

Precision Interferometry with Bose-Einstein Condensates

Alan O. Jamison

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

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Subhadeep Gupta, Chair

J. Nathan Kutz

Blayne Heckel

Program Authorized to Offer Degree:
Physics

©Copyright 2014

Alan O. Jamison

University of Washington

Abstract

Precision Interferometry with Bose-Einstein Condensates

Alan O. Jamison

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

This dissertation describes the creation of the first matter-wave interferometer using ytterbium (Yb) atoms. Most of the experiments focus on a contrast interferometer geometry with a Bose-Einstein condensate (BEC) as source. The recoil frequency of the ^{174}Yb atom is measured with this interferometer. The recoil frequency of an atom is part of a set of precision measurements that together give a value for the fine structure constant.

The experimental results of this dissertation lay the groundwork for a future sub-part-per-billion (ppb) precision measurement of the Yb recoil frequency. The contrast interferometry technique is extended to substantially longer times scales than those achieved in previous experiments. A measurement at the ~ 10 parts-per-million level is made. Systematic effects and statistical scaling are studied and found to be compatible with the desired sub-ppb precision for a future measurement.

Such a measurement requires a detailed theoretical study of possible systematic shifts to the measured value. A substantial portion of this dissertation consists of this analysis, carried out in sufficient generality as to guide future sub-ppb level measurements. In addition to a large number of possible systematic shifts due to well-understood physics, two more complex effects are identified and studied: Diffraction phases and atom-atom interactions.

ACKNOWLEDGMENTS

My time at the University of Washington has been wonderful. And while, in due course, I want to thank many of the people who made it as such, I feel I would be remiss in not reaching back much further into my past.

Ron Bradley, my middle school science teacher, taught me perhaps the most important thing I know about science: That it demands constant skepticism, questioning of authorities and received wisdom, plus just a dash of out-right rebellion. He also gave me my first dizzying glimpse into the strangeness of quantum mechanics. My first science class after Mr. Bradley was Mark Schiffer's Biology class. Mark also had a rebellious streak, perhaps unknown to most of his students. From him I learned to further appreciate the joy and wonder of digging deep into the details of complex systems. Together they impressed upon me the fundamental tension of science: Being willing to tear down any idea or structure, you can build up tremendously powerful structures of ideas.

Dave Pritchard's encouragement helped set me on the path to UW. As Dave once said to me, "I gave you a job (at his online education startup), I wrote you a great letter, and I trained a great advisor for you." All true, and all deeply appreciated.

That advisor was Deep Gupta. Deep is a tremendous physicist. I've learned a great deal from our discussions over the years. He possesses an uncanny ability to diagnose and fix problems in the lab. He's also a great person. I feel quite fortunate to have had such an understanding and amiable advisor. His infectious optimism and love of science have been key to my continued productivity over the years.

Before joining Deep's lab, I spent two years with the particle theory group. In that

time, I learned a lot from Ann Nelson, my advisor, and Larry Yaffe. Each of them taught me a lot about how to think like a physicist. In my work with Ann I gained an appreciation for effective field theory that has continued to shape how I look at physics problems outside of particle theory. She was also kind enough to work with me on some ideas around dark matter BECs long after I'd moved to the basement.

My route from particle theory to Deep's lab ran through the office of Nathan Kutz in the Applied Math Department. Nathan has been a really valuable mentor, in addition to acting as a co-advisor for my work on atomic interactions. He taught me much about math and numerical simulations. He also makes fantastic cappuccinos.

When I joined the lab, two other students from my entering class, Anders Hunnes Helgesen Hansen (a.k.a. "Ahhh") and Alex Khramov (a.k.a. "Alexander Yurivich"), had already been hard at work building the machine for almost two years. They both taught me a great deal about how to make experiments work. We each had fairly different perspectives on physics, which lead to really productive collaboration. It was an honor to be part of the triple A team. We also had a lot of fun both inside and outside of the lab. Anders is a talented beer brewer and unquestioned authority on all things Tolkien. Alex and I had a great number of common interests outside of physics, from linguistics to progressive metal. I especially appreciate the aspects of Russian culture he introduced me to, such as throwing Christmas trees out of second-story windows and *The Master and Margarita*. Anders and Alex both fled to their homelands after graduation and have been very much missed.

