle = \phi(p) \qquad \langle x|p\rangle = \exp(ipx/\hbar)/\sqrt{2\pi\hbar}$$
$$I = \int_{-\infty}^{\infty}|x'\rangle\langle x'|\,\mathrm{d}x' = \int_{-\infty}^{\infty}|p'\rangle\langle p'|\,\mathrm{d}p'$$

where $|p\rangle$ denotes a plane wave with momentum p. Starting with the time-derivative term

$$i\hbar\langle p|\frac{d}{dt}|\psi\rangle = i\hbar\langle p|\frac{d}{dt}I|\psi\rangle \tag{A.17}$$

$$= i\hbar\frac{d}{dt}\int_{-\infty}^{\infty}\langle p|p'\rangle\langle p'|\psi\rangle\,\mathrm{d}p' = i\hbar\dot{\phi}(p) \tag{A.18}$$

The $\exp(\pm i2k\hat{x})$ type terms evaluate to

$$\begin{aligned}\langle p|\exp(\pm i2k\hat{x})|\psi\rangle &= \langle p|\exp(\pm i2k\hat{x})I|\psi\rangle \\ &= \int_{-\infty}^{\infty}\langle p|e^{\pm i2k\hat{x}}|p'\rangle\langle p'|\psi\rangle\,\mathrm{d}p' \\ &= \int_{-\infty}^{\infty}\delta(p'-p\pm 2\hbar k)\phi(p')\,\mathrm{d}p' \\ &= \phi(p\mp 2\hbar k)\end{aligned} \tag{A.19}$$

where in the second to last step we used the following relation

$$\begin{aligned}\langle p|\exp(\pm i2k\hat{x})|p'\rangle =& \langle p|I\exp(\pm i2k\hat{x})I|p'\rangle \\ =& \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\langle p|z\rangle\langle z|\exp(\pm i2k\hat{x})|z'\rangle\langle z'|p'\rangle\,\mathrm{d}z'\,\mathrm{d}z \\ =& \frac{1}{2\pi\hbar}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\exp(-ipz/\hbar)\langle z|z'\rangle\exp(\pm i2kz)\exp(ip'z'/\hbar)\,\mathrm{d}z'\,\mathrm{d}z \\ =& \frac{1}{2\pi\hbar}\int_{-\infty}^{\infty}\exp(i\,(p'-p\pm 2\hbar k)\,z/\hbar)\,\mathrm{d}z \\ =& \delta(p'-p\pm 2\hbar k) \end{aligned} \tag{A.20}$$

□

A.5 Supporting proofs for adiabatic elimination

The proofs in this section derive an identity for the form of an integral of an arbitrary function multiplied by a complex sinusoid. This identity makes it straightforward to adiabatically eliminate the excited state. I am not certain if a proof for Theorem A.3 or Theorem A.4 exists in the literature. A proof for Theorem A.5 exists in [129]. I expect the proofs here already exist somewhere. I derived them from scratch to get a better sense of the exact approximation used in adiabatic elimination.

Theorem A.3. *The integral*

$$\int_0^x e^{i\alpha x'}f(x')\,\mathrm{d}x' \tag{A.21}$$

can be integrated by parts an arbitrary number of times to yield

$$\int_0^x e^{i\alpha x'}f(x')\,\mathrm{d}x' = -\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1}f^{(k)}(x')e^{i\alpha x'}\Big|_0^x + \left(\frac{i}{\alpha}\right)^{m+1}\int_0^x f^{(m+1)}(x')e^{i\alpha x'}\,\mathrm{d}x' \tag{A.22}$$

$$= \lim_{m\to\infty}\left[-\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1}f^{(k)}(x')e^{i\alpha x'}\Big|_0^x + \left(\frac{i}{\alpha}\right)^{m+1}f^{(m)}(0)\frac{e^{x+i\alpha x}-1}{1+i\alpha}\right] \tag{A.23}$$

where in the limit of infinite applications of integration by parts we can get rid of the integral.

Proof. First employ Theorem A.4 where

$$u^{(0)}(x) = f(x) \qquad v^{(n)}(x) = e^{i\alpha x} \tag{A.24}$$

This yields

$$\int_0^x e^{i\alpha x'} f(x')\,\mathrm{d}x' = \sum_{k=0}^{m}(-1)^k f^{(k)}(x')\frac{e^{i\alpha x'}}{(i\alpha)^{k+1}}\Big|_0^x - (-1)^m \int_0^x f^{(m+1)}(x')\frac{e^{i\alpha x'}}{(i\alpha)^{m+1}}\,\mathrm{d}x' \tag{A.25}$$

$$= -\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1} f^{(k)}(x')e^{i\alpha x'}\Big|_0^x + \left(\frac{i}{\alpha}\right)^{m+1}\int_0^x f^{(m+1)}(x')e^{i\alpha x'}\,\mathrm{d}x' \tag{A.26}$$

We are permitted to integrate by parts as many times as we would like, so we take $m \to \infty$

$$\int_0^x e^{i\alpha x'} f(x')\,\mathrm{d}x' = -\lim_{m\to\infty}\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1} f^{(k)}(x')e^{i\alpha x'}\Big|_0^x + \left(\frac{i}{\alpha}\right)^{m+1}\int_0^x f^{(m+1)}(x')e^{i\alpha x'}\,\mathrm{d}x' \tag{A.27}$$

via Theorem A.5 note that $f^{(\infty)}(x) = \lim_{m\to\infty} f^{(m)}(0)e^x$, so

$$\begin{aligned}\int_0^x e^{i\alpha x'} f(x')\,\mathrm{d}x' =& -\lim_{m\to\infty}\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1} f^{(k)}(x')e^{i\alpha x'}\Big|_0^x \\ &+ \lim_{m\to\infty}\left(\frac{i}{\alpha}\right)^{m+1} f^{(m)}(0)\int_0^x e^{x'}e^{i\alpha x'}\,\mathrm{d}x' \qquad \text{(A.28)}\\ =& -\lim_{m\to\infty}\sum_{k=0}^{m}\left(\frac{i}{\alpha}\right)^{k+1} f^{(k)}(x')e^{i\alpha x'}\Big|_0^x \\ &+ \lim_{m\to\infty}\left(\frac{i}{\alpha}\right)^{m+1} f^{(m)}(0)\frac{e^{x+i\alpha x}-1}{1+i\alpha} \qquad \text{(A.29)}\end{aligned}$$

□

Theorem A.4.

$$\int_a^b u^{(0)}(x)v^{(n)}(x)\,\mathrm{d}x = \sum_{k=0}^{m}(-1)^k u^{(k)}(x)v^{(n-k-1)}(x)\Big|_a^b - (-1)^m\int_a^b u^{(m+1)}(x)v^{(n-m-1)}(x)\,\mathrm{d}x \tag{A.30}$$

Proof. Rewrite the standard integration by parts formula, replacing $v^{(1)}(x) \to v^{(n)}(x)$

$$\int_a^b u^{(0)}(x)v^{(n)}(x)\,\mathrm{d}x = u^{(0)}(x)v^{(n-1)}(x)\Big|_a^b - \int_a^b u^{(1)}(x)v^{(n-1)}(x)\,\mathrm{d}x \tag{A.31}$$

Plug the left side into the right hand side

$$\int_a^b u^{(0)}(x)v^{(n)}(x)\,\mathrm{d}x = u^{(0)}(x)v^{(n-1)}(x)\Big|_a^b - \left[u^{(1)}(x)v^{(n-2)}(x)\Big|_a^b - \int_a^b u^{(2)}(x)v^{(n-2)}(x)\,\mathrm{d}x\right] \tag{A.32}$$

$$= \sum_{k=0}^{m}(-1)^k u^{(k)}(x)v^{(n-k-1)}(x)\Big|_a^b - (-1)^m \int_a^b u^{(m+1)}(x)v^{(n-m-1)}(x)\,\mathrm{d}x \tag{A.33}$$

□

We are interested in computing a derivative of form $\lim_{n\to\infty} f^{(n)}(x)$. Surprisingly, it turns out that there is a nice form for this if $\lim_{n\to\infty} f^{(n)}(x) = g(x)$, that is if the limit of the derivative of infinite order approaches some function. This was originally proved in [129], which contains a more in-depth discussion than what is provided here.

