

©Copyright 2018

Benjamin T. Plotkin-Swing

Large Momentum Separation Matter Wave Interferometry

Benjamin T. Plotkin-Swing

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

Doctor of Philosophy

University of Washington

2018

Reading Committee:

Subhadeep Gupta, Chair

Boris B. Blinov

Stephen R. Sharpe

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Large Momentum Separation Matter Wave Interferometry

Benjamin T. Plotkin-Swing

Chair of the Supervisory Committee:
Subhadeep Gupta
Dept. of Physics

This work establishes a new benchmark for momentum separation in a matter wave interferometer with stable, visible fringes. With a path separation of 112 photon recoil momenta ($112\hbar k$), our signal visibility of 30% and phase stability of 0.6rad exceed the performance of earlier free space interferometers. Contributing to this success are the symmetric form of the 3 path contrast interferometer geometry, which rejects phase noise due to vibrations and other systematic errors, the narrow momentum width of the Bose-Einstein condensate (BEC) source, and atom-optics parameters chosen to suppress unwanted diffraction phases. These results can be applied toward a competitive measurement of the fine-structure constant and a test of QED. The described experiments were performed in a new ytterbium BEC apparatus whose design, construction, and operation is documented here.

TABLE OF CONTENTS

	Page
List of Figures	v
Glossary	xi
Chapter 1: Introduction	1
1.1 Interferometry in general	2
1.2 Matter Wave Interferometry	4
1.2.1 Applications and limitations	5
1.2.2 Large momentum transfer	6
1.3 Contrast interferometry	8
1.3.1 Acceleration	10
1.3.2 CI aesthetics	11
1.3.3 CI history	12
1.3.4 CI phase	12
1.4 Recoil frequency and the fine structure constant	13
1.4.1 QED tests	14
1.5 Ytterbium BEC interferometry	15
1.5.1 Ytterbium	15
1.5.2 BEC interferometry	17
Chapter 2: Apparatus	19
2.1 Vacuum chamber	19
2.1.1 The atomic physics machine of Theseus	19
2.1.2 Design	19
2.1.3 Construction	29

2.1.4	Baking	32
2.2	Magnetic Coils	35
2.2.1	Zeeman Slower Coils	35
2.2.2	Magneto Optic Trap Coils	35
2.2.3	Compensation Coils	35
2.3	Electronics	36
2.3.1	Cicero	36
2.3.2	Analog and Digital channels	36
2.3.3	MOT fluorescence Monitor	36
2.3.4	Absorption Imaging CCD Camera	38
2.4	Optics	38
2.4.1	399nm Source	40
2.4.2	556nm Source	41
2.4.3	Injection Locks	41
2.4.4	532nm Source	41
2.4.5	Frequency Stabilization Optics	42
2.4.6	Zeeman Slower Optics	50
2.4.7	Magneto Optic Trap Optics	51
2.4.8	Diffraction Beam Optics	53
2.4.9	Optical Dipole Trap Optics	54
2.4.10	Imaging Optics	56
2.4.11	Readout Light Optics	56
Chapter 3:	Running the machine	61
3.1	Trapping and Cooling	61
3.1.1	Zeeman Slowing	61
3.1.2	Slower viewport clouding	68
3.1.3	Transverse Cooling	70
3.1.4	Crossed beam Slowing	72
3.1.5	Magneto Optic Trapping	73
3.1.6	Optical Dipole Trapping	76

3.2	Bose-Einstein Condensation	79
3.2.1	Achieving a BEC	79
3.2.2	Painting	81
3.2.3	Trap frequencies	82
3.3	Diffraction: building the contrast interferometer	82
3.3.1	Interferometer Geometry in the Lab	82
3.3.2	Diffraction beam heuristics	84
3.3.3	CI Schemes	86
3.3.4	Bloch Oscillations	89
3.3.5	$n = 4$ Contrast Interferometer	89
3.3.6	Adding Acceleration: a lesson on the KISS principle	90
3.3.7	Acceleration with Sequential Bragg Pulses	93
Chapter 4:	Results	97
4.1	Prelude	97
4.2	Slope Data	99
4.3	Visibility Data	101
4.4	Diffraction Phase Data	104
4.5	Interactions Data	104
4.6	Readout Light Data	104
Chapter 5:	Discussion and Outlook	108
5.1	Analysis Methods	108
5.1.1	Extracting the phase from the CI signal	109
5.1.2	Phase extraction with non-linear curve fitting	110
5.1.3	Fourier component phase extraction	112
5.1.4	Signal quality vs n	115
5.2	Current Capabilities	119
5.3	Looking Forward	123
5.3.1	M : Faster BEC production	123
5.3.2	$\delta\Phi$	124
5.3.3	n and T	125

5.3.4	Prospects for updated horizontal diffraction beam CI	125
5.3.5	Two horizontal diffraction beams	127
5.3.6	Vertical diffraction beams	128
5.3.7	Projections for Vertical CI	130
Bibliography		131

LIST OF FIGURES

Figure Number	Page
1.1 Water wave interference in Puget sound.	3
1.2 An illustration of the difference between fringe contrast and visibility	7
1.3 Two path matter wave IFM (left), and 3 path contrast IFM (right)	8
1.4 Space-time schematic of the time dependent atomic density grating that appears at the end of the CI sequence where the paths overlap. The periodicity of the density grating is such that light at the frequency of the diffraction laser reflects off the grating, meeting the Bragg condition with an angle of 90 degrees. As will be described later, this technique can also be used with another available light frequency with a different Bragg angle.	10
1.5 Space-time schematics of the CI with acceleration pulses to increase the momentum separation between paths 1 and 3. We label the momentum separation with n , the number of single photon recoils $\hbar k$ separating the outer paths. Shown are $n = 16$ and $n = 100$	11
1.6 Spins, and abundances of the various stable isotopes of Yb, with the fermionic isotopes shaded. Yb energy level diagram, with arrows representing the two optical transitions used in this work.	16
2.1 Vacuum chamber schematic	20
2.2 Technical drawings for the atomic beam shutter with mirror. The part shown here is a bracket that connects to the shaft of the rotary feedthrough, and holds a 1inch square mirror. It is made from 6061 aluminum, and made by me on the mill in the student shop. It uses set screws to hold the mirror in place. Care was taken to vent the ends of the set screws where they press into the mirror, to avoid trapping a pocket of air creating a virtual leak. When the mirror surface is parallel to the atomic beam it is offset to allow the beam to pass. When It is rotated 45 degrees, the mirror is placed to reflect a beam coming from the main chamber towards a viewport on the differential pumping 6-way cross.	25

2.3	Technical drawings for the custom cross vacuum fitting manufactured by Sharon Vacuum.	26
2.4	Theoretical curves for the viewport AR coating, from Larsen Electronic Glass. The lower curve is for the AR coating applied to glass, for which it is designed. The upper curve is for the same coating applied to a sapphire viewport. The performance is lower because the coating is not designed for sapphire, but we judged it to be good enough to use the coating on our one sapphire viewport.	28
2.5	Main chamber before high temperature bake	33
2.6	The analog and digital inputs and outputs connected to the National Instruments cards controlled by Cicero.	37
2.7	Yb energy level diagram	39
2.8	Optics table schematic overview	40
2.9	Doppler shift of an atom with velocity component v along the propagation axis of a laser with lab frame frequency ω	43
2.10	Optics schematic for the 399nm Saturated Absorption Spectroscopy. The cube on the left hand side of the figure is not present in the real setup, but is included here to help show the relationship between the pump beam and the probe beam.	44
2.11	Schematic for the beat note lock.	47
2.12	Blue and Green readout geometry	57
2.13	Top: photo of the lens tubes containing the readout light collection optics. Bottom: Schematic of the optical elements. The iris size is adjustable, as is its location in the x-y plane (with z being the propagation direction of the readout light). The translation stage iris mount is itself mounted inside a 2" lens tube. This configuration of iris mount and lens tube required some creative modifications of the lens tube, see 2.14	58
2.14	The slightly anthropomorphic iris translation stage mounted in a 2" lens tube, via some band saw modifications to the lens tube. This contraption is part of the readout light collection optics.	59
3.1	The ideal magnetic field profile, along with the predicted design curve	63
3.2	The measured field of the increasing field and reverse slower sections at 3A, small dots, compared to the predicted field, solid line.	64
3.3	The measured field for the three separate slower sections at 3A.	65

3.4	Doppler sensitive fluorescence spectrum of the atomic beam as it passes through the MOT region with no slower light presence. The sharp peak on the left is from a perpendicular, hence Doppler insensitive probe beam. The broad peak is the Doppler sensitive signal, centered at 300m/s.	66
3.5	Spectroscopy signal with slower light on, but no magnetic fields. Atoms are slowed until they go out of resonance with the slower light.	67
3.6	Velocity profile with the slower light and increasing field coils, but no offsets.	67
3.7	Velocity profile for slower light and offset coils, no increasing field.	68
3.8	Velocity profile for the slower light with offset coils and increasing fields. The larger peak near $v = 0$ compared to Fig. 3.6 demonstrates the utility of the offset coils.	69
3.9	Ytterbium coating the inner surface of the slower viewport. The upper half of the Ytterbium smudge has been cleaned by focusing onto the viewport a beam with 10W of 532nm light with a $50\mu\text{m}$ waist. The beam was scanned by hand using a mirror knob, removing the Ytterbium coating in a manner very reminiscent of an Etch A Sketch. The Etch A Sketch pattern is especially clear in the lower half of the smudge, where I made some lines and corners while learning how best to proceed.	71
3.10	Schematic of the crossed beams slower. The crossed beams travel next to the MOT beams, and cross just upstream in the atom beam of the MOT region. Unlike the MOT beams, the crossed slower beams are not retro-reflected. . .	74
3.11	The MOT number vs time with and without the crossed beam slower. The overall slower performance is fully tuned up for each configuration, i.e. the Zeeman slower exit velocity is appropriately set in each case. We observe a factor of 2 higher MOT number and load rate while using the crossed beams.	75
3.12	fluorescence lines split in magnetic field. The two split peaks correspond to different circular polarizations of light.	77
3.13	Absorption images at various TOF and temperatures during the end of evaporative cooling, as the condensate forms. A clear signature of condensation is seen in the reversal of anisotropy of the atomic cloud. The trapped condensate is taller than it is wide. In the right hand column, the thermal atoms expand isotropically. The left column, showing a pure condensate, ends up wider than it is tall. The energy due to atomic interactions in the trap is converted to momentum preferentially into the more tightly trapped directions.	80
3.14	Absorption images at 20ms TOF with and without the decompression step .	83

3.15	Laser schematic for retro-reflected horizontal CI. Referred to in the text as DiffPlan0. In this and the following 3 figures, the double ended arrows represent linear polarizations, while the one ended arrows represent direction of propagation.	86
3.16	Laser schematic for the diffraction beam scheme with a single polarization, but with two fibers instead of a retro mirror, for more flexible frequency control. Referred to in the text as DiffPlan1.	87
3.17	Laser schematic for the diffraction beam scheme with one polarization and two fibers, but with two frequency components in one of the beams. Referred to in the text as DiffPlan2.	87
3.18	Laser schematic for the diffraction beam scheme with two polarization components, and a $\frac{\lambda}{4}$ waveplate to rotate the polarizations to get standing waves with the desired frequency components. Referred to in the text as DiffPlan3.	88
3.19	Early Bloch Oscillation demonstration from 2016	90
3.20	Comparison of a Gaussian and Cosine with similar widths, showing that the Cosine goes to zero more quickly than the Gaussian.	95
3.21	3rd order Bragg pulse efficiency vs Bragg pulse width. In this case, Bragg pulse width is defined as the waist of a Gaussian fit to the recorded diffraction beam intensity during the Bragg pulse.	96
4.1	(a) CI phase Φ vs free evolution time $2T$ with linear fit, for various n . (b) Fit slopes $d\Phi/dT$ vs n , demonstrating the expected quadratic relationship $d\Phi/dT = \frac{1}{2}n^2\omega_{\text{rec}}$ (black curve). (c) Typical fit residuals ($n = 16$).	100
4.2	Averaged signals for $n=4$ through $n=52$. The right column shows the averaged raw data, and the left column shows the data with a diffraction phase correction applied before averaging.	102
4.3	Averaged signals for $n=64$ through $n=122$. The right column shows the averaged raw data, and the left column shows the data with a diffraction phase correction applied before averaging.	103

4.4	Diffraction phase shift vs peak lattice depth (with pulse width held fixed) for the second-order Bragg mirror pulse (black circles) and the third-order acceleration pulse (blue diamonds) from $2\hbar k$ to $8\hbar k$. The dashed vertical lines indicate the peak lattice depths at which the π -pulse condition is met in each case. Overall phase offsets have been removed to zero the diffraction phase at the π -points. Black and blue solid lines are the predictions from the corresponding numerical model. The dotted orange line is the prediction of the model for the diffraction phase for the third-order Bragg pulse for large n .	105
4.5	Top: CI signal vs trap frequency. Higher signals correspond to lower trap frequencies. The vertical axis is voltage, but offsets are added to the signals to keep them from overlapping. Middle: Gaussian width of the fit function envelope vs trap frequency. Bottom: Signal phase vs trap frequency.	106
4.6	Left: the CI signal for a variety of readout light powers. Higher power signals are lower on the plot. The vertical axis is in volts, but offsets have been added so that the signals do not overlap. Upper right: The CI phase vs readout light power. Middle right: Phase standard deviation vs readout light power. Lower left: Signal area vs readout light power.	107
5.1	Example of a single shot CI signal with $n=4$	110
5.2	$n=4$ CI signal with 100 shots averaged together.	111
5.3	The left panel shows 3 signals with different phases, and the right panel shows the average of the three signals. The phase variation leads to a decrease in the amplitude and visibility of the averaged signal, but the fit uncertainty for the averaged signal does not capture the underlying phase variations.	112
5.4	For $n = 100$, the left panel shows a single shot, while the right shows an average of 80 shots.	114
5.5	Histogram of the single shot Fourier phases of the 80 shots shown in the right panel of Fig. 5.4. The blue vertical line is the phase returned by the non-linear fitter applied to the 80 shot average, while the orange line is the mean value of the single shot Fourier phases. The orange shaded region is the standard error of the mean.	115

5.6	The upper panel shows the normalized readout signal amplitude and visibility vs n . The amplitude is based on the Fourier component at $8\omega_{\text{rec}}$, and visibility is extracted from sinusoidal fits. The amplitude decreases for higher n because fewer atoms contribute to the signal due to the finite efficiency of the acceleration pulses. The amplitude model curve is an exponential fit with acceleration efficiency as a fit parameter. Visibility is calculated with (squares) and without (triangles) the diffraction phase correction. The visibility model is based on the 7% rate for spontaneous scattering of readout light photons and the number of atoms illuminated by the readout beam for different n . The lower panel shows the CI phase standard deviation vs n , with (diamonds) and without (circles) the diffraction phase correction.	117
5.7	In order to simulate the reduction in visibility expected due to signals with a spread of phases being averaged together, we plot the visibility of the sum of many sine waves, each with a different phase, versus the standard deviation of the summed phases. We observe that the reduction in visibility due to this effect at the highest phase standard deviation observed in our work is 10%.	119
5.8	Visualization of a simulation of all momentum state trajectories in the CI sequence. We usually only consider the intended paths, neglecting the paths of atoms which are left behind due to inefficiencies of each diffraction pulse. The amplitude along each path is represented by the thickness of the lines. However, a minimum thickness is imposed, so that all paths are visible. We use the simulation that underlies this image to calculate the total number of atoms that arrive in the area that the readout light hits, in order to predict the level of spontaneously scattered light that forms a background to the CI signal.	120
5.9	Left: typical output of the diffraction beam intensity monitor photodiode. This example is for $n = 76$. The diffraction beams enter the main chamber at an angle slightly offset from 90 degrees, and the reflection from the viewport is directed to the diffraction beam monitor photodiodes. This plot shows the signals from the two channels multiplied together. The heights of the acceleration pulses vary over the course of the sequence in order to maintain the π -pulse condition as the atoms pass through different parts of the diffraction beams. Right: the correlation of the total diffraction pulse area vs the corresponding measured CI phase. We implement the diffraction phase correction by fitting a linear model to this correlation, and subtracting the fitted slope from the phase values. As shown in Fig. 5.6, the standard deviations of the corrected phases are substantially lower for n up to 88.	121

GLOSSARY

AOM: Acousto optic modulator

B055: The room in which the CI machine resides

B063: The other Gupta group lab

BEC: Bose-Einstein condensate.

BO: Block Oscillations

CI: Contrast interferometer

CMOT: Compressed magneto optic trap

DDS: Direct digital synthesizer

MOT: Magneto optic trap

ODT: Optical dipole trap

ACKNOWLEDGMENTS

There are many people I want to thank for inspiration and support on my path towards a Ph.D. Since the ultimate motivation for this long and sometimes arduous project is, at its most basic, curiosity and fascination with how the world works, I want to start with acknowledging the people who help instill these traits in me.

First, my family. On top of all the love and support, my mother and father raised me so that asking why, figuring it out, and critically for an experimentalist, building and playing with everything I could get my hands one was the most natural thing in the world.

In physics in particular, I had no shortage of role models. I am a third generation physicist. My grandfather worked with early masers, and helped design and build the lunar range-finding reflectors that were placed on the moon by Apollo astronauts. My father worked in particle physics before moving to the, at the time, new field of computer programming. In such a family, the heroic stories of physics are told around the dinner table. In some ways, I feel I have physics in my blood.

Over my years of education, there are many teachers who nurtured my curiosity and introduced me to valuable facts:

Mikel Matisoo and Danny Greenburg at the Sudbury Valley School.

James Tanton at The Math Circle, who helped me get my first publication, and another young math circle teacher whose name I can't remember, but who showed me how to derive the Lorentz transformation in a way I recall almost 20 years later.

Carl Grossman, who guided my first optics table placements, and John Boccio, who introduced me to my first quantum operator.

Steve Sharpe, Steve Ellis, and Andreas Karch, the professors for my first year graduate physics classes. The rigor and pedagogy of these classes was invigorating. Completing these, and the qualifying exam the next summer, was the milestone where I first felt I could call myself a physicist.

My years of formal education would be near impossible, and far less enjoyable, without a community of fellow students. I learned so much with, and from, my classmates:

Rob, Al, Margaret, Emma, Erin, Eric, and Chris, the 2009 Swarthmore physics graduating class.

The entire 2010 UW graduate physics cohort.

For both of these groups, there is a unique camaraderie formed doing problem sets in the physics lounge at 2AM. Though it sometimes felt like type 2 fun, I treasure the knowledge and friendships formed in these times of my life.

Another type of community is found in an experimental physics lab. In addition to being an invaluable source of experimental knowledge and lab lore, the other members of the Gupta group have a sense of quirky fun and intellectual curiosity that helped attract me to the group in the first place. Though I didn't work with them directly, I have been very glad to know the members of the inter-atomic mixtures group: Alex, Anders, Will, Ricky, Alaina, Jun, and Khang.

The other members of the contrast interferometry team have been my constant companions through years of designing, building, operating, and analyzing. Though table-top atomic physics labs do not require thousand-person collaborations, neither are they one-person projects. I could not have done it without them.

Alan Jamison was the senior student on the experiment for my first years. He has been a mentor and an inspiration, even if we could never completely agree on data-taking-music choices.

Ryan Weh, Brendan Saxberg, and Eric Cooper have all made much-appreciated contributions to the experiment.

Katie and Dan are the next generation of interferometeers. They have contributed instrumentally to the results in this thesis. To them I say, use well this machine that we have built!

Finally, I thank my advisor, Deep Gupta. Deep is a great physicist and lab leader. I am grateful for all the ways Deep supported and assisted with this project, but over the years I've particularly appreciated his enthusiasm for the work. No conversation with Deep is safe from slipping into a discussion of new ideas and next steps for the experiment. Deep is a master in his field, and any sort of problem, from lab equipment mysteries to calculation difficulties, can benefit from his wisdom.

DEDICATION

to Hannah

Chapter 1

INTRODUCTION

That quantum scale objects have some counter-intuitive behaviors is well known, if not well understood, by the public at large. Thought experiments like Schrodinger's cat and Heisenberg's microscope are commonplace, and most have heard phrases such as "wave particle duality". Like many students, learning the truth about concepts underlying the often confused popular understanding of quantum mechanics (QM) was one of my motivations to study physics at a high level.

Interference experiments cut to the heart of the strangeness in QM. Essentially all of the "weird" things about quantum behavior can be traced back to the fact that QM objects have amplitudes that can be negative and imaginary [1]. Amplitudes that cancel or reinforce each other are the source of quantum interference.

In more than a century of controversy and debate among some of the worlds brightest minds, there has been remarkably little progress in making humanity comfortable with the fact that the universe seems to run on quantum mechanics while day to day life is classical. But comfortable or not, the framework of QM is the basis for our stupendously successful standard model of physics.

In deciding to pursue experimental atomic physics, I had a vague hope that by spending some years working with quantum objects in the lab, becoming a true "quantum mechanic", I would gain some practical insight into the nature of QM. Predictably, this has not been the case even after personally creating and interfering tens of thousands of superposition

states. What I have gained is an appreciation for the beauty and power of QM as a tool to understand and manipulate the world. It turns out that truth does not need to be explicable in order to be useful.

1.1 Interferometry in general

Though the fact that quantum objects have amplitudes that lead to interference is a surprising result from QM, the basic operation of an interferometer can be understood classically. The ingredients for interference are:

1. An oscillator, defined as an object or medium with some periodic property.
2. Ability to divide the oscillator along different paths.
3. Ability to recombine the paths so that the oscillator phases add and interfere.
4. Practically, one needs some way of observing the interference.

Some purely classical media have all of these properties, for instance water as shown in Fig 1.1.

In general, one uses an interferometer by varying a parameter that controls the effective path length difference, and hence the relative phase accrued as the oscillators travel along their separate paths. An interferometer with a well defined effective path length difference can be used to measure an oscillator with a less well known frequency. Conversely, an oscillator with a well known frequency can be used in an interferometer to characterize the relative path length difference. In both of these complementary use cases, an interferometer is an excellent tool for precision measurement.

In the case of measuring a distance with a well known oscillator frequency, the interferometer sensitivity scales with the ratio of the oscillator wavelength to the total path length.

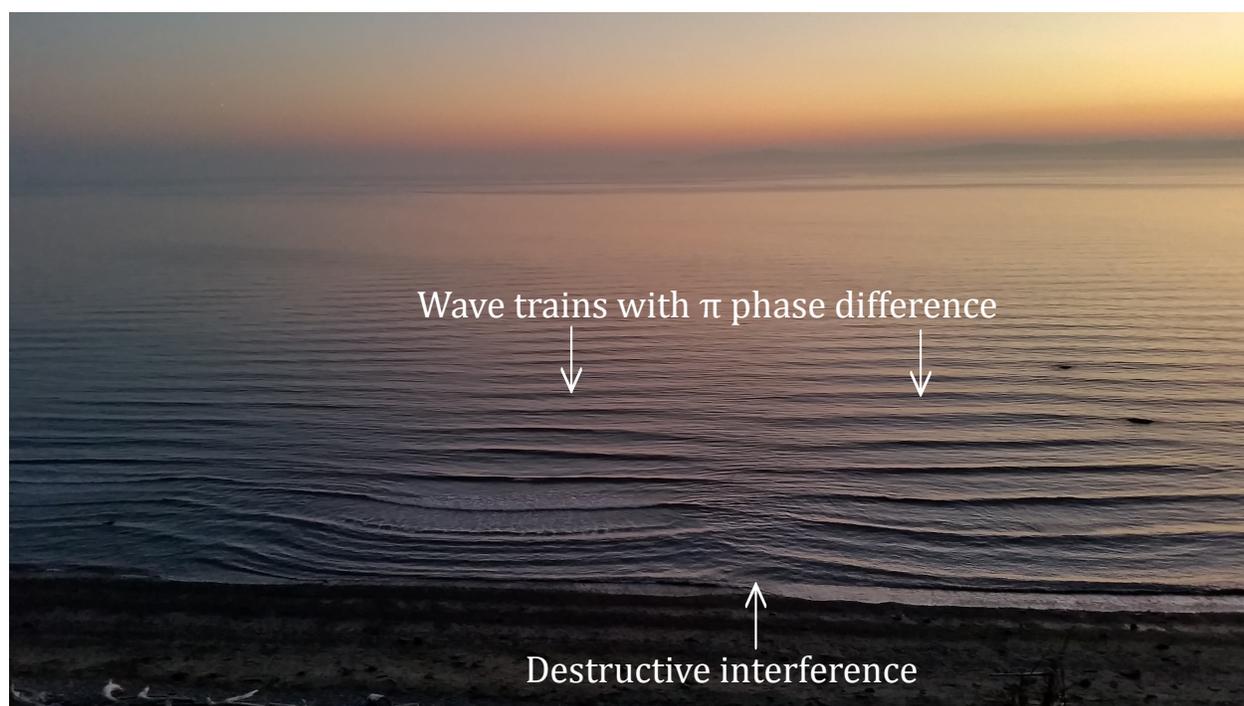


Figure 1.1: Water wave interference in Puget sound.

By making the interfering paths very long, the interferometer is sensitive to small relative path length changes. The extreme precision of these techniques was recently demonstrated with the optical interferometric detection of gravitational waves from colliding black holes [10] and neutron stars [11]. In these measurements the relative path length differences that were detected are on the scale of a proton diameter.

For the purpose of measuring an oscillator frequency given a known path length difference, an interferometer is sensitive, modulo 2π , to the total phase difference between paths at the point of interference. The sensitivity of the interferometric measurement goes as $\frac{\delta\phi}{\phi_{\text{total}}}$, where $\delta\phi$ is the uncertainty in the phase measurement, which often does not depend on the total phase difference between paths ϕ_{total} . Therefore, the measurement can be made very sensitive by increasing ϕ_{total} .

The experiment described here is of the type where a less well known oscillator frequency ω_{rec} is measured using better known path length differences.

1.2 Matter Wave Interferometry

The most commonly used type of interferometer is the optical interferometer. Optical waves have well known frequency properties, and are relatively easy to create, manipulate, and detect. The invention of the laser made optical interferometry even more attractive by providing high intensity and highly coherent optical sources.

With the discovery of quantum mechanics, it became clear that all quantum particles have wavelike properties and could in principle serve as the oscillator in an interferometer. The earliest atom interferometers involved manipulating internal atomic states, for example in Ramsey spectroscopy [31] although this result wasn't construed as interferometry at the time. For the purposes of this work, I will focus on the development of external state matter wave interferometry.

Matter wave interferometry uses the de Broglie wave associated with the center of mass

motion of a particle. The earliest examples of matter wave interferometry used particles like the electron [24] or the neutron [32], but the field started to come into its own with the advent of matter wave atom interferometry [8]. The advantages of atom matter waves over more fundamental particles include: lower cost to produce, stronger couplings to light, and larger range of atomic properties that can be selected and tuned for different applications. Matter wave atom interferometry required the development of a toolbox of atom optics techniques to initialize and control atomic momentum states. Another advantage of matter wave interferometry is that the well known properties of light and light-atom interactions mean that light based atom optics can be very well characterized.

1.2.1 Applications and limitations

Matter wave atom interferometry has been used with great success in a wide range of applications[8]. The typical thermal de Broglie wavelength is much shorter than optical wavelengths, so matter waves can be used to make very sensitive inertial sensors [30, 26, 9]. Matter waves can be used for fundamental physics such as tests of the equivalence principle [13, 34] and of quantum electrodynamics [4].

Matter wave interference experiments that aim to probe fundamental physics with high precision are typically performed with atoms in free fall in a UHV vacuum chamber. UHV pressures are needed to shield the atoms from interactions with background gases, while freely falling atoms are used because any confining potential could be a source of systematic errors on the intended measurement. The use of freely falling atoms in UHV chambers puts some hard limits on the total time available for phase evolution in an interferometer. At some point, the atoms will hit the bottom of the vacuum chamber.

These limits on total interferometer time can be extended a little bit by building a bigger vacuum chamber like the 10m atomic fountain at Stanford [36]. Other groups have gone farther in this direction with experiments in a vacuum chamber freely falling in a multi-storey

drop tower, or flying in a ballistic sounding rocket [28]. The ultimate goal in this direction would be a matter wave interferometer in orbit. A planned Cold Atoms Laboratory user facility on the ISS may provide such an opportunity; in addition, there are various satellite based proposals [28].

That all said, for reasons of cost and practicality most matter wave experiments remain earthbound and in relatively small UHV chambers.

1.2.2 Large momentum transfer

The separate paths in a matter wave interferometer are momentum states, and the rate of differential phase accumulation along the paths depends on how widely separated the paths are in momentum space. This opens up another way to overcome the limitations due to low evolution time: by increasing the magnitude of the momentum space beamsplitters that initially divide the paths, more total phase can be accrued in a fixed amount of time. Large momentum beam splitters could be fruitfully applied to most matter wave experiments, so an outstanding goal in the field is to develop practical techniques for high momentum transfer atom optics. Momentum transfer in these contexts is typically measured in photon momenta $\hbar k$, where k is the wavenumber of the atom-optic diffraction laser. In this work, I use n to indicate the degree of momentum splitting in a matter wave interferometer in units of $\hbar k$.

There have been several previous efforts in this area of large momentum splitting matter waver interferometry. Notable high water marks include n up to 30 in the Muller and Kasevich groups [27, 3].

Higher momentum splittings have been demonstrated, but only in interferometers that have shot-to-shot variations in phase that are greater than 2π , which means that the fringe pattern is not observable, and that the interferometer cannot be used for a direct phase measurement. $n = 102$ [7] has been demonstrated in this manner, where interferometric contrast is inferred based on the population of measured phases being consistent with being

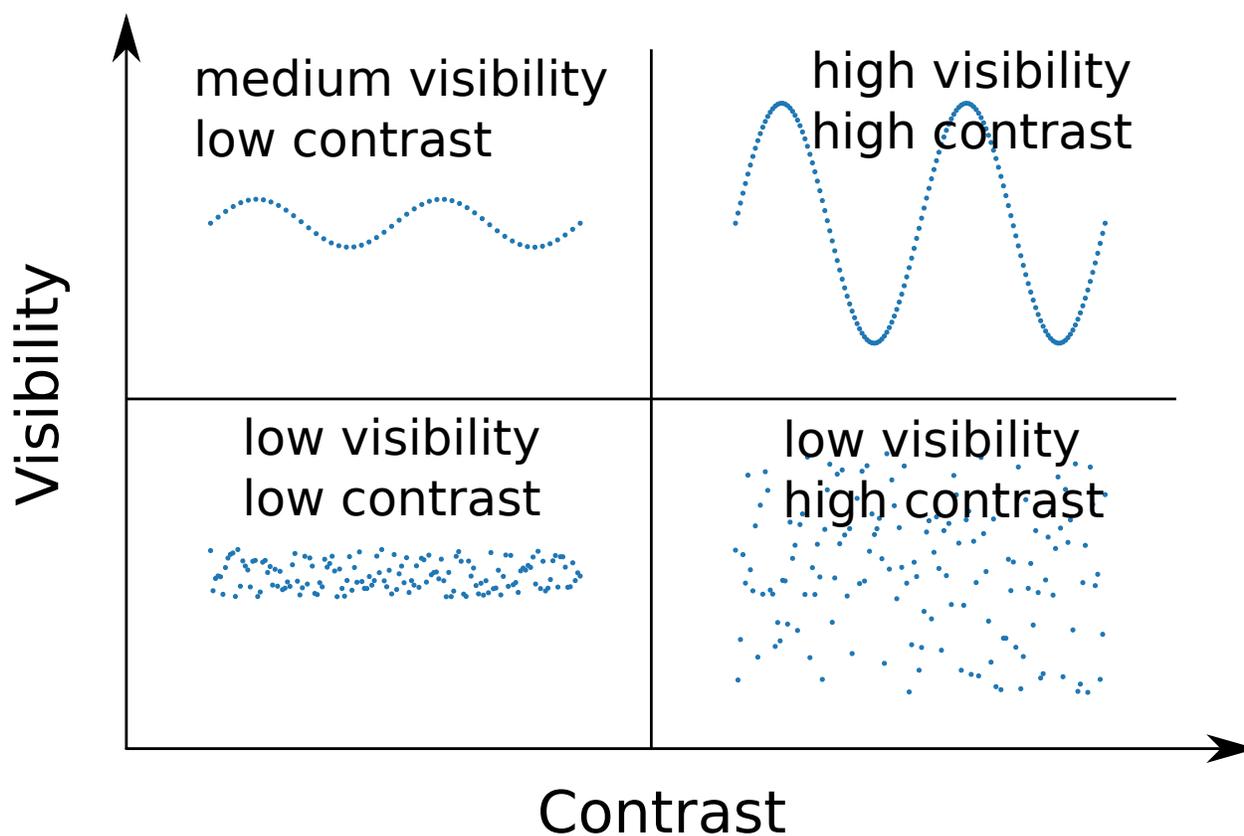


Figure 1.2: An illustration of the difference between fringe contrast and visibility

randomly sampled from a sinusoidal distribution. Such an interferometer can be characterized as having contrast, but no visibility. The Rasel group in Hannover has presented on a conference poster, but not yet published, a demonstration of $n = 400$ with contrast but no visibility. Fig. 1.2 is an illustration of the difference between fringe contrast and visibility.

Visibility of 7% for $n = 80$ has been shown for atoms confined in an optical waveguide potential [25]. This is an impressive result, but for the purposes of precision measurement the confining potential is a difficult to account for systematic error.

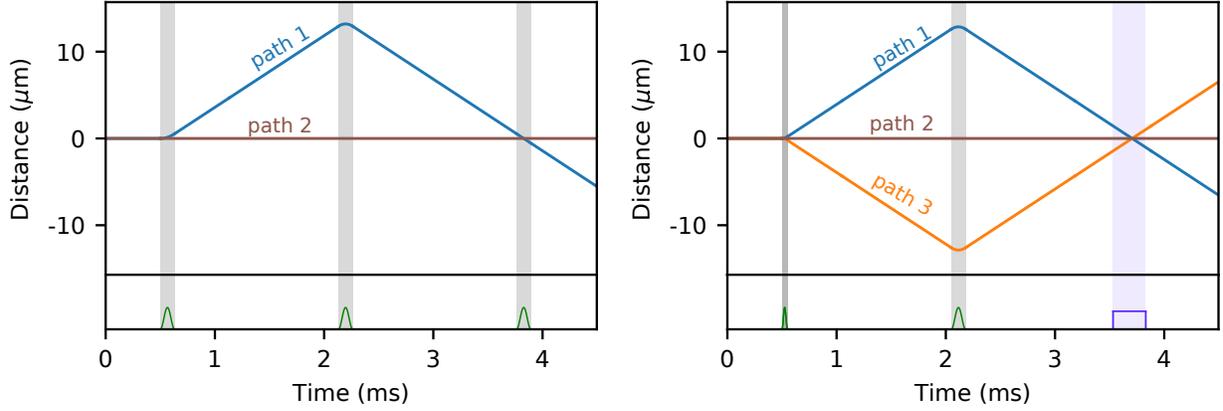


Figure 1.3: Two path matter wave IFM (left), and 3 path contrast IFM (right)

1.3 Contrast interferometry

The overarching goal of my graduate career has been to improve the state of the art in large momentum splitting matter wave interferometry. The tool that gave me the confidence to tackle this goal is the contrast interferometer (CI) [14]. To explain the basics of the CI, I'll first reference a more typical interferometer design.

The left side of Fig. 1.3 shows a space-time schematics of a 2 path atom interferometer:

This two path IFM consists of a $\frac{\pi}{2}$ pulse that puts the atoms into a superposition of the momentum states $|2\hbar k\rangle$ (path 1) and $|0\hbar k\rangle$ (path 2), a π mirror pulse that reverses the momentum of path 1, and a final $\frac{\pi}{2}$ that interferes path 1 and 2 when they overlap. The phase of this interferometer is read out by looking at the relative amplitudes after the final pulse in the two interfering states, or output ports. Typically, in such an interferometer, a single shot does not suffice to determine the phase. Instead, multiple shots are taken while varying one of the interferometer parameters such as the free evolution time T , or the phase of the final pulse of the interferometer (the fact that this type of interferometer is sensitive to the phase of the final pulse is important, since this couples mirror vibrations to the

interferometer phase). As one of these parameters is scanned, the relative amplitudes in the output states change sinusoidally, and the phase of this oscillation is the interferometer phase. This type of interferometer is directly analogous to an optical interferometer, with the atomic wave-function corresponding to EM waves, and the laser diffraction pulses corresponding to mirrors and beam splitters. In both cases, the phase is read out by looking at the relative amplitudes in the output ports while varying an interferometer parameter.

The CI has a different geometry, and takes a different approach to reading out the phase¹. As shown on the right side of Fig. 1.3, the CI has three symmetrical paths. The interferometer consists of a splitting pulse that puts the atoms into an equal superposition of the states $|2\hbar k\rangle$ (path 1), $|0\hbar k\rangle$ (path 2), and $|-2\hbar k\rangle$ (path 3). A mirror pulse after a time T reverses the momenta of paths 1 and 3. After another interval T , the paths overlap.

When the paths overlap, one could apply another splitting pulse to interfere the three paths, and observe the interferometer phase by the relative amplitudes in the output states, as in the 2 path interferometer (such an interferometer was demonstrated in [20] to show the vibration insensitivity of the CI). Instead, we directly detect the phase by reflecting a readout laser beam off the time dependent atomic density grating that appears at the point of overlap, as shown in 1.4. We collect the reflected light on a PMT; the phase of the CI is the phase of this oscillating readout signal.

There are two main advantages to this readout approach:

- The CI phase is not sensitive to the phase of the final pulse, which is in turn sensitive to mirror vibrations. The CI phase is therefore insensitive to optical table vibrations.
- The phase can be obtained in a single shot, as opposed to the phase emerging from combining many output port amplitude measurements.

¹For a two path interferometer “the phase” can be assumed to refer to the phase difference between the two paths. For a three path interferometer like the CI, “the phase” is potentially ambiguous. In this case, I refer to the phase difference between the two outer paths and the inner path, as shown in 1.3.

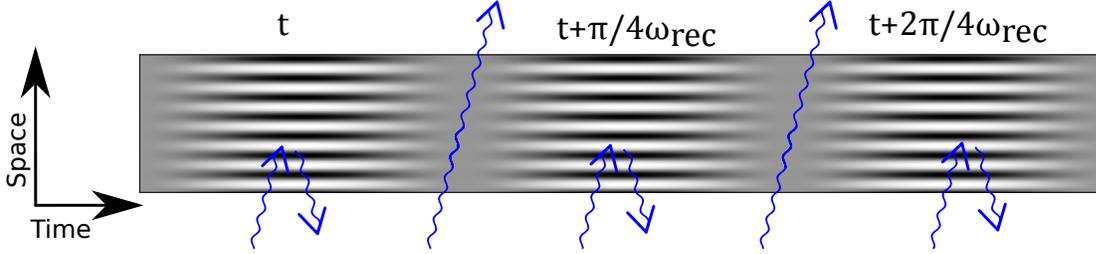


Figure 1.4: Space-time schematic of the time dependent atomic density grating that appears at the end of the CI sequence where the paths overlap. The periodicity of the density grating is such that light at the frequency of the diffraction laser reflects off the grating, meeting the Bragg condition with an angle of 90 degrees. As will be described later, this technique can also be used with another available light frequency with a different Bragg angle.