Will Dowd's up-beat assurance that the experiment will be working soon has helped us push through more than a few hard times in the lab. Will is also a voracious reader, and I've enjoyed talking to him about a variety of books.

Vlad Ivanov was a post doc in our lab who taught me quite a lot about laser cooling, optics, and how to tame a poorly designed Ti:Saph laser. He also set a great

example for how to constructively deal with experimental frustration.

Ben Plotkin-Swing has been the other graduate student enduring interferometry along with me for the past three years. Ben has great breadth as a physicist, being equally at home building things in the lab and talking about the philosophy of science. He's also a fellow connoisseur of bad puns. I've enjoyed our time together, even when that time has come at 3AM on a long data-taking run.

Ricky Roy joined the mixture team after I had largely extracted myself from it to focus on interferometry. Nonetheless, we've had a lot of great discussions and some great hikes. Between Ben and Ricky, I know the Gupta group will be in good hands once Will and I clear out.

Raúl Briceño, Brian Mattern, and John Vinson were good friends and great teachers of physics and programming from our first year at UW right through to their graduations, and on to the present. I could go on for many pages but I should cap these acknowledgments with a generic thank you to all of the others—fellow grad students and professors—who have contributed to my education in physics.

My many great friends outside the physics department have also contributed in keeping me grounded and reminding me there is life outside the lab. My mother Miriam, brother Steve, and sister Michele have all been supportive and encouraged me throughout my life, not just in my grad school career. The memory of my father Pete, who passed away during my time at UW, continues to support and encourage me, as well.

Saving the best for last, I must thank my amazing wife, Kim Rivera. She has been a source of joy and inspiration since the day we met in Prof. Shivamoggi's PDE class at UCF. She keeps me calm and restores my optimism when the inevitable hardships of research have me agitated and morose. We've had more fun together the past eight years than I would have imagined possible. ¡Gracias, mi amor!

TABLE OF CONTENTS

	Page
List of Figures	iv
Chapter 1: Introduction	1
1.1 “Precision” Measurements	1
1.2 Interferometry	4
1.3 The Perfect Quantum System?	5
Part I: Background	7
Chapter 2: Cooling Atoms with Blinding Hot Light	8
2.1 The Physics of Laser Cooling	8
2.2 Ytterbium	14
2.3 Evaporative Cooling	17
2.4 Bose-Einstein Condensation	22
Chapter 3: Atom Interferometry	28
3.1 Interference with Light	28
3.2 Interference with Atoms	36
3.3 BECs in Interferometry	40
Chapter 4: The α and the ω_{rec}	44
4.1 The Fine Structure Constant	44
4.2 History of Recoil Measurements	46
4.3 $g - 2$ and Tests of QED	50
Part II: Contrast Interferometry	56

Chapter 5:	The Light Pulse Contrast Interferometer: Theory	57
5.1	Geometry and Readout Mechanism	57
5.1.1	Basics	57
5.1.2	Advantages of Contrast Interferometry	61
5.2	Optical Standing-Wave Gratings	64
5.2.1	Atom-Light Interactions	65
5.2.2	Modeling Diffraction	68
Chapter 6:	The Light Pulse Contrast Interferometer: Experiments	75
6.1	Building the Interferometer	75
6.1.1	Optical Setup	77
6.1.2	Alignment, Calibration, and the First Recoil Measurement	83
6.2	Finding the Signal	88
6.3	Analysis and Small T Results	91
6.4	Improving the Readout	95
6.4.1	Better Equipment	96
6.4.2	Decoupling Diffraction and Readout	98
6.5	Large T Results	99
6.5.1	Vibrational Immunity	99
6.5.2	Below 10 ppm and Scaling to Better Precision	102
6.5.3	One Detuning, Two Detuning, Red Detuning, Blue Detuning	104
6.6	To BEC, or Not to BEC	105
6.7	Recommendations for Future Experiments	107
Part III:	Systematics	110
Chapter 7:	Systematic Shifts in a Contrast Interferometer	111
7.1	Reference Frame Effects	111
7.2	External Field Effects	115
7.2.1	Gravitational Fields	115
7.2.2	Electromagnetic Fields	117
7.3	Diffraction Beam Effects	118
7.3.1	Beam Geometry	119
7.3.2	Counter-propagation	125