Theorem A.5. *Let $f(x)$ have derivatives of all orders and be an analytic function. If $\lim_{n\to\infty} f^{(n)}(x) = g(x)$ then $g(x) = ae^x$ where $a = \lim_{n\to\infty} f^{(n)}(0)$.*

Two proofs are provided.

Proof. Since $f(x)$ is an analytic function it can be Taylor expanded

$$f(x) = \sum_{m=0}^{\infty} f^{(m)}(0)\frac{x^m}{m!} \tag{A.34}$$

Therefore $f^{(n)}(x)$ is given by

$$f^{(n)}(x) = \sum_{m=0}^{\infty} f^{(m+n)}(0)\frac{x^m}{m!} \tag{A.35}$$

In the limit that $n \to \infty$

$$\lim_{n\to\infty} f^{(n)}(x) = \lim_{n\to\infty}\sum_{m=0}^{\infty} f^{(m+n)}(0)\frac{x^m}{m!} = \sum_{m=0}^{\infty}\left(\lim_{n\to\infty} f^{(m+n)}(0)\right)\frac{x^m}{m!} \tag{A.36}$$

Noting that

$$\lim_{n\to\infty} f^{(m+n)}(0) = \lim_{n\to\infty} f^{(n)}(0) \tag{A.37}$$

we have

$$\lim_{n\to\infty} f^{(n)}(x) = \sum_{m=0}^{\infty}\left(\lim_{n\to\infty} f^{(n)}(0)\right)\frac{x^m}{m!} = \left(\lim_{n\to\infty} f^{(n)}(0)\right)\sum_{m=0}^{\infty}\frac{x^m}{m!} = ae^x \tag{A.38}$$

where

$$a = \lim_{n\to\infty} f^{(n)}(0) \tag{A.39}$$

□

Proof. Let $\lim_{n\to\infty} f^{(n)}(x) = g(x)$. So therefore using the $n-1$ derivative

$$f^{(n-1)}(b) - f^{(n-1)}(c) = \int_c^b g(x)\,\mathrm{d}x \tag{A.40}$$

But in the limit of $n \to \infty$ it must be that $f^{(n-1)}(x) = g(x)$, so the above line becomes

$$g(x) - g(x) = \int_c^b g(x)\,\mathrm{d}x \tag{A.41}$$

Therefore $g(x) = g'(x)$, which means

$$\lim_{n\to\infty} f^{(n)}(x) = g(x) = ae^x \tag{A.42}$$

For $x = 0$ it must be that

$$\lim_{n\to\infty} f^{(n)}(x)\big|_{x=0} = ae^0 = a \tag{A.43}$$

so

$$a = \lim_{n\to\infty} f^{(n)}(0) \tag{A.44}$$

□

A.6 Two-level system in position space

We begin by successively applying $\langle x| \otimes \langle g|$ and $\langle x| \otimes \langle e|$ to the left-hand side of Eq. (2.4). This produces a system of PDEs in $\psi_g(t,z) = (\langle z| \otimes \langle g|)\left|\widetilde{\psi}\right\rangle$ and $\psi_e(t,z) = (\langle z| \otimes \langle e|)\left|\widetilde{\psi}\right\rangle$. Using the logic from Section 2.1 an explicit form for $\psi_e(t,z)$ is determined.

$$\begin{aligned} i\hbar\dot{\psi}_g(t,z) &= -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\psi_g(t,z) + \frac{\hbar}{2}\sum_\ell \Omega_\ell(t)\exp(-i(k_\ell z - \omega_\ell t + \omega_e t))\psi_e(t,z) \\ \psi_e(t,z) &\approx -\sum_\ell \frac{\Omega_\ell(t)}{2(\omega_e - \omega_\ell)}\exp(i(k_\ell z - \omega_\ell t + \omega_e t))\psi_g(t,z) \end{aligned} \tag{A.45}$$

For the case of a bichromatic field with two counterpropagating lasers ($k = k_1 \approx -k_2$) our expression for $\psi_g(t,z)$ becomes

$$i\hbar\dot{\psi}_g(t,z) = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\psi_g(t,z) - \hbar\left[\Omega_1^{\mathrm{AC}}(t) + \Omega_2^{\mathrm{AC}}(t) + \Omega_{\mathrm{eff}}(t)\cos(2kz - \delta t)\right]\psi_g(t,z) \tag{A.46}$$

where we have used the definitions in Eq. (2.13) and the approximation $-\omega_{1,2} + \omega_e \approx \Delta$ in the denominators.

Dropping the AC Stark shift terms we find

$$i\hbar\dot{\psi}_g(t,z) = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial z^2}\psi_g(t,z) - \hbar\Omega_{\text{eff}}(t)\cos\left(2kz-\delta t\right)\psi_g(t,z) \tag{A.47}$$

This PDE can be solved using the split-step method [94] or the Crank-Nicolson method [86, 95–98]. Both methods are discussed in Section 6.2.2.

A.7 Serrodyne phase modulation equations of motion

In Eq. (2.5) and Eq. (2.11) we derived expressions for the ground and excited state populations. We restate them here, requiring that $\Omega_\ell(0) = 0$

$$i\hbar\dot{\tilde{c}}_g(t,p) = \frac{p^2}{2m}\tilde{c}_g(t,p) + \frac{\hbar}{2}\sum_\ell \Omega_\ell(t)\exp(-i(\Delta_\ell t + \theta_\ell(t)))\tilde{c}_e(t,p+\hbar k_\ell) \tag{A.48}$$

$$i\hbar\tilde{c}_e(t,p) = -i\frac{\hbar}{2}\sum_\ell \frac{\Omega_\ell(t)\exp(i(\Delta_\ell t + \theta_\ell(t)))}{\Delta_\ell}\tilde{c}_g(t,p-\hbar k_\ell) \tag{A.49}$$

where we've written the single photon detuning as $\Delta_\ell \approx \omega_e - \omega_\ell$ (dropping the kinetic energy term), and inserted the appropriate dependence on the phase of the field ℓ. Our goal is to show that the equations of motion resulting from these expressions are the same in the case of an atomic fountain with two upward propagating and two downward propagating beams (one of which, in each direction, is phase modulated), and in the case of one upward propagating and one downward propagating beam with phase modulation. Eliminating $\tilde{c}_e(t,p)$ from the expression for $\dot{\tilde{c}}_g(t,p)$ yields

$$\begin{aligned} i\hbar\dot{\tilde{c}}_g(t,p) =& \frac{p^2}{2m}\tilde{c}_g(t,p) \\ &- \frac{\hbar}{4}\sum_{\ell,\ell'}\frac{\Omega_\ell(t)\Omega_{\ell'}(t)}{\Delta_{\ell'}}\exp(i[(\omega_\ell-\omega_{\ell'})t + \theta_{\ell'} - \theta_\ell])\tilde{c}_g(t,p+\hbar(k_\ell - k_{\ell'})) \end{aligned} \tag{A.50}$$

We first consider the two field system, where

$$\begin{gathered} \omega_1 = \omega_2 = \omega \\ k_1 = -k_2 = k \\ \Omega_1(t) = \Omega_2(t) = \sqrt{2\Omega_{\text{eff}}(t)\Delta} \\ \theta_1(t) = 0 \end{gathered}$$

Since both fields have the same frequency the single photon detunings are identical, and we write $\Delta_1 = \Delta_2 = \Delta$. Switching to notation $c_j(t) = \tilde{c}_g(t, 2j\hbar k)$ yields

$$\begin{aligned} i\hbar\dot{c}_j(t) =& 4\hbar j^2\omega_{\rm r}c_j(t) - \hbar\Omega_{\rm eff}(t)c_j(t) \\ &- \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(-i\theta_2(t))c_{j-1}(t) \\ &- \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(i\theta_2(t))c_{j+1}(t) \end{aligned} \tag{A.51}$$

Note that this is Eq. (2.14) with $\delta t \to \theta_2(t)$. We will now show that the same result is found in the four field system with the case of an atom far from rest. We consider the case where

$$\begin{gathered} \omega_3 = \omega_1 \qquad \omega_4 = \omega_2 \\ k_1 = k_2 = -k_3 = -k_4 = k \\ \Omega_1(t) = \Omega_2(t) = \Omega_3(t) = \Omega_4(t) = \sqrt{2\Omega_{\rm eff}(t)\Delta} \\ \theta_1(t) = \theta_3(t) \qquad \theta_2(t) = \theta_4(t) = 0 \end{gathered}$$