In addition, the symmetric 3 path geometry of the CI results in the cancellation of several important systematics. These cancellations along with the vibration insensitive readout make the CI a promising tool for precision measurement, but a final ingredient is needed for the CI to reach its full potential.

1.3.1 Acceleration

The relevant differential phase accumulation between the paths of the CI is due to the difference in kinetic energy between the moving and non-moving paths. Therefore, increasing the momentum separation n between the paths increases the rate of phase accumulation, and hence the sensitivity of the interferometer (for a given free evolution time T). Since the kinetic energy goes as the square of the momentum, the sensitivity of the CI to this phase is proportional to n^2 .

This increased momentum separation can be implemented with a series of acceleration pulses that act on the moving paths after the initial splitting pulse, as shown in Fig. 1.5.

In a nutshell, we use the CI because the vibration insensitivity and systematics rejection of the design allow us to implement large momentum separation while maintaining phase

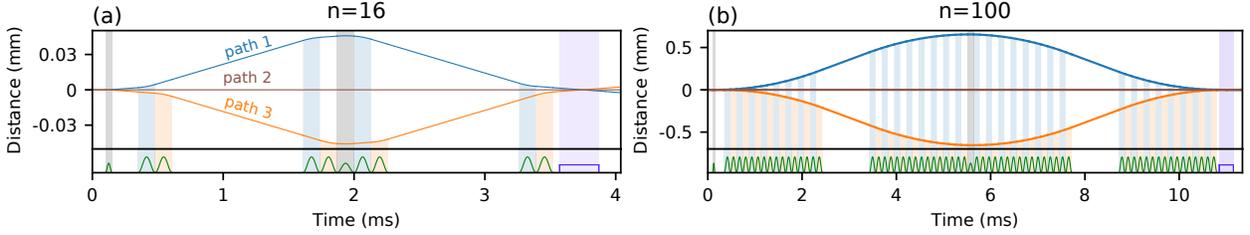


Figure 1.5: Space-time schematics of the CI with acceleration pulses to increase the momentum separation between paths 1 and 3. We label the momentum separation with n , the number of single photon recoils $\hbar k$ separating the outer paths. Shown are $n = 16$ and $n = 100$.

stability, which is otherwise lost for the highest momentum separation shown with other interferometer designs [7].

1.3.2 CI aesthetics

These scientific advantages are of course the main motivation for using this CI design, but I also want to comment on the aesthetics of this type of interferometer. In addition to the pleasing symmetry, the momentum state control and readout scheme have an intriguing series of reversals of the role played by light waves and matter waves.

First, recall that in an optical interferometer, physical gratings can be used to split and recombine light. In comparison to this type of optical interferometer, reversal 1 is that we use an optical standing wave in the role of the physical grating, while replacing light waves with atomic matter waves. Reversal 2 comes when the paths recombine, forming an atomic density grating, which reflects the readout light. Light and matter each serve as both a traveling wave and a Bragg diffraction grating during the CI sequence.

This trick, of using the dual roles of both light and matter, is only possible with a matter wave interferometer. A related interferometer design, the grating echo interferometer[6], uses a similar readout technique.

1.3.3 CI history

The first demonstration of a CI, [14], was led by my advisor, Deep Gupta, under the direction of his advisor, David Pritchard. This experiment showed that a CI could be implemented in the lab, and introduced the application of using a CI for a recoil measurement to determine the fine structure constant.

After Deep became a principal investigator at the University of Washington, he revived the CI project with his student, Alan Jamison. Alan started with theoretical work analyzing one of the main systematics: energy shifts due to atomic interactions [19]. I joined Deep's group just in time to begin assisting with the optics, electronics, and infrastructure to implement the CI in the Gupta group's already existing vacuum chamber. This work led to Alan's thesis, and a publication [20] demonstrating the first Ytterbium atom interferometer, along with experimental and theoretical work to correct for another systematic effect, a phase shift due to the diffraction pulses.

1.3.4 CI phase

Before discussing applications for the CI, here are the relevant phase calculations. Eq. 1.1 gives the atomic wave function at the end of the interferometer sequence, i.e. during the readout pulse, where k_{rec} is the wavevector of the diffraction laser, $\omega_{\text{rec}} = \frac{\hbar k_{\text{rec}}^2}{2m}$ is the 1 photon recoil frequency, and $\theta_1, \theta_2, \theta_3$ are the non-recoil phases accumulated over paths 1, 2 and 3 respectively. The amplitudes for the three paths are A_1, A_2 , and A_3 . The momentum separation between paths 1 and 3 is given by n , where $n=4$ is the smallest value possible.

$$\begin{aligned} \psi(x) = & A_1 \exp\left\{-i(2\vec{k}_{\text{rec}} \cdot \vec{x} + \frac{1}{4}n^2\omega_{\text{rec}}t + \theta_1(t))\right\} \\ & + A_2 \exp\{-i\theta_2(t)\} \\ & + A_3 \exp\left\{-i(-2\vec{k}_{\text{rec}} \cdot \vec{x} + \frac{1}{4}n^2\omega_{\text{rec}}t + \theta_3(t))\right\} \end{aligned} \quad (1.1)$$

The atomic density is then given in Eq. 1.2 . The symmetry of the initial splitting pulse implies that $A_1 = A_3 = A_{13}$:

$$\begin{aligned}
|\psi|^2 &= |A_1|^2 + |A_2|^2 + |A_3|^2 \\
&+ 2A_2A_{13}[\cos\left(2\vec{k}_{\text{rec}} \cdot \vec{x} + \frac{1}{4}n^2\omega_{\text{rec}}t + \theta_1 - \theta_2\right) + \cos\left(2\vec{k}_{\text{rec}} \cdot \vec{x} - \frac{1}{4}n^2\omega_{\text{rec}}t + \theta_3 - \theta_2\right)] \\
&+ 2A_{13}^2[\cos(4k_{\text{rec}}x + \theta_1 - \theta_3)]
\end{aligned} \tag{1.2}$$

We are interested in the second line of this expression, which after the application of some trig identities can be re-written as:

$$2A_2A_{13} \cos\left(2\vec{k}_{\text{rec}}x + \frac{\theta_1 - \theta_3}{2}\right) \cos\left(\frac{1}{4}n^2\omega_{\text{rec}}t + \frac{\theta_1 + \theta_3}{2} - \theta_2\right) \tag{1.3}$$

Some observations about this expression: first, as previously advertised it depends on the energy difference between paths 1 and 3, in the term $\frac{1}{4}n^2\omega_{\text{rec}}t$, which makes clear the quadratic scaling with momentum separation n . Second, the form of the non-recoil phases, $\frac{\theta_1 + \theta_3}{2} - \theta_2$ begins to show how the symmetric form of the CI leads to the cancellation of some systematic effects. A little reflection on this combination of phases should convince you that a constant offset applied to all three phases will cancel out, as will a gradient.

1.4 Recoil frequency and the fine structure constant

The recoil frequency can be connected to the fine structure constant by the expression shown in Eq. 1.4, often called the atomic physics route to α .

$$\alpha^2 = \frac{4\pi R_\infty}{c} \frac{m}{m_e} \frac{\hbar}{m} \tag{1.4}$$

R_∞ is the Rydberg constant, m_e is the mass of the electron, and m is an arbitrary atomic test mass. Since m cancels from the expression, it is clear that any atomic mass can be used. If m is inserted as shown, it can be combined with other constants in advantageous ways.

For several atomic species, the ratio $\frac{\hbar}{m}$ is the least well known term in this expression for α . The Rydberg is known to .008 ppb [37], and the electron mass to .03 ppb [35]. The bare atomic mass m is known to .1ppb for all three of ytterbium[5], rubidium, and cesium, which are the elements currently used for this type of measurement. The ratio $\frac{\hbar}{m}$ is known to a value some ten times less precise than the bare mass, currently from a rubidium recoil measurement[4].

This, then, is the idea of the experiment. By measuring the recoil frequency $\omega_{\text{rec}} = \frac{\hbar k_{\text{laser}}^2}{2m}$, along with the diffraction laser frequency to obtain k_{laser} , we extract $\frac{\hbar}{m}$. Until we surpass one of the other parameters (the first one we would reach is the bare atomic mass m), any increase in the precision of $\frac{\hbar}{m}$ is an increase in the precision of α_{recoil} (α as obtained from a recoil measurement).

Atomic recoil measurements to measure α were pioneered by Chu and Weiss at Stanford [38] using Cesium atoms. This is an active research area, with several groups pushing the state-of-the-art. The best published result is from the Biraben group in Paris, who measure α to .66ppb using rubidium atoms [4], while the Muller group at Berkeley is pending publication.

1.4.1 QED tests

It should be noted that recoil measurements are not currently the best (most precise) way of measuring α . The Gabrielse group at Harvard measures the electron gyro-magnetic ratio g-2, from which α can be extracted at .25ppb [17], with the help of QED theory to obtain α from g-2[2].

Given the success of the g-2 route to α , why also do a recoil measurement? First, there is some hope that a recoil measurement could beat the g-2 value. α_{recoil} is limited to .1ppb by the bare atomic mass measurement, but if the recoil precision was brought to the same level, it would beat the g-2 value by more than a factor of 2, potentially more if the bare

atomic mass is measured more precisely.

In addition, even if one didn't expect to improve the best known value of α , there are good reasons to perform both types of α measurement. Agreement between these two sub-ppb measurements is a powerful confirmation of the methods and systematics analysis involved in both experiments. In particular, comparing the two measured values constitutes a stringent test on a particularly hairy bit of QED theory. $g-2$ is calculated theoretically as a perturbative expansion in powers of α . In order to determine α from $g-2$ at the .1 ppb level, it is necessary to calculate up to 10th order terms in this expansion [2]. The recoil measurement route to α is independent of this theory, so obtaining a consistent value by such an method would validate these difficult calculations.

1.5 Ytterbium BEC interferometry

I've introduced contrast interferometry and recoil measurements in general terms, so far remaining agnostic in the choice of which atom to use, as well as the method of sample preparation. Here is an overview of reasons for working with our particular atomic system.

1.5.1 Ytterbium

In all of the experiments described in this work, we use Ytterbium-174 (^{174}Yb). Ytterbium is named after the town of Ytterby, Sweden, a locale that lends its name to no less than four different elements. It is a rare earth element with atomic number 70, and is the penultimate member of the lanthanide series. At room temperature and pressure, it is a soft, ductile metal that is described by Wikipedia as “somewhat hazardous as an eye and skin irritant”. Luckily for our health and safety, we are protected from this somewhat hazard since our sample of Yb is isolated in our UHV chamber, only to be manipulated with heat and light.

Ytterbium has two valence electrons, and so the physics of the atom are reminiscent of helium. It has two optical transitions that are useful for our purposes: the broad $^1S_0 \rightarrow ^1P_1$

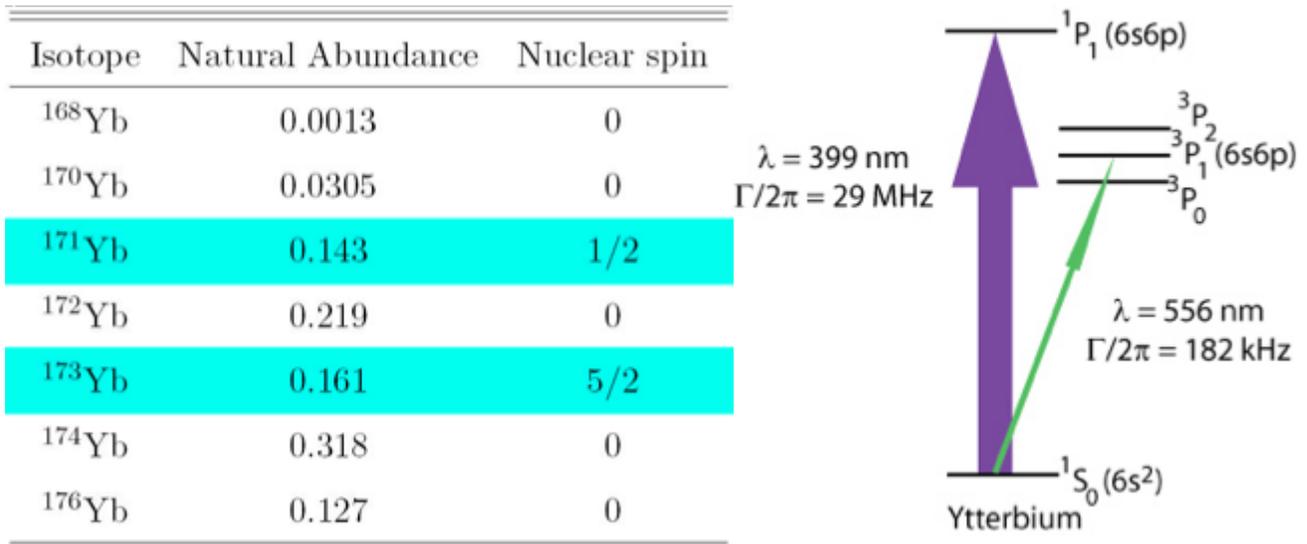


Figure 1.6: Spins, and abundances of the various stable isotopes of Yb, with the fermionic isotopes shaded. Yb energy level diagram, with arrows representing the two optical transitions used in this work.

transition with wavelength $\lambda_b = 399 \text{ nm}$ and width $\Gamma_b = 2\pi \times 28 \text{ MHz}$, and the narrow $^1S_0 \rightarrow ^3P_1$ intercombination line with wavelength $\lambda_g = 556 \text{ nm}$, and width $\Gamma_g = 2\pi \times 182 \text{ kHz}$. The energy level diagram is shown in Fig. 1.6.

Ytterbium has several properties that recommend itself for precision atom interferometry:

- The 1S_0 ground state of ^{174}Yb (and other bosonic isotopes) has no nuclear spin and has $J = 0$, so is practically insensitive to magnetic fields.
- Convenient optical transitions. The broad 399nm line is useful for Zeeman slowing and imaging, while the 556nm line gives a low Doppler temperature MOT, as well as allowing access to a wide range of detunings using standard AOM frequency shifts, needed for generating the diffraction beams.
- The relationship between the two optical transitions is such that the Bragg angle of

the 399nm light reflection from the atomic grating formed by 556nm diffraction beams is 46 degrees, allowing for more convenient CI readout optics.

- Multiple available isotopes. One of our experiment's systematic effects is atomic interactions, which would have different physics for other isotopes. Performing the measurement with different isotopes could be a useful for confirming our ability to model and subtract these effects.
- Can be cooled to quantum degeneracy.

1.5.2 BEC interferometry

The contrast interferometer, as described so far, is single atom physics. Yet we go to quite a bit of trouble to prepare a Bose-Einstein condensate (BEC) interferometer source. Why?

- Signal enhancement from many atoms is necessary for practical measurement. We can't realistically do the experiment one atom at a time.
- Given that we will use many atoms per shot, we'd like them to be as identical as possible. Since the CI signal to noise depends on momentum spread, we desire a narrow momentum spread, which is one of the hallmarks of a BEC.
- The diffraction beams are lasers with a finite waist. We would like to approximate the wavefronts as flat at the atoms, which means that we want the spatial extent of the atomic cloud to be small compared to the waist. A BEC starts small and expands slowly.

However, using a BEC does come with some disadvantages:

- The high in-trap density of a BEC means that there is a significant energy (in precision measurement terms) due to atomic interactions. After the BEC is released from the trap, this interaction energy is converted to kinetic energy, and appears as a systematic in the CI phase [19].
- Preparing the atoms in a BEC state requires a more complicated apparatus than, for instance, laser cooled atoms. This costs both a larger up front investment to build the experiment, as well as increased debugging and down time as there are more things to go wrong.
- Evaporative cooling happens on the order of seconds, which increases the cycle time of the experiment, making it more costly in time to increase precision via statistics ($1/\sqrt{N}$ scaling).

Chapter 2

APPARATUS

2.1 *Vacuum chamber*

2.1.1 *The atomic physics machine of Theseus*

The ship of Theseus is a classic philosophical puzzle about identity. A wooden ship is repaired by replacing various parts until no original components remain. Is it the same ship, and if not, at what point did it change?

The first round of Yb contrast interferometry, led by Alan Jamison, was done on the mixtures apparatus in B063. As Alan's work was wrapping up, it was becoming more and more onerous to switch the mixtures machine back and forth between interferometry and mixtures experiments. Around this time, I began working on the design of what would become the CI Machine.

The CI Machine began life as a ytterbium MOT experiment under the direction of Norval Fortson. The machine had been sitting idle for years. As we contemplated building a new BEC machine for interferometry, the original thought was: here is a vacuum system configured for produced ytterbium MOTs, lets just revive that!

So that is what we did. However, along the way I ended up replacing all but one of the parts.

2.1.2 *Design*

The overarching design goals for the new CI Machine were the following:

- Fast, reliable production of ytterbium BECs.

- Optical Dipole trap geometry to maximize BEC coherence length (painting).
- Optical access for diffraction beams and readout light collection.

Oven design

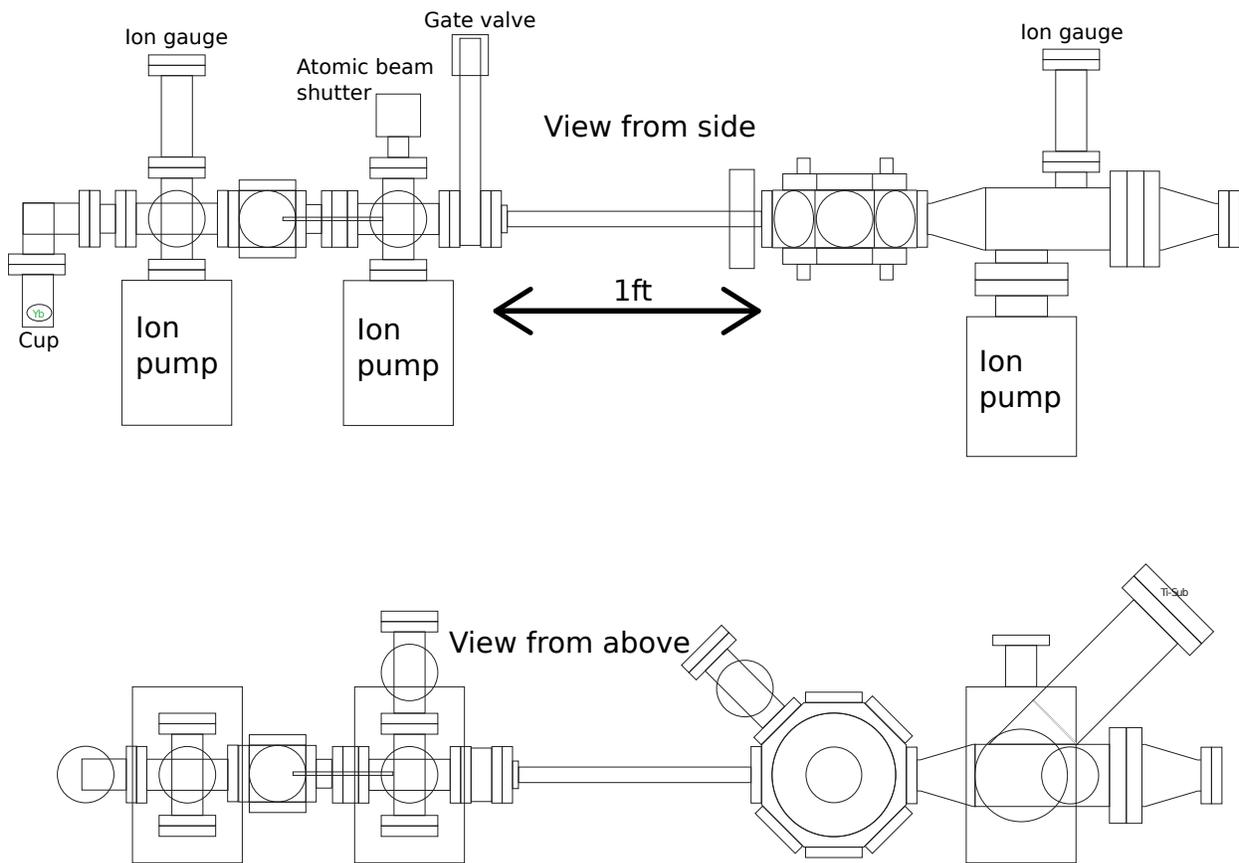


Figure 2.1: Vacuum chamber schematic

The source for the atomic beam is referred to as the oven. I tried to keep the design of the oven as simple as possible. It consists of a “cup”, which is a 2-3/4” conflat half nipple with a welded cap sealing the open end. Metallic ytterbium is placed in the cup, which is

oriented vertically in standard cup fashion so the ytterbium ingots rest at the bottom. The vertical orientation is critical for elements that liquify before evaporating, to make sure the metal does not flow out of the cup. For ytterbium, which sublimates directly from solid to gas, the vertical orientation is still advantageous because it makes it easier to remove and replace the cup if the ytterbium runs out and needs to be replenished.

The cup is attached to a 90 degree elbow, followed by a minimum length nipple, and then the oven region 6-way cross. There are two collimating apertures for the atomic beam, the first is between the elbow and the minimum length nipple, the second is on the opposite side of the 6-way cross. The first aperture is 4mm (#20 drill bit), the second aperture is 5.7mm (#1 drill bit). The second aperture is larger than the first in order to create an atomic beam with some divergence, with the hope that the beam could be further collimated using transverse cooling (more on this later).

The apertures consist of solid disk copper gaskets with holes drilled in them. These gaskets are used in lieu of the usual ring-shaped gaskets which usually go in the junction between two conflat flanges. The holes were drilled by me in the machine shop.

The oven is heated by three band heaters: two on the cup, and one on the first aperture. The band heaters have integrated thermo-couples which are used to feedback on and stabilize the temperatures.

It is important to keep the first aperture at a higher temperature than the cup, so that the atoms do not condense on the aperture and clog it. Our aperture is relatively large so this is not as much of a concern, but it is better to play it safe and not risk needing to break vacuum and rebuild the oven. The cup needs to be maintained at around 380C, so the first aperture needs to stay around 450C. The reason for the minimum length nipple between the elbow and the 6-way cross is to create a little bit of distance between the very hot aperture flange and the 6-way cross.

The four branches of the 6-way cross that are orthogonal to the atomic beam have the

following attachments: Ion pump, titanium sublimation pump (TSP), Ion gauge, and angle valve. The pumps and gauge are so that the oven region can be separately pumped down and monitored, as part of the differential pumping scheme. The angle valve allows the oven region to be attached to the roughing pump and turbo pump.

Transverse cooling and beamline lock chamber

The next module after the oven 6-way cross is a Kimball cube, which along with two ion pumps are the only remaining components of the original CI Machine. This cube has four viewports, allowing optical access to the atomic beam.

Differential pumping region

After the Kimball cube, the next component is a differential pumping tube, followed by another pumping region. The differential pumping tube has dimensions length 4.8", and ID 1.9". It is a stainless steel tube welded to a .75" double sided blank flange. The flange has a clearance hole that the tube passes through, and is counter-sunk so that the clearance hole only touches the tube in the non-counter-sunk region, where the tube and flange are welded together. This counter-sunk design is to avoid regions where gasses may be trapped, leading to internal "virtual leaks".

The double-sided flange holding the differential pumping tube is sandwiched between a minimum length nipple and a 6-way cross. The 6-way cross vertical axis has an ion pump below, and on top a rotating mechanical feedthrough for the atomic beam shutter. The horizontal axis orthogonal to the atomic beam has a viewport on one side, and a T on the other. One branch of the T has a viewport, the other has another T with a TSP and an ion gauge.

After the differential pumping 6-way cross is a gate valve, so that the oven and differential pumping regions can be completely sealed off from the main chamber area. One design

feature that I would add, if I were rebuilding this machine, would be to have a second gate valve on the other side of the differential pumping region. I don't totally trust the gate valve design. We had one gate valve fail during the baking process, and I've heard anecdotally about other failures. It would be a big time-sink to re-bake the main chamber, so in the event that the oven needs to be opened, having two gate valves between the oven and the main chamber would be safer. In addition, having a gate valve between the oven and the differential pumping region would mean that the differential pumping region wouldn't need to be re-baked if the oven needs to be opened.

Rotating mirror atomic beam shutter

It is desirable to have a mechanical shutter to block the atomic beam from reaching the main chamber, to prevent the atomic beam from interfering with atoms during evaporative cooling. One design consideration is to place the atomic beam shutter downstream in the beam from the beamline lock chamber, so that when the shutter is closed it does not block the atomic beam in the beamline chamber. In the CI machine, the atomic beam shutter is mounted on a rotating mechanical feedthrough on the differential pumping 6-way cross.

I designed and fabricated a custom atomic beam shutter that serves a dual purpose: it is also a mirror that is oriented at 45 degrees facing towards the main chamber when the atomic beam shutter is closed. The original reason for creating this mirror/shutter was to allow one of the optical trap beams to be sent in along the slower axis, and exit via the mirror/shutter. The Arduino-based servo controller for the shutter was designed to have an interlock to prevent the high power dipole trap laser from turning on unless the mirror was in place to direct the beam back out of the chamber.

Later, it was decided not to have an optical trap laser on the atomic beam axis, so the mirror/shutter has not been used for that purpose. It has been useful for monitoring the slower light, which reflects from the mirror/shutter and exits the chamber when the shutter

is closed. This slower light diagnostic allows us to monitor the transmission through the slower viewport over time to ensure that the slower viewport is not becoming coated with Yb.

Zeeman slower tube

The vacuum component that connects the gate valve to the main chamber is a custom ordered reducing nipple with 1-1/3" flange on one side and 2-3/4" flange on the other. The Zeeman coils (described in more detail later) are wrapped on brass pipe that can slide on to the nipple. The reasoning behind this design is that the Zeeman coils would be damaged by the temperature of the main chamber bake, so it is desirable to be able to add the coils after the first main chamber bake.

Main chamber

The main chamber is a Kimball Physics spherical octagon, part MCF600-SphOct-F2C8, with 8 2.75" viewports. The octagon geometry is ideal for our application, since it allows three orthogonal axes for the MOT, plus an additional axis for diffraction optics.

Custom cross

One of the challenges of the Spherical octagon is where to attach the pumps. On the one hand, the center of the main chamber is the most critical area to have at high vacuum, so one wants to have high conductance from between the pumps and the main chamber. On the other hand, the main chamber was selected to have as much optical access as possible at the angles needed for our experiment, so one wants to avoid interfering with that.

For pumping speed, the best solution would be to attach the pump to the top or bottom of the octagon. This configuration was considered, but eventually rejected due to another constraint: MOT coil placement. With a large attachment to the top or bottom, the water

mirror holder1.pdf

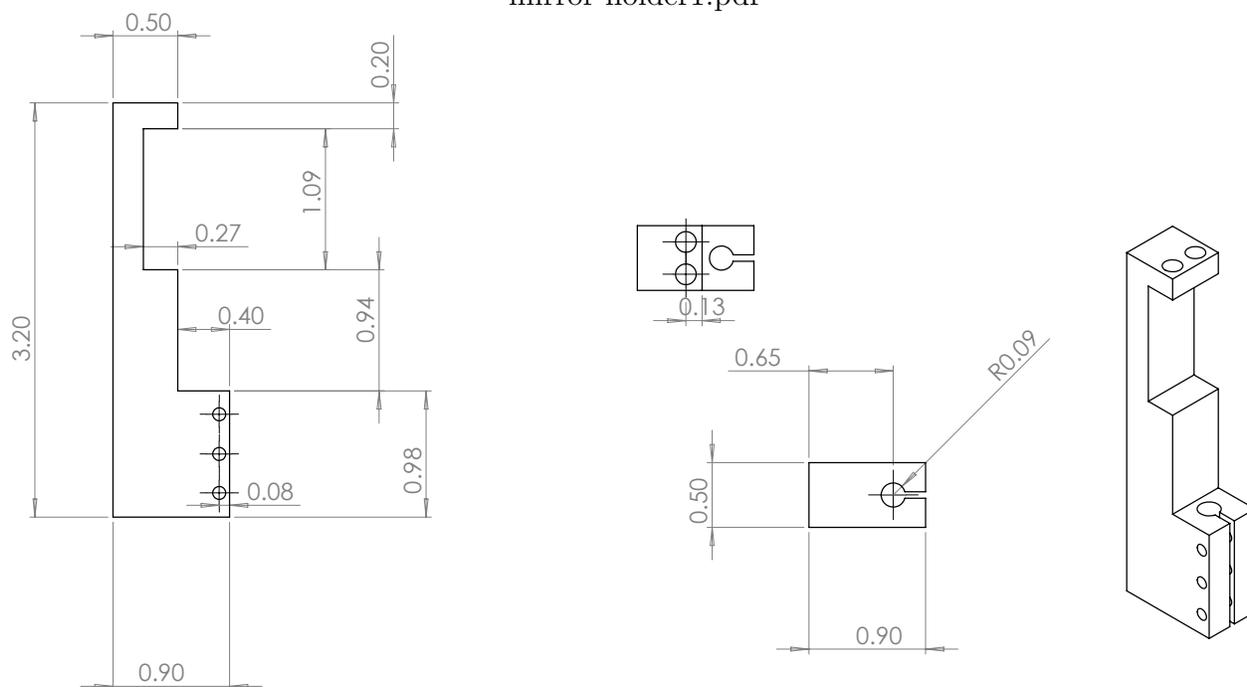


Figure 2.2: Technical drawings for the atomic beam shutter with mirror. The part shown here is a bracket that connects to the shaft of the rotary feedthrough, and holds a 1inch square mirror. It is made from 6061 aluminum, and made by me on the mill in the student shop. It uses set screws to hold the mirror in place. Care was taken to vent the ends of the set screws where they press into the mirror, to avoid trapping a pocket of air creating a virtual leak. When the mirror surface is parallel to the atomic beam it is offset to allow the beam to pass. When It is rotated 45 degrees, the mirror is placed to reflect a beam coming from the main chamber towards a viewport on the differential pumping 6-way cross.

R3 = change ti pump flange to 4-5/8CF
 lengthen "Y" port to 10", note R or NR CF's
 R4 = change y branch port to 4-5/8 x 3

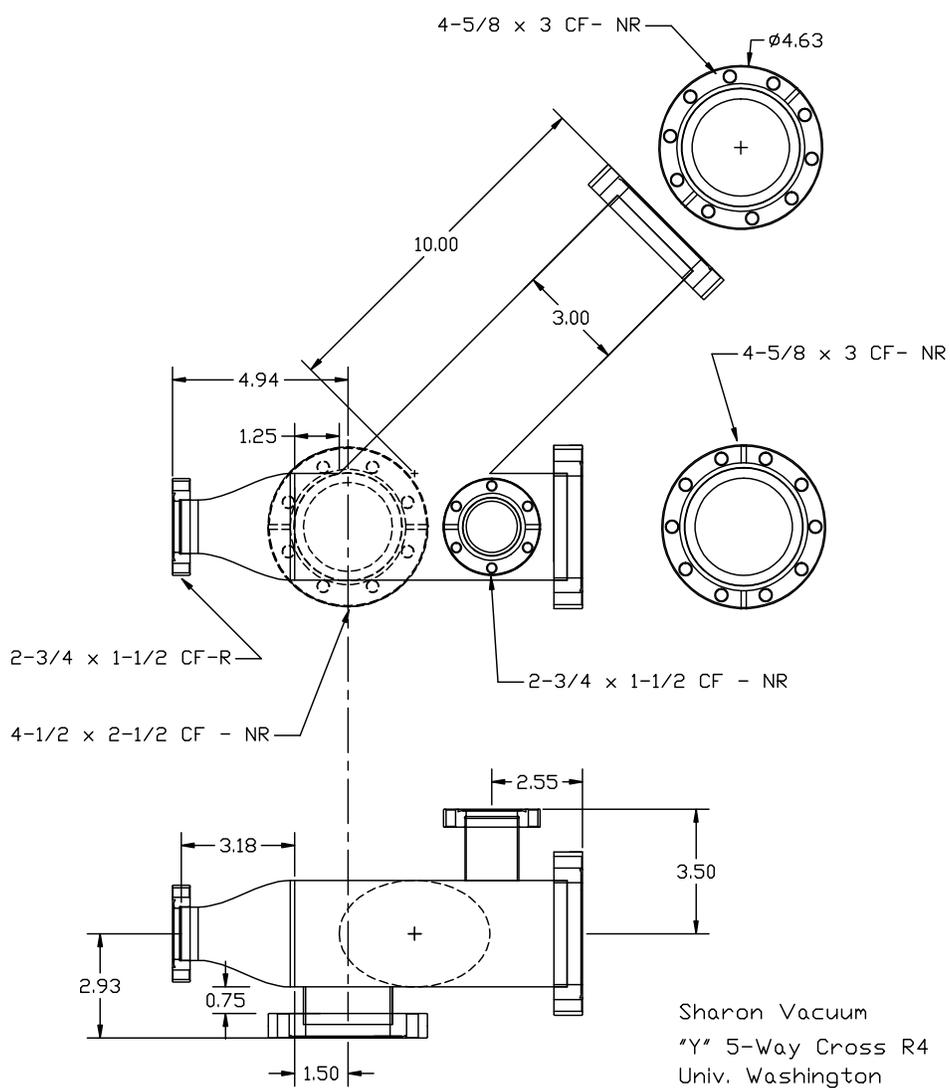


Figure 2.3: Technical drawings for the custom cross vacuum fitting manufactured by Sharon Vacuum.

cooled MOT coils would need to be geometrically entangled with the vacuum system such that they could not be removed after the machine was fully assembled. Since the MOT coils cannot be baked at full temperature, such a machine would not be workable.

The solution we settled on involved a custom manufactured vacuum component we refer to as the “custom cross”, shown in Fig. 2.3. The custom cross attaches to the main chamber opposite to where the atomic beam enters. After the 2.75” flange, the cross expands to a large tube size to increase conductance to the ion pump, and surface area exposed to the TSP. The custom cross has flanges for an ion pump, a TSP, an ion gauge, and a viewport for the slower light to enter along the atomic beam axis.

The TSP wing of the custom cross is attached at a 45 degree angle, to keep it from obstructing the area in front of viewports where optics need to be placed. I took care in specifying the length of the TSP wing so that there would be no line of sight from the viewports in the main chamber to the TSP filaments, to avoid depositing titanium on the viewports.

Slower viewport reducing nipple

The slower viewport is mounted on a modified 4-5/8” to 2-3/4” reducing nipple. The reasoning behind this design is to create some spatial filtering to keep atoms from the atomic beam, which scatter from the heated slower viewport, from bouncing back into the main chamber. The slower viewport is heated to prevent atoms from accumulating on it (with mixed success, see section on Verdi cleaning).

The 4-5/8” flange has a stainless steel ring welded to it. The ring is visible in through the slower viewport in Fig. 3.9. The ring is far enough from the heated viewport that it is close to room temperature, so atoms that bounce off the viewport stick to the cooler ring. The hole in the ring is just big enough to not obstruct the atomic beam or the slower light.

A close inspection of the internal stainless ring will reveal the inscribed initials of the

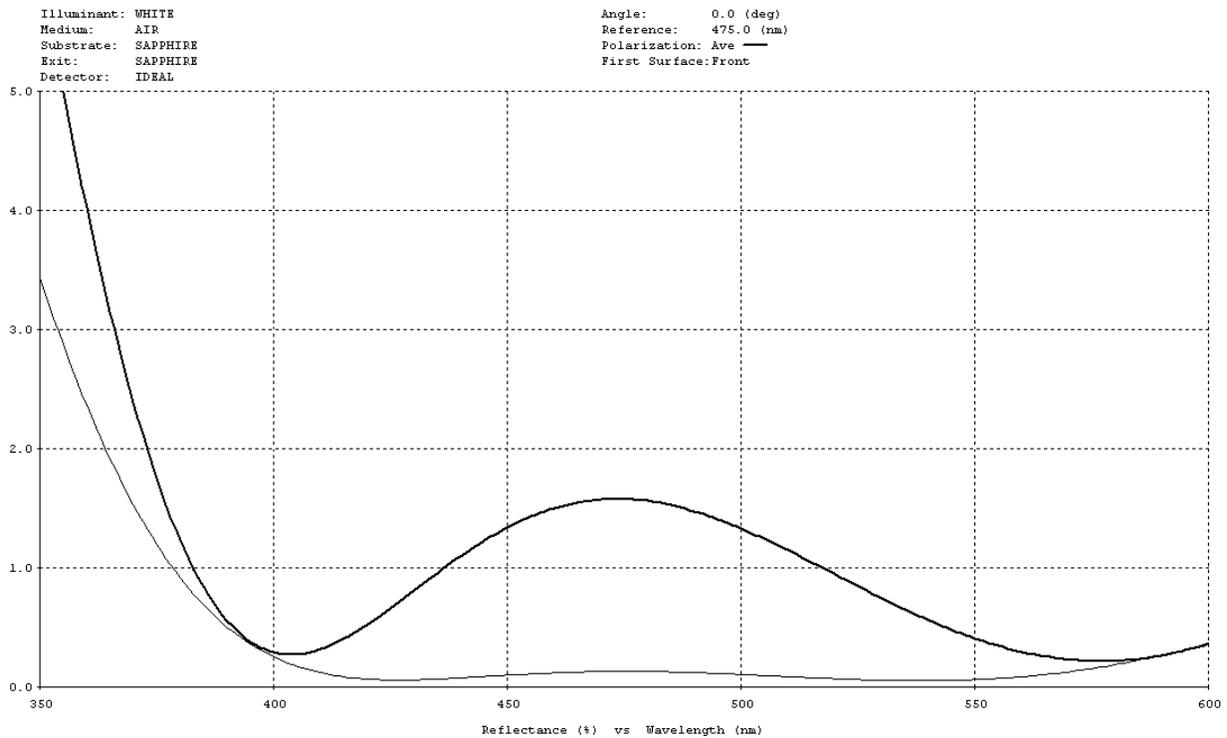


Figure 2.4: Theoretical curves for the viewport AR coating, from Larsen Electronic Glass. The lower curve is for the AR coating applied to glass, for which it is designed. The upper curve is for the same coating applied to a sapphire viewport. The performance is lower because the coating is not designed for sapphire, but we judged it to be good enough to use the coating on our one sapphire viewport.

original CI team (Alan and Ben).

Viewports

The Anti-Reflection (AR) coated viewports were ordered from Larsen Electronic Glass. The theoretical transmission curve for our custom coating is shown in Fig. 2.4. The coating was optimized by Larsen to transmit at 556nm, 532nm, and 399nm, the three laser wavelengths used in our experiment.

All the viewports use standard BK7 glass, except for the slower viewport which is sap-

phire. While the experiment is operating, the slower viewport is maintained at 200C to prevent the atomic beam from condensing on the viewport and changing the transmission properties. The glass-metal junction of the viewport is one of the weak points in the vacuum system, particularly when it is subject to rapid temperature changes. Usually, the slower viewport is maintained at a constant temp, but problems can arise in the event of a power outage. The viewport is heated by a band heater powered by a variac. The viewport is insulated with an aluminum foil wrapping outside the band heater, so if the power goes off the temperature goes down slowly enough that it is not an issue. The problem arises when the power returns, and the variac is still set to bring the temp to 200C, causing the temperature to rise dangerously quickly. This sequence of events led to a vacuum disaster in the B063 mixture. I tried to mitigate this danger in several ways. First, make every effort to turn off the variac after a power outage before the power is restored. Second, power the variac with an uninterruptible power supply (UPS), although I've found that in practice the amount of time the UPS can power the variac is measured in minutes. Finally, the choice of sapphire, which is less susceptible to damage from temperature changes.