7.3.3	AC Stark Effect	126
7.3.4	Index of Refraction	127
Chapter 8:	Diffraction Phases	132
8.1	Basic Framework	132
8.2	Refinements	135
8.2.1	Non-Zero Momentum Width	135
8.2.2	Imbalanced Laser Beams	136
8.3	Calibration and Diffraction Phase Calculation	136
8.4	Results for Current Experiments	138
Chapter 9:	Atomic Interactions in a BEC Interferometer	141
9.1	Ultracold Interactions	141
9.2	Mean-Field Theory	143
9.3	Scaling Solutions	148
9.3.1	Deriving the Scaling Solutions	148
9.3.2	Numerical Simulations	153
9.4	Slowing-Varying Envelope Approximation	157
9.5	Modeling a Complete Experiment: Analytic Results	162
9.5.1	Modified Kinematics and Scaling Behavior	163
9.5.2	Effects of Phase Curvature	167
9.5.3	Summary of the Physics	169
9.6	Modeling a Complete Experiment: Numerical Results	172
Appendix A:	Effects of gravitational confinement on bosonic asymmetric dark matter in stars	179
Bibliography	182

LIST OF FIGURES

Figure Number	Page
1.1 Accuracy vs. Precision	2
2.1 Heuristic diagram of Zeeman slower	10
2.2 Cartoon of a Magneto-Optical Trap	13
2.3 Ytterbium facts	15
2.4 Forced evaporative cooling	19
2.5 The condensation transition in position, phase, and momentum spaces	24
3.1 Young's double slit	29
3.2 Path based calculation of Young's double slit	31
3.3 Mach-Zehnder and Michelson interferometers	32
4.1 Low order contributions to g	53
4.2 Contributions to $g - 2$ by loop content	55
5.1 Schematic of a contrast interferometer	58
5.2 Example contrast signal	62
6.1 My CI	76
6.2 Interferometry breadboard	79
6.3 Calibrating with Kapitza-Dirac	85
6.4 High-order diffraction	85
6.5 Momentum-space Ramsey interferometer	87
6.6 Stray light backgrounds	89
6.7 Images of CI signals	91
6.8 Small T Results	94
6.9 PMT Schematic	95
6.10 Vibration immunity	101
6.11 CI across the BEC transition	106

8.1	Intensity calibration	138
8.2	Mirror pulse	139
8.3	Diffraction phase correction	140
9.1	Accuracy of scaling solutions at center of BEC	154
9.2	Scaling solutions for coherence length	155
9.3	Errors in SVEA from temporal phase matching	161
9.4	Summary of interaction effects	170
9.5	Simulated output of a contrast interferometry experiment	173
9.6	Mean-field shifts versus splitting parameters	174
9.7	Different splittings	177

Chapter 1

INTRODUCTION

The present chapter attempts to introduce the main ideas of this thesis, along with some jargon that is common to precision measurements (foremost amongst them, the dreaded “systematics”) but may be less commonly heard in labs committed to saner pursuits.

1.1 “Precision” Measurements

A common cartoon in grade-school science courses explains precision and accuracy in terms of darts thrown at a dart board. An example is shown in figure 1.1.

One of the key aspects of empirical science is reproducibility. If the same result cannot be seen in different places and times by different experimenters, then it is of dubious value for understanding the world. When applied to a single measurement, the concept of reproducibility is referred to as precision. I would not claim to have a pen that is 10 cm long if I measure it to be 10 cm one time and 8 cm or 12 cm at other times. Precision is also used to refer to the quantitative level of reproducibility. Suppose I have an extremely sensitive measuring instrument that tells me the pen is 10.000 cm long before lunch. After lunch I measure again and get a reading of 10.002 cm. If I then decide to take 100 readings of the same pen with the same instrument, I can develop a statistical sample and report the standard deviation of those measurements. This quantitative accounting for the reproducibility of a measurement is generally what is meant by the precision of a measurement. In this thesis, the mass of a single ytterbium (Yb) atom is measured to a precision of 7 parts per million (ppm), which means that upon repeated runs of the same measurement,