We assume that ω_1 and ω_2 are order 1 kHz to 1 MHz apart. We also assume both of these are optical frequencies (i.e. several hundred THz). This means that we can approximate each $1/\Delta_{\ell'}$ in the sum as the same $1/\Delta$. Letting $\delta = \omega_1 - \omega_2$ and switching to the notation $c_j(t) = \tilde{c}_g(t, 2j\hbar k)$ yields

$$i\hbar\dot{c}_j(t) = 4\hbar j^2\omega_{\rm r}c_j(t) - \hbar\Omega_{\rm eff}(t)\left(1+\cos(\delta t - \theta_1(t))\right)\left(2c_j(t) + c_{j-1}(t) + c_{j+1}(t)\right) \tag{A.52}$$

An atom in the fountain far from rest has high j, and so we are able to make the rotating wave approximation. To do this we transform to the frame $\tilde{c}_j(t) = \exp(i4j^2\omega_{\rm r}t)c_j(t)$ to find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_j(t) =& -\hbar\Omega_{\rm eff}(t)\left(1+\cos(\delta t - \theta_1(t))\right)\left[2\tilde{c}_j(t)\right. \\ &\left. + \exp(i(8j-4)\omega_{\rm r}t)\tilde{c}_{j-1}(t) + \exp(-i(8j+4)\omega_{\rm r}t)\tilde{c}_{j+1}(t)\right] \end{aligned} \tag{A.53}$$

If the cosine term is expanded into exponentials and we choose δ and j to be positive, then we can make the rotating wave approximation to find

$$\begin{aligned} i\hbar\dot{\tilde{c}}_j(t) =& -2\hbar\Omega_{\rm eff}(t)\tilde{c}_j(t) \\ &- \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(i(-\delta t + \theta_1(t) - 4\omega_{\rm r}t + 8j\omega_{\rm r}t))\tilde{c}_{j-1}(t) \\ &- \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(i(\delta t - \theta_1(t) - 4\omega_{\rm r}t - 8j\omega_{\rm r}t))\tilde{c}_{j+1}(t) \end{aligned} \tag{A.54}$$

Reverting the frame transformation (this time including the δ term) leads to

$$\begin{aligned} i\hbar\dot{c}_j(t) =& \hbar\left(4j^2\omega_{\rm r} - j\delta - 2\Omega_{\rm eff}(t)\right)c_j(t) \\ &- \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(i\theta_1(t))c_{j-1}(t) - \hbar\frac{\Omega_{\rm eff}(t)}{2}\exp(-i\theta_1(t))c_{j+1}(t) \end{aligned} \tag{A.55}$$

If $\delta = 4n\omega_{\mathrm{r}}$ then we see this result is equivalent to Eq. (A.51) up to the sign of the phase and a factor of 2 in the AC Stark shift term, which can be dropped as it contributes a common phase to both interferometer arms and therefore does not contribute to the interferometer phase. Therefore the two systems have identical dynamics in the limit we operate them in.

Appendix B

Coil switching electronics

Several custom electrical circuits were developed for this thesis. We discussed the circuit used to stabilize laser intensity in Fig. 4.13, a circuit used to detect the phase of the beatnote between two lasers in Fig. 4.14, and a PI loop capable of automatically relocking in Fig. 4.15. Here we briefly review several other circuits used to switch coil currents in the experiment.

To switch the 3D MOT coils on and off the circuit in Fig. B.1 is used. The 3D MOT coils have an inductance of 6 mH and are driven with 200 A. Therefore, when driven, they store an energy of 120 J. This energy must be rapidly dissipated for the coils to turn off.

When the coils are on, the STMicroelectronics STE145N65M5 MOSFETS Q2-5 act as shorts, and current flows from ground to the −15 V source through the MOT coil and then across Q2-Q5 in parallel. To switch the MOSFETs to closed the BNC connected to the optical isolator U1 is switched and the gate voltage to Q2-5 is changed to close the MOSFETs. In this configuration the coil continues driving current, which now passes across the 4 Ω resistors in parallel and then goes through two IXYS DSA300I200NA Schottky diodes. The energy stored in the inductor is dissipated through resistance in the resistors and diodes. The diodes prevent current from bypassing the coil when current is on. This configuration is called a flyback circuit. The Vishay VOH1016AD-V varistor RV1 shields the MOSFETs from any voltage spike that occurs during the switching process. Circuit components are connected using small gauge (large diameter) wires, or are connected directly with copper strips.

We probe the MOSFET drain voltage at the time of the turn off to determine the time taken for the coils to turn off. The recorded voltage trace is shown in Fig. B.2. We see the turn off takes around 4 ms. Note that the −15 V source is always on. As described in Section 4.1.2, this turn-off time is much faster than the time it takes for magnetic fields due to eddy currents in the chamber to die off.

We use two other basic circuits to set the current of smaller coils in the experiment, for example RSC coils or compensation coils. The first is shown in Fig. B.3, which switches current on and off. It operates using the same principle as the 3D coil switching circuit, but lacks a flyback circuit. The IXYS IXTK32P60P p-channel MOSFET is used for Q1. The other circuit drives current according to a reference voltage, and is shown in Fig. B.4. It uses an op-amp feedback loop to stabilize the current passing through a 1 Ω sense resistor. A

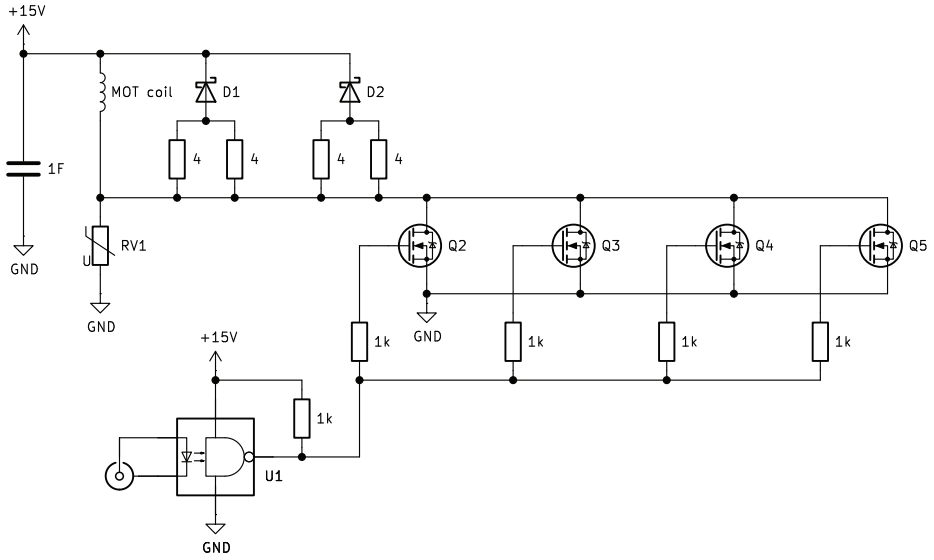


Figure B.1: The 3D MOT coil switching circuit.

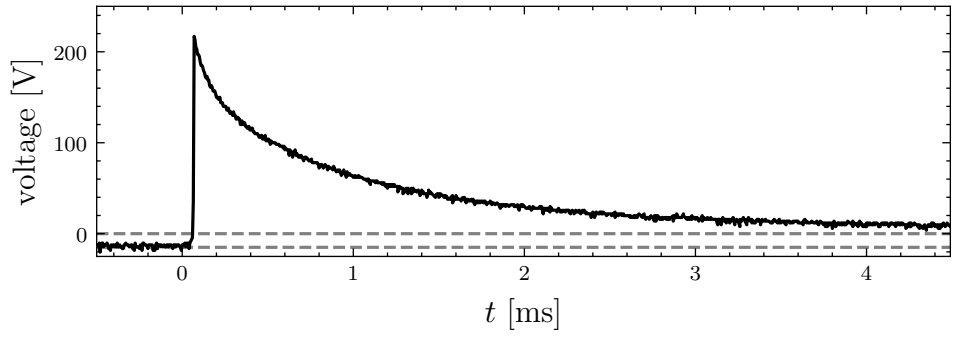


Figure B.2: The 3D MOT coil switching circuit turn off. The lower dashed line is at $-15\,\mathrm{V}$, which we expect the MOSFET drain voltage to be close to before the switch. The upper dashed line is at $0\,\mathrm{V}$, which we expect the probe voltage to decay to once no current flows through the MOT coil.