Unfortunately, the AR coating designed for the glass viewports has different properties when applied to sapphire, and it is economically impractical to do a separate coating run for the single sapphire viewport. The darker curve in Fig. 2.4 shows the transmission for the coating on sapphire, which we found to be good enough for our purposes.

2.1.3 Construction

Here are some general principles of vacuum system construction.

Vacuum Screws

Vacuum screws come in a variety of standard sizes based on the flange size being used. For instance, our most commonly used flange size is 2.75", for which the screw size is 1/4-

28 (1/4" diameter, 28 threads per inch). The thread pitch is only relevant for internally threaded pieces. There are standard sizes depending on whether the connection is between two flanges with through holes, or between one through hole and one internally threaded flange. The Duniway catalogue is a good resource for finding these screw sizes.

For through hole to through hole flange connections, a nut is required. In most cases, the easiest solution is a plate nut, conflats specific piece of hardware with two threaded holes separated by the hole spacing in the flange. The advantage of a plate nut is that once both screws are threaded into it, the plate nut cannot rotate. A standard nut and bolt requires a wrench on both sides, while a plate nut only needs one on the bolt side.

For UHV flanges that are baked at high temperatures, It is a concern that screws may seize, making it difficult to disassemble parts. This is especially a concern for parts of the oven that are both maintained at high temperature and may need to be disassembled and reassembled to replace the atomic source material, but it is something that should be kept in mind for all parts. To minimize the chance of seizing we use silver plated screws. The silver plating is designed to work against a stainless steel surface. This means that only one of the nut and bolt should be silver plated. In most cases we use silver plated bolts and stainless steel nuts (or plate nuts), but there are a few places where for various reasons we use stainless steel threaded rod with silver plated nuts.

Gaskets

Conflat flanges create a vacuum seal by squeezing a soft copper gasket between two "knife edge" circular ridges. It was learned (the hard way) in the B063 mixtures experiment that hot ytterbium reacts with Nickel plated gaskets, so we were careful to use only copper. Silver plated copper gaskets are available, and we used them in parts of the oven that would be maintained at high temperatures.

Creating a seal

The procedure for creating a good vacuum seal:

- Nitrile gloves should be worn while handling clean vacuum fittings, especially while cleaning and handling the knife edge and gasket.
- If possible, a connection should include one rotatable and one non-rotatable flange.
- Visually inspect both knife edges to insure there are no imperfections.
- Clean both flanges and the gasket using methanol and a lint free wipe.
- Put the gasket on one flange, and then bring the flanges together trapping the gasket between them. This can be tricky if one of the flanges is fixed in a vertical orientation. When the gasket is properly seated between the flanges, you should be able to visually verify that the flanges are concentric, and that the gap between the flanges is even all the way around. Further, it should be possible to “feel” that the gasket is in place by, while pushing the flanges together, trying to slide them against each other in a shearing motion. If the gasket is seated properly, it will stop this motion.
- If one or both of the flanges is rotatable (or barring that, if one of the fittings is itself free to rotate with a fixed flange), align the flanges so that the sniffer ports are lined up. If possible, arrange the sniffer ports to be horizontal. Sniffer ports are used to allow methanol or helium gas to flow to the gasket seal for leak testing.
- While holding the flanges together, insert the vacuum bolts and plate nuts. There is a small washer that comes with the bolts that should go on the bolt side, not the nut side. An extra pair of hands is often useful if not required for getting the bolts in place. Finger tighten the bolts, but be careful not to over tighten any one bolt before

all bolts are in place. Once all bolts are finger tight, again verify that the gap between the flanges is even all the way around.

- Begin tightening the bolts with a wrench. Tighten the bolts in a star pattern, being careful not to tighten two bolts next to each other. Tighten each of the bolts before returning to the first one. During each round of tightening, turn each bolt by about a 1/2 turn, reducing to a 1/4 each near the end of the process
- Continue tightening the bolts in the star pattern as needed. This is something that you'll have to be shown to get a feel for. It is not necessary to tighten to the point that the flanges touch each other.

2.1.4 Baking

In order to achieve Ultra High Vacuum (UHV), it is necessary to “bake” the apparatus to speed up the desorption of contaminants on the inner walls of the chamber. The vacuum fittings are all coated with trace amounts of water from the atmosphere, and at room temperature, it would take years for this water to desorb and be pumped out of the system. Baking accelerates this process to a matter of days.

We did the bake of the CI Machine in three stages. First, we baked the main chamber at high temperatures before installing any of the viewports or magnetic coils. Next, we baked the oven region separately from the main chamber. Finally, we put everything together, and did a third lower temperature bake.

Main Chamber High Temperature Bake

For the high temperature bake, the main chamber was assembled up to the Zeeman slower tube. Instead of connecting to a gate valve, the slower tube instead had a blank. All viewports were replaced with blank flanges. As an extra measure, we put a TSP sticking into

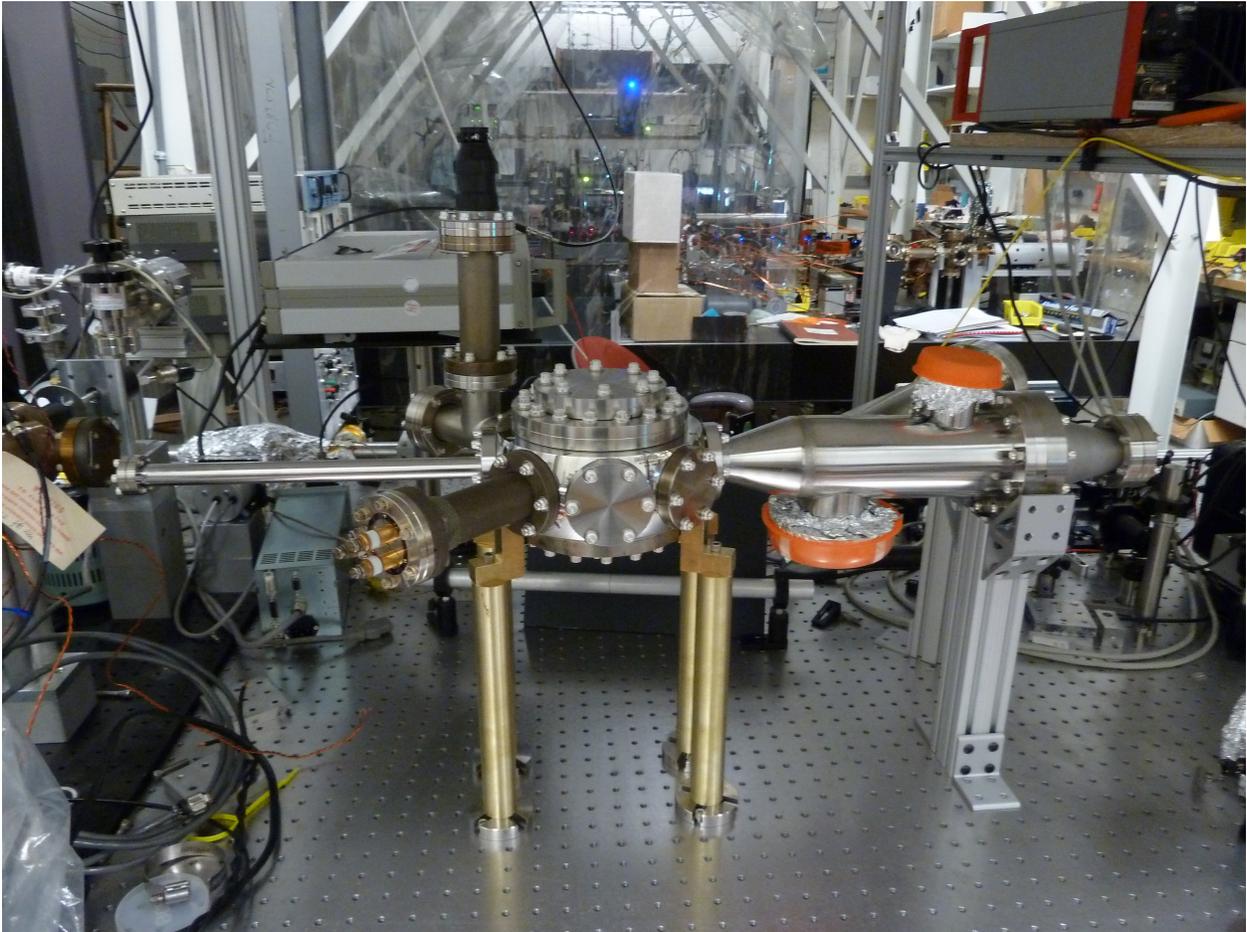


Figure 2.5: Main chamber before high temperature bake

the center of the main chamber. We fired this TSP during the bake, coating the inside of the main chamber with titanium. Although it is likely that this titanium layer was effectively nullified when the main chamber was brought to atmospheric pressure under argon, it is possible that after the second low temperature bake it still played some useful role.

Procedure for baking the chamber:

- Apply heating elements to the chamber. The goal of this step is to apply heat as uniformly as possible. We used a combination of heating tapes and band heaters.

While applying heating tapes, it is best practice to avoid crossing the tape over itself to create a double layer, since this can cause a hot spot. It is especially important to have well controlled and uniform heating of the viewports; band heaters are ideal for this purpose.

- Attach thermo-couples for monitoring the temperature. The thermo-couples should be placed first of all in areas that may get very hot, such as under band heaters. After accounting for hot spots, thermo-couples should be distributed to try to sample the temperature in as many places as possible. More thermo-couples are better, up to the point where it becomes impractical to monitor them all. It is important that thermo-couples are anchored securely in contact with the vacuum fittings. Kapton tape can be useful for this, but it is usually wise to “strain relieve” the thermo-couple wire by winding it around the fitting or a support bracket so that it is not just tape holding it in place.
- Insulate the whole assembly with aluminum foil. Good insulation is needed to attain a close to uniform temperature over the whole chamber. It is best to insulate in onion-like layers. Take care while foiling the chamber not to dislodge any of the thermo-couples, and make sure that the thermo-couple wires and the power cords for the heating elements don’t get trapped in the insulation.
- Turn up heat *slowly*, while monitoring the thermo-couples to make sure that the temperatures are as uniform as possible over the whole chamber. We would bring everything up to the target temperature over a day or two, with an equal time to bring the temperature back down.
- For the secondary, lower temperature bakes: fire the TSPs and turn on the ion pumps while the chamber is at the full target temperature.

- While the temperature is going down, monitor the composition of the remaining gasses with the residual gas analyzer.

2.2 Magnetic Coils

2.2.1 Zeeman Slower Coils

The Zeeman coils were designed by Dan as his REU project the summer before he joined UW as a grad student. Dan and I worked together to wind the coils on the lathe in the machine shop (turned by hand!). For details on this process, see Dan's REU final paper.

2.2.2 Magneto Optic Trap Coils

The MOT coils were wound using the same materials and similar techniques as the Zeeman slower coils. The MOT coils are used in an anti-Helmholtz configuration, with a gradient at the center of 0.64 Gauss/cm per amp.

The MOT coils are attached to the main chamber with custom machine aluminum clamps, attached by screws to threaded holes in the main chamber.

2.2.3 Compensation Coils

The compensation coils are used to zero out any stray fields (for instance due to the ion pump magnets), and to make fine adjustments to the exact location of the zero point of the MOT gradient. This is important for overlapping the location of the MOT with the ODT.

The compensation coils are wound directly on the Kimball spherical octagon main chamber. They are wired up in Helmholtz configuration. The two pairs of horizontal coils each generate 0.55 Gauss/amp, while the vertical coils generate 0.43 Gauss/amp.

2.3 Electronics

2.3.1 Cicero

To control the experiment, we use the computer program Cicero, a standard tool for cold atom experiments [22]. This program has a graphical user interface that allows us to configure experimental sequences consisting of a series of time steps.

2.3.2 Analog and Digital channels

Digital TTL channels are output via National Instruments PCIe 6535B card, which is connected to a lab-built (thanks to Ryan Weh) digital isolator box.

Analog channels are output via a pair of National Instruments PCI 6713 cards, each connected to an NI BNC 2110 breakout box.

2.3.3 MOT fluorescence Monitor

In a typical MOT based experiment, the MOT is loaded for a fixed amount of time during each experimental cycle. If the MOT loading rate changes from run to run, the total number of atoms loaded during this fixed time can vary. In an effort to reduce atom number variations, I wanted to find a way to “trigger” the experiment based on the MOT fluorescence.

I accomplished this goal by adding an Opal Kelly Field Programmable Gate Array (FPGA) to the Cicero/Atticus system. The FPGA is mainly intended to allow the system to run on a variable timebase clock, but the feature of interest to me was the ability to pause a Cicero sequence midway through until a TTL trigger is received.

I set up a Thorlabs PMM02 PMT to collect MOT fluorescence. The PMT signal is sent through an SRS 560 low noise pre-amplifier where it is low-pass filtered, and then to a comparator circuit. The comparator circuit has a TTL output that goes from low to high when the MOT fluorescence signal climbs above an adjustable voltage threshold. This TTL

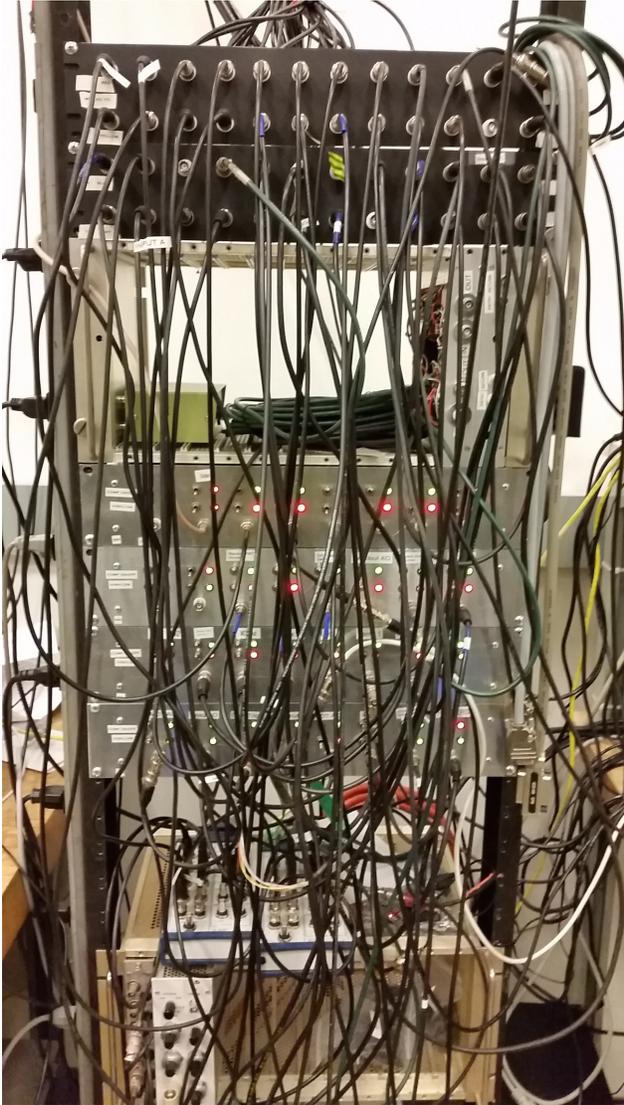


Figure 2.6: The analog and digital inputs and outputs connected to the National Instruments cards controlled by Cicero.

is sent to release the Cicero timestep hold.

This functionality has been extremely nice to have, especially on days when the injection locks are misbehaving. When a laser unlocks, instead of running a whole sequence only to see that no atoms were loaded, the sequence hangs at the MOT load step until the lasers are relocked, after which the sequence can continue and finish. This also improves atom number stability, since it prevents the sequence from running in cases where the MOT did not load for the entire MOT load step.

2.3.4 Absorption Imaging CCD Camera

To start with, the absorption imaging for the CI machine was done with an ancient Princeton Instruments camera. Late in 2016, at an experimentally inopportune moment, the Princeton camera died. We considered replacing the camera with something similar to the Andor camera used in B063, to the tune of around \$20,000. In order to try to save some money, we instead chose to try out a cheaper option.

We purchased a Mako G-32B CCD camera from Allied Vision Technology for less than \$300. The camera control and absorption imaging software had to be re-written in Matlab. After a couple weeks, we were back up and running. The performance of the Mako camera is sufficient for our purposes.

2.4 Optics

As mentioned in the introduction, ytterbium has two optical transitions that are useful for our purposes: the broad $^1S_0 \rightarrow ^1P_1$ transition with wavelength $\lambda_b = 399$ nm and width $\Gamma_b = 2\pi \times 28$ MHz, and the narrow $^1S_0 \rightarrow ^3P_1$ intercombination line with wavelength $\lambda_g = 556$ nm, and width $\Gamma_g = 2\pi \times 182$ kHz.

The energy level diagram is shown in fig2.7

The strength of the 399nm transition is most useful in cases where we need a high

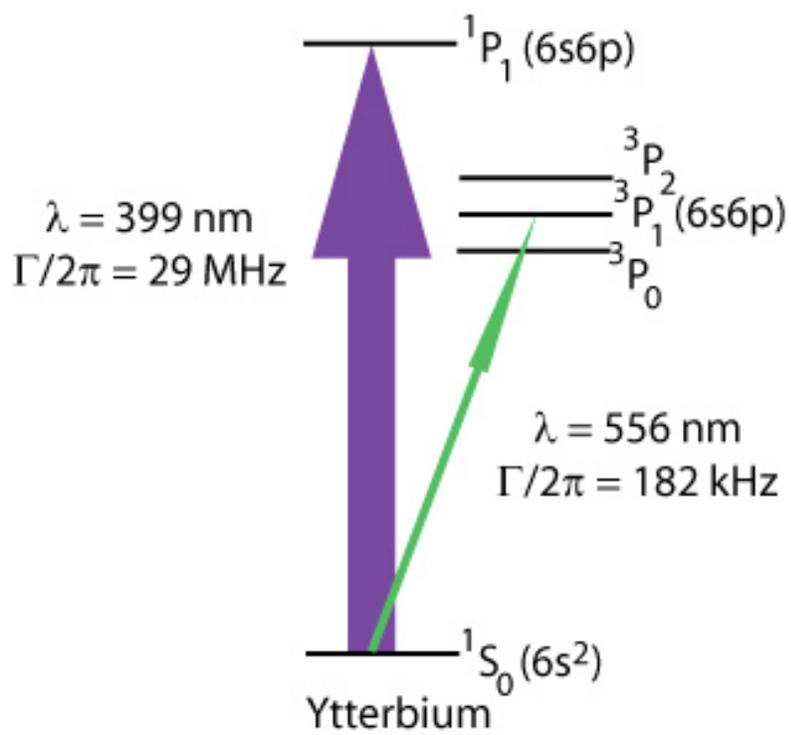


Figure 2.7: Yb energy level diagram

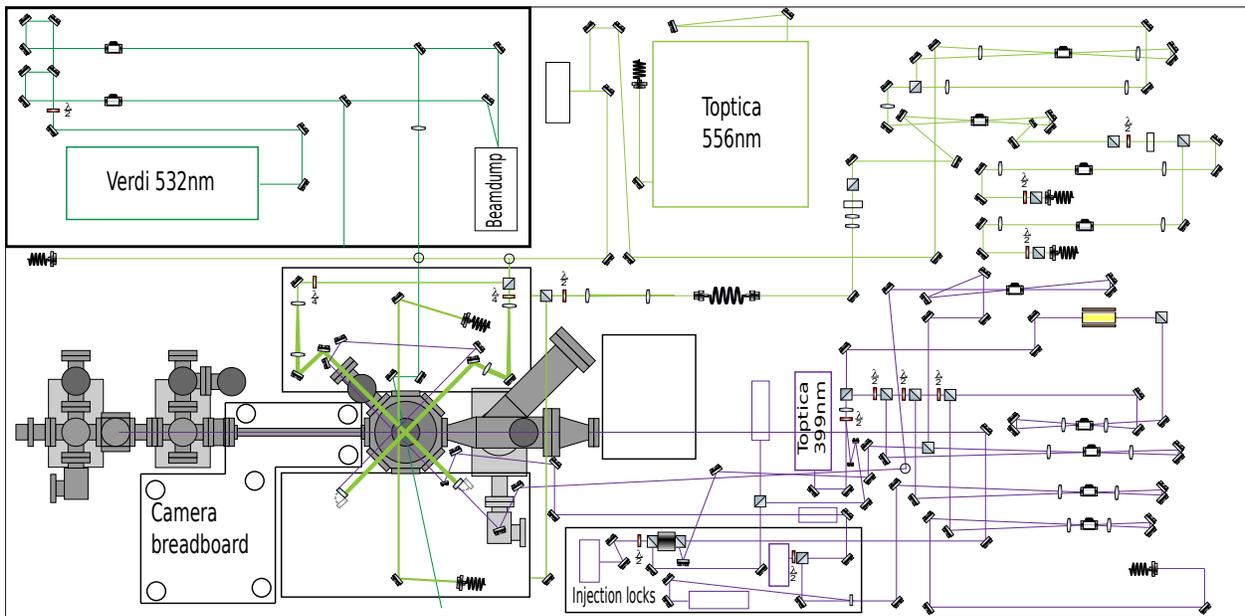


Figure 2.8: Optics table schematic overview

scattering rate. For us, this describes our initial laser cooling in the Zeeman slower and our absorption imaging. We also use 399nm for the readout light due to a fortuitous relationship between the 399nm wavelength and the periodicity of the atomic density grating created by the 556nm diffraction beams. More details on this in a bit.

For applications where a narrower linewidth is desired, we use the 556nm transition. These applications include the Magneto-Optic Trap (MOT), where the narrow width leads to a lower Doppler temperature, and the diffraction beams, where the narrow width allows us to reach the necessary detuning using a 200MHz AOM.

2.4.1 399nm Source

The 399nm source is a Toptica DL Pro module outputting 35 mW. This is an External Cavity Diode Laser (ECDL). This source has functioned well, my only complaint is that

it would be nice if it produced more optical power, so that the injection locks were not necessary.

The one hiccup in the operation of this laser was that the optical isolator (contained in the body of the laser) developed an imperfection after about a year off use. We've seen similar "spots" develop in polarizing beam cubes used with 399nm light. It seems that these relatively high energy photons can damage optical surfaces with long term use.

2.4.2 556nm Source

We derive our 556nm light from a Toptica TA/SHA Pro laser system. This integrated system consists of a DL Pro (diode laser) module outputting 35mW at 1112nm, a Tapered Amplifier (TA) chip that increases the power to 900mW, and finally a Second Harmonic Generation (SHG) cavity that doubles the frequency to the desired 556nm. The output power at 556nm has been about 390mW during the time where the key data for this thesis was taken. The power has been slowly decreasing since the laser was installed, see the laser power log in the lab notebook for details.

The TA/SHA Pro is powered by a control rack with several modular components. The definitive guide to these electronics is of course the manual.

2.4.3 Injection Locks

The 35mW output by the main 399nm laser is not enough power for the Zeeman slower. In order to avoid the cost of a more powerful 399nm laser system, we generate more optical power by the method of injection locking. [16]

2.4.4 532nm Source

We use a Coherent Verdi V18 laser outputting 18W at 532nm for our Optical Dipole Trap (ODT). The requirements for the ODT laser are high power (at least 10s of Watts), and red

detuning from the strong optical transition at 399nm. Since 532nm is blue of the 556nm transition there is a small (12%) repulsive component to the interaction, but the net effect is attractive.

532nm was chosen mostly due to the availability of the 18W Verdi laser, which was sitting unused when we started building our experiment. It had previously been used to pump a Titanium-Sapphire laser which was part of a now-replaced 399nm source in the B063 mixture experiment. Compared to the 1064nm optical trap in B063, the Verdi was nice to work with since it is in the visible frequency range. Although the power is much higher than the limit where the blink reflex can protect you, it is still comforting to know where and when the beam is on.

2.4.5 Frequency Stabilization Optics

Our experiment relies critically on controlling frequency difference, or detuning, between the laser frequency and the atomic resonance frequency. This section describes how we reference the 399nm and 556nm lasers to their respective atomic resonances

399nm lock

We lock the 399nm laser using Saturated Absorption Spectroscopy in a Hamamatsu L2783-70NE-YB hollow cathode lamp (HCL). Described in detail in section 8.3 of [12], saturated absorption is one of several techniques to overcome the problem of Doppler broadening of an atomic line during spectroscopy of a thermal vapor.

An atom with velocity \mathbf{v} relative to a laser with lab frame wavevector \mathbf{k} sees the laser frequency shifted by $\mathbf{k} \cdot \mathbf{v}$, see Fig. 2.9. When one looks at the absorption spectrum of a laser passing through an atomic vapor, this Doppler shift “smears out” the natural linewidth of the transition if $k\delta v > \Gamma$ where δv is the velocity spread of the vapor. Since we need to control the laser frequency at the level of the linewidth, this Doppler broadened spectrum

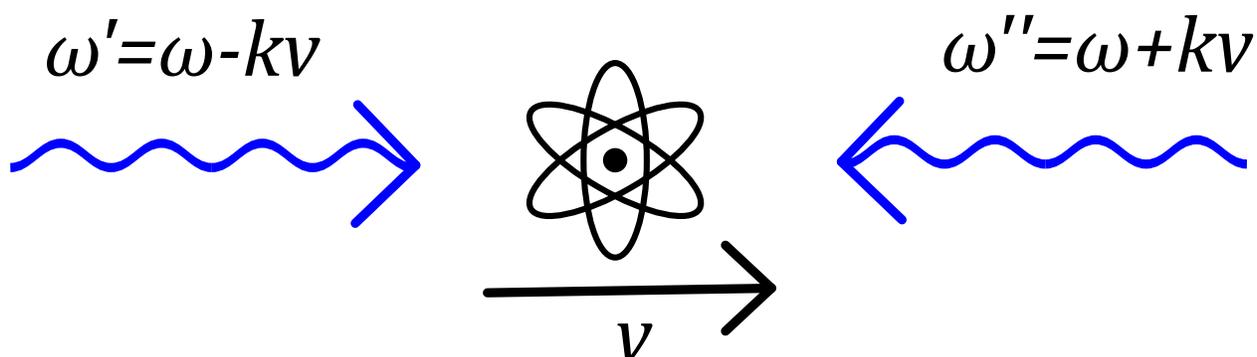


Figure 2.9: Doppler shift of an atom with velocity component v along the propagation axis of a laser with lab frame frequency ω

does not provide a narrow enough feature to lock the laser to.

Saturated absorption is one type of “Doppler free” spectroscopy. As shown in Fig. 2.10, we use two beams counter-propagating through an atomic vapor in the HCL. The weak probe beam is well below the saturation intensity of the transition, while the pump beam is above the saturation intensity. Recall that above the saturation intensity, the scattering rate does not increase with more intensity because the atom spends 50% of the time in the excited state. The scattering rate is limited by the time it takes the atom to decay back to the ground state, where it is available to absorb another photon.

First, consider the signal on the photodiode as the frequency of the probe is swept in the absence of the pump beam. In this case, we’d see the Doppler broadened transition. As the frequency is swept, the probe beam addresses different velocity classes in the vapor. When the probe is red detuned from the transition frequency of the lab frame atoms, due to the Doppler shift it interacts with atoms in the vapor that are moving counter to the beam. When it is blue detuned, it interacts with atoms moving with the beam.

Next, we’ll consider the effect of the pump beam. For the moment, we’ll say that the pump and probe have the same frequency ω . As before, when the frequency is swept, the beams interact with different velocity classes, but since the pump and probe are counter-

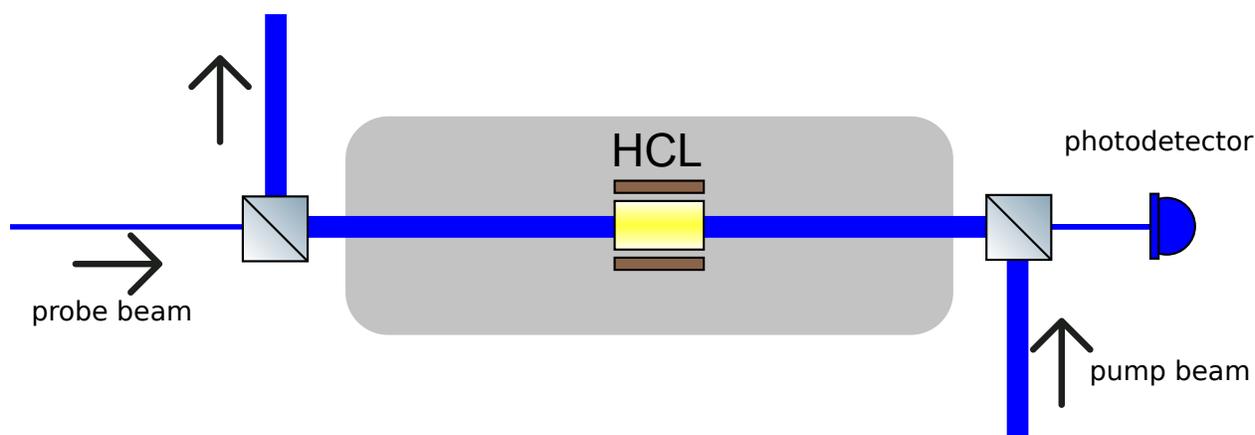


Figure 2.10: Optics schematic for the 399nm Saturated Absorption Spectroscopy. The cube on the left hand side of the figure is not present in the real setup, but is included here to help show the relationship between the pump beam and the probe beam.

propagating, in most cases they interact with different velocity classes. Referring to Fig. 2.10, when the laser is red detuned, the pump beam interacts with atoms moving left while the probe beam interacts with atoms moving right. When the laser is blue detuned, the directions above are reversed.

Now, consider what happens when the laser is on resonance. In this case, the pump and probe both address the atoms with no velocity component along the laser axis. Since the pump beam saturates the transition, the probe light that would otherwise be absorbed is allowed to pass. This effect appears as a dip in the absorption spectrum.

Our actual spectroscopy is slightly different in that the pump and probe are at different frequencies. In the example described above, the laser would end up being locked exactly at the resonance frequency for the lab frame atoms. For our purposes, it is more practical to have the laser frequency at some fixed offset from the transition.

To accomplish this, we frequency shift the pump beam with an AOM. The AOM is also needed as the means to modulate the frequency as part of the locking process, which we'll come to in a moment. When the pump and probe are at different frequencies, the velocity

class that they both interact with is shifted from 0 by half the frequency difference. Our spectroscopy AOM is typically set to shift the pump beam by +840MHz, which means that the laser frequency ends up -420MHz red detuned from the transition.

556nm lock

Over the life of the CI machine we used three different locking schemes for our 556nm laser, and are now preparing to move to a fourth.

The first scheme was to do saturated absorption, similar in principle to the 399nm lock described above. Because of the relatively weaker 556nm transition, the optical density of the atomic vapor in the HCL is not sufficient to do saturated absorption at 556nm. To get a higher density atomic vapor, we used a simple vacuum system referred to as the “heat pipe”. Working with Anupriya Jayakumar, a postdoc who worked for a time in our group, we designed and built a novel heat pipe designed to simultaneously lock the 399nm and 556nm transitions [21]. Unfortunately, after demonstrating the basic functionality of the design, the viewports of the heat pipe became coated with Yb, rendering the device non-functional. It is possible that the dual-axis heatpipe could be reconfigured to function, but we found it more expedient at the time to move on to other locking schemes. We moved the 399nm lock to the HCL as described above, and moved the 556nm to a beamline lock.

Crossed Beam Spectroscopy

Crossed beam spectroscopy is another technique to mitigate Doppler broadening. The key insight is that in a transversely collimated atomic beam there is a large spread of velocities along the axis of propagation, but the velocity spread transverse to the propagation is much smaller. By setting the spectroscopy laser beam perpendicular to the atomic beam, the spectroscopy is only sensitive to this narrowed velocity spread. For an ideally collimated atomic beam with zero transverse velocity, the linewidth would have no Doppler broadening.

In the context of the CI machine, a beamline lock is possible in the Kimball cube between the oven and the differential pumping region. In fact, one of the CI machine design considerations was to put the atomic beam shutter downstream from this chamber so that the atomic beam could close without blocking the atomic beam in the Kimball cube.

As discussed in the section on vacuum system design, the atomic beam is formed by two apertures. The second aperture is intentionally made bigger than the first, leading to a somewhat diverging atomic beam. This was done because at the time the machine was designed, we were contemplating doing transverse cooling in the Kimball cube region. We intentionally made a diverging atomic beam so that the transverse cooling could potentially “focus” the atomic beam to increase the flux in the MOT capture region.

Unfortunately the design configurations for increasing flux in the MOT region with transverse cooling are counter to the ideal design for beamline spectroscopy, which would be to narrow the transverse velocity of the atomic beam as much as possible. Nevertheless, we did set up and operate a beamline lock.

The optics for this locking scheme are relatively simple. A spectroscopy beam is derived from the spectroscopy AOM, shifted from the laser frequency by 400MHz. This beam is brought to the spectroscopy chamber, and intersected with the atomic beam at right angles. The fluorescence from the beam crossing is collected on a Thorlabs PMM02 PMT. The laser frequency is scanned, resulting in a fluorescence spectroscopy line shape. By adding modulation to the spectroscopy AOM and demodulating with the locking electronics described in the electronics section, an error signal is derived and locked to.

This is the lock with which we achieved our first MOT and BEC. However, we were plagued by instability in the number of atoms loaded into the ODT while using this lock. All parts of the new CI machine were in flux during this time period and so there were many suspects, but we eventually narrowed things down to the green lock. We found that the Doppler broadening from the diverging atomic beam was still too much to achieve the

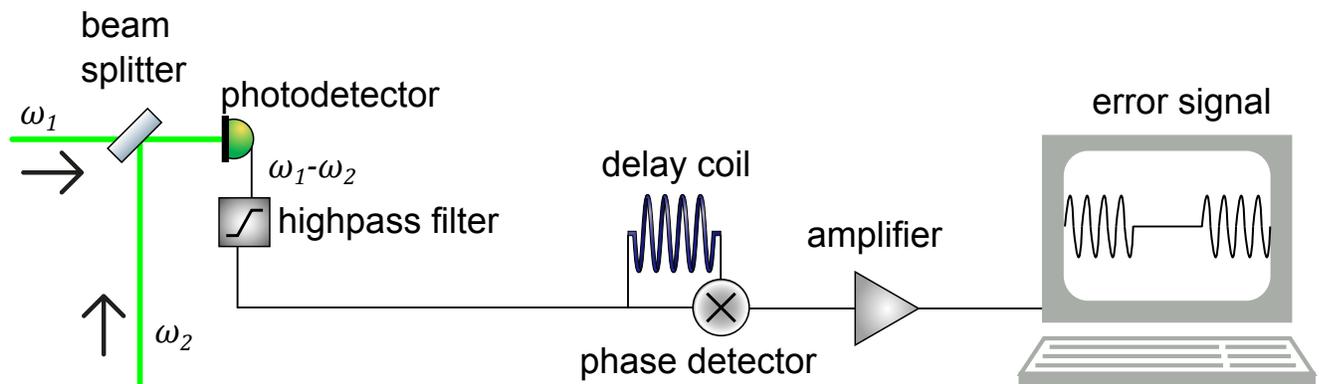


Figure 2.11: Schematic for the beat note lock.

locking stability that we required. Adjusting the atomic beam properties once the vacuum system is sealed up, baked, and at UHV is supremely impractical and undesirable. Therefore, we moved on to other solutions.

Beat Note Lock

At this point we again were faced with the option to revive the heat pipe apparatus and set up a saturated absorption lock, but since we were already at the point of making BECs and were impatient to continue working towards our interferometry goals, we wanted a faster solution. Since a well functioning saturated absorption setup for 556nm was already in place in B063 down the hall from us, we decided to reference the lasers to each other using a beat note lock.

A beat note lock is a way of referencing one laser to another. As the name implies, it is based on the beat note frequency $\omega_1 - \omega_2$ that arises when lasers with those frequencies are overlapped. If you have experience tuning a musical instrument, you are familiar with the sound of the beat note getting lower and disappearing as one string is brought perfectly in tune with another.

The beat note is created by combining two lasers on a beamsplitter and observing the

intensity of the combined beam on a photodiode, as shown in Fig. 2.11. The photodiode signal oscillates at the difference frequency $\omega_1 - \omega_2$. This oscillating signal is split into two parts using an RF power splitter. One part travels through a delay coil while the other takes a direct path, and the two parts are recombined on a phase detector. The speed of propagation through the BNC cables is about 8 inches/nanosecond, or $\frac{2}{3}c$. This means that, for example, if $\omega_1 - \omega_2 = 100\text{MHz}$, then a 40 inch path difference would give a $\frac{\pi}{2}$ phase shift. This phase shift depends on both the length of the delay coil and the frequency $\omega_1 - \omega_2$. So for a given fixed delay coil length, there is a frequency dependent phase shift between the two paths that feed into the phase detector. The phase detector output varies depending on the phase difference between its inputs. So if $\omega_1 - \omega_2$ is sweeping linearly, the phase detector oscillates sinusoidally. The frequency (for a given sweep rate of $\omega_1 - \omega_2$) of the phase detector signal depends on the length of the delay coil.

The phase detector signal can be used directly as an error signal to lock the laser frequency. Say that ω_1 is the fixed frequency of the B063 556nm laser that is locked to the transition using a saturated absorption setup. ω_2 is the frequency of the B055 556nm laser. If we zoom in on a single period of the phase detector signal, we can stabilize $\omega_1 - \omega_2$ by feeding back on the slope of the signal ω_2 . We can set the sign of the feedback to lock to positive or negative slopes. Because there are many such slopes, we can fix $\omega_1 - \omega_2$ to many different values. Practically, we'd like to have the same value of $\omega_1 - \omega_2$ each time we lock the laser, which means we need a way of picking the same slope each time. To aid this process, we add a high pass filter between the photodetector and the RF power splitter, which suppresses the signal until $\omega_1 - \omega_2$ reaches some finite value. We use a 50MHz high pass filter. The photodetector has a bandwidth of 200MHz, so the signal is attenuated unless the absolute value of $\omega_1 - \omega_2$ is between 50 and 200MHz. The example error signal in Fig. 2.11 shows the flat region where the high pass filter suppresses the signal. The edge of the suppressed region serves as a reference point for picking which slope to lock to.

A slight complication is that the B063 laser is 400MHz red detuned, while we'd like our laser to be 400MHz blue detuned, since it is desirable to have the diffraction beams blue detuned. In order to facilitate this, we use a double pass 400MHz AO to shift our laser frequency before beating it with the B063 light. This is the lock with which our main data sets were collected.

New beamline lock

As I write this thesis, a new beamline lock on a separate atomic beam apparatus is being built and tested by Katie McAlpine. This atomic beam is designed to be more collimated, to avoid the problems that we had with the CI machine atomic beam lock.

This new locking apparatus is desirable for two reasons. First, it is better not to be coupled to the B063 experiment. Although it has been about as problem free as could be expected, we from time to time have had issues where the fiber that couples the light between the labs gets bumped on the other side, leaving us suddenly lockless. There is also the concern that if there was any downtime of the B063 556nm laser, it would affect us as well.