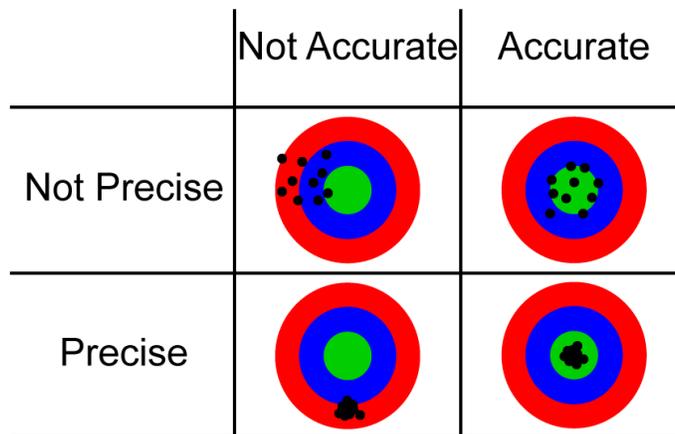


Figure 1.1: Accuracy vs. Precision. The four dart boards demonstrate the ideas of accuracy and precision. Striking the bull's eye corresponds to measuring the quantity you set out to measure. Tight clusters correspond to high precision or reproducibility. Points near the bull's eye correspond to high accuracy. The radius of the black dots represents the uncertainty in the value from a single measurement.

the value should fall within 7 ppm of our reported value roughly 68% of the time, i.e., the standard deviation of a large set of such measurements would be 7 ppm.

A distinction is made between relative and absolute precision. For example, 7 ppm is the precision of the mass measurement relative to the actual measured mass. An absolute precision would be something along the lines of “The mass was measured to $\pm 2 \times 10^{-30}$ kg.”

Accuracy is a more subtle issue and presents the real difficulty in precision measurements. As such, they should probably be called “accuracy measurements,” though somehow that sounds less appealing. In setting out to measure the mass of a single Yb atom, one must have an idea what this phrase means. The actual quantity that comes out of the experiment is the phase of a sine wave fitted to data that looks like a piece of a sine wave. An operationalist account of this measurement would include a list of all of the steps going into generating this data. Indeed, chapter 6 will give just such an account. After extracting the phase of this sine wave, the phase from one experiment is compared to those of slightly different experiments to extract a frequency from the set of phases. Then using freshman physics, namely $E = p^2/(2m)$, and the connection between frequency and energy in quantum mechanics we arrive finally at a value for the mass of an Yb atom.

One important point to be taken from this is that a tremendous number of assumptions go into the arrival at a value for mass. Under an operationalist account, the term “mass” would be defined in terms of experiments having nothing to do with measuring phases of sine waves. These assumptions must all be tested before the final measurement can be believed to be accurate. The word “believed” in the previous sentence points out an essential difference between the dart board cartoon and the real world. Precision is always something that may be known: take enough measurements and it is purely a question of statistics. Accuracy is never a sure thing. That is to say, in the real world we cannot see a bull’s eye to know whether we have measured the thing we wish to measure. In a measurement such as that described in this thesis,

the use of $E = p^2/(2m)$ assumes that the frequency we have measured corresponds purely to kinetic energy rather than kinetic energy plus a tiny potential energy due to stray magnetic fields, the earth's gravitational field, or, most troubling, some effect we have not considered at all in our stated assumptions.

Effects other than the one we set out to measure that can nonetheless change the value of our measurement are referred to as systematic effects or just systematics for short. The design of our experiment, detailed in chapter 5, is largely conceived to avoid or reduce potential systematic effects. The third main section of this thesis is devoted purely to the analysis of systematic effects.

1.2 Interferometry

Having spent several pages on the first word of the thesis title, we'll now move more quickly through the rest. Interferometry describes the use of quantum interference effects for purposes of measurement. Quantum interference effects are very sensitive to the relative phase of the two quantum amplitudes that are interfering with one another (e.g., the amplitude for an atom to pass through the left or the right slit in the classic two-slit interference experiment). To use this sensitivity for measurement, we must create a situation wherein the effect we wish to measure creates such a relative phase shift. As described above, to measure the mass of an Yb atom, we first measure its kinetic energy. Thus, if we can make two interfering amplitudes for the atom with one in motion and one static, we have introduced just the sort of phase shift we need.