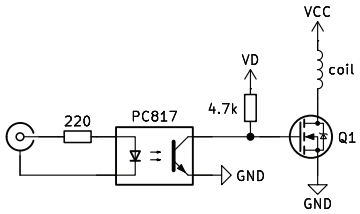


Figure B.3: Basic coil switching circuit.

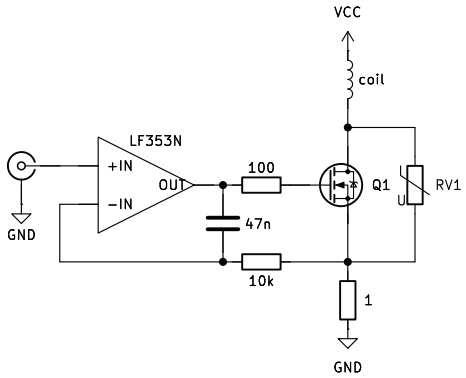


Figure B.4: Coil driver circuit.

varistor protects the MOSFET during turn off. A power resistor is used for the sense resistor, and the MOSFET is mounted to a large heatsink to enable the use of large currents.

Bibliography

[1] C. Davisson and L. H. Germer. "Diffraction of Electrons by a Crystal of Nickel". In: *Physical Review* 30.6 (Dec. 1927), pp. 705–740. ISSN: 0031-899X. DOI: 10.1103/physrev.30.705.

[2] R. Colella, A. W. Overhauser, and S. A. Werner. "Observation of Gravitationally Induced Quantum Interference". In: *Physical Review Letters* 34.23 (June 1975), pp. 1472–1474. ISSN: 0031-9007. DOI: 10.1103/physrevlett.34.1472.

[3] Richard H. Parker, Chenghui Yu, Weicheng Zhong, Brian Estey, and Holger Müller. "Measurement of the fine-structure constant as a test of the Standard Model". In: *Science* 360.6385 (2018), pp. 191–195. DOI: 10.1126/science.aap7706.

[4] Albert A. Michelson and Edward W. Morley. "LIX. On a method of making the wavelength of sodium light the actual and practical standard of length". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 24.151 (Dec. 1887), pp. 463–466. ISSN: 1941-5990. DOI: 10.1080/14786448708628131.

[5] Robert S. Shankland. "Michelson and his interferometer". In: *Physics Today* 27.4 (Apr. 1974), pp. 37–43. ISSN: 1945-0699. DOI: 10.1063/1.3128534.

[6] J. von Fraunhofer. *Prismatic and Diffraction Spectra: Memoirs*. Harper's scientific memoirs. Harper & Brothers, 1984. URL: https://books.google.com/books?id=flZCAQAAMAAJ.

[7] Niels Bohr. "IV. On the Series Spectrum of Hydrogen and the Structure of the Atom: (Phil. Mag. 29 (1915) 332)". In: *Work on Atomic Physics (1912–1917)*. Ed. by L. Rosenfeld and Ulrich Hoyer. Elsevier, 1981, pp. 376–380. DOI: 10.1016/s1876-0503(08)70065-3.

[8] A. Sommerfeld. "Zur Quantentheorie der Spektrallinien". In: *Annalen der Physik* 356.17 (Jan. 1916), pp. 1–94. ISSN: 1521-3889. DOI: 10.1002/andp.19163561702.

[9] U.D. Jentschura and I. Nándori. "Attempts at a determination of the fine-structure constant from first principles: a brief historical overview". In: *The European Physical Journal H* 39.5 (Dec. 2014), pp. 591–613. ISSN: 2102-6467. DOI: 10.1140/epjh/e2014-50044-7.

[10] Helge Kragh. "Magic Number: A Partial History of the Fine-Structure Constant". In: *Archive for History of Exact Sciences* 57.5 (July 2003), pp. 395–431. ISSN: 1432-0657. DOI: 10.1007/s00407-002-0065-7.

[11] Peter J. Mohr, David B. Newell, Barry N. Taylor, and Eite Tiesinga. "CODATA recommended values of the fundamental physical constants: 2022". In: *Reviews of Modern Physics* 97.2 (Apr. 2025). ISSN: 1539-0756. DOI: 10.1103/revmodphys.97.025002.

[12] Axel Beyer, Lothar Maisenbacher, Arthur Matveev, Randolf Pohl, Ksenia Khabarova, Alexey Grinin, Tobias Lamour, Dylan C. Yost, Theodor W. Hänsch, Nikolai Kolachevsky, and Thomas Udem. "The Rydberg constant and proton size from atomic hydrogen". In: *Science* 358.6359 (Oct. 2017), pp. 79–85. ISSN: 1095-9203. DOI: 10.1126/science.aah6677.

[13] S. Sturm, F. Köhler, J. Zatorski, A. Wagner, Z. Harman, G. Werth, W. Quint, C. H. Keitel, and K. Blaum. "High-precision measurement of the atomic mass of the electron". In: *Nature* 506.7489 (Feb. 2014), pp. 467–470. ISSN: 1476-4687. DOI: 10.1038/nature13026.

[14] Meng Wang, G. Audi, F. G. Kondev, W.J. Huang, S. Naimi, and Xing Xu. "The AME2016 atomic mass evaluation (II). Tables, graphs and references". In: *Chinese Physics C* 41.3 (Mar. 2017), p. 030003. ISSN: 1674-1137. DOI: 10.1088/1674-1137/41/3/030003.

[15] D. S. Weiss, B. C. Young, and S. Chu. "Precision measurement of h/m Cs based on photon recoil using laser-cooled atoms and atomic interferometry". In: *Applied Physics B Lasers and Optics* 59.3 (Sept. 1994), pp. 217–256. ISSN: 1432-0649. DOI: 10.1007/bf01081393.

[16] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. "Measurement of the Electron Magnetic Moment". In: *Physical Review Letters* 130.7 (Feb. 2023). ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.130.071801.

[17] Rym Bouchendira, Pierre Cladé, Saïda Guellati-Khélifa, François Nez, and François Biraben. "New Determination of the Fine Structure Constant and Test of the Quantum Electrodynamics". In: *Physical Review Letters* 106.8 (Feb. 2011). ISSN: 1079-7114. DOI: 10.1103/physrevlett.106.080801.

[18] Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa. "Determination of the fine-structure constant with an accuracy of 81 parts per trillion". In: *Nature* 588.7836 (Dec. 2020), pp. 61–65. ISSN: 1476-4687. DOI: 10.1038/s41586-020-2964-7.

[19] D. Hanneke, S. Fogwell, and G. Gabrielse. "New Measurement of the Electron Magnetic Moment and the Fine Structure Constant". In: *Physical Review Letters* 100.12 (Mar. 2008). ISSN: 1079-7114. DOI: 10.1103/physrevlett.100.120801.

[20] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. "Theory of the Anomalous Magnetic Moment of the Electron". In: *Atoms* 7.1 (Feb. 2019), p. 28. ISSN: 2218-2004. DOI: 10.3390/atoms7010028.

[21] Brian Vincent Estey. "Precision Measurement in Atom Interferometry Using Bragg Diffraction". PhD thesis. University of California, Berkeley, 2016. URL: https://escholarship.org/content/qt4t19t317/qt4t19t317_noSplash_80c87dcf7c8d1757dfe21b0e35c034bf.pdf.

[22] Chenghui Yu. "Measuring the Fine Structure Constant with a State-of-the-art Atom Interferometer". PhD thesis. UC Berkeley, 2018, p. 186. ISBN: 978-0-438-64251-5. URL: https://www.proquest.com/dissertations-theses/measuring-fine-structure-constant-with-state-art/docview/2135778785/se-2.

[23] Weicheng Zhong. "Precision Measurement of the Fine-Structure Constant with Atom Interferometry". English. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2026-05-06. PhD thesis. 2020, p. 109. ISBN: 9798597063386. URL: https://www.proquest.com/dissertations-theses/precision-measurement-fine-structure-constant/docview/2511291027/se-2.

[24] Zachary Pagel. "Towards a next-generation measurement of the fine-structure constant". PhD thesis. UC Berkeley, 2023. URL: https://escholarship.org/uc/item/5qm678c2.

[25] Jack Roth, Andrew Christensen, Madeline Bernstein, Yuno Iwasaki, Hana Lampson, and Holger Mueller. "Linewidth narrowing and wideband frequency modulation of a DBR laser". In: *Optics Express* 33.20 (Sept. 2025), p. 42240. ISSN: 1094-4087. DOI: 10.1364/oe.572129.