More importantly, the new lock is designed to give an absolutely reliable offset from the 556nm transition frequency. While saturated absorption spectroscopy does give a lock close to the natural linewidth, the absolute value of the frequency that is locked to is not guaranteed to be the transition frequency due to the unknown magnetic fields in the locking region. It is difficult to precisely control these fields in a heat pipe because of the variac powered band heaters required to keep the pipe hot.

The new beamline has 3-axis Helmholtz coils to, on two axis cancel out stray magnetic fields, and on the third axis create a bias field to separate the 556nm excited state energy levels. By shifting the magnetically sensitive lines using the bias field, and choosing the light polarization to excite the magnetically insensitive π transition, the beamline lock can be

made insensitive to magnetic fields.

This is important for the precision measurement of α , since we need to know the exact wave vector k to determine $\frac{\hbar}{m}$ from ω_{rec} . To know k in absolute terms, we need a lock that gives us a reliable offset from the true 556nm transition frequency, combined with a value for the 556nm transition provided by the Ludlow Yb clock group[29].

2.4.6 Zeeman Slower Optics

The light for our Zeeman slower is derived from one of our two injection locks. The seed light for the injection lock is shifted -490MHz from the laser frequency by the slower AOM, giving a total detuning of -910MHz. The output from the slower injection lock passes through a pair of cylindrical lenses, to fix up the elliptical output from the injected diode, and a pair of plano-convex lenses to expand the beam. The lenses are arranged so that the slower beam is converging to a focus somewhere slightly behind the oven nozzle. The goal in shaping the slower beam is to match the shape of the atomic beam in order to maximize the number of atoms affected by the slower. Since the atomic beam is slightly diverging from the oven nozzle, this is the shape we aim for.

In earlier versions of the slower I made the initial slower beam spot size comparable to the size of the MOT, which is about 1cm diameter. The problem with this arrangement is that it becomes impossible to avoid having the slower beam pass through the MOT itself, which can adversely affect the trap. I experimented with making a hole in the slower beam by placing a small dot of paper supported by narrow wires in the path of the expanded slower beam, creating a ring like beam profile. The idea was that the slower beam would surround, but not hit, the MOT. I did have some success with this approach, but I found that the alignment was tricky and difficult to maintain. Eventually, I settled on a smaller initial slower beam size .3cm diameter with no hole, taking care to miss the MOT with the beam. This is the configuration with which our main datasets were taken.

2.4.7 Magneto Optic Trap Optics

The light from the 556nm laser is used for three main purposes: locking the laser, MOT light, and diffraction light. The small amount of light needed for the beat note lock is split from the main beam path with a polarizing beam cube. The remainder goes to the MOT AOM which is set up in double pass configuration. The zeroth-order diffraction path from the MOT AOM, i.e. the light that passes straight through, is picked off after the first pass and sent to the diffraction beam AOMs. The diffraction beams are never used at the same time as the MOT, so this allows us to use the entire laser power (minus spectroscopy needs) for either of the two.

The MOT light needs to be tunable in the range of 50 to 1 linewidths red detuned. Since the 556nm laser is blue detuned by 400MHz, the MOT light needs to be shifted down in frequency by around 410MHz. The second pass from the MOT AOM is picked off and directed towards the main chamber. It passes through a $\frac{\lambda}{2}$ waveplate and polarizing beam cube to stabilize the polarization, since we have observed some temperature sensitive birefringence of the MOT AOM. It passes through a Uniblitz shutter, so that the MOT light can be entirely blocked from reaching the main chamber during evaporative cooling and interferometry. This is a critical consideration during the final stages of evaporative cooling towards the BEC transition, where even microwatts of near resonant light can heat the atoms in the ODT.

The beam passes through a pair of cylindrical lenses to circularize a slight ellipticity in the laser output. After the cylindrical telescope, the beam profile is very close to 1mm $\frac{1}{e^2}$ radius waist.

The beam is then coupled into Thorlabs PM-S405-XP single mode, polarization maintaining optical fiber. Earlier iterations of the experiment had the MOT free space coupled to the main chamber region. We moved away from that setup because alignment of the MOT beams with respect to the main chamber had to be redone every time the MOT AO was

adjusted. This realignment procedure proved to be unreliable, leading to a large amount of frustration as the performance of the MOT would vary wildly from week to week.

By fiber coupling the MOT, we decouple the AO alignment from the MOT beam alignment. A further advantage is that the fiber outputs a clean Gaussian mode, leading to more uniform MOT beams. After switching to the fiberized MOT, I did not find it necessary to realign the MOT beams for nearly a year while taking nearly all the data central to this thesis, which is a huge improvement over the weekly realignments that used to be necessary.

I have to admit here that I was a little stubborn in how long it took me to implement this solution. I knew that B063 had fiberized their MOT beams previously, to great success. I had achieved excellent MOT performance in the past without the fiber, and I overestimated how disruptive adding a fiber to our setup would be, while underestimating how much time we were losing to constantly realigning the MOT. I add this as a caution to future students: if a known better solution exists, don't be afraid to use it even if it means a bit of construction and tinkering (within reason! Don't break vacuum or anything).

After the MOT fiber, the beam goes through a 2:1 telescope, and then a $\frac{\lambda}{2}$ waveplate and cube divide the vertical MOT beam from the horizontal beams. A telescope brings the horizontal beams up to the main chamber breadboard, where a second cube again divides the beam. All three beams pass through a $\frac{\lambda}{4}$ waveplate, changing the beams to their appropriate circular polarizations, and then a 3:1 telescope. I used individual telescopes for each beam because I didn't want the expanded beam to be clipped while passing through the beam-splitting cubes. After the beam is expanded, 2" optics are used to prevent aperturing. The retro mirrors have $\frac{\lambda}{4}$ waveplates integrated into the optomechanic mirror holder, to reverse the circular polarization as necessary to create a MOT.

2.4.8 Diffraction Beam Optics

The diffraction beam AOM paths went through several iterations as we tested out different diffraction schemes, as described in the next chapter. I'm only going to describe in detail the optical setup used for our main datasets.

As mentioned above, the light for the diffraction beams is taken from the zeroth order of the MOT AOM. This beam passes through a polarizing cube to turn any polarization fluctuation (due to the temperature dependent birefringence of the MOT AOM crystal) into intensity fluctuation, which can be removed downstream with intensity stabilization. After the cleanup cube, it passes through $\frac{\lambda}{2}$ waveplate and then a splitting cube, to divide the beam into the two diffraction paths.

Each diffraction path has a single pass through a Brimrose 200 MHz AOM. We used a single pass because an earlier diffraction scheme called for adding two frequency components to a single beam. This was the last setup we used before moving to our most successful scheme, so once things were working with the single pass AOs we stuck with them. In the future, it would be advisable to move to double pass, to avoid a frequency dependent pointing change of the 1st order diffracted beam. We tested that the pointing change does not affect the post-AOM fiber coupling (described below) by more than 10% until the frequency is changed from its central value by 1 MHz, which is within the range that can be compensated for with our intensity stabilization.

After the AOMs, each diffraction beam is coupled into a single mode polarization maintaining optical fiber, which brings the light to the main chamber breadboards. One fiber goes to each of the two breadboards. The fiber output couplers are Thorlabs TC25APC-543 FC/APC Triplet Fiber Collimation Packages with EFL = 25 mm. These collimators output a very clean Gaussian mode at a beam waist that is acceptably large for our diffraction beams, avoiding any additional beam shaping optics after the fiber. The diffraction beams pass through the center of the main chamber, counterpropagating from either side, and are

cross coupled into each other. We achieve cross coupling efficiency of 25%. Cross coupling the fibers in this manner insures a high degree of beam alignment.

The original diffraction fibers were FC/PC on the AOM side, and FC/APC on the main chamber side. Once these fibers were cross coupled, it was discovered that the two FC/PC sides formed an optical cavity. The cavity interference seriously degraded the diffraction beam performance. Our diffraction schemes assumed that a given frequency and polarization would come from one fiber through the chamber, and not reflect back. In the presence of these back reflections from the FC/PC side of the far fiber, this assumption broke down, leading to diffraction order chaos. We temporarily fixed the problem by de-cross-coupling the fibers. Later, we switched to fibers that are FC/APC on both sides. This allows the fibers to be cross coupled without leading to the optical cavity effect.

The diffraction beams enter the main chamber at a slight angle to the normal vector from the viewport. This is to prevent viewport reflections from impinging on the atoms, leading to the kind of unintentional cavitation described above. We collect the small viewport reflection on a photodiode to monitor and stabilize the diffraction beam power. By using the viewport reflection, the last optical element before the laser enters the vacuum, we are monitoring the diffraction beam power at the closest possible spot to the atoms.

2.4.9 Optical Dipole Trap Optics

The optical dipole trap optics are an area where we desired the utmost in stability. Additionally, the 18W Verdi laser requires a planned and cautious approach to optics placements. In other parts of the table, we indulge in tightly packed optical elements, and at times creative beam paths in order to fit everything together. In the Verdi region, I marked off a large area so that optics could be placed far apart and in their most logical locations. The Verdi laser head and AOM optics are all in a solid plastic box with 80-20 edges. The box, built by Dan, is designed to support the weight of a moderately sized graduate student. This has

proved to be a huge boon, as it allows access to the far side of the main chamber breadboard by sitting atop the box. Although it can be very useful, please exercise extreme caution if attempting this maneuver.

Most of the mirrors in the Verdi box are 532nm laser line mirrors from Melles-Griot which are designed for 45 degree angle of incidence, which necessitates the laser path to always turn at right angles. The beam passes through a $\frac{\lambda}{2}$ waveplate and a polarizing beam splitter, and is divided into two paths, one for the vertical and one for the horizontal ODT beam. Each beam passes through an 80MHz AOM. One beam takes the +1 order, and the other the -1 order, so that the two beams are at different frequencies. This prevents the two ODT beams from interfering with each other at the crossing point where atoms are trapped. One AO is oriented vertically, the other horizontally, to facilitate the “painting” of the optical trap, described later.

The zeroth order from the AOMs is directed into a high power beam dump. When the 1st order diffraction is maximized, there is also enough power in the 2nd order to melt the walls of the plastic box, so the 2nd orders are also picked off and directed to the beam dump.

The horizontal ODT beam is expanded with a two lens telescope and brought via a periscope to the main chamber breadboard. It is then focused to a $22\mu m$ waist at the atoms. The horizontal ODT beam passes through the same viewports as the diffraction beams, but it goes through at as much of an angle as possible in order to keep the viewport as clear as possible for other beams and lenses.

The vertical ODT beam comes up from below, crossing the horizontal beam near the MOT location. I made an effort to align the painting axis of the vertical beam with the angle of the horizontal beam, so that as the vertical beam is painted it moves along the axial direction of the horizontal beam, rather than moving radial and uncrossing the beams. The focus of the vertical beam is calculated to be about $60\mu m$, but it has proven difficult to experimentally confirm this number. Since the role of the vertical beam is to provide

confinement along the axial direction of the horizontal beam, the exact value of the waist is not critical.

2.4.10 Imaging Optics

The 399nm laser is about 420MHz red detuned, so the imaging AOM needs to shift the frequency up by that amount to create the resonant light needed for absorption imaging. After the double pass 200MHz AOM, the imaging light is coupled into a single mode fiber. At one point, we achieved 50% coupling efficiency into this fiber, which is the highest I've seen for this wavelength of light. Unfortunately, sometime after that the fiber tip, which was AR coated from Thorlabs, became terribly scratched, and the coupling efficiency went way down to 25%.

The fiber output coupler has an attached lens tube with integrated lenses to expand and collimate the beam. The imaging light can be directed to a vertical or horizontal imaging path with the placement or removal of a mirror on a magnetic mount. Horizontal imaging conflicts with the diffraction beams, so while doing interferometry only vertical imaging is used.

After passing through the chamber, the vertical and horizontal paths converge on an upper breadboard above the Zeeman slower. A 150mm achromat lens images the atoms, and a second lens re-images the atomic shadow onto the CCD camera. For vertical imaging, there are two choices for the second lens leading to two different magnifications.

2.4.11 Readout Light Optics

The optics for collecting the readout light reflected from the atoms were a difficult and low performing aspect of the first Yb contrast interferometer. However, lessons learned in the old setup helped to guide the design of the current apparatus. Some of the issues with the old setup are related to the geometry needed for green readout, and so are not immediately

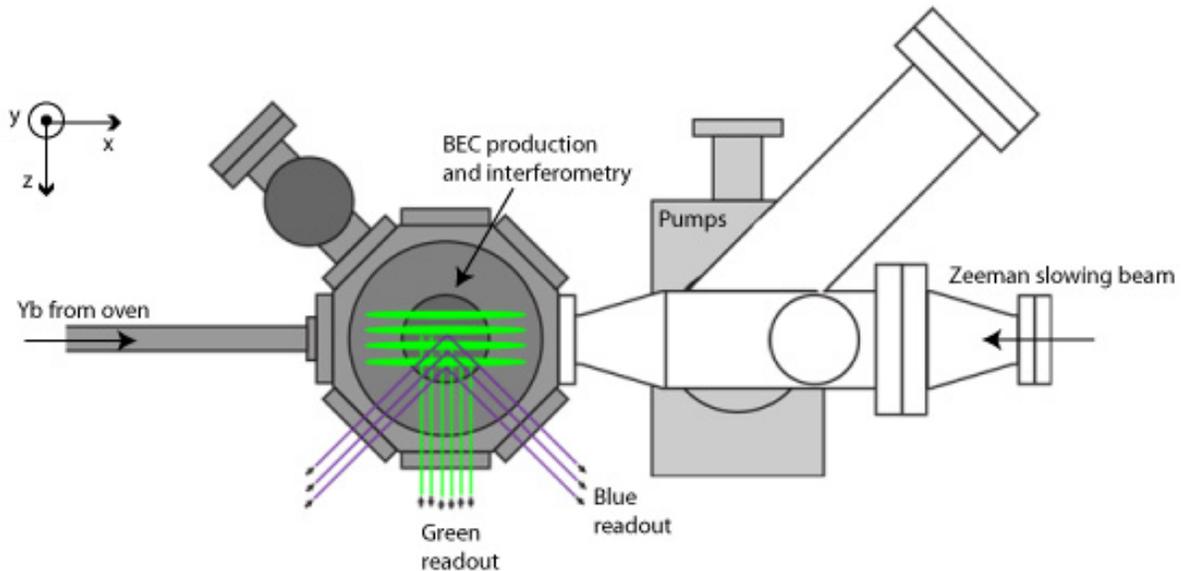


Figure 2.12: Blue and Green readout geometry

relevant to the current experiment which uses blue readout.

When using the green readout, which has a 180 degree back reflection off the atoms, great care must be taken to separate the returning signal from the reflections and scatter from every optic that the incoming beam passes through. When using blue readout, which reflects at a 44 degree angle, this issue is removed entirely.

The CI machine blue readout optics are fully integrated inside a lens tube system, as shown in Fig. 2.13. The readout light is collected by a 2" diameter $f=75\text{mm}$ achromat lens located a distance $2f$ from the atoms. Light from the atoms forms an image $2f$ on the other side of the lens. At the imaging plane, there is an iris on an x-y translation stage inside the lens tube. The translation stage had to be modified, and the lens tube a slit cut into it, to make this possible, see Fig. 2.14. After the iris, a 30mm lens collimates the rays from the image. Next comes a lens to integrated adjustable mirror assembly. The mirror is elliptical,

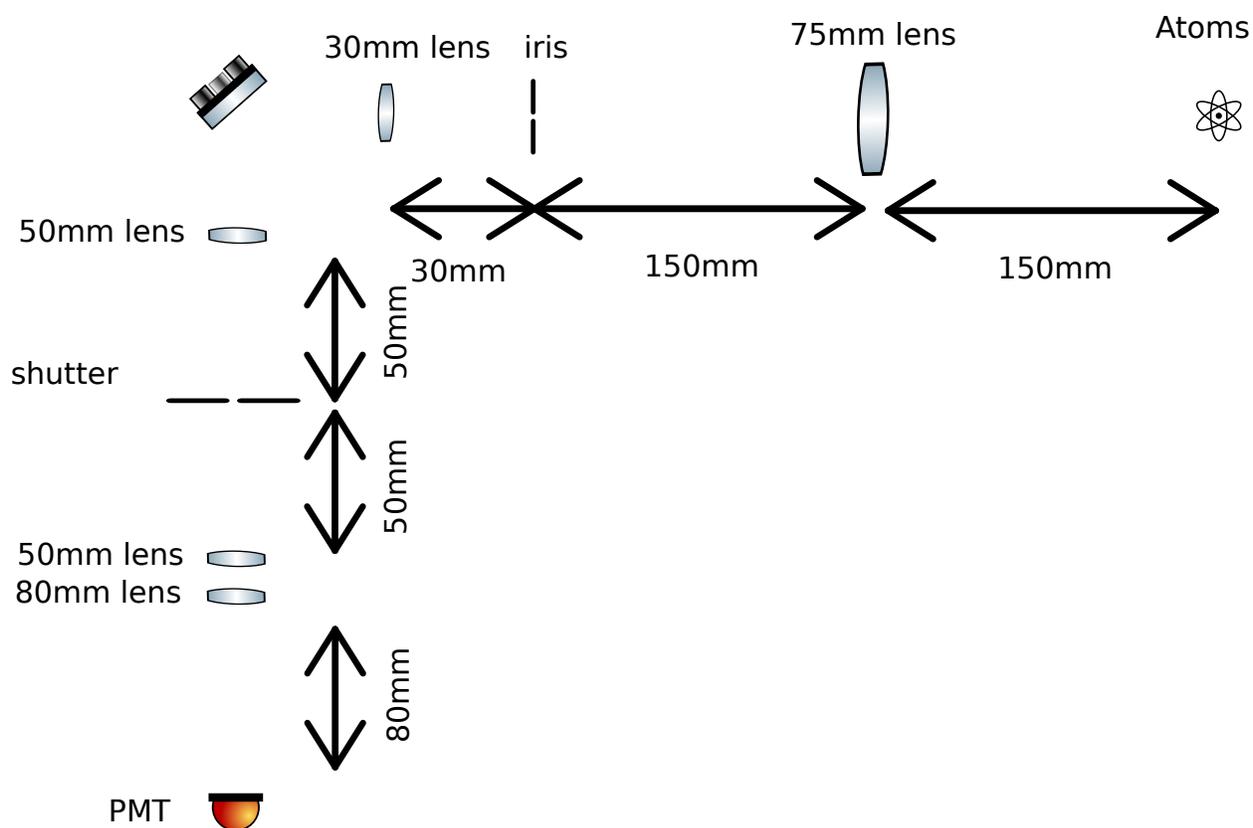
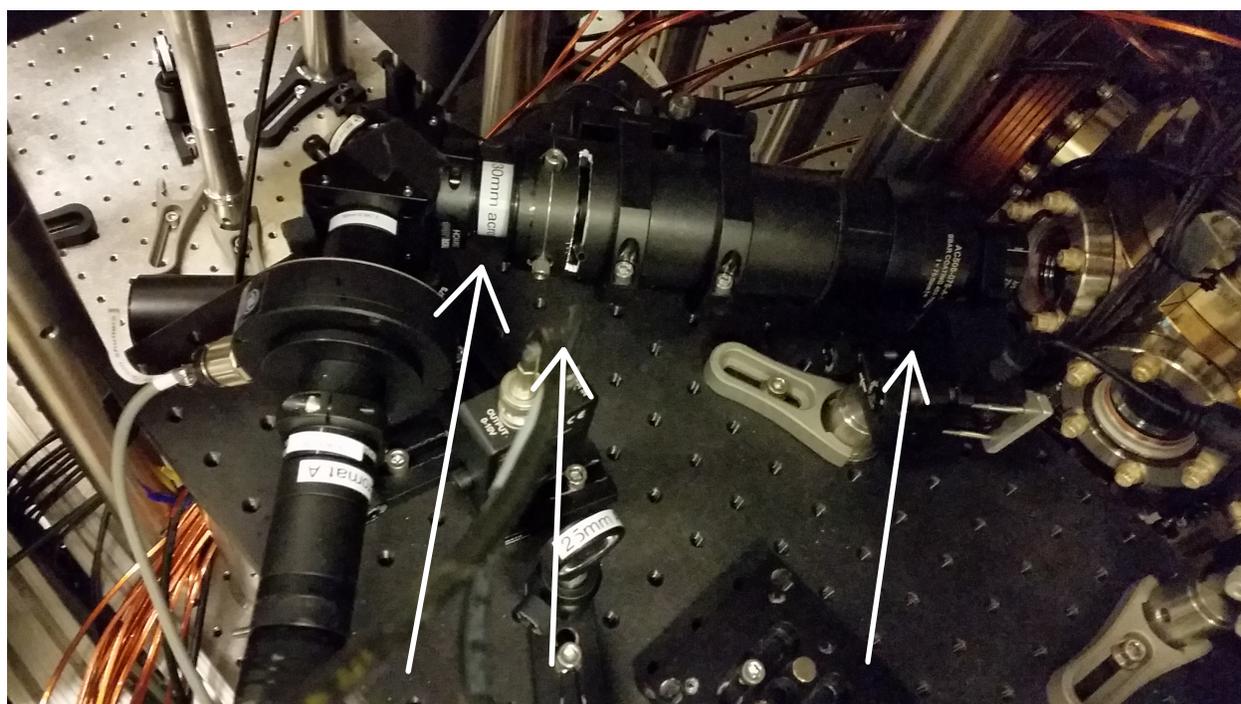


Figure 2.13: Top: photo of the lens tubes containing the readout light collection optics. Bottom: Schematic of the optical elements. The iris size is adjustable, as is its location in the x-y plane (with z being the propagation direction of the readout light). The translation stage iris mount is itself mounted inside a 2" lens tube. This configuration of iris mount and lens tube required some creative modifications of the lens tube, see 2.14



Figure 2.14: The slightly anthropomorphic iris translation stage mounted in a 2" lens tube, via some band saw modifications to the lens tube. This contraption is part of the readout light collection optics.

to avoid creating a smaller than needed aperture. After the mirror, a 1-to-1 telescope made with 50mm lenses brings the rays to a focus so that they can pass through a Uniblitz shutter. Finally, an 80mm lens images the rays onto the PMT.

For the purposes of collecting readout light, it is not important that the light be imaged onto the PMT, the light just needs to get there somehow whether it is an image or not. However, the easiest way to align the readout path is do absorption imaging. We swap the PMT for our CCD camera (another advantage of the Mako camera is that it is small and easy to move around), and use a MOT beam for green imaging along the readout axis.

The entire readout optics assembly is mounted on a linear stage that adjusts the distance from the 75mm lens to the atoms. The height and horizontal displacement are adjustable by loosening screws attaching the assembly to the stage. The process of the initial alignment of the absorption imaging while using these adjustment requires a very fine touch, and likely some trial and error. After the image is found, the gross adjustment screws are tightened down, and the translation stage is used to focus the image. The elliptical mirror is used to center the image, and the iris is used as a confocal filter to block any source of light other than the atoms.

While adjusting the translation stage, the elliptical mirror optomechanic base is left to slide along the table. After the translation adjustment is finished, the system is locked down by screwing the elliptical mirror to the breadboard.

Chapter 3

RUNNING THE MACHINE

In the previous chapter, I've described in detail the design and construction of the CI machine. In this chapter, I'll go through the process of bringing the experiment online: trapping and cooling the atoms to form BEC, and configuring the laser pulse atom optics needed to create the Contrast Interferometer.

3.1 Trapping and Cooling

3.1.1 Zeeman Slowing

Ytterbium is metallic at room temperature and pressure, and remains so as the pressure is reduced to UHV. In order to create a vapor of Yb atoms of sufficient density, the metal must be heated to around 360C. The design of the oven, where the Yb is heated, and the nozzle apertures that allow the Yb vapor to escape in the form of an atomic beam, are described in 2.1.2.

The longitudinal velocity distribution in the atomic beam is given by $v^3 \exp\left(\frac{-v^2}{2k_b T/M}\right)$ [12] table 8.1. This expression has an extra factor of v compared to the Boltzmann velocity distribution in a gas. This is because higher velocity atoms in the oven have more collisions with the wall, and so are likelier to hit the exit aperture and become part of the atomic beam. For our typical operating oven temperatures, the peak of this distribution is around 300m/s.

The MOT has a capture velocity of around 5 m/s (calculated by integrating the MOT scattering force over the distance of the capture region). The atomic beam velocity distri-

bution has very few atoms below 5m/s, which makes it difficult to load a MOT from the un-modified atomic beam.

In order to slow the atoms in the atomic beam, we want to take advantage of the scattering force of the 399nm light.

$$F_{\text{scatt}} = \hbar k \frac{\Gamma}{2} \frac{I/I_{\text{sat}}}{1 + I/I_{\text{sat}} + 4\delta^2/\Gamma^2} \quad (3.1)$$

Where I is the intensity of the light, I_{sat} is the saturation intensity of the transition, and $\delta = \omega - \omega_0 + kv$ is the detuning from resonance, taking into account the Doppler shift.

In order to slow the atoms most effectively, we'd like to have $\delta = 0$, keeping the scattering force at its maximum value $F_{\text{max}} = \hbar k \frac{\Gamma}{2}$. Clearly, the velocity dependance in δ means that we can't maintain $\delta = 0$ as the atoms change velocity without taking further measures. How much of a problem is this? Recall that for the 399nm transition, $\Gamma_b = 2\pi \times 28$ MHz. $kv = \Gamma_b$ when $v = 11.2m/s$, so reducing the velocity from 300 to 0 m/s is changing the detuning by over 20 linewidths.

We compensate for the Doppler shift by applying a position depending magnetic field. The Zeeman effect adds another term to the detuning, so $\delta = \omega - \omega_0 + kv - \frac{\mu_b B(z)}{\hbar}$. Considering an initial velocity v_0 and a distance L_0 over which the field is applied, we work backwards from the assumption that $F_{\text{scat}} = \frac{F_{\text{max}}}{2}$ at all times to find the desired magnetic field profile:

$$B(z) = \frac{\hbar k v_0}{\mu_b} (\sqrt{1 - z/L_0} - 1) \quad (3.2)$$

For our experiment, we use a so called increasing field slower. This means that in the field profile above, we set the slower laser frequency so that $\delta = 0$ for zero magnetic field for a velocity v_0 , called the capture velocity. We set v_0 to be near or slightly above the peak velocity class in the atomic beam. We used the increasing field design so that the slower light would be as far detuned from the zero field lab frame resonance as possible, in order to minimize the effect of the slower beam on the MOT. In order to address the correct velocity

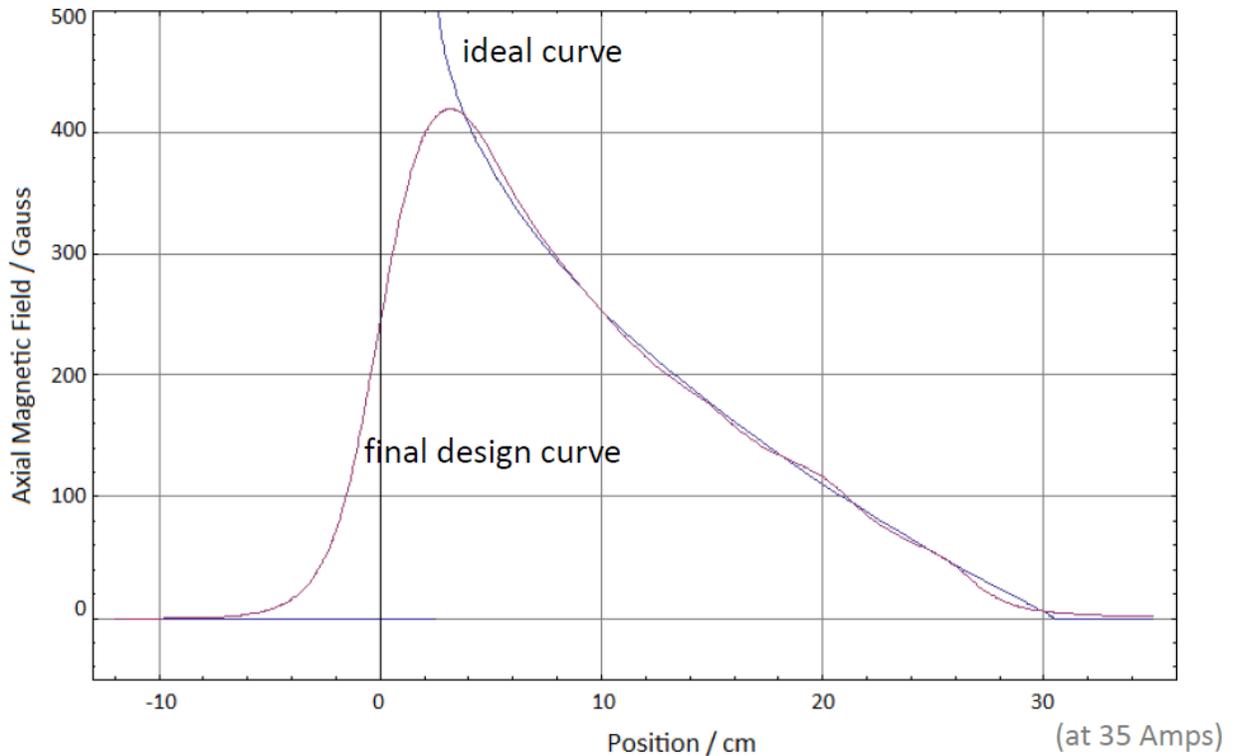


Figure 3.1: The ideal magnetic field profile, along with the predicted design curve

class, we use $\sigma+$ light for the slower beam.

The ideal field profile is shown in Fig. 3.1, along with the actual field based on the design model for our slower coils. The design model is compared to measured field values in Fig. 3.2.

Our Zeeman coils were designed by Dan Gochner as his REU summer project the year before he joined our lab as a graduate student. Dan and I worked together building the coils. The figures and measurements of the completed coils are from Dan's final REU presentation.

The Zeeman slower is composed of 3 separate coils: Offset coils, increasing field coils, and reverse coils. The use of the offset coils is described below. The reverse coils are used

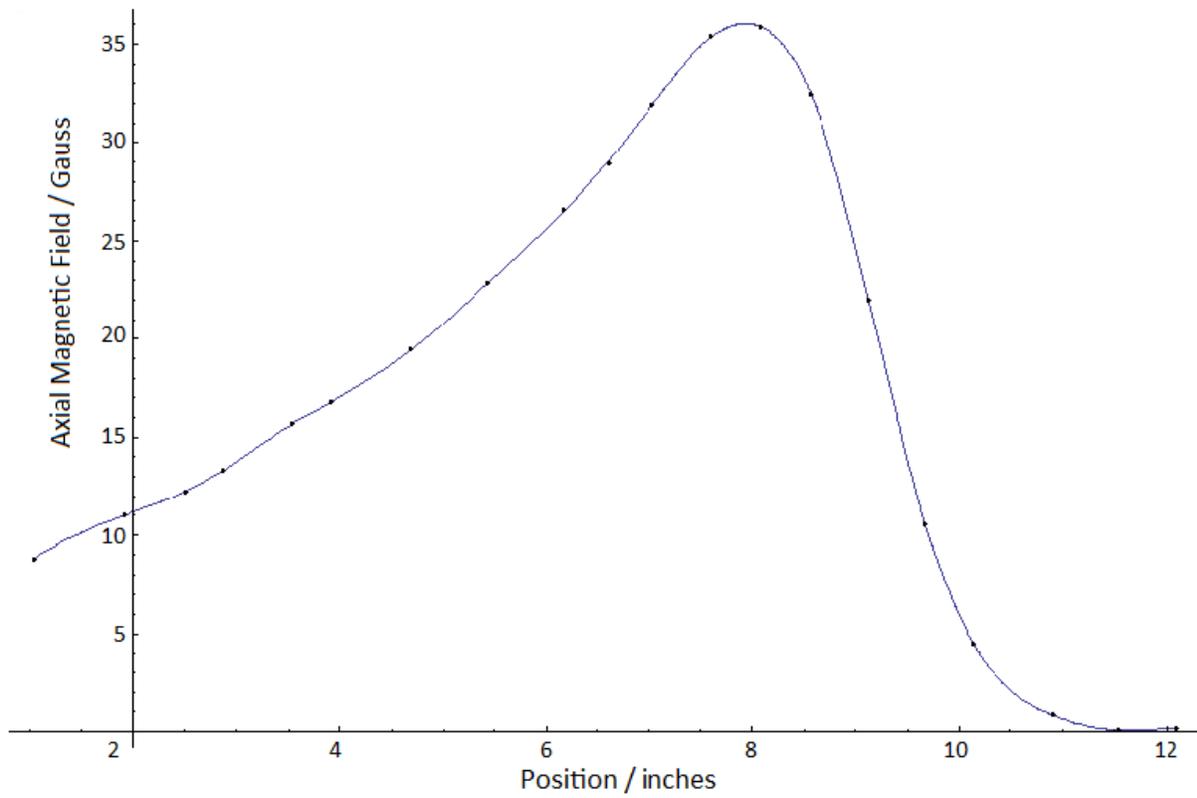


Figure 3.2: The measured field of the increasing field and reverse slower sections at 3A, small dots, compared to the predicted field, solid line.

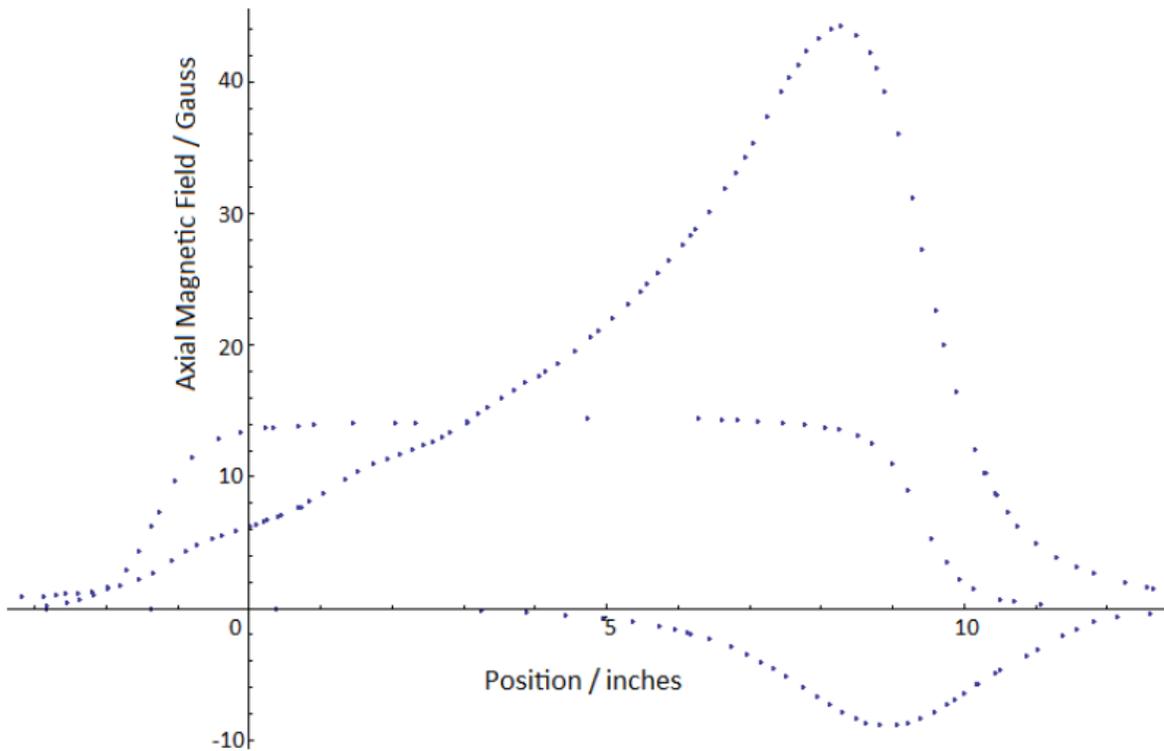


Figure 3.3: The measured field for the three separate slower sections at 3A.

to minimize both the magnetic field and its first derivative at the location of the MOT. In practice, increasing field and reverse coils are connected in series, since the reverse coils are designed to cancel the field when at the same current as the increasing field coils. The current of the offset coils is independently adjustable.

In order to quantify the performance of the slower, we set up a Doppler sensitive fluorescence spectroscopy. The frequency of the 556nm probe beam is swept near the atomic resonance, while the slower beam is held at a fixed frequency. The principle is similar to the crossed beam spectroscopy described in 2.4.5, but instead of a probe beam perpendicular to the atomic beam, we use a probe beam at 45 degree angle. In this case, the effective detuning of the laser beam depends on the atomic beam longitudinal velocity.

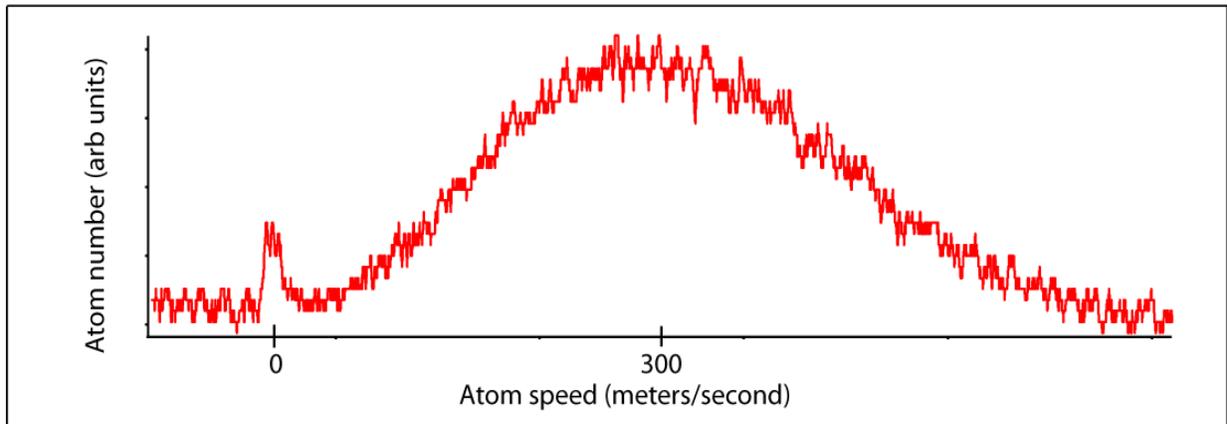


Figure 3.4: Doppler sensitive fluorescence spectrum of the atomic beam as it passes through the MOT region with no slower light presence. The sharp peak on the left is from a perpendicular, hence Doppler insensitive probe beam. The broad peak is the Doppler sensitive signal, centered at 300m/s.

Fig. 3.4 shows the Doppler sensitive fluorescence spectrum of the atomic beam with no slower light present. There is also a Doppler insensitive peak from a probe beam perpendicular to the atomic beam, hence insensitive to the longitudinal velocity of the atomic beam. The Doppler insensitive peak indicates where on the horizontal axis corresponds to lab frame atoms. A correctly operating slower would increase the number of atoms near the lab frame velocity.

Fig. 3.5 shows what happens when the slower light is turned on without any B fields present. The atoms near 300m/s are slowed, but quickly go out of resonance with the slower beam, leading to a bump in the velocity distribution below 300/ms.

Next, we turn on the increasing field coils, with the result shown in Fig. 3.6. Our goal is to get as many atoms as possible in the range 0 to 5m/s, so this figure demonstrates a successful slower.