Interferometry is advantageous for measurement because phase can only be measured modulo 2π . Suppose that you know, based on previous measurements, that the total phase evolution should be 1000 ± 1 rad. If the exact phase evolution is $1000.000 = 2\pi \times 159 + 0.974$, the interferometer is only measuring the 0.974. To find the strength of the effect to one part in one million, you need only measure the interference effect to a precision of one part in one thousand, which is a dramatic

increase in ease of measurement.

As a concrete example, the distance between two objects may be measured through interferometry with light to precisions much smaller than the wavelength of the light itself. A change in distance of one thousandth of a wavelength can make a measurable change in the brightness of an interference fringe. The distance between two objects separated by 1 mm or 1 km may be measured in this way to the same absolute precision. However, the relative precision in the latter case is one million times larger, without any greater complication to the actual measuring process. In fact, the LIGO experiment does just such a measurement, though to far better precision than one thousandth of a wavelength, in its search for gravitational waves.

1.3 The Perfect Quantum System?

Rather than light or individual atoms, the experiments described in this thesis leverage interference in Bose-Einstein condensates (BECs). A BEC is a peculiar state of matter consisting of a large number of atoms all in identical quantum states. Thus, quantum effects that might be expected to manifest only at the level of single atoms, such as interference, can be seen on a macroscopic level. In fact, we can photograph the interference fringes in certain types of BEC interferometers described in chapter 6. With a single-atom source this would be impossible. In a sense, one can think of a BEC as an amplifier of “quantumness,” bringing the counter-intuitive behavior of single atoms into the macroscopic realm. In fact, BEC is at the heart of such strange, macroscopic manifestations of quantum mechanics as superfluidity (the flow of liquids, e.g. liquid helium, without viscosity) and superconductivity (certain materials ability, at cryogenic temperatures, to carry electrical current with zero resistance).

The use of BECs allows for large signal-to-noise in atom interferometry experiments. Knowing that every atom is in the same quantum state makes it much easier to suppress noise due to interactions with the environment. This extreme level of quantum control also facilitates study of potential systematics. On the other hand,

it introduces a large new systematic not present in single-atom interferometry: The atom-atom interactions discussed in chapter 9. Fear of this systematic has delayed the adoption of BECs as sources for precision measurements. I believe the techniques and results of chapter 9 allow us to stop worrying and love the BEC.

Chapter 2 discusses how we generate BECs in the lab, as well as elaborating on the nature and usefulness of BECs. Two of the stars of our experiments, Yb and BECs, will be explained in more depth. Then, chapter 3 delves into the details of interferometry, both as a manifestation of fundamental aspects of quantum mechanics and as a powerful tool for precision measurement. The background section is completed in chapter 4 where the connection between an atom's mass and the fine structure constant is explained, along with the larger context of fine structure constant measurements and tests of Quantum Electrodynamics (QED).

Section II contains the central story. The theory of a contrast interferometer, the workhorse of this thesis, is described in chapter 5. The main experimental results are presented in chapter 6. While Section II contains the central story, the gritty details of systematic effects are to be found in Section III. Chapter 7 is devoted to understanding a variety of potential systematic effects, dealing with physics that is generally well-understood. Chapters 8 and 9 each examine a more complex/problematic systematic in detail. Chapter 8 details a systematic effect arising from the laser diffraction pulses, common to many types of atom interferometer. Chapter 9 shows that atomic interaction systematics, a serious concern for any BEC interferometer, may be controlled and subtracted.

Part I
BACKGROUND

Chapter 2

COOLING ATOMS WITH BLINDING HOT LIGHT

Laser cooling of atomic gases began in the early 1980s. The field moved quickly, overcoming some initial barriers, such as the Doppler limit, accidentally. By the time the Nobel prize was awarded for laser cooling in 1997, the techniques had been refined and combined with evaporative cooling in magnetic traps to achieve Bose-Einstein condensates (BECs) in these dilute atomic gases. A BEC is a new state of matter achieved at extremely low temperatures in dilute gases, which has properties useful to precision measurement. Achievement of BEC was awarded the Nobel prize in 2001.