[26] Jack Roth, Andrew Christensen, Madeline Bernstein, Yuno Iwasaki, and Holger Mueller. *A fast and accurate method for simulating Bragg atom interferometers*. 2026. DOI: 10.48550/ARXIV.2601.17236.

[27] Jack Roth, Madeline Bernstein, Andrew Christensen, Yuno Iwasaki, Holger Mueller, Zaire Sprowal, and Robert Compton. *Serrodyne Matterwave Optics*. 2026. DOI: 10.48550/ARXIV.2607.24511.

[28] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. "Optics and interferometry with atoms and molecules". In: *Rev. Mod. Phys.* 81 (3 July 2009), pp. 1051–1129. DOI: 10.1103/RevModPhys.81.1051.

[29] Liang Hu, Nicola Poli, Leonardo Salvi, and Guglielmo M. Tino. "Atom Interferometry with the Sr Optical Clock Transition". In: *Physical Review Letters* 119.26 (Dec. 2017). ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.119.263601.

[30] Liang Hu, Enlong Wang, Leonardo Salvi, Jonathan N. Tinsley, Guglielmo M. Tino, and Nicola Poli. "Sr atom interferometry with the optical clock transition as a gravimeter and a gravity gradiometer". In: *Classical and Quantum Gravity* 37.1 (Nov. 2019), p. 014001. ISSN: 1361-6382. DOI: 10.1088/1361-6382/ab4d18.

[31] Jesse S. Schelfhout, Thomas M. Hird, Kenneth M. Hughes, and Christopher J. Foot. "Single-photon large-momentum-transfer atom interferometry scheme for Sr or Yb atoms with application to determining the fine-structure constant". In: *Physical Review A* 110.5 (Nov. 2024). ISSN: 2469-9934. DOI: 10.1103/physreva.110.053309.

[32] Daniel A. Steck. *Cesium D Line Data*. Tech. rep. revision 1.6. Los Alamos National Laboratory, Oct. 2003. URL: https://steck.us/alkalidata/cesiumnumbers.1.6.pdf.

[33] Holger Müller, Sheng-wey Chiow, and Steven Chu. "Atom-wave diffraction between the Raman-Nath and the Bragg regime: Effective Rabi frequency, losses, and phase shifts". In: *Physical Review A* 77.2 (Feb. 2008). ISSN: 1094-1622. DOI: 10.1103/PhysRevA.77.023609.

[34] Tim Kovachy, Jason M. Hogan, David M. S. Johnson, and Mark A. Kasevich. "Optical lattices as waveguides and beam splitters for atom interferometry: An analytical treatment and proposal of applications". In: *Physical Review A* 82.1 (July 2010). ISSN: 1094-1622. DOI: 10.1103/PhysRevA.82.013638.

[35] Tim Kovachy, Sheng-wey Chiow, and Mark A. Kasevich. "Adiabatic-rapid-passage multiphoton Bragg atom optics". In: *Physical Review A* 86.1 (July 2012). ISSN: 1094-1622. DOI: 10.1103/PhysRevA.86.011606.

[36] Ekkehard Peik, Maxime Ben Dahan, Isabelle Bouchoule, Yvan Castin, and Christophe Salomon. "Bloch oscillations of atoms, adiabatic rapid passage, and monokinetic atomic beams". In: *Physical Review A* 55.4 (Apr. 1997), pp. 2989–3001. ISSN: 1094-1622. DOI: 10.1103/PhysRevA.55.2989.

[37] Pierre Cladé, Saïda Guellati-Khélifa, François Nez, and François Biraben. "Large Momentum Beam Splitter Using Bloch Oscillations". In: *Physical Review Letters* 102.24 (June 2009). ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.102.240402.

[38] Pierre Cladé, Manuel Andia, and Saïda Guellati-Khélifa. "Improving efficiency of Bloch oscillations in the tight-binding limit". In: *Physical Review A* 95.6 (June 2017). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.95.063604.

[39] Zachary Pagel, Weicheng Zhong, Richard H. Parker, Christopher T. Olund, Norman Y. Yao, and Holger Müller. "Symmetric Bloch oscillations of matter waves". In: *Physical Review A* 102.5 (Nov. 2020). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.102.053312.

[40] Mark Kasevich and Steven Chu. "Atomic interferometry using stimulated Raman transitions". In: *Physical Review Letters* 67.2 (July 1991), pp. 181–184. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.67.181.

[41] Mark Kasevich, David S. Weiss, Erling Riis, Kathryn Moler, Steven Kasapi, and Steven Chu. "Atomic velocity selection using stimulated Raman transitions". In: *Phys. Rev. Lett.* 66 (18 May 1991), pp. 2297–2300. DOI: 10.1103/PhysRevLett.66.2297.

[42] R. Grimm, Yu. B. Ovchinnikov, A. I. Sidorov, and V. S. Letokhov. "Observation of a strong rectified dipole force in a bichromatic standing light wave". In: *Phys. Rev. Lett.* 65 (12 Sept. 1990), pp. 1415–1418. DOI: 10.1103/PhysRevLett.65.1415.

[43] J. Söding, R. Grimm, Yu. B. Ovchinnikov, Ph. Bouyer, and Ch. Salomon. "Short-Distance Atomic Beam Deceleration with a Stimulated Light Force". In: *Phys. Rev. Lett.* 78 (8 Feb. 1997), pp. 1420–1423. DOI: 10.1103/PhysRevLett.78.1420.

[44] S. E. Galica, L. Aldridge, and E. E. Eyler. "Four-color stimulated optical forces for atomic and molecular slowing". In: *Phys. Rev. A* 88 (4 Oct. 2013), p. 043418. DOI: 10.1103/PhysRevA.88.043418.

[45] Jan-Niclas Siemß, Florian Fitzek, Sven Abend, Ernst M. Rasel, Naceur Gaaloul, and Klemens Hammerer. "Analytic theory for Bragg atom interferometry based on the adiabatic theorem". In: *Physical Review A* 102.3 (Sept. 2020). ISSN: 2469-9934. DOI: 10.1103/physreva.102.033709.

[46] A. Béguin, T. Rodzinka, J. Vigué, B. Allard, and A. Gauguet. "Characterization of an atom interferometer in the quasi-Bragg regime". In: *Physical Review A* 105.3 (Mar. 2022). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.105.033302.

[47] Brian Estey, Chenghui Yu, Holger Müller, Pei-Chen Kuan, and Shau-Yu Lan. "High-Resolution Atom Interferometers with Suppressed Diffraction Phases". In: *Physical Review Letters* 115.8 (Aug. 2015). ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.115.083002.

[48] Víctor José Martínez-Lahuerta, Jan-Niclas Kirsten-Siemß, Klemens Hammerer, and Naceur Gaaloul. "Diffraction-phase-free Bragg atom interferometry". In: *AVS Quantum Science* 8.1 (Mar. 2026), p. 014404. ISSN: 2639-0213. DOI: 10.1116/5.0307499. eprint: https://pubs.aip.org/avs/aqs/article-pdf/doi/10.1116/5.0307499/20943413/014404_1_5.0307499.pdf.

[49] Bruce W. Shore. *The Theory of Coherent Atomic Excitation*. Wiley, 1990.

[50] Daniel A. Steck. *Quantum and Atom Optics*. revision 0.13.10, 22 September 2021. 2021. URL: http://steck.us/teaching.

[51] Xuejian Wu, Fei Zi, Jordan Dudley, Ryan J. Bilotta, Philip Canoza, and Holger Müller. "Multiaxis atom interferometry with a single-diode laser and a pyramidal magneto-optical trap". In: *Optica* 4.12 (Dec. 2017), p. 1545. ISSN: 2334-2536. DOI: 10.1364/optica.4.001545.

[52] Norman F. Ramsey. "A New Molecular Beam Resonance Method". In: *Phys. Rev.* 76 (7 Oct. 1949), p. 996. DOI: 10.1103/PhysRev.76.996.

[53] Norman F. Ramsey. "A Molecular Beam Resonance Method with Separated Oscillating Fields". In: *Phys. Rev.* 78 (6 June 1950), pp. 695–699. DOI: 10.1103/PhysRev.78.695.