However, we can do a little bit better. By using the offset coils, we can take advantage

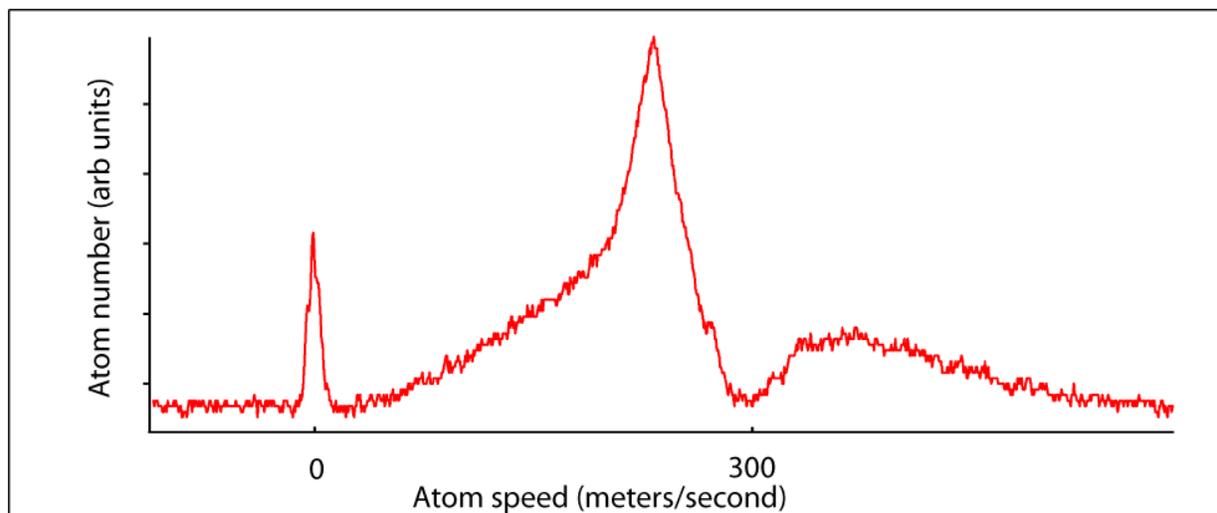


Figure 3.5: Spectroscopy signal with slower light on, but no magnetic fields. Atoms are slowed until they go out of resonance with the slower light.

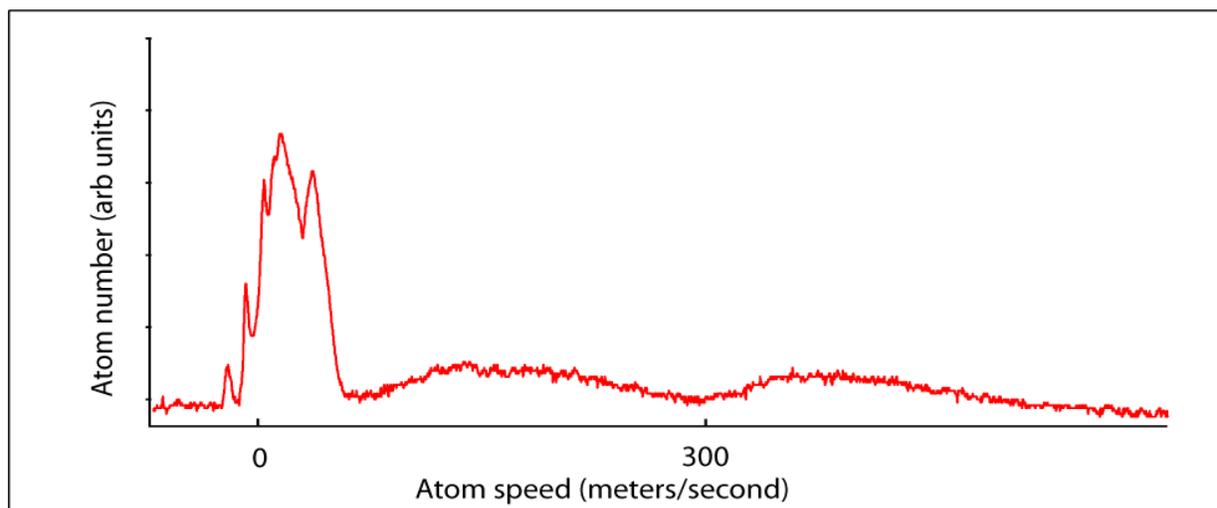


Figure 3.6: Velocity profile with the slower light and increasing field coils, but no offsets.

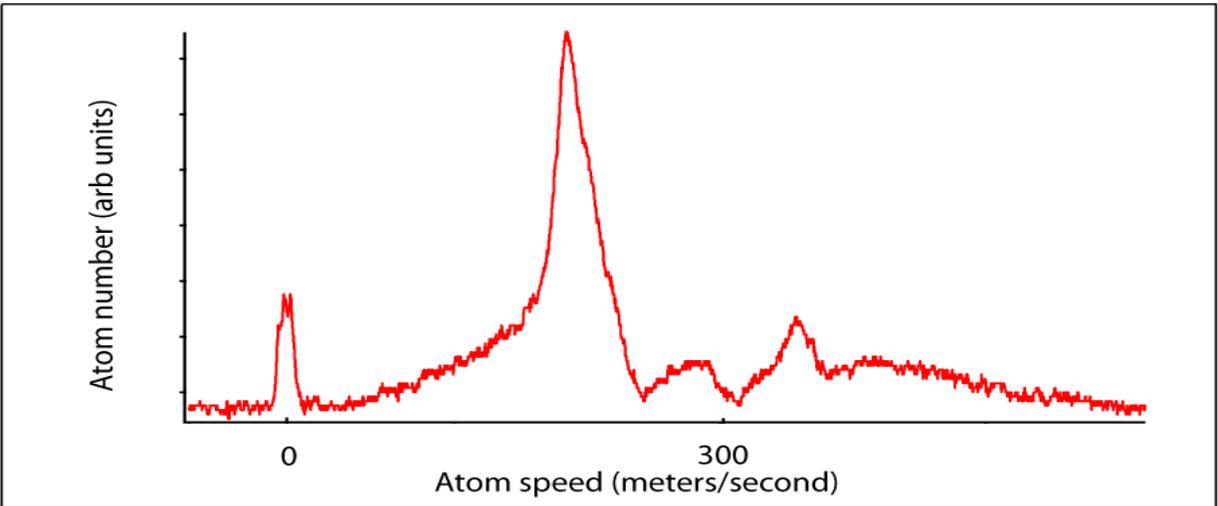


Figure 3.7: Velocity profile for slower light and offset coils, no increasing field.

of the region between the oven nozzle and the beginning of the Zeeman slower to do a first round of slowing that is unassisted by the magnetic fields. Fig. 3.7 shows the velocity profile with slower light on and current in the offset coils, but not in the increasing field coils. The atoms are slowed once in the region before the slower, and again in the constant B field of the offset coils.

Finally, we can turn on the increasing field coils in addition to the offset coils, as shown in Fig. 3.8. We observe a larger peak near $v = 0$ compared to Fig. 3.6, indicating better slower performance with the offset coils active, compared to without them.

3.1.2 Slower viewport clouding

As mentioned in the Apparatus chapter, the viewport through which the slower light enters the chamber is prone to becoming coated on the inside surface by Ytterbium from the atomic beam. To mitigate this, we keep the slower viewport heated so the atomic beam mostly bounces off. The viewport AR coating has a max temperature of 200C, so we initially

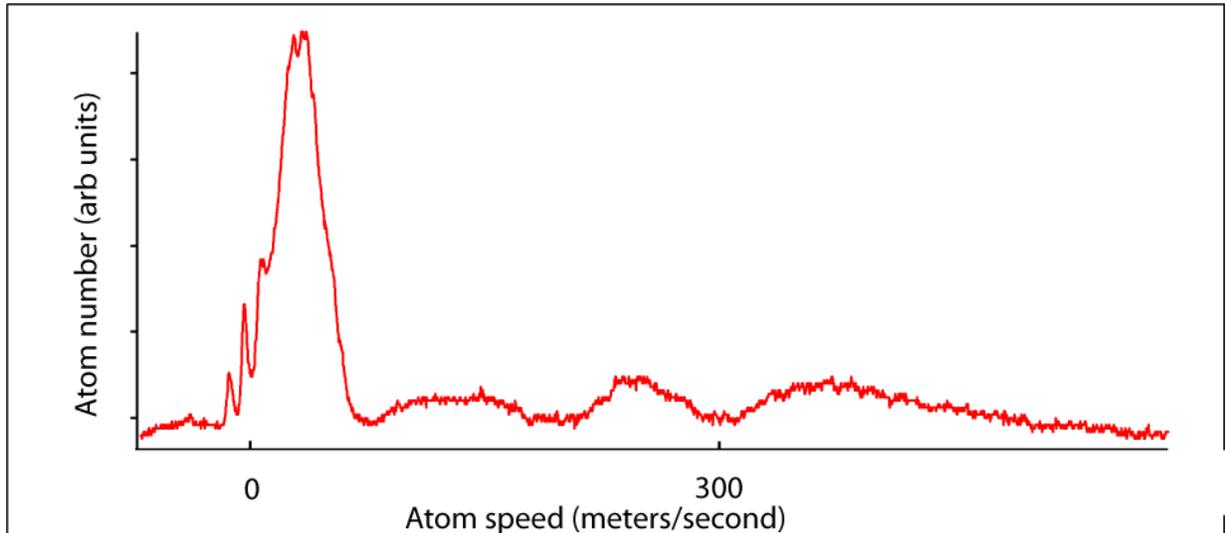


Figure 3.8: Velocity profile for the slower light with offset coils and increasing fields. The larger peak near $v = 0$ compared to Fig. 3.6 demonstrates the utility of the offset coils.

kept the viewport around 150C to keep the coating safe. This turned out to be too low a temperature to keep Ytterbium from collecting on the viewport. When I realized that this had become an issue, with the symptom being reduced slower performance, I put into use a piece of lab lore: cleaning a viewport using the high power Verdi 532nm laser.

This is a technique which was brought to our attention by the Takahashi lab in Japan, by way of Deep. Whenever this method was discussed, it would be described as using the laser “like a firehose”, accompanied by a hand movement as if spraying the laser in little circles.

I decided that it would be unwise to wield the Verdi like a firehose, so I set up a more controlled system. I redirected the horizontal ODT path onto the slower viewport bread-board, and set up the final mirror before the beam entered the slower viewport so that I could aim the beam by adjusting the mirror knobs. I experimented with focusing the beam more or less tightly on the viewport, and found that a quite small waist was required in order to remove the coating, something on the order of $50\mu\text{m}$.

Cleaning the viewport with this technique was like using a very dangerous and expensive Etch A Sketch. The results of the cleaning are shown in Fig. 3.9. As the laser was scanned, it cleared a skinny path through the coating, so to clean the viewport I had to tediously scan the laser back and forth in a raster pattern. I only cleared the upper half of the smudge, since that is the area that the slower beam passes through (due to irregularities in the custom welding job of the cross, it angles slightly downward, so the slower beam enters high on the slower viewport when the beam centered on the rest of the CI machine).

A word of caution! This process came to an abrupt end when, despite extensive cleaning before I started, I hit a dust spot, or a pre-existing imperfection in the AR coating. Whatever it was, the laser quickly heated it, creating a larger spot, visible in Fig. 3.9 just above the uncleaned smudge, left of center. Luckily, this happened towards the bottom of the area that I intended to clean, below the area that the slower beam hits, so this new spot does not affect performance.

After this cleaning, the slower viewport temp was raised to 200C, and since then the viewport has not been re-clouding. If the Verdi firehose trick is ever needed again, great care should be taken to make sure the viewport is clean.

3.1.3 Transverse Cooling

Before we moved to the first version of the beamline lock (see section 2.4.5), we tested out a transverse cooling scheme in the Kimball cube between the oven and the differential pumping region. As described in section 2.1.2, the atomic beam is designed to be slightly diverging. The goal of transverse cooling is to focus the atomic beam to increase the flux through the MOT region.

We directed some red detuned near resonant 399nm light to the Kimball cube, and created a 2D optical molasses. The blue beam was “recycled” to make many passes through the atomic beam. We did observe an increase in atomic flux of around 50%. However,

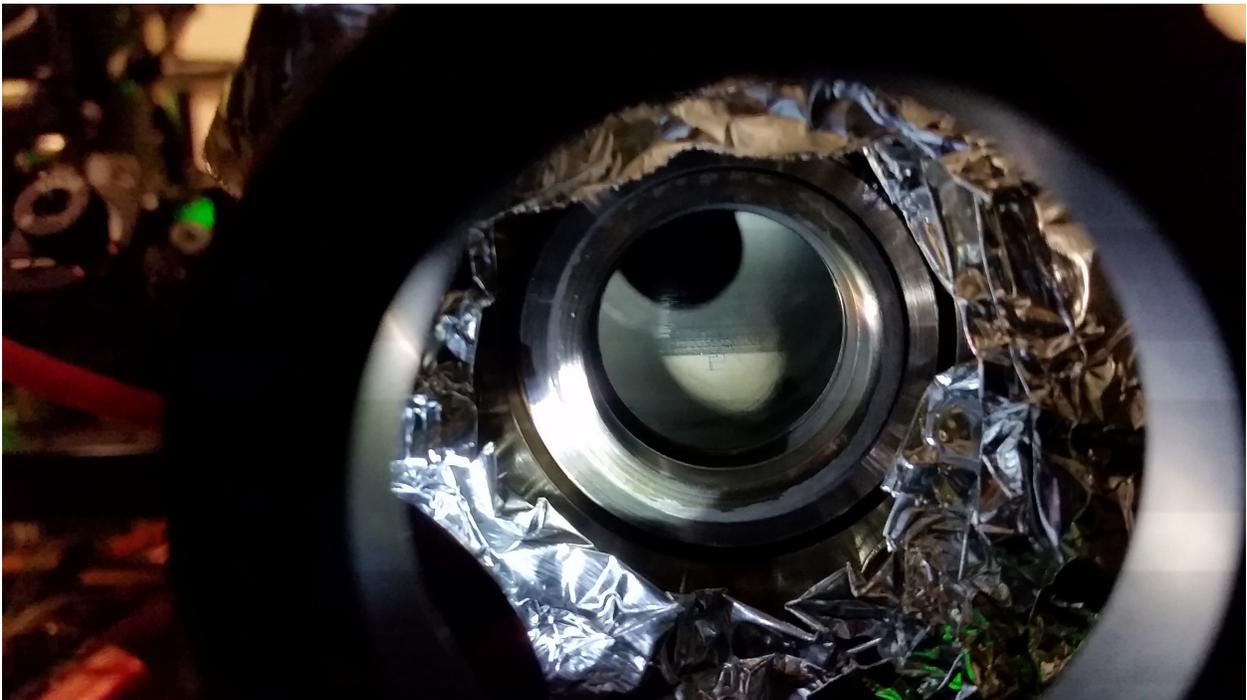


Figure 3.9: Ytterbium coating the inner surface of the slower viewport. The upper half of the Ytterbium smudge has been cleaned by focusing onto the viewport a beam with 10W of 532nm light with a $50\mu\text{m}$ waist. The beam was scanned by hand using a mirror knob, removing the Ytterbium coating in a manner very reminiscent of an Etch A Sketch. The Etch A Sketch pattern is especially clear in the lower half of the smudge, where I made some lines and corners while learning how best to proceed.

when we decided to start using a beamline lock for the 556nm laser, we had to remove the transverse cooling light to make room. We have since moved on to other 556nm locking methods, so we could reconstitute the transverse cooling setup. However, we have since found another use for the former transverse cooling light, described in the next section.

3.1.4 *Crossed beam Slowing*

A major limitation of the 556nm MOT is the low capture velocity. A 399nm MOT would have a much higher capture velocity, allowing many more atoms to be captured, but the Doppler temperature ($k_b T_D = \frac{\hbar\Gamma}{2}$) would be too high to load the ODT with sufficiently high phase space density to facilitate evaporation to BEC. Jeongwon et al [23] demonstrate a scheme that combines the advantages of the blue and green MOTs. Their “core-shell” MOT has standard 556nm MOT beams surrounded by a “shell” of 399nm light. The MOT beam profile is a green dot surrounded by a blue ring. Atoms are captured in the blue region, but as they cool they eventually are slow enough to be confined to the central green region, where they can be cooled to the 556nm Doppler temperature.

We considered re-creating the core-shell MOT for the CI machine, but eventually made the judgement that the technical challenges and costs were not worth it. For instance, obtaining quarter waveplates active at both 399nm and 556nm is a significant expense.

While pondering the core-shell MOT, I asked myself “in what region of the ‘shell’ is the blue light most effective?” My conclusion: the region immediate upstream from the MOT in the atomic beam. This was the spark of the idea for what eventually became Crossed Beam Slowing.

As shown in Fig. 3.10, the crossed beam slower consists of 399nm beams that cross each other just upstream of the MOT region. The transverse component of the scattering force from these beams cancels, leaving a longitudinal component.

When using only the standard longitudinal slower beam, we tune the slower beam fre-

quency and the Zeeman slower current to maximize the number of atoms near 5m/s. There are approximately 3 inches between the end of the Zeeman slower field and the MOT capture region, so atoms leaving the slower region at 0m/s do not reach the MOT, while atoms faster than 5 m/s are not captured. In addition, the atoms have some transverse velocity at the end of the slower. At 5m/s longitudinal velocity, the transverse velocity of some atoms is sufficient to make them miss the MOT capture region.

When using the crossed beam slower, we tune the Zeeman slower so the exit velocity is a bit higher. This reduces the effect of the transverse component of the beam, so that more of the atoms reach the MOT region, though they are slightly too fast to be captured. The crossed beams do the last bit of slowing directly before the MOT region.

The crossed beam slowing idea is a rare fruit: an experimental concept that worked almost immediately. After I configured the optics for the crossed beam slower, I was quickly able to observe an increase in the near-zero-velocity peak of the Doppler sensitive spectroscopy. This success is also apparent where it matters most: in the MOT number and loading rate. Fig. 3.11 shows the data quantifying the performance boost with the crossed beam slower.

3.1.5 *Magneto Optic Trapping*

The principles of the MOT are well known, see [12] section 9.4. We've seen how a near resonant beam can slow and cool atoms. A 3D optical molasses is formed when red detuned beams are counter-propagating on three orthogonal axes, slowing the atoms no matter which direction they may be going. However, such a molasses does not prevent atoms from eventually deffusing out of the area. In order to trap the atoms, a position dependent force is needed. This is obtained by adding a magnetic field gradient with a zero point at the center of the optical molasses. The incoming beams are arranged with the appropriate circular polarizations so that if the atom is displace from the center of the trap in a given direction, the Zeeman shift brings the beam coming from that direction into resonance so the atom

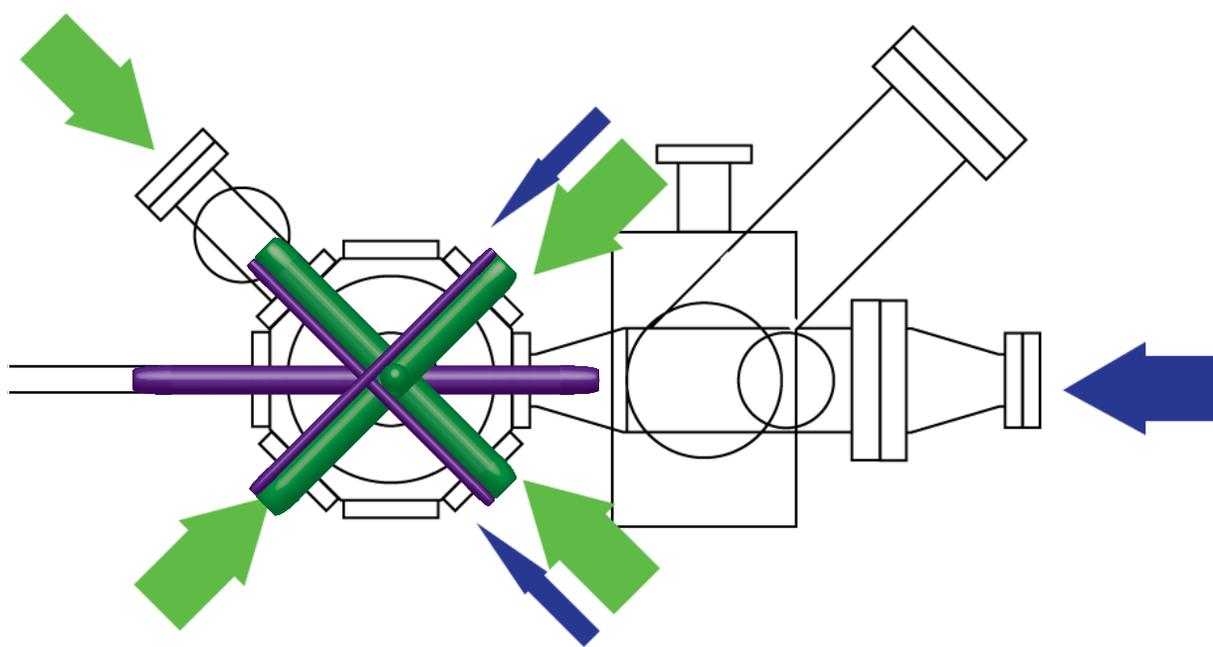


Figure 3.10: Schematic of the crossed beams slower. The crossed beams travel next to the MOT beams, and cross just upstream in the atom beam of the MOT region. Unlike the MOT beams, the crossed slower beams are not retro-reflected.

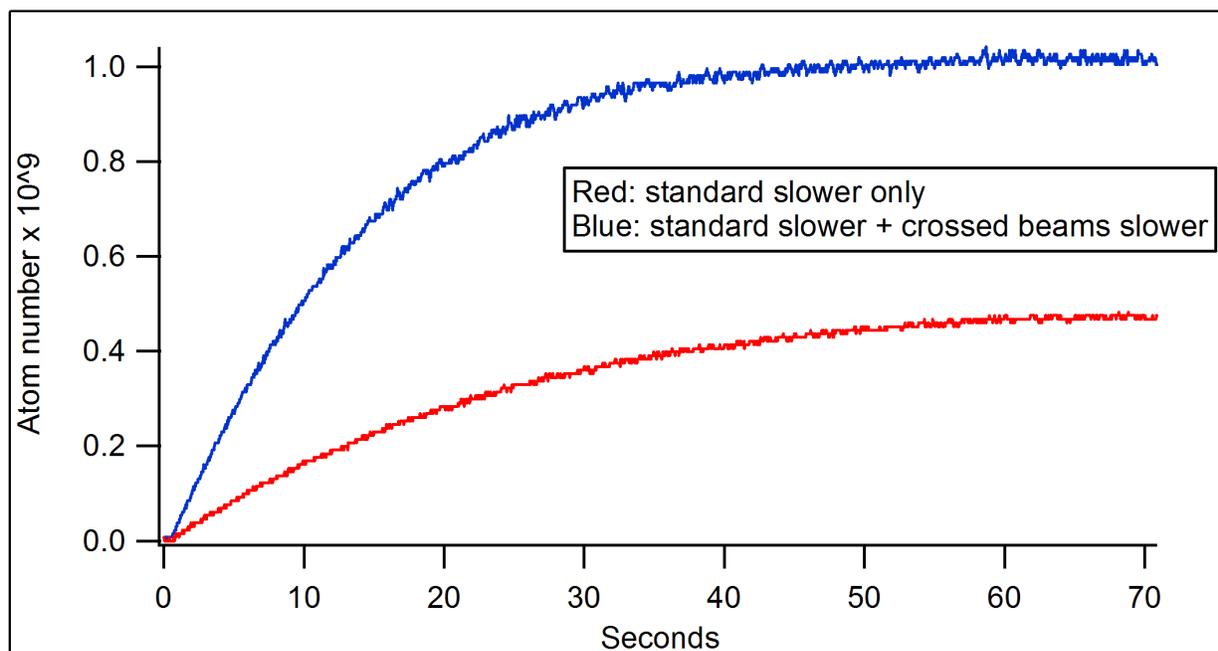


Figure 3.11: The MOT number vs time with and without the crossed beam slower. The overall slower performance is fully tuned up for each configuration, i.e. the Zeeman slower exit velocity is appropriately set in each case. We observe a factor of 2 higher MOT number and load rate while using the crossed beams.

experiences a restoring force.

The most difficult part of producing a MOT for the first time was getting all of the circular polarizations set up correctly. In principle, this should be easy to figure out. In practice, there are many degrees of freedom and it is hard to get them all right at once. In particular, some of the waveplates we used did not have well labeled fast and slow axes, which made it difficult to know which circular polarization was being produced. The technique that we eventually had success with was to “ask the atoms”. This involved applying a bias B field while looking at the crossed beam spectroscopy signal. We used the vertical comp coils to apply the bias field, and the vertical MOT beam for the spectroscopy light. With this geometry, the $\sigma+$ and $\sigma-$ transitions are split as shown in Fig. 3.12. While looking at this signal, adjust the quarter waveplate for the vertical MOT beam until the $\sigma+$ or $\sigma-$ line is maximized (and the other minimized or gone), as appropriate depending on the sign of the B field gradient. Once the vertical beam polarization is determined in this way, the horizontal polarizations can be set to match using a “polarization checker” consisting of a quarter waveplate and then a polarizing beam cube. Put the polarization checking in the vertical beam, and observe which port of the cube the light exits from. Then put the polarization checker in one of the horizontal beams, and turn the horizontal beam quarter waveplate (not the polarization checker quarter waveplate) until the light exits from the port opposite from the one that the vertical beam exits.

3.1.6 Optical Dipole Trapping

Our optical dipole trap (ODT) consists of a horizontal beam (11W, $20\mu\text{m}$ waist) and a vertical beam (5W, $60\mu\text{m}$ waist). The horizontal beam is the main trap, while vertical beam helps to localize the atoms into a smaller volume at the intersection of the beams. The smaller volume crossed beam trap is needed for evaporative cooling to degeneracy. So when setting up the ODT, we first trap atoms in the horizontal beam, and then set up the vertical

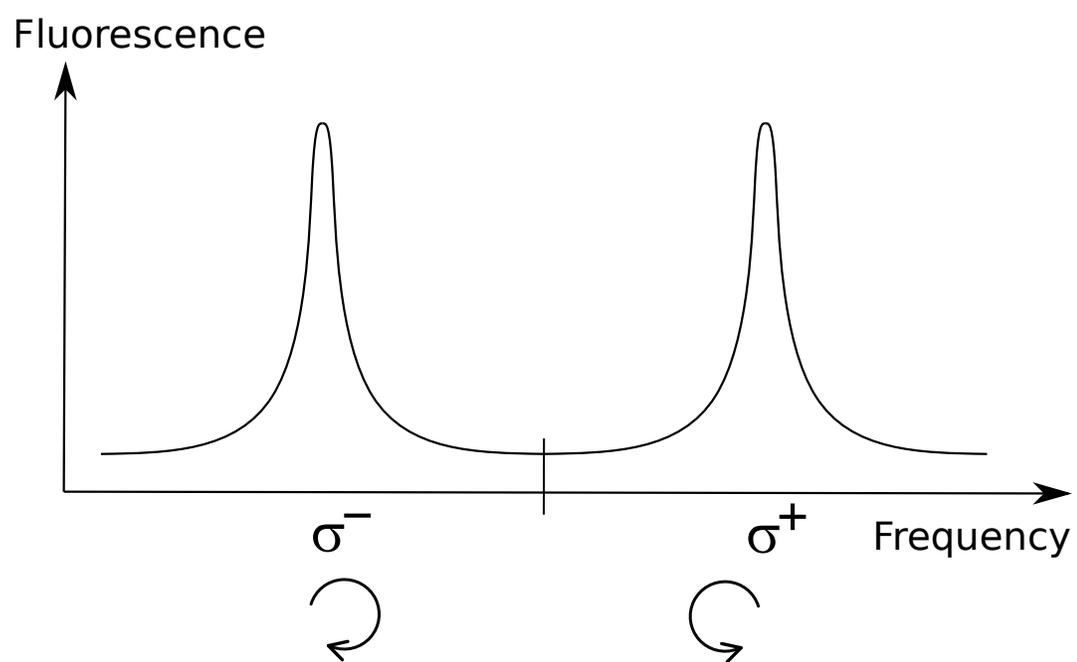


Figure 3.12: fluorescence lines split in magnetic field. The two split peaks correspond to different circular polarizations of light.

beam.

Experimentally, the crux of setting up an optical dipole trap is aligning the tightly focused beam with the atoms in the middle of a vacuum chamber. We accomplished this task by setting up an imaging path along the ODT axis. Conveniently, the Melles-Griot 532nm mirrors have good transmission at 399nm, so we were able to insert 399nm light directly onto the ODT path. We set up an imaging path bringing the horizontal beam up to the camera. After imaging the CMOT on the camera and carefully noting the position of the image, we switched to looking at the ODT light on the camera.

It is worthwhile to pause for a moment to carefully consider what is being done here. The output of an 18W laser is being focused directly onto the sensor of an expensive and delicate scientific CCD camera. The beam power should start at zero, and get turned up slowly while monitoring the camera response. It is good to do this procedure working with another person, and triple check each other's work.

After the ODT light is visible on the camera, adjust it so that the light is at the same location as the atoms were in the camera image.

Next, go back to absorption imaging on the regular horizontal imaging axis. The ODT axis imaging optics must be removed at this point, and the ODT beam dump put back in place. If the procedure has been done correctly, dipole trapped atoms should be visible.

Once there are atoms visibly trapped in the ODT, it is a relatively simple, if time consuming, optimization job to get the focus of the horizontal beam overlapped with the CMOT location. In a coordinate system where z is the propagation direction of the horizontal beam, the x and y alignment is easy to do by looking at the side images. The z axis position of the beam waist is more difficult. The waist position can be adjusted using the final ODT lens before the chamber, which is on a translation stage. It is not uncommon to need to move the translation stage on the order of a cm to get the waist lined up with the CMOT. Switching to vertical imaging can help this process.

Finally, align the vertical beam to cross the horizontal beam where it overlaps the CMOT. We were preparing to set up absorption imaging on the vertical axis, but found that we had a faint signature of the vertical beam visible in vertical imaging. This was enough to start the optimization procedure to get the beams overlapped.

3.2 *Bose-Einstein Condensation*

3.2.1 *Achieving a BEC*

Typical benchmarks for BEC production in the CI machine:

- We begin with 10^8 atoms in the compressed MOT, at a temperature between 20 and $30 \mu\text{K}$.
- After the ODT load, we have 10^7 atoms in the ODT after a 100ms hold. At this point, the atoms fill the horizontal ODT beam, the vertical beam is too weak to contain the atoms in the area where the beams cross.
- As the evaporation sequence proceeds, the atoms eventually cool to the point where they are contained in the crossed trap. At this point in the sequence, we have 10^6 atoms at $1\mu\text{K}$. The rightmost column of Fig. 3.13 shows a series of absorption images of atoms at this temperature.
- At around 200 nK, the atoms begin to condense, as shown in the middle column of Fig. 3.13. For the atoms in the condensate, thermal temperature is no longer a meaningful parameter. Instead, we observe the chemical potential of the condensate converted to kinetic energy. I'll refer to this as the kinetic temperature. Thermal atoms expand isotropically, while one of the hallmarks of the BEC state that the kinetic temperature can differ along different axes, related to the geometry of the trapping potential in

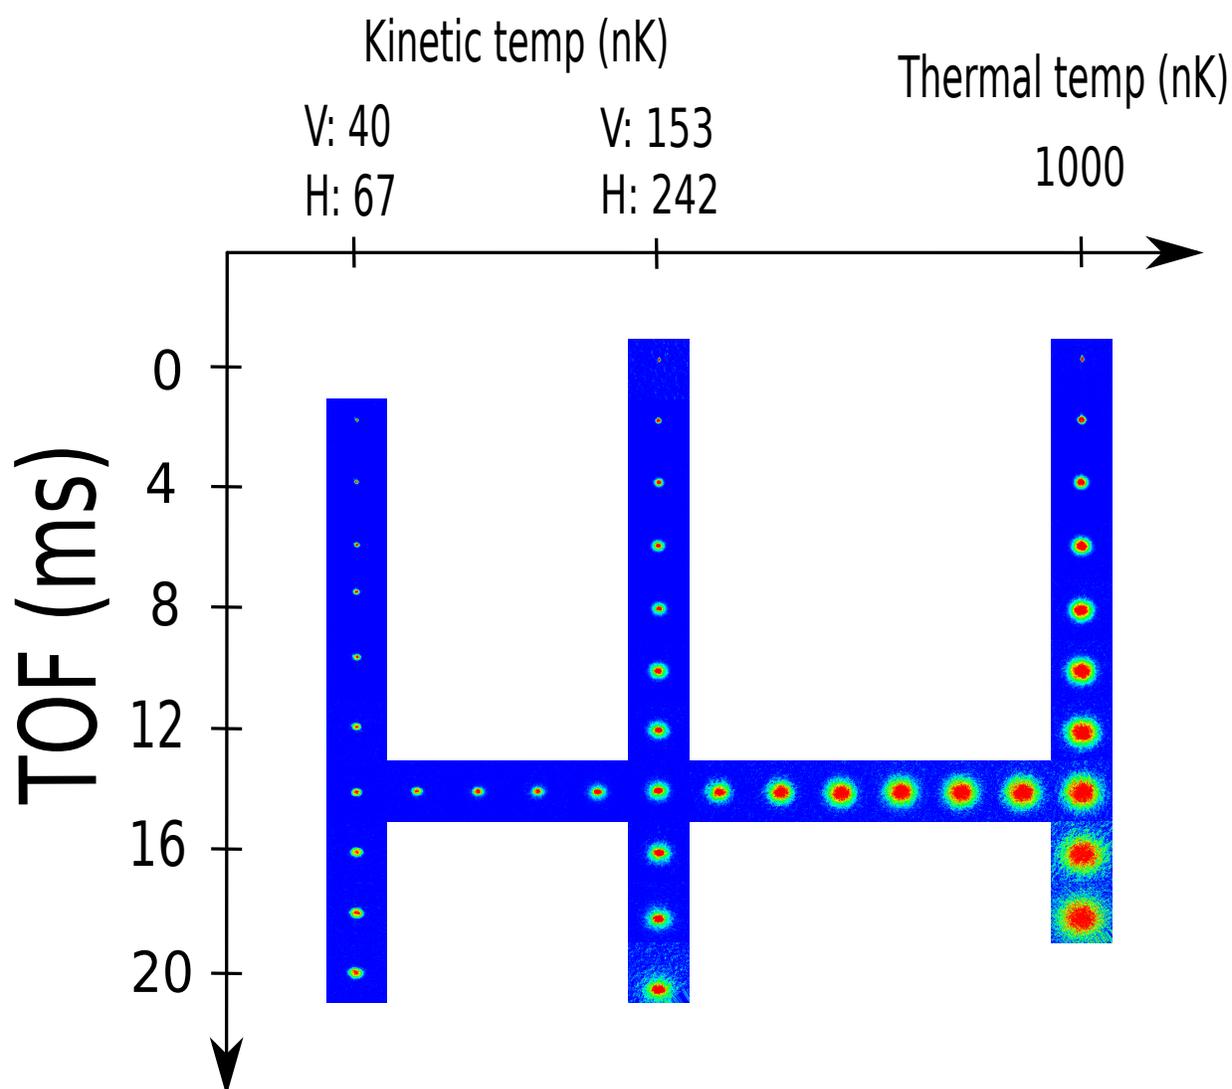


Figure 3.13: Absorption images at various TOF and temperatures during the end of evaporative cooling, as the condensate forms. A clear signature of condensation is seen in the reversal of anisotropy of the atomic cloud. The trapped condensate is taller than it is wide. In the right hand column, the thermal atoms expand isotropically. The left column, showing a pure condensate, ends up wider than it is tall. The energy due to atomic interactions in the trap is converted to momentum preferentially into the more tightly trapped directions.

which the BEC is formed. Higher kinetic temperature on a given axis corresponds to a tighter trap along that axis.

- The sequence ends with 10^5 atoms with a kinetic temperature (chemical potential) of 50 nK, with nearly all atoms in the condensate, shown in the left column of Fig. 3.13.

3.2.2 *Painting*

The condensate shown in Fig. 3.13 is produced by an evaporation sequence consisting only of varying the optical power in the horizontal ODT beam. We can gain additional control over the trap parameters by “painting” the horizontal ODT beam by modulating the ODT AOM frequency, creating a time-averaged potential. Ricky Roy in our neighboring lab used this technique to create condensates with 10^6 atoms, and to make smaller condensates in 1.6 seconds[33]. Some combination of the following factors currently prevent us from achieving results similar to this work in terms of condensate number: less power in our trapping laser, different optical trap geometry, less available Zeeman slower laser power, and the spontaneous scattering rate of our 532nm ODT (compared to 1064nm in[33]).

During evaporative cooling, one needs to keep the atomic density in a sweet spot: dense enough so that the atoms re-thermalize at a sufficient rate, but not so dense that three-body loss becomes significant. Achieving this balance is difficult with an un-painted ODT beam because the trap depth and trap frequency are not independently controllable. Painting can be used to decouple these parameters, and keep the density in the sweet spot throughout evaporation. With no painting, the trap depth and frequency both increase with optical power. Painting allows one to decrease the trap depth while lowering the trap frequency at a given optical power.

However, in the current CI machine ODT, we have never observed three-body loss playing a significant role during our evaporation sequence. This means that we would like to make the trap tighter, which is not something that painting can help with. One factor contributing

to this is that there is more spontaneous scattering (ie, 1-body) loss than in B063. We have not yet fully characterized this issue. A potential future upgrade to the ODT would be to tighten the waists of both the vertical and horizontal beams.

That said, we have found painting to be useful for adjusting the trap frequency after the condensate is already formed. We use a “decompress” step which involves increasing the painting amplitude, which relaxes the trap in the horizontal beam radial direction, turning down the vertical beam power, which relaxes the trap in the vertical beam radial direction. During this, we compensate for the reduced trap depth by turning up the horizontal beam power. As shown in Fig 3.14, we can substantially reduce the diffraction axis kinetic temperature using this technique. As shown in the results chapter, the increased coherence length gained by this reduction in momentum spread results in a significantly strong CI signal using our most decompressed trap.

3.2.3 *Trap frequencies*

For the main datasets of this thesis, the trap frequencies, measured at the end of the decompression step, are the following: $\omega_x = 2\pi \times 200\text{Hz}$, $\omega_y = 2\pi \times 16\text{Hz}$, $\omega_z = 2\pi \times 80\text{Hz}$, with a geometric mean $\omega_{\text{mean}} = 2\pi \times 63\text{Hz}$.

Here x is the diffraction axis, y is the vertical axis, and z is the slower axis.