2.1 The Physics of Laser Cooling

We begin with a broad overview of how a laser cooling and trapping experiment works, including a heuristic picture of the Magneto-Optical Trap (MOT), one of the key components of any laser cooling experiment. We will quote well-known mathematical results on laser cooling and trapping. A more thorough account may be found in many textbooks, for example that of Foot[31]. A particularly nice derivation using the Heisenberg equation of motion for the density matrix and master equations for the probability density functions may be found in the problem book of Basdevant and Dalibard[8]. For more details on the machine we use for laser cooling, see the dissertations of Hansen and Khramov[40, 52].

Most of the atoms which are amenable to laser cooling are solids at room temperature. This is true of ytterbium (Yb), the element focused on in this dissertation. Therefore, the first step of many laser-cooling experiments is to heat the material. For

Yb, we must raise its temperature above 400°C to achieve sufficient vapor pressure. At these temperatures, enough gas phase Yb exists in our oven that a beam emerges from a small aperture in the front of the oven. The beam passes through several collimating apertures, which reduce the transverse momentum spread.

In this way, we achieve our first goal. The Yb atoms are in the gas phase and all moving in roughly the same direction through an ultra-high vacuum (UHV) system. UHV will be key to keeping them isolated from outside influence once they are trapped. While they're all headed the same way, they are moving quite quickly, with average speed approximately 340 m/s. For light elements this may be over 1000 m/s. If we are to trap them with laser light, they must be slowed down. This is achieved through the use of a Zeeman slower, first demonstrated by Phillips and Metcalf in 1982[82].

The Zeeman slower makes use of radiation pressure to slow the atomic beam. When an atom encounters a photon resonant with one of its electronic transitions, it will absorb the photon and then re-emit it in a random direction. After absorbing many photons from the same direction and re-emitting each into a random direction, the atom has a net momentum change in the direction the photons were originally going. So, shining resonant light directly toward the atomic beam should eventually slow it down, as depicted in figure 2.1 a).

However, as the atom slows down, the frequency it sees changes. For light to be on resonance when propagating toward high-speed atoms, the light must be below the resonant frequency in the lab frame, so the Doppler shift can bring it up to resonance. As the atoms slow, the Doppler shift is weakened, taking the light out of resonance. For a slowing laser beam to continue to work as the atoms are slowed, something must compensate this frequency shift. In a Zeeman slower, a magnetic field shifts the value of the resonance frequency. The field is designed to keep the atoms on resonance as they slow down, so that a constant radiation pressure may be applied. Figure 2.1 b) shows a typical profile. As the atoms slow, they cover less ground during the same

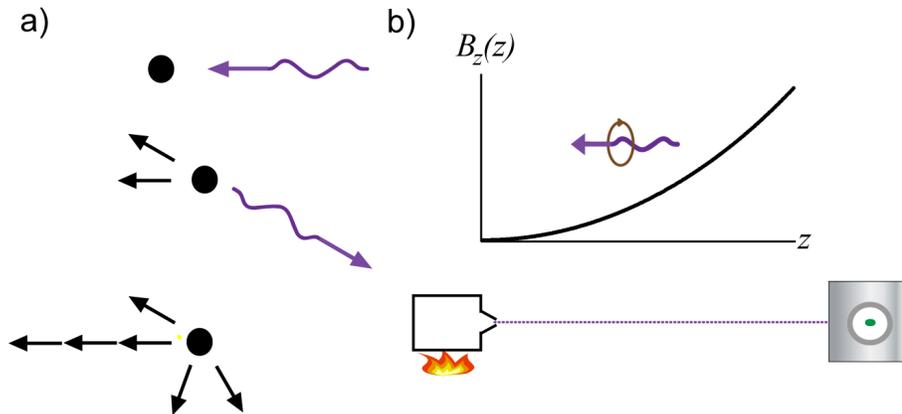


Figure 2.1: Heuristic diagram of Zeeman slower. In a) the idea of slowing by radiation pressure is shown. Photons are symbolized by wavy violet lines and momentum kicks are symbolized by black arrows. Each absorption gives a momentum kick to the left. Each emission gives a momentum kick in a random direction. The two processes are shown in the top two cartoons. The bottom cartoon shows the result of three absorption/re-emission events, showing how the net momentum kick is largely along the axis of light propagation. In b) the magnetic field profile used to compensate Doppler shifts with Zeeman shifts is shown with the positions of the oven and main trapping chamber indicated below. The main chamber contains a green spot to symbolize the MOT, discussed below. The violet dotted line shows the path of the atomic beam from the oven to the main chamber.

change in velocity, so the field must change more quickly as it approaches the end of the slower.