[54] Norman F. Ramsey and Henry B. Silsbee. "Phase Shifts in the Molecular Beam Method of Separated Oscillating Fields". In: *Phys. Rev.* 84 (3 Nov. 1951), pp. 506–507. DOI: 10.1103/PhysRev.84.506.

[55] Robert G. Littlejohn. "Rotations in Quantum Mechanics, and Rotations of Spin 1/2 Systems". Physics 221A, Academic Year 2020–21, Notes 12. 2020. URL: http://bohr.physics.berkeley.edu/classes/221/2021/221.html.

[56] D.B. Sullivan. "Time and frequency measurement at NIST: the first 100 years". In: *Proceedings of the 2001 IEEE International Frequency Control Symposium and PDA Exhibition (Cat. No.01CH37218)*. 2001, pp. 4–17. DOI: 10.1109/FREQ.2001.956152.

[57] Robert G. Littlejohn. "The Propagator and the Path Integral". Physics 221A, Academic Year 2020–21, Notes 9. 2020. URL: http://bohr.physics.berkeley.edu/classes/221/2021/221.html.

[58] Sheng-wey Chiow, Sven Herrmann, Steven Chu, and Holger Müller. "Noise-Immune Conjugate Large-Area Atom Interferometers". In: *Physical Review Letters* 103.5 (July 2009). ISSN: 1079-7114. DOI: 10.1103/physrevlett.103.050402.

[59] Pippa Storey and Claude Cohen-Tannoudji. "The Feynman path integral approach to atomic interferometry. A tutorial". In: *Journal de Physique II* 4.11 (Nov. 1994), pp. 1999–2027. ISSN: 1286-4870. DOI: 10.1051/jp2:1994103.

[60] Albert Roura, Wolfgang Zeller, and Wolfgang P Schleich. "Overcoming loss of contrast in atom interferometry due to gravity gradients". In: *New Journal of Physics* 16.12 (Dec. 2014), p. 123012. ISSN: 1367-2630. DOI: 10.1088/1367-2630/16/12/123012.

[61] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. "Trapping of Neutral Sodium Atoms with Radiation Pressure". In: *Physical Review Letters* 59.23 (Dec. 1987), pp. 2631–2634. ISSN: 0031-9007. DOI: 10.1103/physrevlett.59.2631.

[62] Aden Zhenhao Lam, Claire Warner, Niccolò Bigagli, Stephan Roschinski, Weijun Yuan, Ian Stevenson, and Sebastian Will. *Compact Two-Dimensional Magneto-Optical Trap for Ultracold Atom Setups*. 2020. DOI: 10.48550/ARXIV.2012.06688.

[63] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Yu. B. Ovchinnikov, and T. Pfau. "Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap". In: *Physical Review A* 66.2 (Aug. 2002). ISSN: 1094-1622. DOI: 10.1103/physreva.66.023410.

[64] E. B. Norrgard, N. Sitaraman, J. F. Barry, D. J. McCarron, M. H. Steinecker, and D. DeMille. "In-vacuum scattered light reduction with black cupric oxide surfaces for sensitive fluorescence detection". In: *Review of Scientific Instruments* 87.5 (May 2016). ISSN: 1089-7623. DOI: 10.1063/1.4949503.

[65] Enrico Maccallini, Fabrizio Siviero, Antonio Bonucci, Andrea Conte, Peeyush Srivastava, and Manini Paolo. "Non evaporable getter (NEG) technology: A powerful tool for UHV-XHV systems". In: *AIP Conference Proceedings*. AIP, 2012, pp. 24–27. DOI: 10.1063/1.4732360.

[66] Gary C. Bjorklund. "Frequency-modulation spectroscopy: a new method for measuring weak absorptions and dispersions". In: *Optics Letters* 5.1 (Jan. 1980), p. 15. ISSN: 1539-4794. DOI: 10.1364/ol.5.000015.

[67] Fei Zi, Xuejian Wu, Weicheng Zhong, Richard H. Parker, Chenghui Yu, Simon Budker, Xuanhui Lu, and Holger Müller. "Laser frequency stabilization by combining modulation transfer and frequency modulation spectroscopy". In: *Applied Optics* 56.10 (Mar. 2017), p. 2649. ISSN: 1539-4522. DOI: 10.1364/ao.56.002649.

[68] John F. Barry. "Laser cooling and slowing of a diatomic molecule". PhD thesis. Yale University, 2013, p. 307. ISBN: 978-1-303-71588-4. URL: https://www.proquest.com/dissertations-theses/laser-cooling-slowing-diatomic-molecule/docview/1496011189/se-2.

[69] F.-X. Esnault, E. Blanshan, E. N. Ivanov, R. E. Scholten, J. Kitching, and E. A. Donley. "Cold-atom double-Λ coherent population trapping clock". In: *Physical Review A* 88.4 (Oct. 2013), p. 042120. ISSN: 1094-1622. DOI: 10.1103/physreva.88.042120.

[70] Xiaochi Liu, Eugene Ivanov, Valeriy I. Yudin, John Kitching, and Elizabeth A. Donley. "Low-Drift Coherent Population Trapping Clock Based on Laser-Cooled Atoms and High-Coherence Excitation Fields". In: *Physical Review Applied* 8.5 (Nov. 2017), p. 054001. ISSN: 2331-7019. DOI: 10.1103/physrevapplied.8.054001.

[71] Gilbert Grynberg, Alain Aspect, Claude Fabre, and Claude Cohen-Tannoudji. *Introduction to Quantum Optics: From the Semi-classical Approach to Quantized Light*. Cambridge University Press, Sept. 2010. ISBN: 9780511778261. DOI: 10.1017/cbo9780511778261.

[72] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Springer New York, 1999. DOI: 10.1007/978-1-4612-1470-0.

[73] Paul D. Lett, Richard N. Watts, Christoph I. Westbrook, William D. Phillips, Phillip L. Gould, and Harold J. Metcalf. "Observation of Atoms Laser Cooled below the Doppler Limit". In: *Physical Review Letters* 61.2 (July 1988), pp. 169–172. ISSN: 0031-9007. DOI: 10.1103/physrevlett.61.169.

[74] J. Dalibard and C. Cohen-Tannoudji. "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models". In: *Journal of the Optical Society of America B* 6.11 (Nov. 1989), p. 2023. ISSN: 1520-8540. DOI: 10.1364/josab.6.002023.

[75] Claude N. Cohen-Tannoudji and William D. Phillips. "New Mechanisms for Laser Cooling". In: *Physics Today* 43.10 (Oct. 1990), pp. 33–40. ISSN: 1945-0699. DOI: 10.1063/1.881239.

[76] Vladan Vuletić, Cheng Chin, Andrew J. Kerman, and Steven Chu. "Degenerate Raman Sideband Cooling of Trapped Cesium Atoms at Very High Atomic Densities". In: *Physical Review Letters* 81.26 (Dec. 1998), pp. 5768–5771. ISSN: 1079-7114. DOI: 10.1103/physrevlett.81.5768.

[77] Andrew J. Kerman, Vladan Vuletić, Cheng Chin, and Steven Chu. "Beyond Optical Molasses: 3D Raman Sideband Cooling of Atomic Cesium to High Phase-Space Density". In: *Physical Review Letters* 84.3 (Jan. 2000), pp. 439–442. ISSN: 1079-7114. DOI: 10.1103/physrevlett.84.439.

[78] Dian-Jiun Han, Steffen Wolf, Steven Oliver, Colin McCormick, Marshall T. DePue, and David S. Weiss. "3D Raman Sideband Cooling of Cesium Atoms at High Density". In: *Physical Review Letters* 85.4 (July 2000), pp. 724–727. ISSN: 1079-7114. DOI: 10.1103/physrevlett.85.724.

[79] Philipp Treutlein, Keng Yeow Chung, and Steven Chu. "High-brightness atom source for atomic fountains". In: *Physical Review A* 63.5 (Apr. 2001). ISSN: 1094-1622. DOI: 10.1103/physreva.63.051401.

[80] Andrew J. Kerman. "Raman sideband cooling and cold atomic collisions in optical lattices". English. PhD thesis. Stanford University, 2002, p. 285. ISBN: 978-0-493-87501-9. URL: https://www.proquest.com/dissertations-theses/raman-sideband-cooling-cold-atomic-collisions/docview/305520254/se-2.