3.3 *Diffraction: building the contrast interferometer*

3.3.1 *Interferometer Geometry in the Lab*

The CI is implemented by putting the atoms in a BEC into a superposition of momentum states using a series of coherent photon recoils from a laser beam. The CI phase is determined by the difference in energy between the interfering states. We want to maximize the energy splitting in order to increase the sensitivity of the CI. We do this by adding “acceleration pulses” to the diffraction sequence, increasing the momentum of the moving paths of the

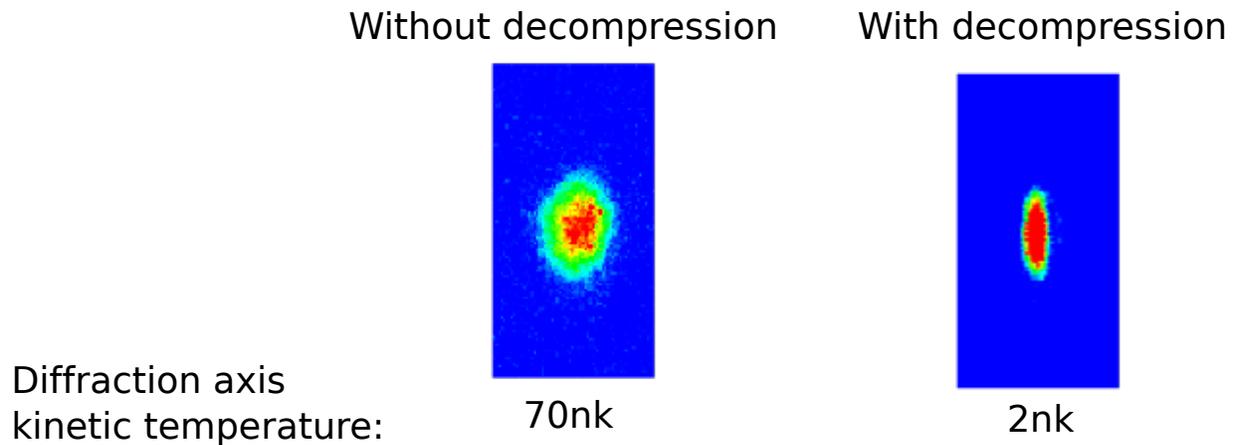


Figure 3.14: Absorption images at 20ms TOF with and without the decompression step

interferometer.

In order to minimize all energy shifts except those due to the photon recoils, the CI sequence is applied to atoms in free fall. The motion of the freely falling atoms with respect to the apparatus in the lab frame needs to be accounted for during the diffraction sequence.

The diffraction beams could be oriented either parallel or perpendicular to the gravitational field (anything in between would be the worst of both worlds), and either choice has pros and cons. The advantage of horizontal beams is that there is no Doppler shift from gravity induced motion along the diffraction axis. The disadvantage is that the atoms fall out of the diffraction beams, limiting the amount of phase evolution possible. Vertically oriented beams allow for longer evolution times, but at the expense of needing to compensate for the Doppler shift due to gravity.

Our experimental strategy has been to start with horizontal diffraction beams. The technology needed to implement acceleration pulses is a pre-requisite for doing the Doppler corrections that would be needed for vertical diffraction beams. The logical progression is to demonstrate acceleration in the horizontal case first, before attempting vertical.

Further, it was not clear from the outset whether going vertical would be necessary. It was plausible, given some (in retrospect) highly optimistic assumptions about the acceleration we could achieve, that we could reach our desired precision in a horizontal geometry.

3.3.2 Diffraction beam heuristics

Before launching into the detailed discussion of acceleration schemes, I want to share some heuristics for understanding how the diffraction beams affect the atoms. For a detailed mathematical treatment of the light-matter interactions that we use, see [15]. Here, I aim to give a quick guide to the intuitive picture that we use day to day in the lab while operating the experiment. In particular, I'm ignoring pulse shape and area (intensity vs time profile), and detuning from the atomic resonance, which are important parameters but are not needed for this discussion of interferometer design.

The coherent light-atom interactions that we are considering all involve a standing wave of light. In fact, the reason we call them diffraction beams is that there is an analogy between the physics of light diffracting from matter grating, and atoms interacting with an optical grating, or standing wave.

There are three types of standing wave pulses that we use: The splitting pulses, Bragg pulse, and Bloch Oscillations.

Splitting pulse

The splitting pulse takes atoms in the state $|0\hbar k\rangle$ and puts them into a superposition of the three states $|0\hbar k\rangle$ and $|\pm 2\hbar k\rangle$, where the states are labeled with their momentum along the direction of propagation of the diffraction beam.

Key facts about the splitting pulse:

- It is a standing wave in the frame of the $|0\hbar k\rangle$ atoms.

- It should only take atoms to $|\pm 2\hbar k\rangle$, not to higher momentum states.

Bragg pulse

A Bragg pulse can connect any two momentum states separated by an even number of photon recoils. For instance, a 1st order Bragg pulse connects $|0\hbar k\rangle$ to $|2\hbar k\rangle$, while a 3rd order Bragg pulse connects $|0\hbar k\rangle$ to $|6\hbar k\rangle$.

- A Bragg pulse is a standing wave in the frame halfway between the momentum states it connects. For example, a 3rd order Bragg pulse connecting $|0\hbar k\rangle$ to $|6\hbar k\rangle$ is a standing wave in the $|3\hbar k\rangle$ frame.
- In practice, higher order Bragg pulses have lower diffraction efficiencies. We've found experimentally that the per $\hbar k$ efficiency for Bragg π pulses is maximized for 3rd order pulses, at least for our experimental parameters of source distribution.

Bloch oscillations

Like Bragg pulses, Bloch oscillations can connect any two momentum states. A Bloch oscillation sequence involves turning on a standing wave in the initial frame of the atoms, and then smoothly accelerating the standing wave. While the optical lattice is on, the atoms receive a series of 2 photon kicks that keep them at rest with respect to the accelerating standing wave. So for instance, Bloch oscillations to take atoms from $|0\hbar k\rangle$ to $|100\hbar k\rangle$ would involve turning on a standing wave in the $|0\hbar k\rangle$ frame, smoothly accelerating the standing wave until it is in the $|100\hbar k\rangle$ frame, and then turning it off. The atoms come along for the ride, and are transferred to $|100\hbar k\rangle$ with high efficiency.

With this toolkit in mind, here are some of the schemes we have considered for increasing the momentum splitting in a CI.

3.3.3 CI Schemes

The simplest instance of a CI can be implemented using a linearly polarized retro-reflected laser beam, as shown in fig3.15

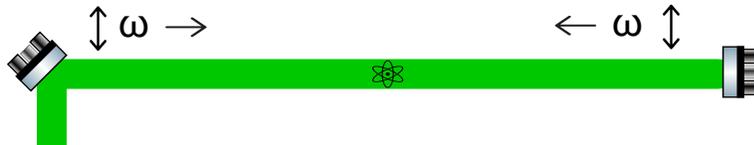


Figure 3.15: Laser schematic for retro-reflected horizontal CI. Referred to in the text as DiffPlan0. In this and the following 3 figures, the double ended arrows represent linear polarizations, while the one ended arrows represent direction of propagation.

No matter the laser frequency ω , this retro reflected beam forms a standing wave in the lab frame (or more precisely, the frame of the retro mirror). This configuration is simple and reliable, but it cannot be used to accelerate the atoms after the initial splitting pulse. Bloch oscillations are not possible, and the only Bragg pulse that can be done with this configuration connects states that are symmetrical around $|0\hbar k\rangle$, for instance taking $| -2\hbar k\rangle$ to $| +2\hbar k\rangle$ and vis versa. This latter pulse is referred to as the mirror pulse, and is a key component of any CI configuration. I'll return later to this point, that any CI requires a lab frame standing wave for the mirror pulse.

In order to implement acceleration pulses, more flexible frequency control is required. One such setup is shown in Fig. 3.16.

In this scheme, the retro mirror is replaced with a second counter propagating beam. The two beams have are sourced from the same laser with frequency ω , but are shifted by separate AOMs giving frequency shifts δ_1 and δ_2 . This configuration is very flexible in that, by choosing appropriate values for δ_1 and δ_2 , a standing wave can be created in any frame, and linearly sweeping δ_1 or δ_2 creates an accelerating lattice as required for Bloch oscillations. This scheme has one big limitation, which is that it can only create one standing wave at

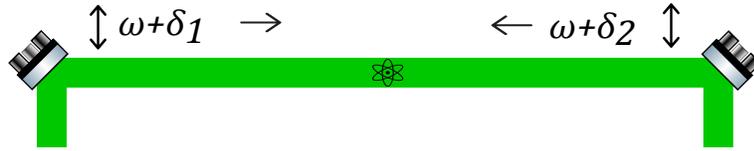


Figure 3.16: Laser schematic for the diffraction beam scheme with a single polarization, but with two fibers instead of a retro mirror, for more flexible frequency control. Referred to in the text as DiffPlan1.

a time. This means that it is not possible to implement acceleration pulses that maintain complete symmetry between paths 1 and 3 of the CI.

In order to have fully symmetric acceleration, we need a way of generating two standing waves in two different frames. One such method is shown in Fig. 3.17.

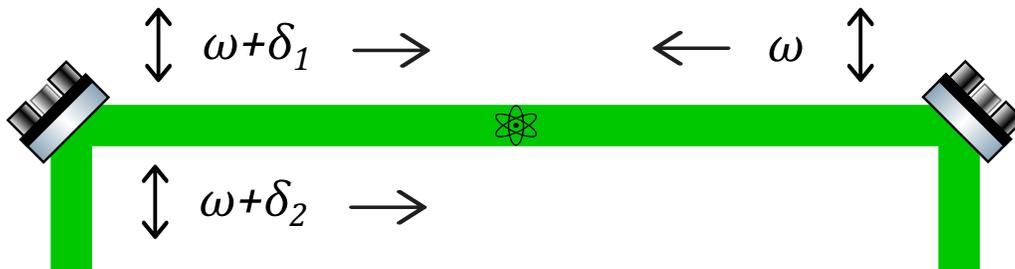


Figure 3.17: Laser schematic for the diffraction beam scheme with one polarization and two fibers, but with two frequency components in one of the beams. Referred to in the text as DiffPlan2.

Like the previous one, this configuration consists of two counter-propagating beams with separate frequency control. One beam has the laser frequency ω , while the other beam has two separately controllable frequency components. Both beams have the same linear polarization. By choosing different values for δ_1 and δ_2 , standing waves can be created in any two frames, and frequency sweeps of δ_1 and δ_2 can create two accelerating standing waves.

This scheme seems to give us everything we need, but it turns out to have significant drawbacks. The two standing waves have the same polarization, which means processes involving one standing wave can be coupled to processes involving the other. Further, the two frequency components in one beam create a beat note at the difference frequency between δ_1 and δ_2 . When the difference between δ_1 and δ_2 is less than the bandwidth of the diffraction beam intensity stabilization circuit, the beat note interferes with the stabilization. At the start of the acceleration sequence, $\delta_1 - \delta_2$ needs to be less than the 100kHz intensity stabilization bandwidth, which means that intensity stabilization is effectively impossible using this scheme.

Ideally, what we would like is two separately controllable standing waves in orthogonal linear polarizations. Fig. 3.18 shows one way to do this.

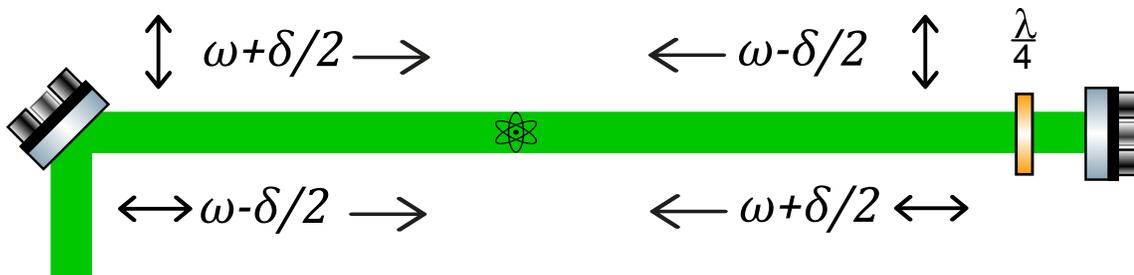


Figure 3.18: Laser schematic for the diffraction beam scheme with two polarization components, and a $\frac{\lambda}{4}$ waveplate to rotate the polarizations to get standing waves with the desired frequency components. Referred to in the text as DiffPlan3.

In this scheme, we return to a retro-reflected diffraction beam. The diffraction beam is formed by combining two beams, each with a controllable frequency offset, on a polarizing beam cube. Right before the retro mirror is a $\frac{\lambda}{4}$ waveplate. The waveplate is oriented so that after the beam passes through it twice, the linear polarizations in the beam are rotated by 90 degrees. In effect, the polarizations of the two frequency components are swapped.

The end result is two standing waves in orthogonal linear polarizations. A standing wave is only formed when the counter-propagating beams have the same polarization, so the two standing waves have different frequencies as shown in Fig. 3.18.

3.3.4 Bloch Oscillations

Our assumption at the outset was the Bloch Oscillations BO would be the best performing large momentum transfer atom optics with which to add acceleration to the CI. Since BO had not been demonstrated previously in the Gupta group, doing so was our first order of business with the new BEC machine.

The simplest diffraction beam setup that can do BO is DiffPlan1 as shown in Fig. 3.16, so this was the first thing we set up. We initially struggled with poor performance until we figured out how to avoid forming an optical cavity with the FC/PC ends of our diffraction fibers (see section 2.4.8 for more details). After this, we were able to demonstrate BO with 98.5% efficiency per $\hbar k$, as shown in Fig. 3.19. This was a little worse than expected, which turned out to be due to a glitch in the DDS system due to a switching power supply connected to one of the DDS voltage inputs. After moving to a linear power supply, we were able to demonstrate BO with 99.9% efficiency per $\hbar k$, at 30kHz/8 μ s sweep rate, with 45mW per diffraction beam.

3.3.5 $n = 4$ Contrast Interferometer

After our experimentation with acceleration tools using DiffPlan1, we regressed to DiffPlan0 (Fig. 3.15) in order to set up the simplest possible contrast interferometer. We used this configuration to align the readout light collection optics, as described in section 2.4.11

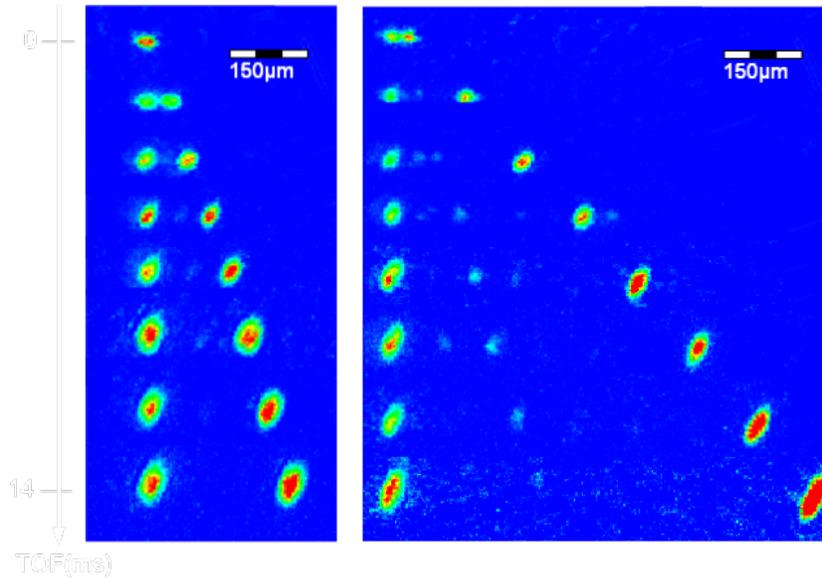


Figure 3.19: Early Bloch Oscillation demonstration from 2016

3.3.6 Adding Acceleration: a lesson on the KISS principle

Finally, it came time to put these pieces together to create the accelerated CI. Adding acceleration pulses with the DiffPlan0 diffraction beam configuration is impossible, so we had to choose which of the more complicated configurations to test first. DiffPlan3 is the only configuration that allows us to symmetrically accelerate the paths of the CI, so this is what we tried first.

The first thing we tried was doing BO symmetrically on two arms at the same time. We

had success with this, but the efficiencies were lower than expected.

Next, we tried setting up the $n=4$ CI. Here is where we really started running in to problems. We found that the mirror pulse had inconsistent diffraction efficiency. This is not behavior we ever noticed with the DiffPlan0 configuration. We spent some time trying to observe a CI signal for $n=16$, but were unsuccessful.

To see what was causing this problem, refer again to Fig 3.18, and consider what happens when $\delta = 0$. In the ideal case, we'd end up with two standing waves at the same frequency and with orthogonal polarizations. There is nothing to distinguish these standing waves besides polarization, so the total field is a sum of these two components. The resulting field depends on the relative phase of the components. Because the two diffraction beam paths are split and take different routes through different AOMs and fibers, time dependent phase differences can result.

If the two diffraction beams were perfectly in phase, a standing wave with polarization at a 45 degree angle would result. If the beams are out of phase, we end up with circular polarizations that do not form a standing wave. Hence the lattice depth of the resulting standing wave depends on the unstable phase difference between the two paths.

One way to solve this problem would be to remove the quarter waveplate for the diffraction pulses where $\delta = 0$. The timing requirements for removing and replacing the waveplate are difficult: the switch needs to be made on the scale of microseconds, which is impossible for any scheme involving physically removing a glass waveplate.

There is a way to accomplish this: a Pockels cell. These devices are effectively voltage controlled waveplates. Preliminary discussions with Thorlabs technical support suggest that the Thorlabs EO-PC-550 Pockels cell would be able to switch between a zero waveplate and a quarter waveplate at the timescales needed for our experiment. It would be a significant project to set up the high voltage control electronics for the Pockels cell and we were looking for quicker results, so we kept this possibility in reserve and moved on to other diffraction

beam configurations that we could test right away with the equipment on hand.

We wanted to maintain our ability to symmetrically accelerate the paths, so the next configuration we tested was DiffPlan2 (Fig. 3.17). This setup required a reconfiguration of the diffraction beam AOMs to single pass, so that two frequency components could be applied to one AOM.

This configuration fixed the problem of unstable diffraction efficiency for pulses where $\delta = 0$ like the mirror pulse and splitting pulse. Using this setup with simultaneous Bragg pulses, we did obtain the first ever signals for an accelerated CI. However, we struggled with the efficiency of the acceleration pulses, and had a hard time accelerating very much beyond $n=16$. We also had no success with observing a CI signal with BO acceleration. For both Bragg pulses and BO, we found that the acceleration efficiency in one path was hurt by the presence of the optical lattice accelerating the other path. We also found that the CI phase fluctuated more than we would like when we could observe it.

At this point, I finally started questioning my assumption that the complete symmetry of the CI must not be broken. If we can't accelerate the two moving paths simultaneously, maybe we can do it sequentially as long as we switch back and forth very quickly? As long as the moving paths have the same momentum during the free evolution time, the interferometer should work as normal.

So we moved back to DiffPlan1, which is what we'd started with. We used Bragg pulses sequentially applied to one moving arm and then the other, and almost immediately starting having much more success observing the CI signal for higher n , and with lower phase fluctuations.

I have a few thoughts I want to share about the experimental path we followed. Our original thinking was to test the new parts of the experiment first. We assumed that BO would be the best acceleration method, and that we would use the diffraction beam configuration DiffPlan3, which on paper had the best properties. So we spent a significant amount

of time testing these techniques before establishing the $n=4$ CI on the new machine. Once we did set up the CI, it became clear very quickly once we tried to combine the CI with BO and with DiffPlan3 that something was amiss. Especially with regard to the excessive phase fluctuations seen while attempting the BO accelerated CI, this is not something we could have known about testing BO on their own. We should have had access to the CI signal as a diagnostic before testing new acceleration technology.

So my advice to future students on this project is to “close the loop” faster, i.e. make sure to test as quickly as possible whether a new experimental technique is working for its actual purpose.

3.3.7 Acceleration with Sequential Bragg Pulses

After we settled on acceleration with sequential Bragg pulses, we spent some time finding the optimal pulse parameters. We eventually settled on 3rd order Bragg pulses with a single oscillation sinusoid profile with a $130\mu s$ total pulse time.

To be clear, the pulse described above is what is output by Cicero on an analog channel that controls the diffraction beam intensity feedback setpoint. The actual intensity vs time profile applied to the atoms determined by the response of the feedback circuit to this input. We monitor and record the diffraction beam powers. The recorded Bragg pulse profile is slightly asymmetric, and is fit reasonably well by a gaussian with $1/e$ width of $29\mu s$.

We tested several different Bragg pulse profiles. In addition to acceleration efficiency, we were also concerned with the amount of time it takes to accelerate the atoms. The diffraction beams are horizontally oriented, and there are only around 13ms of total interferometer time available before the atoms fall out of the beams.

The standard Bragg pulse profile is a Gaussian. The problem with this profile for our purposes is the relatively long “tail” of the gaussian function. The intensity feedback does not respond well to square edged intensity profiles, so we need the pulse profile to go nearly

to zero before we begin and end the time step. This led the time steps for the Gaussian pulses to be longer than seemed necessary.

We first tried an inverted parabola profile. Although this pulse profile allowed very short time steps for a given peak width, we found that the Bragg efficiency was significantly worse than for Gaussian shaped pulses. It seems that some inflection point is needed.

Our next idea (credit goes to Dan for coming up with this) was to try a pulse profile based on a single peak of a sinusoidal function. The thought was that this is a function with an inflection point like a Gaussian, but that goes to zero more quickly for a given peak width, as shown in Fig. 3.20.

We found that this “cosbragg” profile gave us similar Bragg efficiency as the Gaussian profile for a given peak width, while the cosbragg would fit in a shorter time step as desired.

Once we settled on the cosbragg profile, we optimized the pulse width as shown in Fig. 3.21. For each width we tested, we optimized the pulse height to form a Bragg π pulse. We did a series of 6 pulses, and looked at the number of atoms in the final target momentum state divided by the total number of atoms.

We also tested 4th order Bragg pulses, but found that they had a lower per $\hbar k$ efficiency than 3rd order for our source distribution. We did not test 2nd order Bragg pulses, because the acceleration rate with such pulses would be too low to do a reasonable amount of acceleration with horizontally oriented diffraction beams. In a vertical diffraction beam interferometer geometry where time is not so limited, 2nd order Bragg pulses would be a reasonable option if the Bragg efficiency is found to be higher.

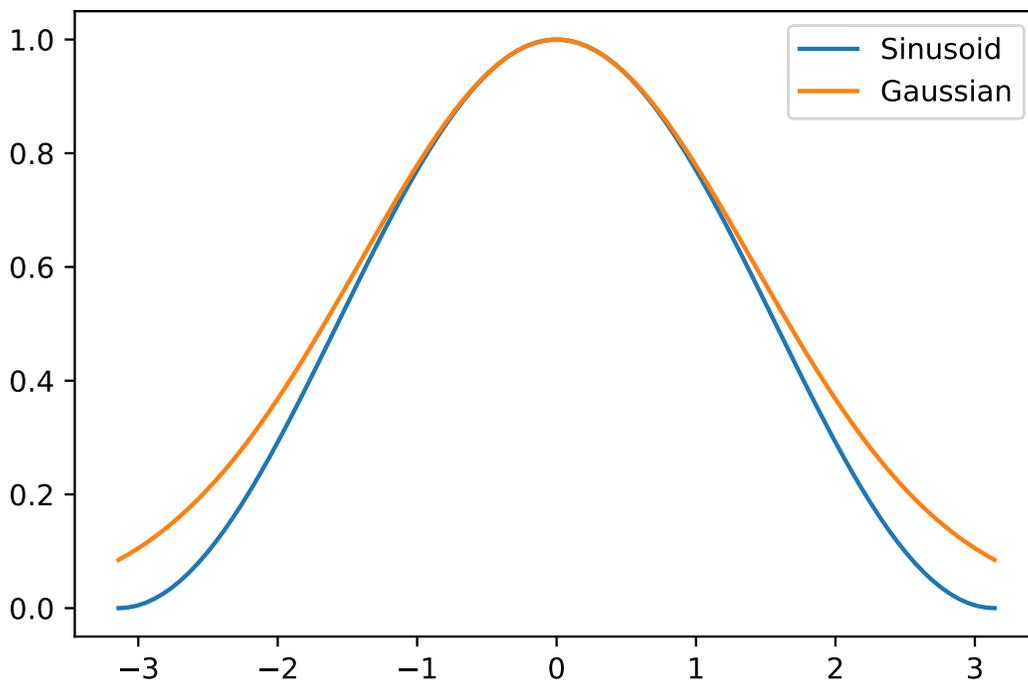


Figure 3.20: Comparison of a Gaussian and Cosine with similar widths, showing that the Cosine goes to zero more quickly than the Gaussian.

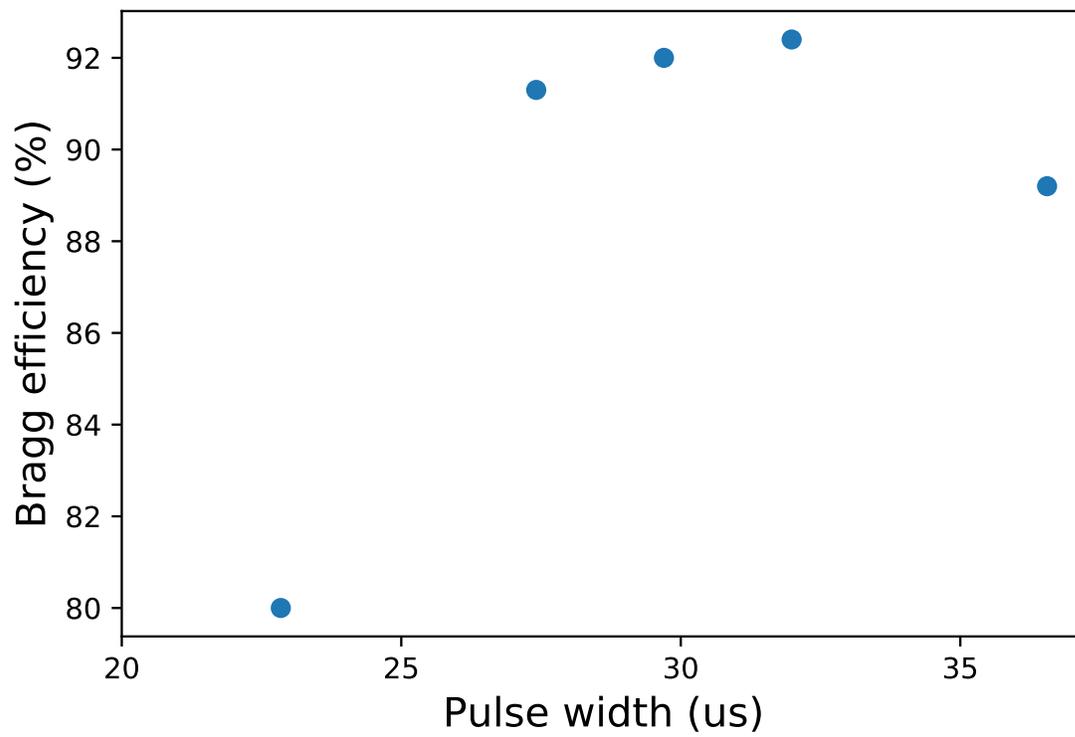


Figure 3.21: 3rd order Bragg pulse efficiency vs Bragg pulse width. In this case, Bragg pulse width is defined as the waist of a Gaussian fit to the recorded diffraction beam intensity during the Bragg pulse.

Chapter 4

RESULTS

4.1 *Prelude*

It took years to build the CI machine and get it to the point where it was producing BECs. By late 2015, we were starting to work towards demonstrating the new atom optics techniques, like Bloch Oscillations, that would be needed to add acceleration to the CI. After demonstrating Bloch oscillations using a single polarization lattice, we moved on to doing $n=4$ interferometry in order to align our readout optics and demonstrate blue readout for the first time. After this, we moved to the two polarization retro reflected setup, which we anticipated would be the best configuration for the accelerated CI. We did some proof of concept work (or so we thought) showing acceleration in two directions with this diffraction beam setup.

So by late summer 2016, it seemed like all the pieces were in place. We had demonstrated the $n=4$ CI on the new machine using blue readout. We had demonstrated Bloch oscillations in a single polarization lattice, as well as acceleration in two directions using the two polarization lattice. All that remained was to put these pieces together.

What followed was a very frustrating few months where it seemed that nearly every aspect of the machine had to be retuned. MOT loading instabilities lead us to the following upgrades:

- Cleaning Yb building on the slower viewport by blasting it with the focused Verdi.
- Re-optimizing the slower by going back to Doppler sensitive spectroscopy.

- Fiberizing the MOT light, which turned out to be the final missing piece to get a stable MOT.

Although we had been making BECs for some time, the number of condensed atoms had not been satisfactory. It was enough for our testing up till that point, but in order to do high n interferometry, where one loses significant fraction of the initial atoms to acceleration inefficiency, we needed more than we were getting. Our evaporation sequence seemed to be less efficient than it could be. The ODT optics are not something that we touch or adjust lightly, but I made the decision to tighten the focus of the horizontal ODT beam. This quickly led to more efficient evaporation, and larger BECs.

Also, during this time the Princeton camera broke, and a new camera from Allied Vision Technology had to be purchased and integrated into the experiment.

After we had a stable MOT, large BECs, and a working camera, we were finally able to pursue interferometry again, where it quickly became clear that the readout optics were no longer aligned. We had to go back to and regain the readout from scratch.

That was the breakthrough point where we finally started to attempt contrast interferometry with acceleration. However, we quickly learned that the strategy of testing Bloch Oscillations on their own before applying them to the CI was not ideal. The CI signal with BO acceleration proved elusive. This led to a period of experimenting with different diffraction optics configurations.

Finally, after months of experimentation, we settled on sequential 3rd order Bragg acceleration in a single polarization lattice. With this method, we started to have success with the accelerated CI. In late winter and early spring of 2017, we often had weekly breakthroughs in our acceleration abilities. As those breakthroughs started to slow down, we shifted to taking data that we hoped would be publication worthy. Sometimes we took a few versions of these datasets. What I'm presenting here are the stopping points of these processes, the datasets that were deemed to be good enough. Most of these "good enough" datasets were

taken in a single day, but those single days are balanced precariously on the months and years of designing, building, and tuning the experiment.

4.2 Slope Data

Our first goal with the accelerated CI was to demonstrate the n^2 scaling of the phase accumulation rate. The CI has two free evolution intervals of length T , and for a given T the CI phase is $\Phi(2T) = \frac{1}{2}n^2\omega_{\text{rec}}T + \phi_{\text{offset}}$. We take data at several different T , and extract the slope of $\Phi(2T)$ vs $2T$.

The resolution of Cicero timestep length is $2\mu\text{s}$. In general, the two free evolution intervals should be the same length, but a $2\mu\text{s}$ difference is acceptable. This means that the smallest interval by which we can change $2T$ is $2\mu\text{s}$. This is no problem for $n=4$, where the phase accumulation of $2\mu\text{s}$ is much less than π , but for higher n the phase accumulation over $2\mu\text{s}$ can be significant. For our highest n of $112\hbar k$, the CI phase changes by more than 100π over $2\mu\text{s}$. I'll go into more detail about how we fit and analyze this data in the Discussion section.

Our slope data taking procedure was to start by taking 5 data points each $2\mu\text{s}$ apart. After that, we would increase the time separation between subsequent points by a factor of 2, so the next point would be $4\mu\text{s}$, the next $8\mu\text{s}$ after that and so on until these jumps are too large to accommodate in the total interferometer time available.

We took slope data for n between 4 and 52 on March 21st 2017. We stopped at $n = 52$ because this was the highest n for which we could reliably fit individual shots. After some further tuning of the machine, and the decision that we didn't need to fit individual shots, but could just look at the averaged signal, we pushed to higher n . In quick succession we achieved $n = 64$ on May 18th, $n = 76$ on May 19th, $n = 88$ on May 22nd, and $n = 100$ on May 26th. Finally, on June 19th we pushed it to $n = 112$. It was clear that we could not get any further with the current machine configuration. For $n = 100$ and $n = 112$, we only

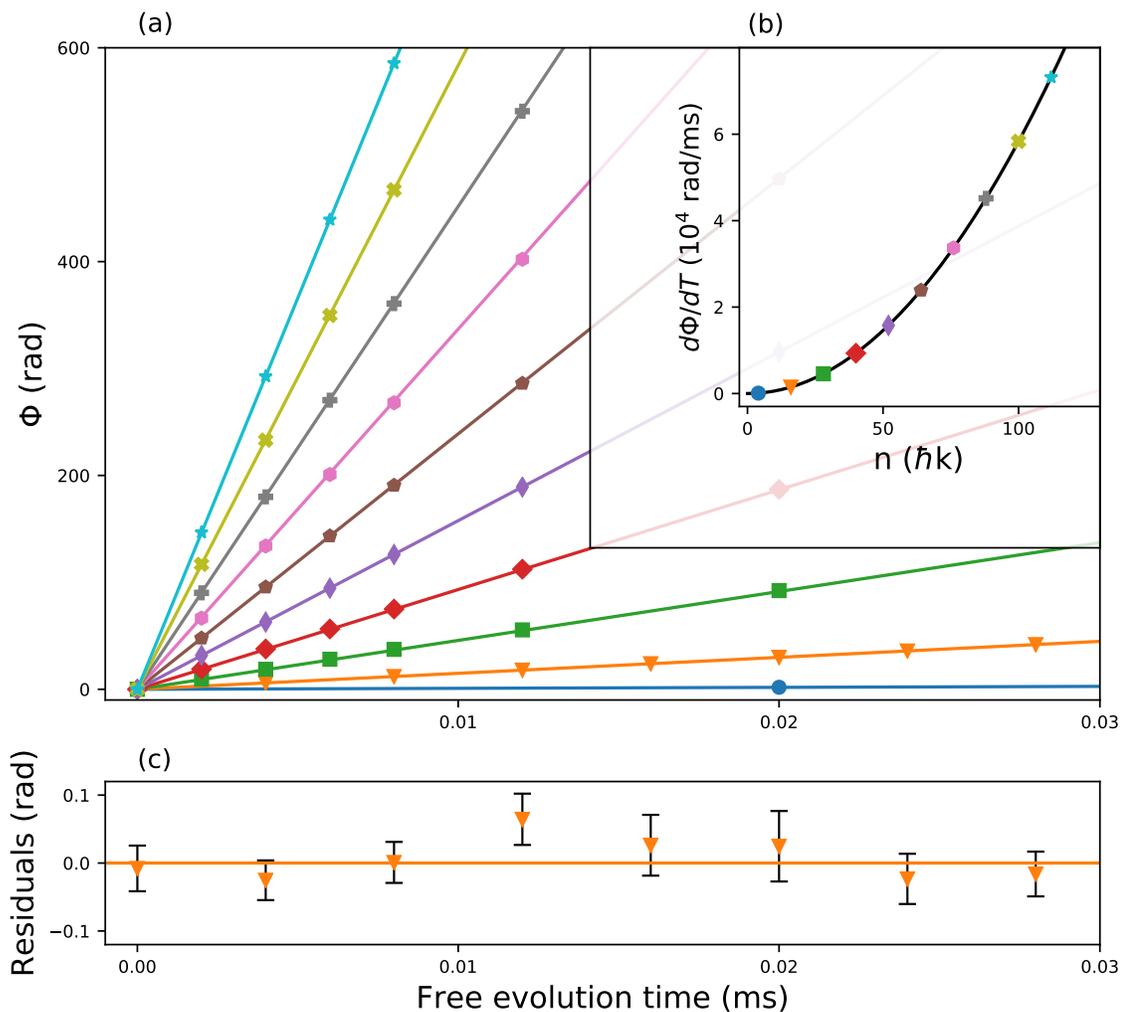


Figure 4.1: (a) CI phase Φ vs free evolution time $2T$ with linear fit, for various n . (b) Fit slopes $d\Phi/dT$ vs n , demonstrating the expected quadratic relationship $d\Phi/dT = \frac{1}{2}n^2\omega_{\text{rec}}$ (black curve). (c) Typical fit residuals ($n = 16$).

took data at 5 different $2T$ values. For each $2T$ we had to take 50 or more shots.

Slope data for $n = 4$ through $n = 112$ are shown in Fig. 4.1(a). Fig. 4.1(b) shows the value of the slope for each n demonstrating the n^2 scaling, and Fig. 4.1(c) shows typical residuals for the $n = 16$ data.

4.3 Visibility Data

We wanted to study in detail how the CI signal changes as a function of n . In principle these data can be extracted from the slope data mentioned in the previous section. In practice, there are two problems with the slope data for this purpose. First, the data were taken on many different days, with gaps of weeks between different n . Second, the data were taken with different total interferometer times. The CI signal quality depends on the total time of flight.¹ The length in time of the acceleration sequence varies with n , so identical $2T$ times have different total interferometer times for different n . To compare the signal quality between different n , we want to compare signals not with the same $2T$, but with the same total interferometer time.

On June 16th 2017, we took a large dataset focused on signal quality vs n . We started with the $n = 100$ CI (our highest at the time) and took data with $2T = 20\mu s$. Next we took data for $n = 88$ with $2T = 20\mu s + AccTime$, where $AccTime$ is the length of the acceleration pulses removed. Basically, we held the total interferometer time fixed while reducing n step by step down to 4. A few days later we achieved $n = 112$. Unfortunately those data are at a longer interferometer time than the rest of the signal quality data, but we judged it to be close enough and include the $n = 112$ with the rest.

These data are shown in Fig. 4.2 and Fig. 4.3. The right column shows the averaged raw data, and the left column shows the data with a diffraction phase correction applied before averaging.

¹In the previous generation Yb CI, the signal quality did not change with TOF up to 22ms. We are still investigating why the current CI signal strength depends on TOF

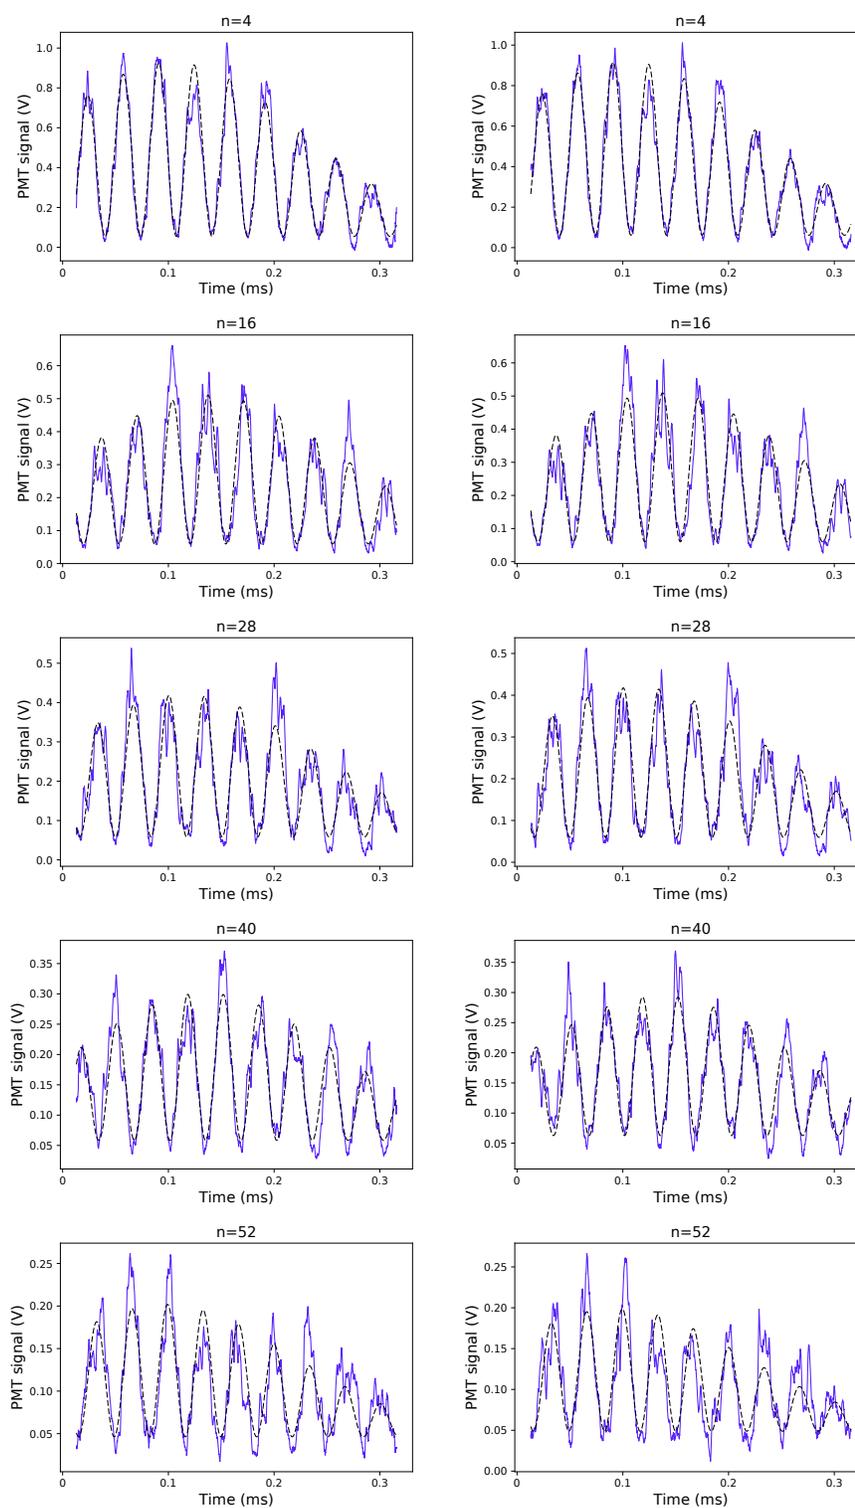


Figure 4.2: Averaged signals for $n=4$ through $n=52$. The right column shows the averaged raw data, and the left column shows the data with a diffraction phase correction applied before averaging.