To understand this more quantitatively, we need to know the force applied to an atom by the laser. The effective smooth force generated by many discrete photon scattering events, called the scattering force, is given by

$$\vec{F}_{\text{sc}} = \hbar \vec{k} \frac{\Gamma}{2} \frac{\omega_{\text{R}}^2/2}{\delta^2 + \omega_{\text{R}}^2/2 + \Gamma^2/4}, \quad (2.1)$$

where Γ is the natural linewidth of the transition, ω_{R} is the Rabi frequency, and δ is the detuning of the light from atomic resonance in the atom's frame of reference. The linewidth can also be understood as the inverse of the excited state lifetime, in accord with the time-energy indeterminacy relation¹. The Rabi frequency is the coupling strength between the ground and excited states. The Rabi frequency is related to the size of the transition matrix element, which itself relates to Γ , and the square root of the intensity of the laser light².

The detuning merits careful consideration. Laser cooling essentially depends on manipulating this detuning, making it a function of position and velocity. For an atom at rest in empty space, with an energy difference between ground and excited state of $\hbar\omega_0$, light with angular frequency ω has a detuning $\delta = \Delta \equiv \omega - \omega_0$. If the atom is moving with velocity \vec{v} , it sees a different value for ω than is seen in the laboratory reference frame. Namely, it is Doppler shifted by $-\vec{v} \cdot \vec{k}$, where the sign agrees with the common experience of an observer moving toward a source ($\vec{v} \cdot \vec{k} < 0$) hearing a sound of higher frequency (pitch) than a stationary observer. Thus, the detuning becomes $\delta = \Delta - \vec{v} \cdot \vec{k}$. Just as ω may be changed by velocity, a space dependent magnetic

¹More explicitly, the natural linewidth is the indeterminacy in the excited state energy ΔE divided by \hbar . As such, the expected lifetime comes from the Heisenberg indeterminacy relation: $\Delta E \Delta t \geq \hbar$. If this inequality is saturated, then we find the lifetime $\Delta t = \hbar/\Delta E = 1/\Gamma$. That the equality should be saturated comes from following this argument backward: If the lifetime is fixed, then the natural (i.e., unperturbed) energy uncertainty should saturate the uncertainty principle, as it has nothing to increase it beyond the Fourier transform limit.

²When spontaneous emission is negligible (i.e., when $\omega_{\text{R}} \gg \Gamma$) the state will oscillate between ground and excited states with angular frequency ω_{R} .

field can give ω_0 space dependence. Suppose for the moment that the ground state is insensitive to magnetic fields (this is true for Yb). The Zeeman shift of the excited state is $-\vec{\mu} \cdot \vec{B}$. For a z directed magnetic field, this becomes $m_z g \mu_B B(\vec{x})$, where m_z is the projection of the angular momentum on the z axis, g is the gyromagnetic ratio for the excited state (generally $g < 0$), and μ_B is the Bohr magneton. Thus, finally, $\delta = \Delta - \vec{v} \cdot \vec{k} - m_z g \mu_B B(\vec{x})$.

For designing a Zeeman slower, we want maximum acceleration and so want $\delta = 0$, for a particular initial velocity, the entire length of the slower. Atoms with initially lower speed will tend to feel little acceleration until the target class has slowed to equal their speed, at which point they follow the same path as the target class. For constant acceleration a , $z = (v_0^2 - v^2)/(2a)$, with v_0 the initial speed of the target class. This gives $v = \sqrt{v_0^2 - (2a)z}$. To compensate for the changing Doppler shift of the target class then requires a magnetic field with a square root profile as a function of z , as reflected in figure 2.1 b). A well-tuned Zeeman slower can deliver a sizable flux of atoms with speeds of 1 m/s or so.