[81] G. Grynberg, B. Lounis, P. Verkerk, J.-Y. Courtois, and C. Salomon. "Quantized motion of cold cesium atoms in two- and three-dimensional optical potentials". In: *Physical Review Letters* 70.15 (Apr. 1993), pp. 2249–2252. ISSN: 0031-9007. DOI: 10.1103/physrevlett.70.2249.

[82] Matthew Jaffe. "Atom Interferometry in an Optical Cavity". English. PhD thesis. UC Berkeley, 2018, p. 229. ISBN: 9798597063911. URL: https://www.proquest.com/dissertations-theses/atom-interferometry-optical-cavity/docview/2511368679/se-2.

[83] Radim Halıř and Jan Flusser. "Numerically stable direct least squares fitting of ellipses". In: *Proc. 6th International Conference in Central Europe on Computer Graphics and Visualization. WSCG*. Vol. 98. 1998, pp. 125–132.

[84] Ying-Cheng Chen, Yean-An Liao, Long Hsu, and Ite A. Yu. "Simple technique for directly and accurately measuring the number of atoms in a magneto-optical trap". In: *Physical Review A* 64.3 (Aug. 2001). ISSN: 1094-1622. DOI: 10.1103/physreva.64.031401.

[85] Robert G. Littlejohn. "Spatial Degrees of Freedom". Physics 221A, Academic Year 2020–21, Notes 4. 2020. URL: http://bohr.physics.berkeley.edu/classes/221/2021/221.html.

[86] Toshiaki Iitaka. "Solving the time-dependent Schrödinger equation numerically". In: *Physical Review E* 49.5 (May 1994), pp. 4684–4690. ISSN: 1095-3787. DOI: 10.1103/PhysRevE.49.4684.

[87] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. 3rd ed. USA: Cambridge University Press, 2007. ISBN: 0521880688.

[88] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[89] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. "Numba: a LLVM-based Python JIT compiler". In: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*. LLVM '15. Austin, Texas: Association for Computing Machinery, 2015. ISBN: 9781450340052. DOI: 10.1145/2833157.2833162.

[90] J.R. Dormand and P.J. Prince. "A family of embedded Runge-Kutta formulae". In: *Journal of Computational and Applied Mathematics* 6.1 (Mar. 1980), pp. 19–26. ISSN: 0377-0427. DOI: 10.1016/0771-050x(80)90013-3.

[91] Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Pushmeet Kohli, and Matej Balog. *AlphaEvolve: A coding agent for scientific and algorithmic discovery*. 2025. DOI: 10.48550/ARXIV.2506.13131.

[92] Utkarsh Utkarsh, Valentin Churavy, Yingbo Ma, Tim Besard, Prakitr Srisuma, Tim Gymnich, Adam R. Gerlach, Alan Edelman, George Barbastathis, Richard D. Braatz, and Christopher Rackauckas. "Automated translation and accelerated solving of differential equations on multiple GPU platforms". In: *Computer Methods in Applied Mechanics and Engineering* 419 (Feb. 2024), p. 116591. ISSN: 0045-7825. DOI: 10.1016/j.cma.2023.116591.

[93] M.D Feit, J.A Fleck, and A Steiger. "Solution of the Schrödinger equation by a spectral method". In: *Journal of Computational Physics* 47.3 (Sept. 1982), pp. 412–433. ISSN: 0021-9991. DOI: 10.1016/0021-9991(82)90091-2.

[94] Florian Fitzek, Jan-Niclas Siemß, Stefan Seckmeyer, Holger Ahlers, Ernst M. Rasel, Klemens Hammerer, and Naceur Gaaloul. "Universal atom interferometer simulation of elastic scattering processes". In: *Scientific Reports* 10.1 (Dec. 2020). ISSN: 2045-2322. DOI: 10.1038/s41598-020-78859-1.

[95] W. van Dijk and F. M. Toyama. "Accurate numerical solutions of the time-dependent Schrödinger equation". In: *Physical Review E* 75.3 (Mar. 2007). ISSN: 1550-2376. DOI: 10.1103/PhysRevE.75.036707.

[96] W. van Dijk, J. Brown, and K. Spyksma. "Efficiency and accuracy of numerical solutions to the time-dependent Schrödinger equation". In: *Physical Review E* 84.5 (Nov. 2011). ISSN: 1550-2376. DOI: 10.1103/PhysRevE.84.056703.

[97] W. van Dijk and F. M. Toyama. "Numerical solutions of the Schrödinger equation with source terms or time-dependent potentials". In: *Physical Review E* 90.6 (Dec. 2014). ISSN: 1550-2376. DOI: 10.1103/PhysRevE.90.063309.

[98] Wytse van Dijk. "On numerical solutions of the time-dependent Schrödinger equation". In: *American Journal of Physics* 91.10 (Oct. 2023), pp. 826–839. ISSN: 1943-2909. DOI: 10.1119/5.0159866.

[99] 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". In: *New Journal of Physics* 14.2 (Feb. 2012), p. 023009. ISSN: 1367-2630. DOI: 10.1088/1367-2630/14/2/023009.

[100] Richard H. Parker, Chenghui Yu, Brian Estey, Weicheng Zhong, Eric Huang, and Holger Müller. "Controlling the multiport nature of Bragg diffraction in atom interferometry". In: *Physical Review A* 94.5 (Nov. 2016). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.94.053618.

[101] P. Arora, S. B. Purnapatra, A. Acharya, R. Kumar, and A. Sen Gupta. "Measurement of Temperature of Atomic Cloud Using Time-of-Flight Technique". In: *MAPAN* 27.1 (Mar. 2012), pp. 31–39. ISSN: 0974-9853. DOI: 10.1007/s12647-012-0003-3.

[102] Mathieu Beau, Timothey Szczepanski, Rafael Martellini, and Lionel Martellini. "Quantum arrival times in free fall". In: *Journal of Physics: Conference Series* 3017.1 (June 2025), p. 012022. ISSN: 1742-6596. DOI: 10.1088/1742-6596/3017/1/012022.

[103] A. Medina-Herrera, C. A. Ortiz, E. de Carlos-López, and J. M. López-Romero. "Accurate temperature measurement of cold atoms in cesium fountain clocks". In: *Optics Express* 30.14 (June 2022), p. 24261. ISSN: 1094-4087. DOI: 10.1364/oe.461360.

[104] A. Béguin, T. Rodzinka, L. Calmels, B. Allard, and A. Gauguet. "Atom Interferometry with Coherent Enhancement of Bragg Pulse Sequences". In: *Physical Review Letters* 131.14 (Oct. 2023). ISSN: 1079-7114. DOI: 10.1103/physrevlett.131.143401.

[105] 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". In: *Physical Review Letters* 114.6 (Feb. 2015). ISSN: 1079-7114. DOI: 10.1103/physrevlett.114.063002.

[106] Sheng-wey Chiow, Tim Kovachy, Hui-Chun Chien, and Mark A. Kasevich. "102$\hbar k$ Large Area Atom Interferometers". In: *Physical Review Letters* 107.13 (Sept. 2011). ISSN: 1079-7114. DOI: 10.1103/physrevlett.107.130403.

[107] Raymond Cumming. "The Serrodyne Frequency Translator". In: *Proceedings of the IRE* 45.2 (1957), pp. 175–186. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1957.278387.

[108] Rachel Houtz, Cheong Chan, and Holger Müller. "Wideband, efficient optical serrodyne frequency shifting with a phase modulator and a nonlinear transmission line". In: *Optics Express* 17.21 (Oct. 2009), p. 19235. ISSN: 1094-4087. DOI: 10.1364/OE.17.019235.

[109] D. M. S. Johnson, J. M. Hogan, S.-w. Chiow, and M. A. Kasevich. "Broadband optical serrodyne frequency shifting". In: *Optics Letters* 35.5 (Feb. 2010), p. 745. ISSN: 1539-4794. DOI: 10.1364/OL.35.000745.

[110] T. Lévèque, A. Gauguet, F. Michaud, F. Pereira Dos Santos, and A. Landragin. "Enhancing the Area of a Raman Atom Interferometer Using a Versatile Double-Diffraction Technique". In: *Physical Review Letters* 103.8 (Aug. 2009). ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.103.080405.

[111] Sabrina Hartmann, Jens Jenewein, Sven Abend, Albert Roura, and Enno Giese. "Atomic Raman scattering: Third-order diffraction in a double geometry". In: *Physical Review A* 102.6 (Dec. 2020). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.102.063326.