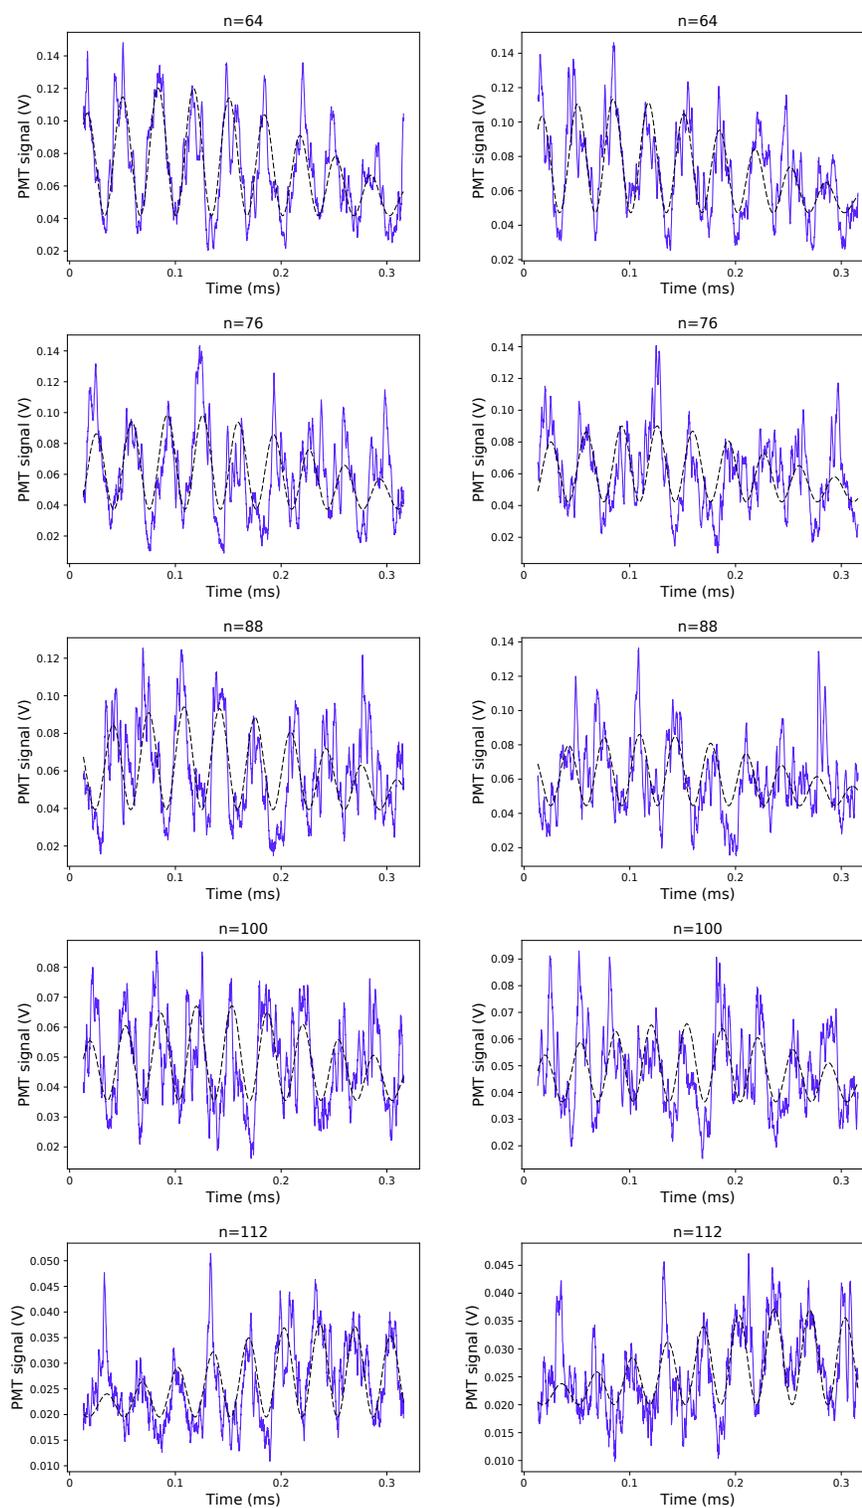


Figure 4.3: Averaged signals for $n=64$ through $n=122$. The right column shows the averaged raw data, and the left column shows the data with a diffraction phase correction applied before averaging.

4.4 Diffraction Phase Data

Diffraction phases are an intensity dependent phase that accrues during standing wave laser pulses. We first studied diffraction phases in the context of the contrast interferometer in [20]; they are also discussed in Alan's thesis [18]. During the normal operation of the CI, we try to minimize the shot-to-shot variation in diffraction phase by stabilizing the diffraction beam intensity. In order to study diffraction phase, we also took data where we intentionally vary the Bragg pulse area. These results are shown in Fig 4.4.

4.5 Interactions Data

As mentioned in Section 3.2.2, we can control the condensate density, and hence momentum spread, by varying the amount of ODT decompression to adjust the final trap frequency before release. The data in Fig. 4.5 show the CI signal vs trap frequency. It is clear that the coherence length of our lowest density condensates is significantly longer. In addition, we observe that the CI phase has little, if any, change while the trap frequency is varied by a factor of 3.

4.6 Readout Light Data

In order to study the effect of the readout light intensity on the CI signal, we took a dataset while varying the readout light intensity. The data are shown in Fig 4.6. For reference, the main CI datasets were taken with readout power of $6\mu\text{w}$.

We note that the CI phase seems to have no dependence on the readout light power. The amplitude of the signal grows, but less quickly for higher readout power, as if there is a saturation mechanism. The envelope becomes asymmetric at higher readout power, with the signal getting larger towards the beginning of the readout pulse. This is consistent with the idea that the readout pulse begins to destroy the atomic grating. A full theoretical description of this process could be an interesting project for the future.

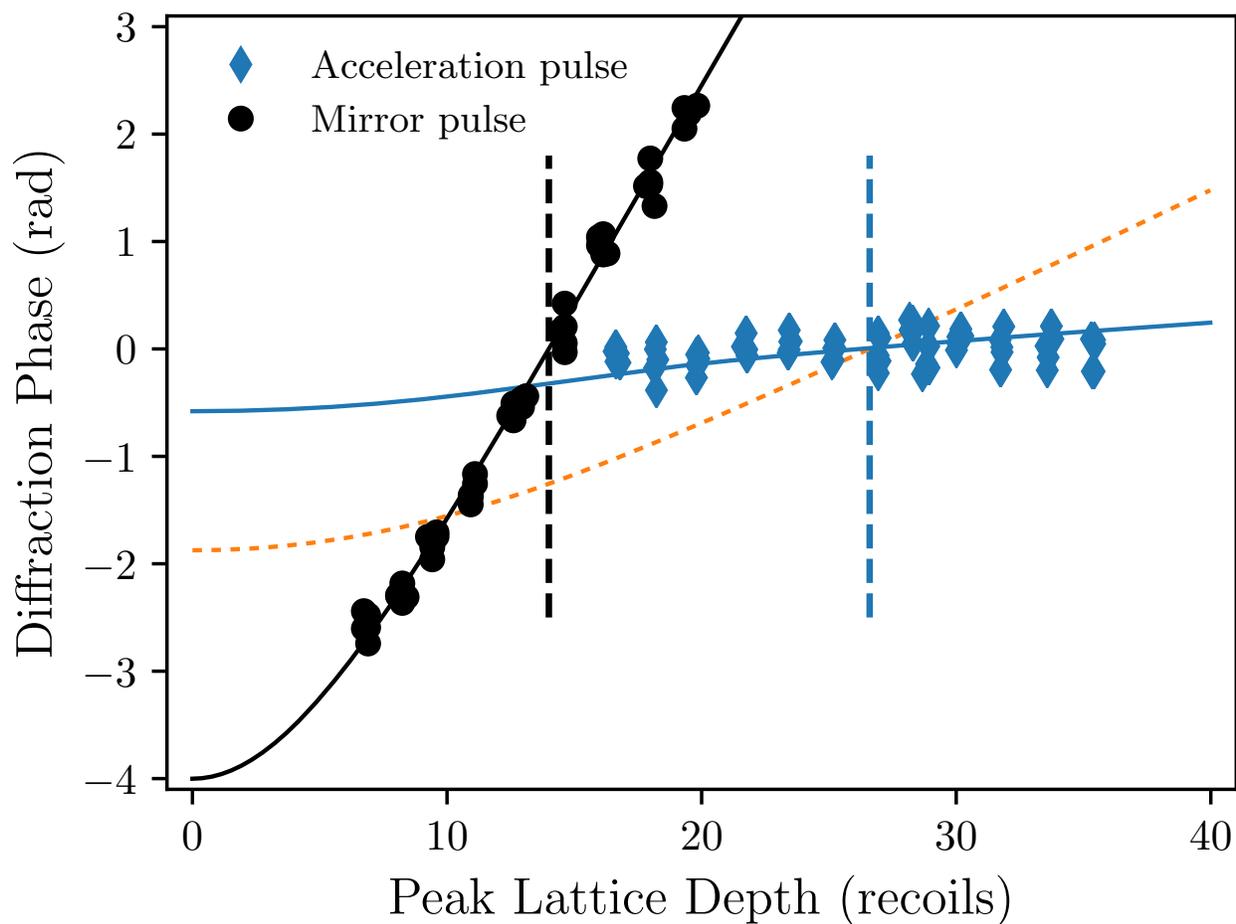


Figure 4.4: Diffraction phase shift vs peak lattice depth (with pulse width held fixed) for the second-order Bragg mirror pulse (black circles) and the third-order acceleration pulse (blue diamonds) from $2\hbar k$ to $8\hbar k$. The dashed vertical lines indicate the peak lattice depths at which the π -pulse condition is met in each case. Overall phase offsets have been removed to zero the diffraction phase at the π -points. Black and blue solid lines are the predictions from the corresponding numerical model. The dotted orange line is the prediction of the model for the diffraction phase for the third-order Bragg pulse for large n .

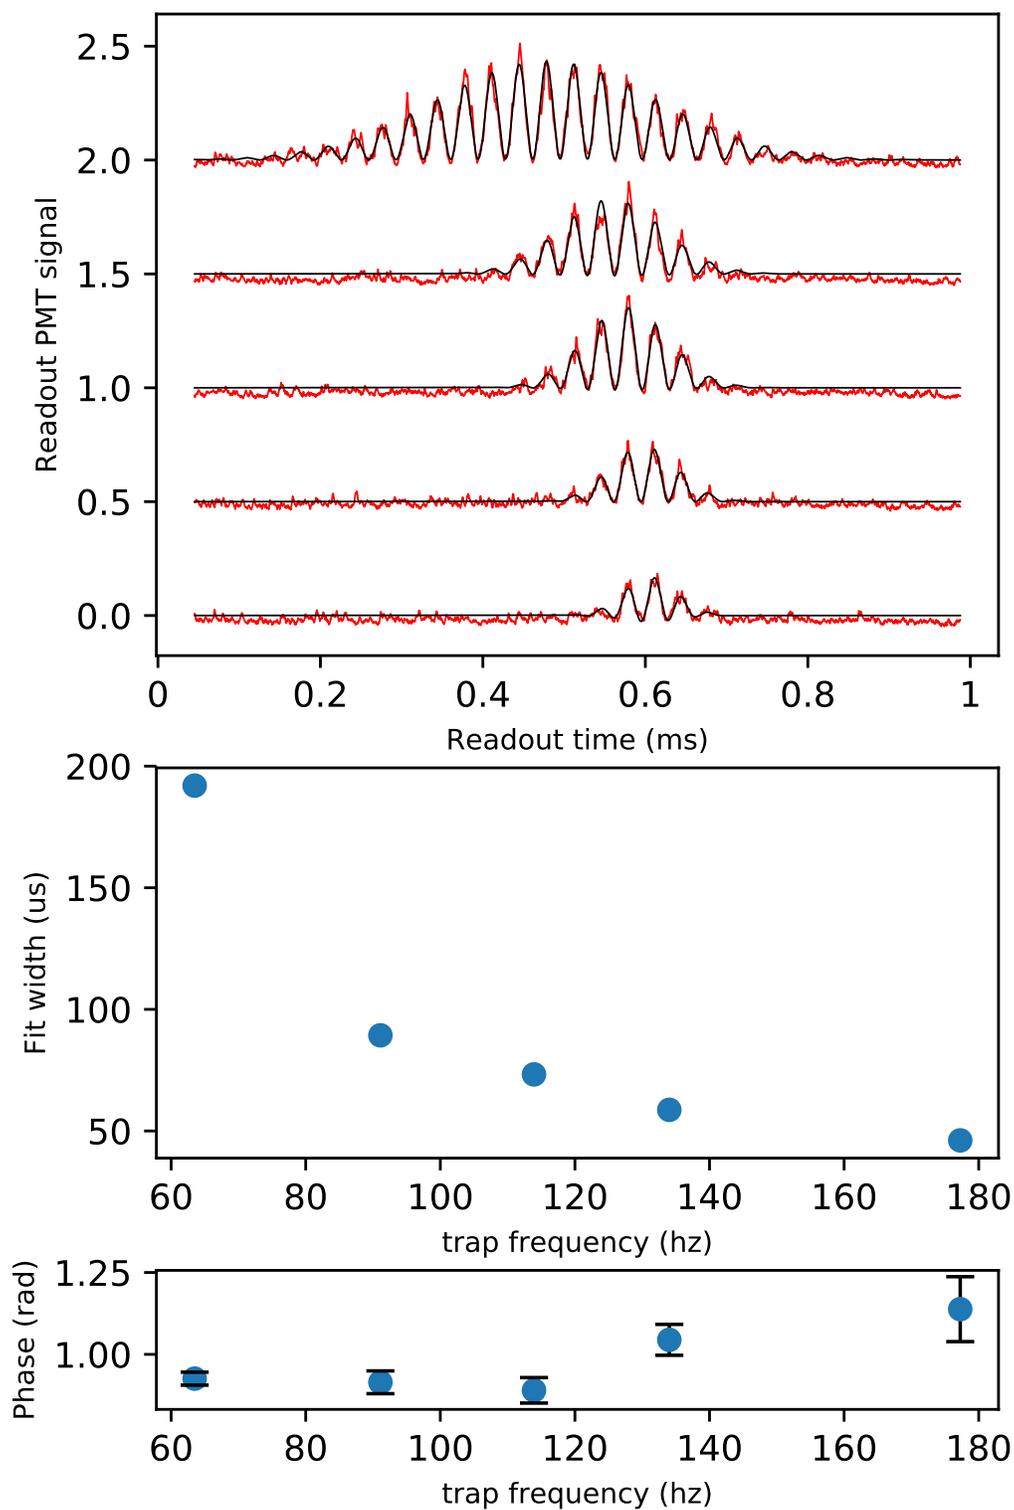


Figure 4.5: Top: CI signal vs trap frequency. Higher signals correspond to lower trap frequencies. The vertical axis is voltage, but offsets are added to the signals to keep them from overlapping. Middle: Gaussian width of the fit function envelope vs trap frequency. Bottom: Signal phase vs trap frequency.

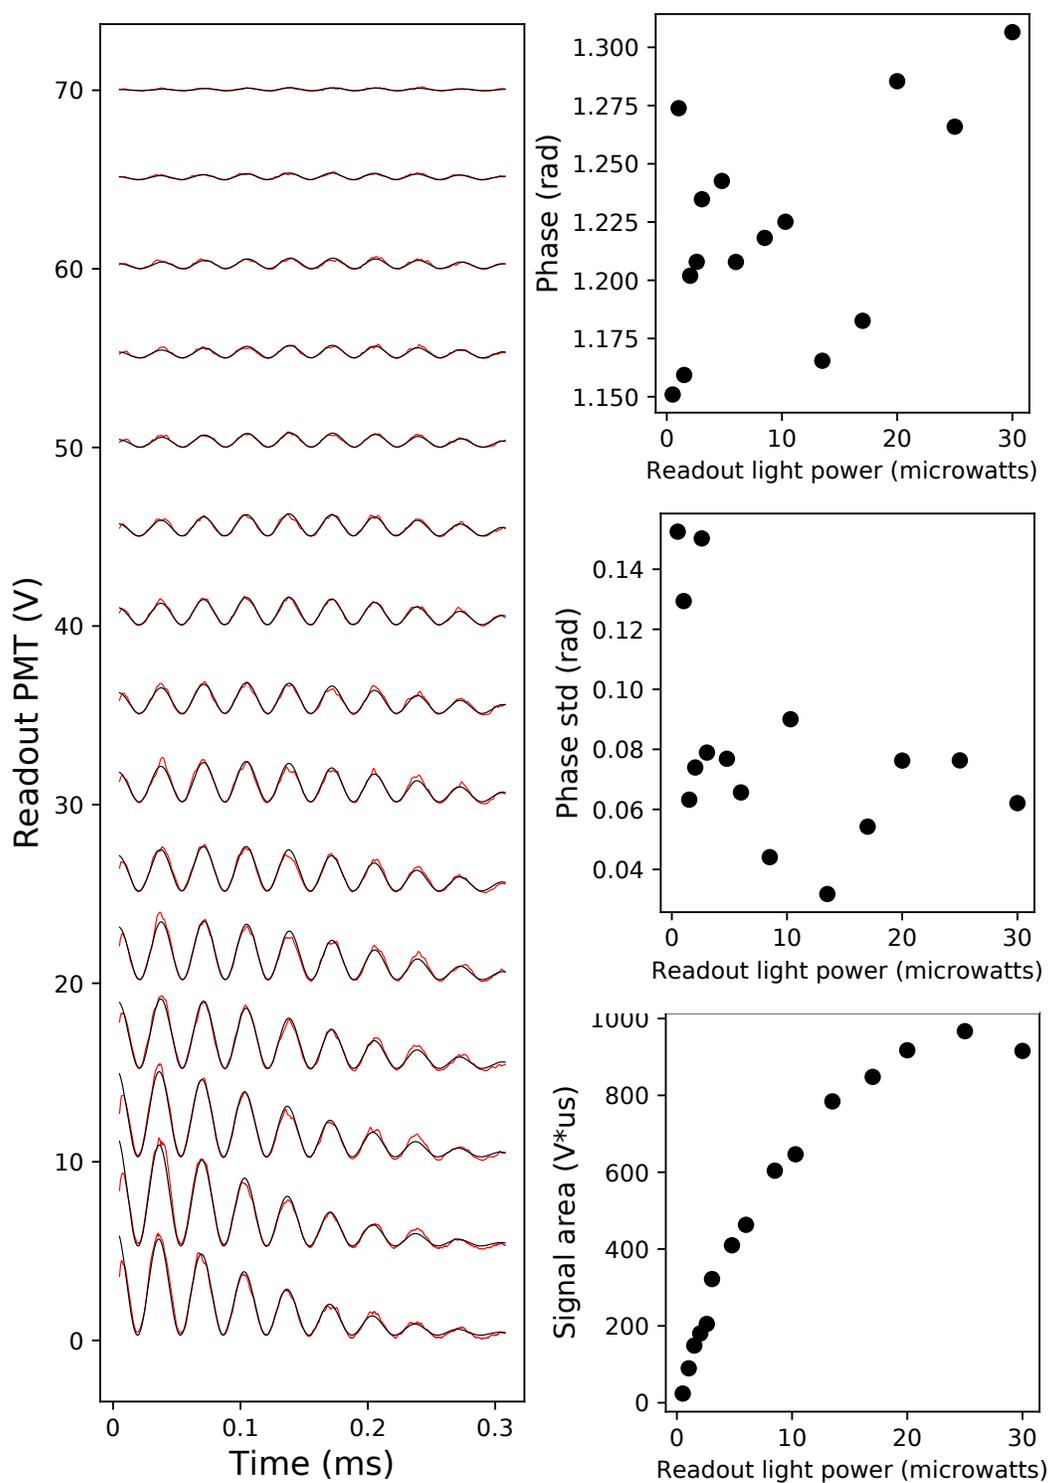


Figure 4.6: Left: the CI signal for a variety of readout light powers. Higher power signals are lower on the plot. The vertical axis is in volts, but offsets have been added so that the signals do not overlap. Upper right: The CI phase vs readout light power. Middle right: Phase standard deviation vs readout light power. Lower left: Signal area vs readout light power.

Chapter 5

DISCUSSION AND OUTLOOK

In the previous chapter, I described the data that we gathered over the course of running the CI. Here, I go into more detail about how the data was processed and analyzed, and what conclusions we draw from it.

5.1 Analysis Methods

The readout PMT voltage and the diffraction beam intensity photodiode signals are recorded on oscilloscopes. For each experimental run, the oscilloscope traces are saved as .csv files.

I analyzed these data using Python 3.5, particularly the Scipy, Numpy and Matplotlib libraries. I used the Anaconda Python distribution, which automatically installs all the Python packages I needed including the Spyder Interactive Development Environment (IDE), which I used to develop the analysis code. The very active Python open source community is an excellent resource. For nearly any problem I faced working with Python, a quick search would return a plethora of example code and discussions.

Most of the previous analysis work in the Gupta lab has been done with Igor Pro or Mathematica. For specific tasks these programs have some advantages, Mathematica for symbolic calculations and graphing, and Igor Pro for graphing and because some of our established data collection processes are based on Igor Pro.

Despite some extra effort associated with moving to a new language, I found that working with Python was advantageous. Although ground level analysis tasks such as making a single fit or plot are easier in Mathematica, more complicated data processing schemes with layers of abstraction are easier to handle in Python.

My analysis suite consists of five main classes (in the sense of Object Oriented Programming). The Data class, the Phase class, the Bragg class, and the Slope class. The Data class has methods to import the readout signal and diffraction monitor .csv files into Python. The Phase class has methods to fit and extract the phase from the readout signals. The Bragg class has methods to process the diffraction beam intensity profiles, and apply corrections to the signal data based on the shot-by-shot diffraction phase. The Slope class takes the processed Phase data and does a bootstrapping routine to fit the phase points to a slope.

5.1.1 *Extracting the phase from the CI signal*

A single shot of the CI returns a full fringe pattern. The fringes are bounded by a roughly Gaussian envelope function (determined by the time dependent spatial overlap of the three BEC components), and oscillate with frequency $8\omega_{\text{rec}}$. For reference, this corresponds to a period of about $34\mu\text{s}$ ¹. Example signals for $n=4$ are shown in Fig. 5.1 (single shot) and Fig. 5.2 (100 shot average).

Our goal is to extract the phase of the $8\omega_{\text{rec}}$ oscillation. We reference the phase to the left edge of the readout window. The oscilloscope is triggered by the same digital channel that triggers the readout pulse. As long as the phase reference point is fixed with respect to the readout pulse trigger it doesn't matter what the reference point is, so the left edge of the scope window is as good a choice as any.²

¹Years of working with the CI signal have given me a keen eye for $34\mu\text{s}$ oscillations, wherever they may appear.

²A consideration for future high precision data taking is that changes in the position of the envelope function with respect to the phase reference point do affect the measured phase if the frequency of the fit function is different from the signal oscillation frequency.

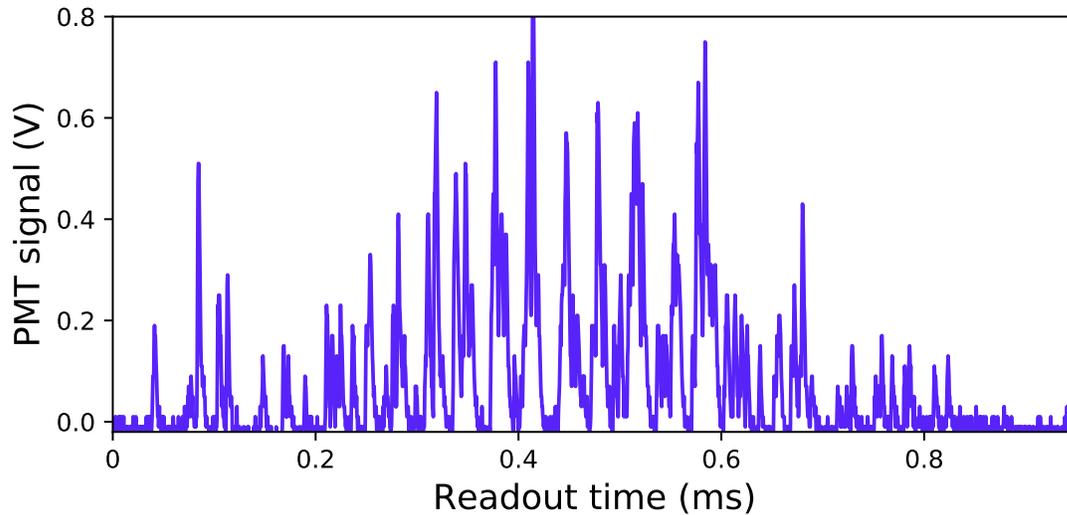


Figure 5.1: Example of a single shot CI signal with $n=4$.

5.1.2 Phase extraction with non-linear curve fitting

Our first strategy for phase extraction was to use a least squares minimizing fitting routine, such as `curve_fit` from the `scipy.optimize` Python module, to fit a model function to the data. I used a variety of model functions for different purposes, but the most commonly used fit function is³:

$$FitFunc(t) = amp \times \exp\left(-\left(\frac{t - shift}{width}\right)^2\right) \times \sin^2(t * 4\omega_{rec} + phase) + offset \quad (5.1)$$

In general, I tried to minimize the number of free parameters in the fit, which makes the fitting routine more likely to converge. The floated parameters in Eq. 5.1 are *phase*, *amp*,

³We use a sine squared function because the PMT is sensitive to the intensity of the readout beam electric field, while the CI contrast oscillation frequency maps to the electric field amplitude. Our use of the sine squared means that our measured phase has a π rather than a 2π ambiguity. Our choice to use this fit function is purely a matter of convention. One could use an un-squared sine and multiply the phase by 2, as long as one remembers to compensate appropriately when backing out physical values.

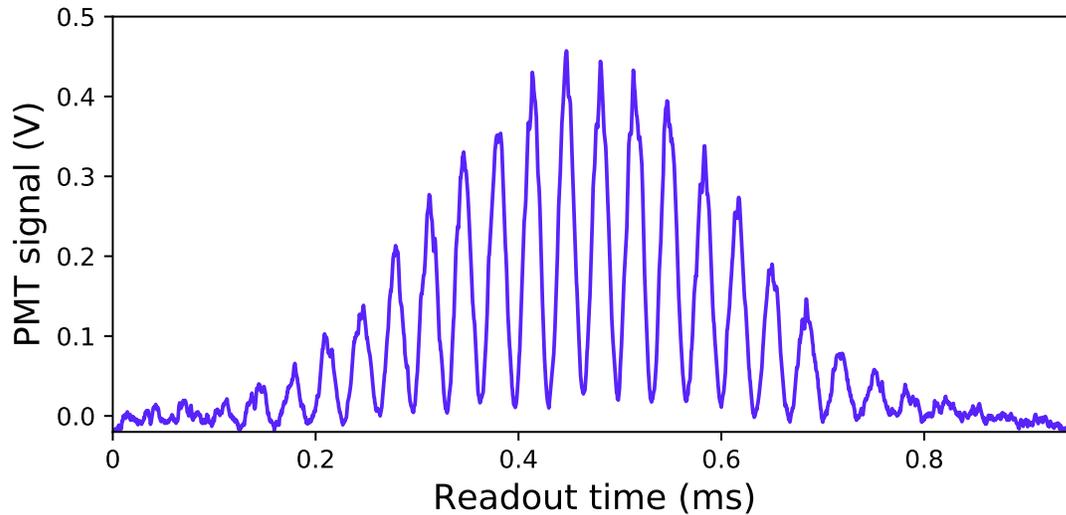


Figure 5.2: $n=4$ CI signal with 100 shots averaged together.

offset, and *shift*. I hardcoded the average width for a given dataset, and the oscillation frequency is always fixed to the best currently known value of ω_{rec} .

This method works fine for relatively strong readout signals, especially averaged signals. However, the readout signal strength depends strongly on n . Around $n = 52$, this fitting procedure becomes unreliable as the routine sometimes fails to converge for single shots. It is possible to run the fitter on averaged signals even up to our highest n .

For the purposes of measuring the average CI phase at a given $2T$, one might think that fitting the average signal is sufficient. This is the case as far as it goes, but a problem arises when one considers what should be used as the phase error bar. The fitting routine does return an uncertainty for each floated fit parameter, but this fit uncertainty does not capture the information about shot-to-shot phase variations. Fig. 5.3 is an illustration of why this is. The left panel shows three (simulated) signals with different phases, while the right shows the average of those three signals. The averaged signal has lower amplitude and visibility, but the fit uncertainty for the averaged signal would not be any different than the

fit uncertainty for a single simulated signal. In other words, the sum of a bunch of sines with small phase variations is still a sine as far as the fitter is concerned, so the fit uncertainty of an averaged signal does not capture small phase variations.

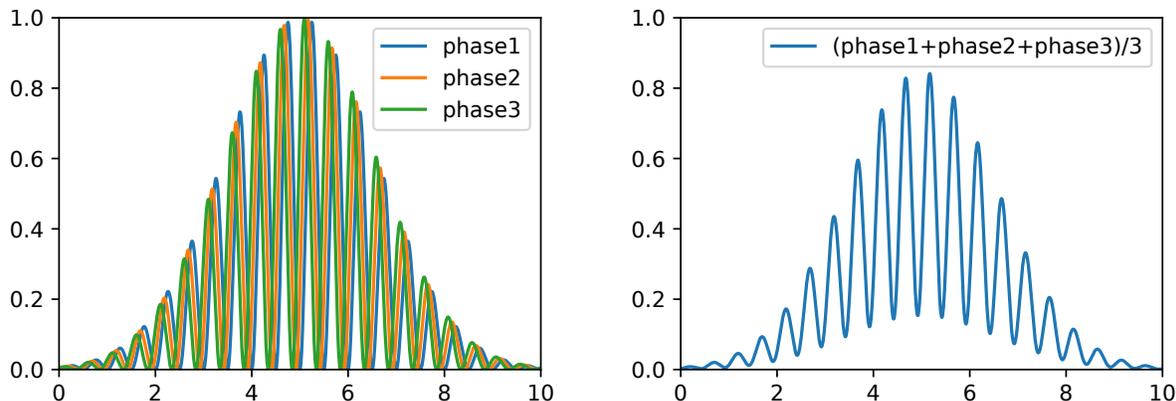


Figure 5.3: The left panel shows 3 signals with different phases, and the right panel shows the average of the three signals. The phase variation leads to a decrease in the amplitude and visibility of the averaged signal, but the fit uncertainty for the averaged signal does not capture the underlying phase variations.

The fit uncertainty does scale with $\frac{1}{\sqrt{N_{\text{ph}}}}$ where N_{ph} is the number of photons contributing to a given signal, but it does not capture shot-to-shot phase variation. For the purposes of precision measurement, we need to know about the shot-to-shot variation in the underlying phase. For $n \geq 64$, we can't get this information with the non-linear fitting method because the fits often fail to converge for single shots.

5.1.3 Fourier component phase extraction

During summer the 2017, we had an excellent REU ⁴ student named Eric Cooper. Eric was a pleasure to work with, and he managed to make significant contributions to our group in

⁴Research Experience for Undergraduates

the short time he was with us. One of those contributions is the Fourier component phase extraction method described below.

The basic idea of this method is to extract the phase by looking at the Fourier transform of the CI signal $S(t)$ evaluated at $8\omega_{\text{rec}}$:

$$\hat{S}(8\omega_{\text{rec}}) = \int_0^{t_0} S(t) \exp(-i * 8\omega_{\text{rec}} * t) dt \quad (5.2)$$

Where t_0 is the length of the signal. For our discrete data, the procedure is as follows:

- List 1 is the readout signal data consists of a list of voltage values vs time.
- Create List 2 with length equal to List 1, consisting of the imaginary exponential $\exp(-i * 8\omega_{\text{rec}} * t)$ evaluated for each t index of List1.
- Calculate the dot product of List 1 and List 2 (component wise multiplication followed by summation).
- The result is a complex number. The CI phase is the angle in the complex plane, while the amplitude is a metric for signal strength.

An advantage of the Fourier phase method that is immediately apparent is that it can be applied to any signal without concern that a fitting routine may fail to converge. The question we need to answer is, for these weak signals, does the Fourier phase return a sensible result?

Fig. 5.4 can give you an idea of the signals we are working with. The left plot shows a single shot for the $n = 100$ CI, while the right panel shows an average of 80 shots. To my eye at least, the single shot has no discernible oscillation at $8\omega_{\text{rec}}$. The non-linear fitter is very unlikely to converge while fitting a single shot signal such as this, while it has no problem with the 80 shot average.

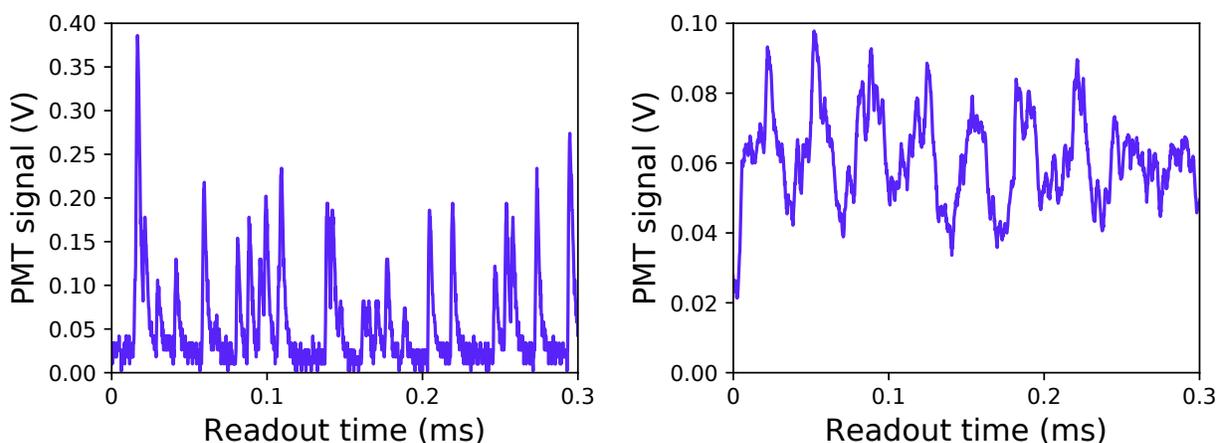


Figure 5.4: For $n = 100$, the left panel shows a single shot, while the right shows an average of 80 shots.

We can apply the Fourier phase technique to the 80 shots at $n = 100$ that make up the averaged signal in Fig. 5.4. A histogram of the single shot Fourier phases is shown in Fig. 5.5, along with the mean value of the single shot Fourier phases. Also shown is the phase returned by the non-linear fitter applied to the averaged signal in Fig. 5.4. We observe that the histogram is centered around both average phase values, which gives us confidence that the Fourier phase technique returns a reasonable value even for weak signals.

The second big advantage of the Fourier phase technique is that it gives us sensible error bars. For the uncertainty on the phase, we use the standard error⁵ on the mean of the single-shot Fourier phases. The orange shaded region in Fig. 5.5 represents the standard error, which overlaps with the non-linear fit phase from the averaged signal, showing that the two methods are consistent.

⁵Standard error=standard deviation/sqrt(number of datapoints).

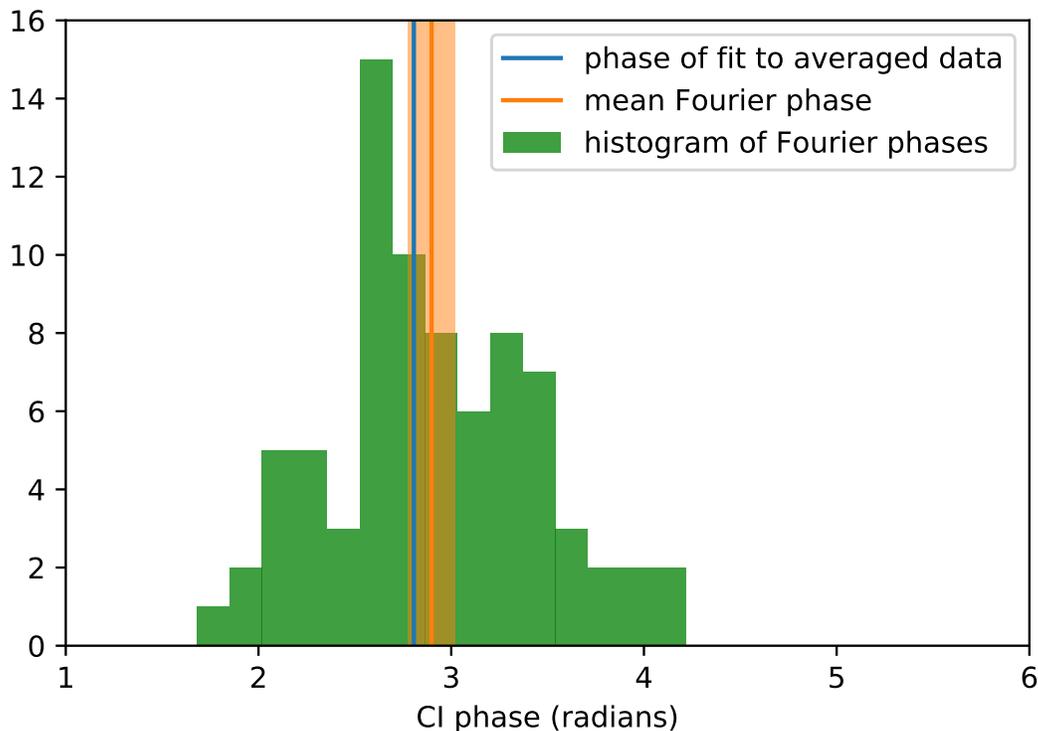


Figure 5.5: Histogram of the single shot Fourier phases of the 80 shots shown in the right panel of Fig. 5.4. The blue vertical line is the phase returned by the non-linear fitter applied to the 80 shot average, while the orange line is the mean value of the single shot Fourier phases. The orange shaded region is the standard error of the mean.

5.1.4 *Signal quality vs n*

Armed with these signal fitting and phase extraction techniques, we can analyze in more detail the signal quality vs n data shown in Fig. 4.2 and Fig. 4.3.

Amplitude vs n

We observe that the signal amplitude decreases with n . Why might this be? The first candidate that comes to mind is that fewer atoms contribute to the contrast signal for higher n , because the acceleration pulses are not 100% efficient. The more acceleration pulses, the fewer atoms there are on paths 1 and 3 by the end of the CI sequence.

Our two signal analysis techniques, non-linear fitting and Fourier analysis, both give a signal amplitude metric, so we need to decide which one to use while studying amplitude vs n . We saw in Fig. 5.3 that averaging signals with shot-to-shot phase variations can reduce the amplitude. For our current purpose we'd like to look at the average single shot amplitude, rather than the amplitude of the averaged signal⁶. So we use the Fourier amplitude for our amplitude vs n analysis.

We want to compare the amplitude vs n data to a model based on the number of atoms contributing to the CI signal: $\#atoms = A^\zeta$, where A is the per photon recoil efficiency of our acceleration pulses, and $\zeta = \frac{4(n-4)}{2}$ is the total number of acceleration photon recoils for path 1 or 3 accounting for the four steps of acceleration or deceleration in a complete experiment and removing the initial splitting and mirror pulses, which are common to all experiments.

The amplitude theory curve in Fig. 5.6 is a fit of this model to the amplitude data with A as a fit parameter, returning $A = .984$, or 91% per 3rd order Bragg pulse. This efficiency is consistent with the 92% 3rd order Bragg efficiency that we found in Fig. 3.21. Given that it is difficult to maintain peak efficiency across all acceleration pulses in the CI, 91% is a reasonable average efficiency over all the pulses.

The model fits the data well, which suggests that the reduction in atoms due to the acceleration pulse efficiency is the primary cause of the signal amplitude decreasing with n .

⁶This sentence is tricky, but stick with it. It does make sense.

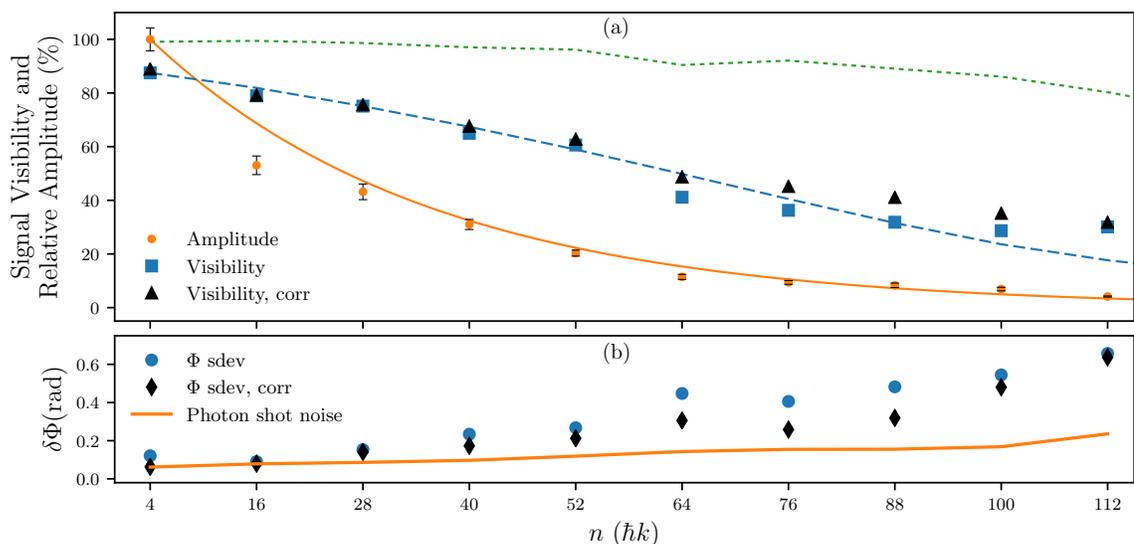


Figure 5.6: The upper panel shows the normalized readout signal amplitude and visibility vs n . The amplitude is based on the Fourier component at $8\omega_{\text{rec}}$, and visibility is extracted from sinusoidal fits. The amplitude decreases for higher n because fewer atoms contribute to the signal due to the finite efficiency of the acceleration pulses. The amplitude model curve is an exponential fit with acceleration efficiency as a fit parameter. Visibility is calculated with (squares) and without (triangles) the diffraction phase correction. The visibility model is based on the 7% rate for spontaneous scattering of readout light photons and the number of atoms illuminated by the readout beam for different n . The lower panel shows the CI phase standard deviation vs n , with (diamonds) and without (circles) the diffraction phase correction.

Visibility vs n

Interferometric visibility is defined as $100 \cdot (\text{Max} - \text{Min}) / (\text{Max} + \text{Min})$, where Max and Min are the highest and lowest points in the fringe pattern. The CI gives a complete fringe pattern in a single shot making it possible to define a single shot visibility, as opposed to other interferometers where a fringe pattern is only observable by combining many shots. However, if we want our visibility metric to be comparable to other interferometers, we should only consider the visibility of the CI signal averaged over many shots.

There are two obvious mechanisms by which visibility can be reduced: shot-to-shot variation in the interferometer phase, and decreased signal/noise ratio. Fig 5.7 illustrates how phase variation affects visibility. The right panel shows the amplitude of the average of many sinusoids with phases drawn from a Gaussian population with a given width. The Gaussian width is directly comparable to the standard deviation of the CI phase, shown in the lower panel of Fig 5.6. From this we can conclude that even for our highest n with $\text{std} = .06$ radians, the amplitude reduction is on the order of 10%. For our first pass at analyzing visibility vs n , we neglect this relatively small effect and focus on signal to noise considerations.

We find that the behavior of our visibility vs n data is captured with a model based only on signal/noise ratio. As shown in Fig. 5.8, we model the expected noise based on the number of atoms in position to spontaneously scatter a readout light photon (which decreases with n) and the 7% spontaneous scattering rate. This model is shown as the dashed blue curve in Fig. 5.6

Diffraction phase correction

We monitor the diffraction beam intensities during the CI sequence, and record the pulse profile for each shot. As shown in Fig. 5.9, we observe a correlation between the total area of the diffraction pulse sequence and the recorded CI phase. We apply a shot-by-shot diffraction phase correction based on this correlation, as shown in Fig. 5.6.

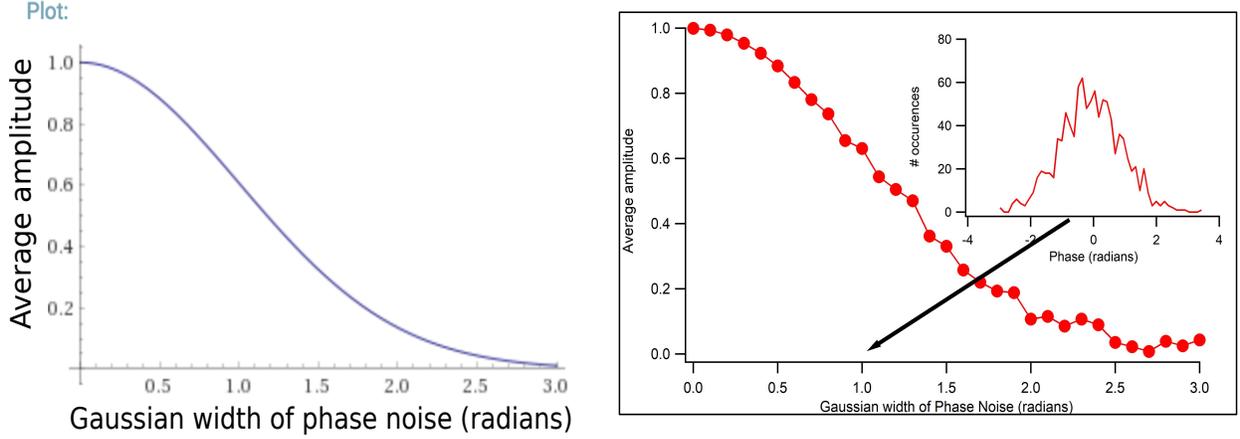


Figure 5.7: In order to simulate the reduction in visibility expected due to signals with a spread of phases being averaged together, we plot the visibility of the sum of many sine waves, each with a different phase, versus the standard deviation of the summed phases. We observe that the reduction in visibility due to this effect at the highest phase standard deviation observed in our work is 10%.

5.2 Current Capabilities

We now turn to the application of the large n CI technique towards a photon recoil and α measurement. The precision in ω_{rec} can be written as:

$$\frac{\delta\omega_{\text{rec}}}{\omega_{\text{rec}}} = \frac{\delta\Phi}{\Phi} = \frac{\delta\Phi}{\frac{1}{2}n^2\omega_{\text{rec}}\Delta T\sqrt{M}} \quad (5.3)$$

where $2\Delta T$ is the range of free evolution times over which the slope of $\Phi(2T)$ is measured, $\delta\Phi$ is the uncertainty in the measured CI phase at a given value of T , and M is the number of experimental shots. The best (lowest) uncertainty is obtained when the quantity $n^2\omega_{\text{rec}}T/2$ is maximized. In our current CI setup, the total interferometer time is constrained by the atoms falling out of the horizontally oriented diffraction beams. This means that there is a tradeoff between n and T , since the acceleration pulses that increase n take a significant amount of time. Thus, increasing n means reducing T . Since n helps us quadratically while

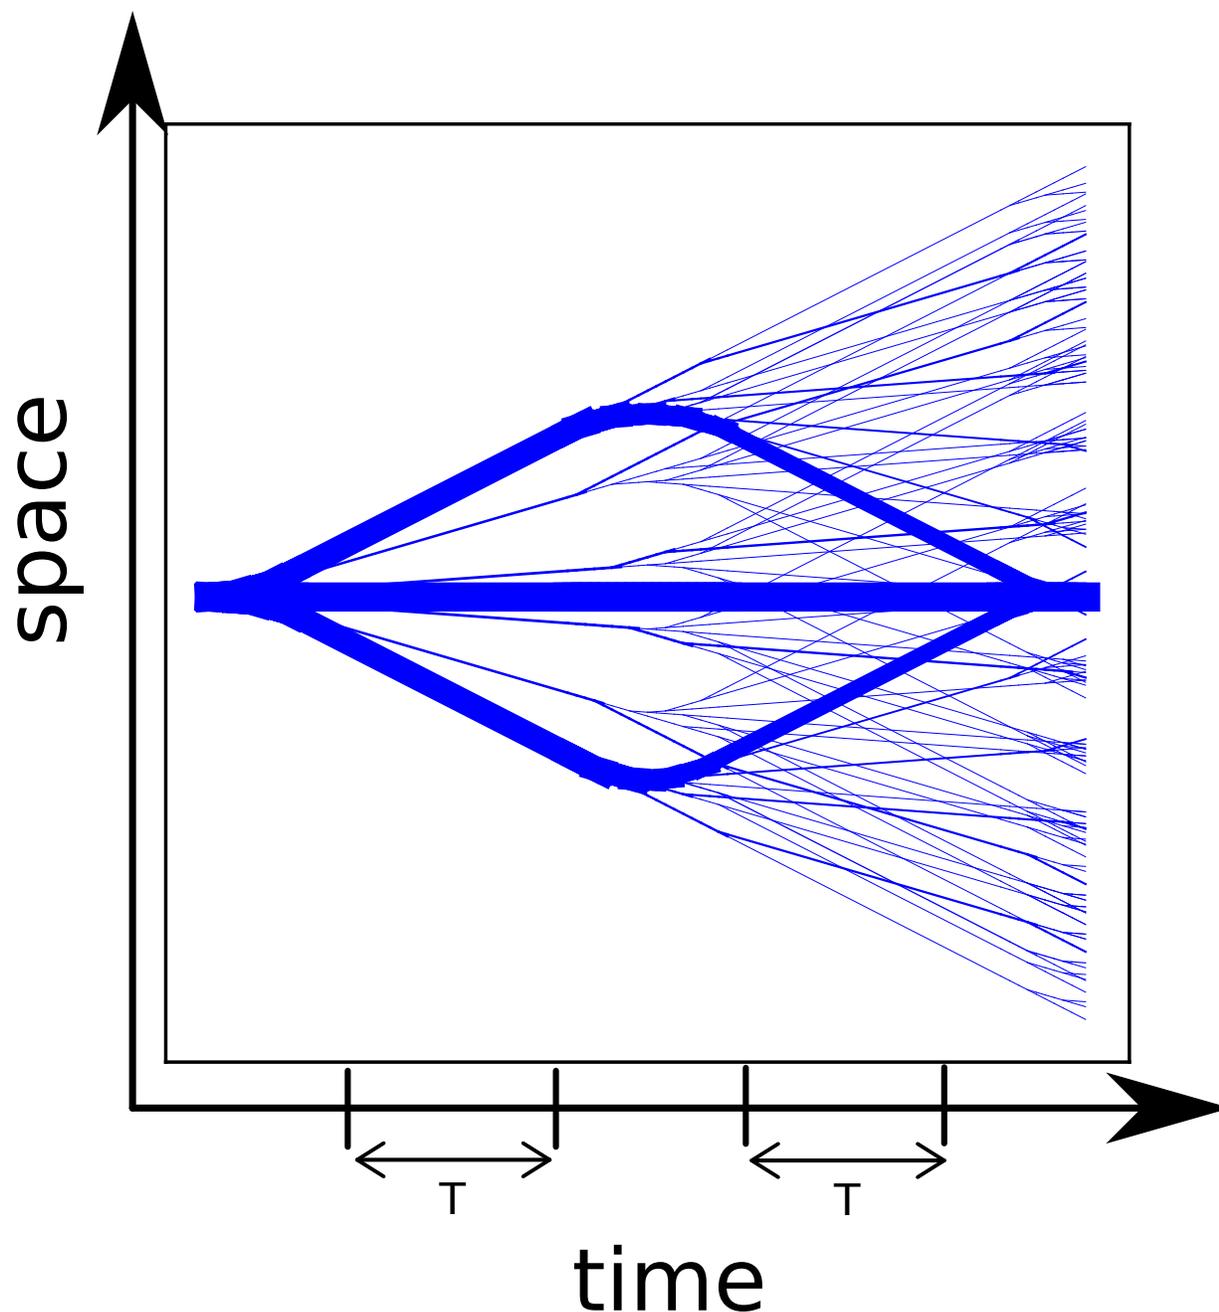


Figure 5.8: Visualization of a simulation of all momentum state trajectories in the CI sequence. We usually only consider the intended paths, neglecting the paths of atoms which are left behind due to inefficiencies of each diffraction pulse. The amplitude along each path is represented by the thickness of the lines. However, a minimum thickness is imposed, so that all paths are visible. We use the simulation that underlies this image to calculate the total number of atoms that arrive in the area that the readout light hits, in order to predict the level of spontaneously scattered light that forms a background to the CI signal.

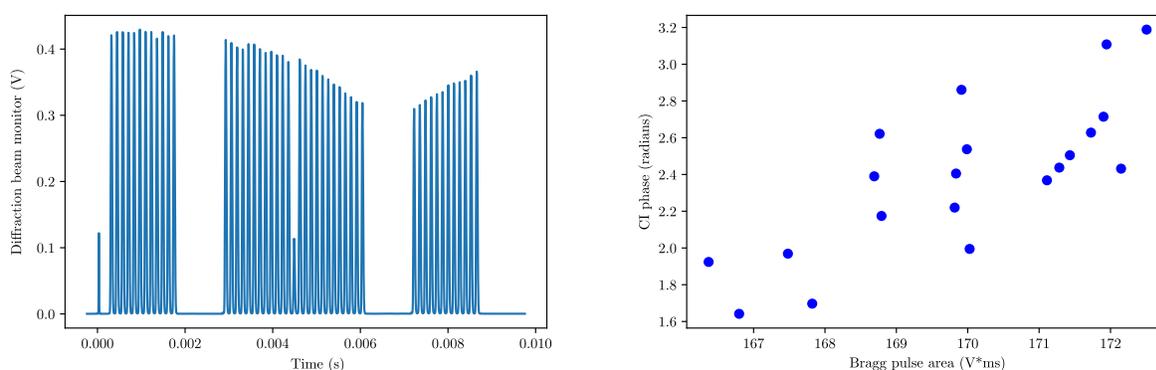


Figure 5.9: Left: typical output of the diffraction beam intensity monitor photodiode. This example is for $n = 76$. The diffraction beams enter the main chamber at an angle slightly offset from 90 degrees, and the reflection from the viewport is directed to the diffraction beam monitor photodiodes. This plot shows the signals from the two channels multiplied together. The heights of the acceleration pulses vary over the course of the sequence in order to maintain the π -pulse condition as the atoms pass through different parts of the diffraction beams. Right: the correlation of the total diffraction pulse area vs the corresponding measured CI phase. We implement the diffraction phase correction by fitting a linear model to this correlation, and subtracting the fitted slope from the phase values. As shown in Fig. 5.6, the standard deviations of the corrected phases are substantially lower for n up to 88.

T only helps linearly, we'd prefer to increase n at the expense of T . However, we do need some T for the experiment to function. We can calculate the ideal ratio of n to T given our experimental parameters.

As discussed in section 3.3.6, we accelerate the atoms using 3rd order Bragg pulses with a length of $130\mu\text{s}$. We accelerate the paths sequentially, so each pair of 3rd order pulses takes $260\mu\text{s}$, and increases n by 12 (six photon recoils in each direction). The CI sequence has four instances of acceleration, so the marginal time increase per set of acceleration pulses is $t_{\text{acc}} = 1040\mu\text{s}$. Also, recall that the lowest possible value of n is 4.

The total time of the interferometer T_{tot} can be broken down into components, where $t_{n=4}$ is the time taken by diffraction pulses to make the $n=4$ CI, and $t_{\text{acc}}(\frac{n-4}{12})$ is the time taken by acceleration pulses for a given n . ΔT (the parameter in Eq. 5.3) is bounded from above by T .

$$T_{\text{tot}} = 2T + t_{n=4} + t_{\text{acc}}\left(\frac{n-4}{12}\right) \quad (5.4)$$

Assuming that $t_{n=4} \ll T_{\text{tot}}$, and $n \gg 4$, we simplify and re-arrange and get an expression for T :

$$T = \frac{T_{\text{tot}}}{2} - \frac{t_{\text{acc}}n}{24} \quad (5.5)$$

We can then calculate the total phase evolution as a function of n :

$$\phi_{\text{tot}} = \frac{1}{4}n^2\omega_{\text{rec}}\left(T_{\text{tot}} - t_{\text{acc}}\frac{n}{12}\right) \quad (5.6)$$

ϕ_{tot} is maximized when the derivative with respect to n is zero:

$$0 = n\omega_{\text{rec}}T_{\text{tot}} - \omega_{\text{rec}}t_{\text{acc}}\frac{n^2}{8} \quad (5.7)$$

$$n = \frac{8T_{\text{tot}}}{t_{\text{acc}}} \quad (5.8)$$

This maximum occurs when the total time spent accelerating $t_{\text{acc}}(\frac{n}{12})$ is 2/3 of T_{tot} . Applying this ratio to our currently available parameters, we find that the quantity $n^2\omega_{\text{rec}}T/2$ is optimized to 2.1×10^5 for $n = 76$, and $T = 3$ ms. For these parameters, the maximum separation of interfering states is 1.5mm. The observed $\delta\Phi=.25$ radians at $n = 76$ then gives a precision of 8.7×10^{-8} in ω_{rec} in 200 shots.

5.3 Looking Forward

As can be seen from Eq. 5.3, we can increase the precision of our measurement in four ways: Decrease $\delta\Phi$, Increase the number of shots M , increase the free evolution time T , or increase the momentum splitting n . I'll address the prospects for improving each of these parameters.

5.3.1 M : Faster BEC production

To paraphrase Gerry Gabrielse, precision measurement via statistical scaling is an act of desperation, i.e. what you do when you don't have any better ideas. Anyone with patience can keep taking more and more data. That said, there is some room for cleverness in decreasing the cycle time of the experiment, so that more shots can be taken in a given amount of time, hence increasing M .

The interferometer cycle time is dominated by the BEC production time. While this is 10s for this work, we have demonstrated Yb BEC cycle times as low as 1.6s in our group [33]. We anticipate that we can get closer to this result by modifying our optical trap to take better advantage of the painting technique mentioned in section 3.2.2. Using 3s as a reasonable benchmark for longterm measurements, the above numbers scale to 1.1×10^{-8} in ω_{rec} in 10 hrs of integration time.

5.3.2 $\delta\Phi$

We see in Fig. 5.6 that $\delta\Phi$ grows with n . The number of atoms contributing to the signal, and hence the number of photons at the detector per shot, decreases with higher n due to the finite efficiency of the acceleration pulses. One consequence of this is that the expected shot noise increases with n . Fig. 5.6 shows that, while we are at the shot noise limit for low n , our measured $\delta\Phi$ grows faster than the shot noise limit with n .

We don't have a quantitative model for why the observed $\delta\Phi$ is higher than shot noise, but we can make some educated guesses.

First, we note that until around $n = 88$, the diffraction phase correction is significant, reducing $\delta\Phi$ to near the shot noise limit. In this work, the diffraction phase correction is based on the correlation between the CI phase and the total area of all the diffraction pulses in the sequence. Ideally, the correction would be done on a per-pulse basis. This kind of per-pulse correction was attempted, but found to be ineffective with the current setup. The likely reason for this is that the time resolution of the diffraction beam monitor data is insufficient to capture the detailed variation of each individual pulse. The oscilloscope that records these data takes 2500 voltage samples per shot, with a spacing determined by the time axis settings of the scope. Since the total interferometer time is on the order of 10ms, there are around 30 samples per $130\mu\text{s}$ pulse. This is fairly sparse data at the peak of the pulse.

This is all to say that there is reason to believe that with more fine grained diffraction pulse monitor data, it may be possible to do a per-pulse diffraction phase correction, which could be more effective than the total area correction that is currently implemented. This more detailed correction could reduce $\delta\Phi$ to the shot noise limit for n up to 88.

For n above 88, the diffraction phase correction becomes less effective, and $\delta\Phi$ gets significantly higher than the shot noise limit. One factor that likely contributes to this is that the sensitivity of the CI phase to diffraction pulse timing increases for higher n . The

higher the rate of phase accumulation, the more significant are variations in pulse timing. In the current setup, the stability of the timing of the diffraction pulses can only be confirmed at the 100ns level, due to the above mentioned sparseness of the diffraction pulse monitor data. Combining a 100ns timing uncertainty with the phase accumulation rate for various n suggests that this timing uncertainty starts to be significant around $n = 88$. This then is a prime suspect for why $\delta\Phi$ begins to increase for this n .

We have confirmed the stability of the analog and digital outputs from Cicero at the 1ns level, so if the pulse timing is fluctuating at the 100ns level, the variation is due to something else. Suspicion immediately falls on the diffraction beam intensity feedback circuit. Changes in the response of this circuit, likely due to thermal effects, could lead to shot-to-shot variations in the diffraction pulse profiles that would constitute variations in the effective times of the pulses.

In conclusion, better diffraction beam intensity stabilization electronics could reduce pulse timing variations, which could improve $\delta\Phi$ for higher n .

5.3.3 n and T

Due to the limitation of the atoms falling out of the horizontally oriented diffraction beams, it is clear that the diffraction beam configuration needs to be updated in order to increase n and T . There are several different ways forward from this point, but one key decision is whether to stick with horizontally oriented diffraction beams, or move to vertical oriented beams. I'll discuss some of the advantages and challenges of both choices.

5.3.4 *Prospects for updated horizontal diffraction beam CI*

As discussed in section 5.2, for our current acceleration rate and total interferometer time, we are already operating with the best possible parameters, where 2/3 of our total time is spent accelerating the atoms. In order to do better, we need to either accelerate faster, or

increase the total interferometer time.

A simple strategy to increase total interferometer time would be to use diffraction beams with a larger waist. The amount that can be gained in this way is limited by available laser power. Increased waist size means a decrease in optical intensity at the atoms. With our current waist size, the overhead in available intensity is not large. In a pinch, we may be able to increase the waist by 50% and just barely maintain sufficient intensity. In addition, we may be possible to decrease detuning, in order to increase the Rabi frequency for a given optical power. At our current detuning of 1300 linewidths, we observed some atom loss due to spontaneous scatter while doing BO, while spontaneous scatter did not seem to be an issue for 3rd order Bragg pulses with our parameters. For now, I'll assume we stick with the current waist and detuning, which allows, conservatively, 10ms of interferometer time.

In the current CI sequence, the atoms start near the top of the diffraction beam, and fall to the bottom. The available time can be doubled by starting the sequence with an upward launch.

We currently accelerate the atoms with Bragg pulses addressed sequentially to paths 1 and 3. As discussed in 3.3.1, it is possible to address paths 1 and 3 simultaneously with what I called DiffPlan3, which uses a Pockels cell as a rapidly adjustable waveplate. If this scheme could be made to work, it would increase the acceleration rate by a factor of 2.

At this point, it must be faced that the efficiency of the 3rd order Bragg acceleration pulses is too low to go much beyond our current benchmark of $n = 112$. We simply run out of atoms at that point. We have demonstrated much higher acceleration efficiency, and rate, using Bloch oscillations, but found unacceptably high diffraction phase fluctuations when Bloch oscillations were applied to the CI. Given the 99.9% efficiency per $\hbar k$ seen for BO in section 3.3.4, compared to 98.5% per $\hbar k$ for 3rd order Bragg pulses with our parameters, if the diffraction phases accrued during BO could be tamed then much higher n could be attained, and our precision goals could be met in the small total interferometer times allowed

by the horizontal beam geometry. Analyzing and controlling BO diffraction phases is thus one future research direction. One concern with this approach is that the requirements for time measurement resolution increase with the phase accumulation rate, which increases with n^2 . So for approaches that rely on very high n , care must be taken that it is experimentally possible to measure the free evolution intervals to the required precision.

5.3.5 *Two horizontal diffraction beams*

There is another potential way that a horizontal diffraction beam CI could be used given our current acceleration abilities, which is to use a second parallel diffraction beam. The basic insight that leads to this approach is that the final deceleration sequence of the CI is part of the fixed offset phase. It is critical that the first three accelerations sequences impart the same momentum change, because they determine the kinetic energy during the two free evolution intervals. The final deceleration sequence needs to slow the atoms to the correct final momentum states for the CI to function, but as long as it is the same shot-to-shot, it need not be exactly the same as the first three.

The “heart shaped” CI takes advantage of this by using two parallel horizontal diffraction beams, and one vertical “elevator” beam for vertical motion. High precision requires not just large T , but large ΔT , the difference between two available T configurations with sufficiently identical acceleration pulses. The heart shaped CI relies on the elevator beam to move the atoms vertically so that the final deceleration pulse always happens in the same beam.

It would work like this:

For the longer T : The atoms start in the upper beam with an upward launch. The initial splitting takes place in the upper beam. The first free evolution time takes place while the atoms move up, halt, and then start moving downward under gravity. As the atoms pass back through the upper beam, they are decelerated, mirrored, and re-accelerated. The second free evolution interval takes place as the fall below the upper beam. Then the final

deceleration happens in the lower beam.

For the shorter T : The sequence begins as above, with a slower initial launch so that the atoms return to the upper beam after a shorter free evolution time. With no intervention, the atoms would not fall far enough to reach the lower beam during the second free evolution time. In order for the final deceleration to happen in the lower beam, the elevator beam is used to quickly move the atoms to the lower beam, where they are decelerated using the same sequence as in the longer T .

The weakness in this scheme is that the short and the long T configuration differ in that one uses the elevator beam mid sequence while the other does not. In principle, the elevator beam is perpendicular to the diffraction axis, so it should not impart momentum in the axis relevant to the measurement. It is the case, however, that the elevation happens while the three paths are separated in space. If the elevator beam imparts any differential phase to the three paths due to the paths occupying different positions in the beam, this could be a systematic that would affect the measurement.

The advantage of this scheme is that it uses horizontal diffraction beams that have already been demonstrated in the lab. In addition, the CI machine is currently optimized for readout light access on the horizontal plane.

5.3.6 *Vertical diffraction beams*

In contrast to horizontal diffraction beams, vertically oriented beams:

- Allow for much longer T , limited only by the size of the vacuum chamber, since gravity does not move the atoms out of the diffraction beams.
- Simplify diffraction phase calculations, because the atoms remain in the same part of the beam for the duration of the experiment.

- Require that the relative frequencies of the diffraction beams be swept to account for the motion of the atoms due to gravity, which is now along the diffraction axis.
- Have some optics complications due to the CI machine being primarily optimized for the horizontal geometry. In particular, the 399nm readout scheme may be difficult to implement with vertical beams.

The great advantage of vertical beams is that, as calculated below, it allows us to reach our precision goals by increasing T , while using the n that has already been demonstrated. Given that the n shown in this work is novel, while much higher T has been seen in other experiments, this seems like a more reliable path forward. It is less certain how much higher n can be raised, compared to the relative certainty that T can be increased.

The requirement to compensate for the motion due to gravity is a new technical complication, but not an insurmountable one. We have planned out the electronics required for performing this sweep in tandem with the established CI sequence.

The biggest immediate obstacle to the vertical scheme is difficulty with CI machine vacuum system. The main chamber we designed with the horizontal CI in mind. There are two issues: one is that the 45 degree angle needed for the 399nm readout is difficult or impossible to access for the vertical CI. We are reluctant to return to the 556nm readout because the 399nm readout was easier to set up and has better signal to noise. However, we did eventually get good results with 556nm readout in the first round CI, so it would not be the end of the world to return to it.

A potentially more difficult issue is that the top and bottom viewports of the CI machine main chamber are oriented close to perfectly horizontally. The vertical diffraction beams must be parallel with gravity, which means they will be hitting the viewports at a right angle. Such an angle of incidence is undesirable, since viewport reflections are directed back onto the beam. Further, since the top and bottom viewport are parallel with each other, they

may form a weak optical cavity. We have seen bad effects from an optical cavity forming between flat ends of optical fibers in the diffraction path. It is possible that the weaker reflections from the AR coated viewports will not have this effect, but if it is there, then there is no easy solution short of breaking vacuum and reconfiguring the main chamber.

5.3.7 Projections for Vertical CI

If vertical diffraction beams are implemented, the limitation on the interferometer free evolution time is lifted and $2T = 240$ ms is possible in the 10 cm vertical region available in the current main chamber. Even larger $2T$ would be possible in a slightly modified main chamber. The precision scaling equation then indicates a precision of 2.8×10^{-10} in ω_{rec} in 10 hrs. The corresponding precision in α , which can be determined by combining ω_{rec} and measurements of other fundamental constants [38] is a factor of two lower. Together with potential improvements in $\delta\Phi$ from better pulse timing control, n from more efficient Bragg pulses, or T from a modified vacuum chamber, this approach holds promise for a competitive measurement of α [17, 4, 2].

BIBLIOGRAPHY

- [1] S. Aaronson. *Quantum Computing Since Democritus*. Quantum Computing Since Democritus. Cambridge University Press, 2013.
- [2] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio. Tenth-Order QED Contribution to the Electron $g-2$ and an Improved Value of the Fine Structure Constant. *Phys. Rev. Lett.*, 109:111807, 2012.
- [3] P. Asenbaum, C. Overstreet, T. Kovachy, D.D. Brown, J.M. Hogan, and M.A. Kasevich. Phase Shift in an Atom Interferometer due to Spacetime Curvature across its Wave Function. *Phys. Rev. Lett.*, 118:183602, 2017.
- [4] R. Bouchendira, P. Clade, S. Guellati-Khelifa, F. Nez, and F. Biraben. New determination of the fine structure constant and test of the quantum electrodynamics. *Phys. Rev. Lett.*, 106:080801, 2011.
- [5] M. P. Bradley, J. V. Porto, S. Rainville, J. K. Thompson, and D. E. Pritchard. Penning Trap Measurements of the Masses of ^{133}Cs , $^{87,85}\text{Rb}$ and ^{23}Na with Uncertainties ≤ 0.2 ppb. *Phys. Rev. Lett.*, 83:4510, 1999.
- [6] S. B. Cahn, A. Kumarakrishnan, U. Shim, T. Sleator, P. R. Berman, and B. Dubetsky. Time-domain de broglie wave interferometry. *Phys. Rev. Lett.*, 79:784–787, Aug 1997.
- [7] S. Chiow, T. Kovachy, H. Chien, and M. A. Kasevich. $102\hbar k$ large area atom interferometers. *Phys. Rev. Lett.*, 107:130403, 2011.
- [8] A.D. Cronin, J. Schmiedmayer, and D.E. Pritchard. *Rev. Mod. Phys.*, 81:1051, 2009.
- [9] D.S. Durfee, Y.K. Shaham, and M.A. Kasevich. . *Phys. Rev. Lett.*, 97:240801, 2006.
- [10] Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.
- [11] Abbott et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.

- [12] C.J. Foot. *Atomic physics*. Oxford master series in physics. Oxford University Press, 2005.
- [13] S. Fray, C.A. Diez, T.W. Hänsch, and M. Weitz. Atomic Interferometer with Amplitude Gratings of Light and Its Applications to Atom Based Tests of the Equivalence Principle. *Phys. Rev. Lett.*, 93:240404, 2004.
- [14] S. Gupta, K. Dieckmann, Z. Hadzibabic, and D. E. Pritchard. Contrast Interferometry using Bose-Einstein Condensates to Measure h/m and α . *Phys. Rev. Lett.*, 89:140401, 2002.
- [15] S. Gupta, A. E. Leanhardt, A. D. Cronin, and D. E. Pritchard. Coherent Manipulation of Atoms with Standing Light Waves. *Cr. Acad. Sci. IV-Phys*, 2:479, 2001.
- [16] G. Hadley. Injection locking of diode lasers. *IEEE Journal of Quantum Electronics*, 22(3):419–426, Mar 1986.
- [17] D. Hanneke, S. Fogwell, and G. Gabrielse. New measurement of the electron magnetic moment and the fine structure constant. *Phys. Rev. Lett.*, 100:120801, 2008.
- [18] A. O. Jamison. Precision interferometry with bose-einstein condensates. *University of Washington PhD Thesis*, 2014.
- [19] A. O. Jamison, J. N. Kutz, and S. Gupta. Atomic interactions in precision interferometry using bose-einstein condensates. *Phys. Rev. A.*, 84:043643, 2011.
- [20] A. O. Jamison, B. Plotkin-Swing, and S. Gupta. Advances in precision contrast interferometry with yb bose-einstein condensates. *Phys. Rev. A*, 90:063606, 2014.
- [21] A. Jayakumar, B. Plotkin-Swing, A.O. Jamison, and S. Gupta. Dual-axis vapor cell for simultaneous laser frequency stabilization on disparate optical transitions. *Rev. Sci. Instr.*, 86:073115, 2015.
- [22] Aviv Keshet and Wolfgang Ketterle. A distributed, graphical user interface based, computer control system for atomic physics experiments. *Review of Scientific Instruments*, 84(1):015105, 2013.
- [23] Jeongwon Lee, Jae Hoon Lee, Jiho Noh, and Jongchul Mun. Core-shell magneto-optical trap for alkaline-earth-metal-like atoms. *Phys. Rev. A*, 91:053405, May 2015.

- [24] L. Marton. Electron interferometer. *Phys. Rev.*, 85:1057–1058, Mar 1952.
- [25] G.D. McDonald, C.C.N. Kuhn, S. Bennetts, J.E. Debs, K.S. Hardman, M. Johnsson, J.D. Close, and N.P. Robins. $80\hbar k$ momentum separation with Bloch oscillations in an optically guided atom interferometer. *Phys. Rev. A*, 88:053620, 2013.
- [26] J.M. McGuirk, G.T. Foster, J.B. Fixler, M.J. Snadden, and M.A. Kasevich. Sensitive absolute-gravity gradiometry using atom interferometry. *Phys. Rev. A*, 65:033608, 2002.
- [27] H. Muller, S. Chiow, Q. Long, S. Herrmann, and S. Chu. Atom Interferometers with Scalable Enclosed Area. *Phys. Rev. Lett.*, 100:180405, 2008.
- [28] H. Müntinga, H. Ahlers, M. Krutzik, A. Wenzlawski, S. Arnold, D. Becker, K. Bongs, H. Dittus, H. Duncker, N. Gaaloul, C. Gherasim, E. Giese, C. Grzeschik, T. W. Hänsch, O. Hellmig, W. Herr, S. Herrmann, E. Kajari, S. Kleinert, C. Lämmerzahl, W. Lewoczko-Adamczyk, J. Malcolm, N. Meyer, R. Nolte, A. Peters, M. Popp, J. Reichel, A. Roura, J. Rudolph, M. Schiemangk, M. Schneider, S. T. Seidel, K. Sengstock, V. Tamma, T. Valenzuela, A. Vogel, R. Walser, T. Wendrich, P. Windpassinger, W. Zeller, T. van Zoest, W. Ertmer, W. P. Schleich, and E. M. Rasel. Interferometry with bose-einstein condensates in microgravity. *Phys. Rev. Lett.*, 110:093602, Feb 2013.
- [29] L. Nenadovic and J.J. McFerran. Clock and inter-combination line frequency separation in ^{171}Yb . *J. Phys. B.*, 49:065004, 2016.
- [30] A. Peters, K.Y. Chung, and S. Chu. High-precision gravity measurements using atom interferometry. *Metrologia*, 38:25, 2001.
- [31] Norman F. Ramsey. A new molecular beam resonance method. *Phys. Rev.*, 76:996–996, Oct 1949.
- [32] H. Rauch, W. Treimer, and U. Bonse. Test of a single crystal neutron interferometer. *Physics Letters A*, 47:369–371, April 1974.
- [33] R. Roy, A. Green, R. Bowler, and S. Gupta. Rapid cooling to quantum degeneracy in dynamically shaped atom traps. *Phys. Rev. A.*, 93:043403, 2016.
- [34] D. Schlippert, J. Hartwig, H. Albers, L.L. Richardson, C. Schubert, A. Roura, W.P. Schleich, W. Ertmer, and E.M. Rasel. Quantum Test of the Universality of Free Fall. *Phys. Rev. Lett.*, 112:203002, 2014.

- [35] S. Sturm, F. Kohler, 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. *Nature*, 506:467, 2014.
- [36] Alex Sugarbaker. *Atom interferometry in a 10 m fountain*. PhD thesis, Stanford, 2014.
- [37] T. Udem, A. Huber, B. Gross, J. Reichert, M. Prevedelli, M. Weitz, and T. W. Hansch. Phase-Coherent Measurement of the Hydrogen $1S \rightarrow 2S$ Transition Frequency with an Optical Frequency Interval Divider Chain. *Phys. Rev. Lett.*, 79:2646, 1997.
- [38] D. S. Weiss, B. C. Young, and S. Chu. Precision Measurement of the Photon Recoil of an Atom Using Atomic Interferometry. *Phys. Rev. Lett.*, 70:2706, 1993.