[112] Jens Jenewein, Sabrina Hartmann, Albert Roura, and Enno Giese. "Bragg-diffraction-induced imperfections of the signal in retroreflective atom interferometers". In: *Physical Review A* 105.6 (June 2022). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.105.063316.

[113] E. Giese, A. Roura, G. Tackmann, E. M. Rasel, and W. P. Schleich. "Double Bragg diffraction: A tool for atom optics". In: *Physical Review A* 88.5 (Nov. 2013). ISSN: 1094-1622. DOI: 10.1103/PhysRevA.88.053608.

[114] Fiodor Sorrentino, Kai Bongs, Philippe Bouyer, Luigi Cacciapuoti, Marella de Angelis, Hansjoerg Dittus, Wolfgang Ertmer, A. Giorgini, J. Hartwig, Matthias Hauth, Sven Herrmann, Massimo Inguscio, Endre Kajari, Thorben T. Könemann, Claus Lämmerzahl, Arnaud Landragin, Giovanni Modugno, Frank Pereira dos Santos, Achmin Peters, Marco Prevedelli, Ernst M. Rasel, Wolfgang P. Schleich, Malte Schmidt, Alexander Senger, Klaus Sengstock, Guillaume Stern, Guglielmo Maria Tino, and Reinhold Walser. "A Compact Atom Interferometer for Future Space Missions". In:

Microgravity Science and Technology 22.4 (Sept. 2010), pp. 551–561. ISSN: 1875-0494. DOI: 10.1007/s12217-010-9240-7.

[115] Kai Frye, Sven Abend, Wolfgang Bartosch, Ahmad Bawamia, Dennis Becker, Holger Blume, Claus Braxmaier, Sheng-Wey Chiow, Maxim A. Efremov, Wolfgang Ertmer, Peter Fierlinger, Tobias Franz, Naceur Gaaloul, Jens Grosse, Christoph Grzeschik, Ortwin Hellmig, Victoria A. Henderson, Waldemar Herr, Ulf Israelsson, James Kohel, Markus Krutzik, Christian Kürbis, Claus Lämmerzahl, Meike List, Daniel Lüdtke, Nathan Lundblad, J. Pierre Marburger, Matthias Meister, Moritz Mihm, Holger Müller, Hauke Müntinga, Ayush M. Nepal, Tim Oberschulte, Alexandros Papakonstantinou, Jaka Perovšek, Achim Peters, Arnau Prat, Ernst M. Rasel, Albert Roura, Matteo Sbroscia, Wolfgang P. Schleich, Christian Schubert, Stephan T. Seidel, Jan Sommer, Christian Spindeldreier, Dan Stamper-Kurn, Benjamin K. Stuhl, Marvin Warner, Thijs Wendrich, André Wenzlawski, Andreas Wicht, Patrick Windpassinger, Nan Yu, and Lisa Wörner. "The Bose-Einstein Condensate and Cold Atom Laboratory". In: *EPJ Quantum Technology* 8.1 (Jan. 2021). ISSN: 2196-0763. DOI: 10.1140/epjqt/s40507-020-00090-8.

[116] Sheng-wey Chiow, Jason Williams, Nan Yu, and Holger Müller. "Gravity-gradient suppression in spaceborne atomic tests of the equivalence principle". In: *Physical Review A* 95.2 (Feb. 2017). ISSN: 2469-9934. DOI: 10.1103/PhysRevA.95.021603.

[117] Xuejian Wu, Zachary Pagel, Bola S. Malek, Timothy H. Nguyen, Fei Zi, Daniel S. Scheirer, and Holger Müller. "Gravity surveys using a mobile atom interferometer". In: *Science Advances* 5.9 (Sept. 2019). ISSN: 2375-2548. DOI: 10.1126/sciadv.aax0800.

[118] Jongmin Lee, Roger Ding, Justin Christensen, Randy R. Rosenthal, Aaron Ison, Daniel P. Gillund, David Bossert, Kyle H. Fuerschbach, William Kindel, Patrick S. Finnegan, Joel R. Wendt, Michael Gehl, Ashok Kodigala, Hayden McGuinness, Charles A. Walker, Shanalyn A. Kemme, Anthony Lentine, Grant Biedermann, and Peter D. D. Schwindt. "A compact cold-atom interferometer with a high data-rate grating magneto-optical trap and a photonic-integrated-circuit-compatible laser system". In: *Nature Communications* 13.1 (Sept. 2022). ISSN: 2041-1723. DOI: 10.1038/s41467-022-31410-4.

[119] Shau-Yu Lan, Pei-Chen Kuan, Brian Estey, Philipp Haslinger, and Holger Müller. "Influence of the Coriolis Force in Atom Interferometry". In: *Physical Review Letters* 108.9 (Feb. 2012). ISSN: 1079-7114. DOI: 10.1103/physrevlett.108.090402.

[120] Maximilian Beyer, Jack C. Roth, Eustace Edwards, and David DeMille. "Frequency-doubled Nd:YAG MOPA laser system with programmable rectangular pulses up to 200 microseconds". In: *Optics Express* 29.13 (June 2021), p. 20370. ISSN: 1094-4087. DOI: 10.1364/oe.427832.

[121] Chang Xu, Chuan Guo, Hai-Bo Yu, Zhi-Min Wang, Jun-Wei Zuo, Yuan-Qin Xia, Qi Bian, Yong Bo, Hong-Wei Gao, Ya-Ding Guo, Sheng Zhang, Da-Fu Cui, Qin-Jun Peng, and Zu-Yan Xu. "High pulse energy, high beam quality microsecond-pulse Ti:sapphire laser at 819.7 nm". In: *Applied Physics B* 123.3 (Mar. 2017). ISSN: 1432-0649. DOI: 10.1007/s00340-016-6627-1.

[122] G. Schimmel, I. Doyen-Moldovan, S. Janicot, M. Hanna, J. Decker, P. Crump, G. Blume, G. Erbert, P. Georges, and G. Lucas-Leclin. "Rear-side resonator architecture for the passive coherent combining of high-brightness laser diodes". In: *Optics Letters* 41.5 (Feb. 2016), p. 950. ISSN: 1539-4794. DOI: 10.1364/ol.41.000950.

[123] P. Albrodt, M. Niemeyer, P. Crump, J. Hamperl, F. Moron, P. Georges, and G. Lucas-Leclin. "Coherent beam combining of high power quasi continuous wave tapered amplifiers". In: *Optics Express* 27.20 (Sept. 2019), p. 27891. ISSN: 1094-4087. DOI: 10.1364/oe.27.027891.

[124] P. Albrodt, M. T. Jamal, A. K. Hansen, O. B. Jensen, G. Blume, K. Paschke, P. Crump, P. Georges, and G. Lucas-Leclin. "Coherent combining of high brightness tapered amplifiers for efficient non-linear conversion". In: *Optics Express* 27.2 (Jan. 2019), p. 928. ISSN: 1094-4087. DOI: 10.1364/oe.27.000928.

[125] J. Hamperl, P. Albrodt, P. Georges, and G. Lucas-Leclin. "Compact module for high power coherent beam combining of tapered amplifiers". In: *2019 IEEE High Power Diode Lasers and Systems Conference (HPD)*. IEEE, Oct. 2019, pp. 33–34. DOI: 10.1109/hpd48113.2019.8938676.

[126] Minjeong Kim, Remy Notermans, Chris Overstreet, Joseph Curti, Peter Asenbaum, and Mark A. Kasevich. "40 W, 780 nm laser system with compensated dual beam splitters for atom interferometry". In: *Optics Letters* 45.23 (Nov. 2020), p. 6555. ISSN: 1539-4794. DOI: 10.1364/ol.404430.

[127] Robert G. Littlejohn. "The Mathematical Formalism of Quantum Mechanics". Physics 221A, Academic Year 2020–21, Notes 1. 2020. URL: http://bohr.physics.berkeley.edu/classes/221/2021/221.html.

[128] Robert G. Littlejohn. "Time Evolution in Quantum Mechanics". Physics 221A, Academic Year 2020–21, Notes 5. 2020. URL: http://bohr.physics.berkeley.edu/classes/221/2021/221.html.

[129] R. P. Boas and K. Chandrasekharan. "Derivatives of infinite order". In: *Bulletin of the American Mathematical Society* 54.6 (1948), pp. 523–526. ISSN: 0273-0979. DOI: 10.1090/s0002-9904-1948-09031-0.