At the end of the slower is the main chamber. Here, the atoms need to be trapped in a single spot. This is the purpose of the Magneto-Optical Trap (MOT). First demonstrated in 1987 by Chu and Pritchard, the idea was originally conceived by Dalibard[86]. Consider first an atom confined to move in only one dimension. A one-dimensional MOT, for instance in the z direction, would consist of a magnetic field gradient, passing through zero. Then lights, tuned below resonance in the lab frame, with opposite projections of angular momentum on the z axis (m_z) shine from each side. The angular momentum projection is equivalent to the handedness of the circularly polarized light. This setup is shown in figure 2.2 a). The energy levels for an excited state with angular momentum $F = 1$ are shown in figure 2.2 b), along with a green stripe to show the energy of a photon in the lab frame. The width of the stripe indicates the natural linewidth of the transition. We'll assume for simplicity a transition from $F = 0$ to $F = 1$. Conveniently, this is the structure of the cooling

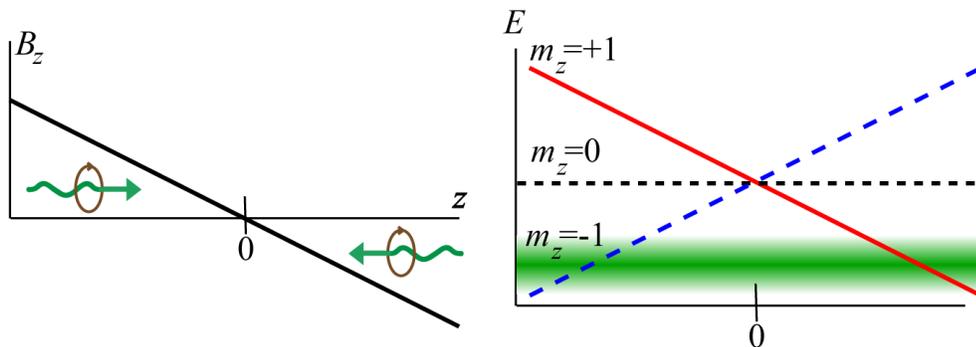


Figure 2.2: Cartoon of a Magneto-Optical Trap. The MOT may be visualized as a pair of Zeeman slower each aimed at the same point. An atom moving to the right will be slowed and stopped to the right of the magnetic field zero. It will receive a small kick back toward the field zero and eventually be slowed and turned around to the left of the zero. This allows a sample of cold atoms to accumulate near the field zero.

transitions for ^{174}Yb , the isotope the experiments described in this dissertation were carried out with.

A fast-moving atom entering from the left would see the cooling light Doppler shifted too high in frequency (energy) to transition to the $m_z = -1$ excited state. However, as it moves to the right the energy of the $m_z = +1$ will drop due to the changing magnetic field until it is on resonance. The light from the right has the correct polarization to reach the $m_z = +1$ state, and so the atom will begin scattering this light. The atom responds the same way that an atom moving through a Zeeman slower would: As it slows it may move out of resonance, but as the field changes it will come back into resonance, with this sequence repeating until it finally stops. Stopping in a place where it is in resonance would cause continued scattering of the leftward moving light until the atom begins moving to the left itself, quickly coming out of resonance through the combination of the Zeeman shift and the Doppler shift.

as it moves to the left, the energy of the $m_z = -1$ state will continue to lower until it is on resonance with the light shining from left to right (including the small Doppler shift). This light can address the $m_z = -1$ state and so will slow and stop the atom again. This process repeats, cooling the atom in each pass and eventually ending with the cold atom confined to a small region near the point where $B_z = 0$, denoted $z = 0$ for convenience in the figure. The final temperature achieved in a MOT comes from a competition between the dissipative cooling that absorption of red-detuned light creates and the diffusion in momentum space due to the spontaneous scatter of those absorbed photons. The minimum attainable temperature of $\hbar\Gamma/(2k_B)$, achieved with a detuning of $\Delta = \Gamma/2$, is known as the Doppler limit.

The MOT in three dimensions implements this scheme in all three dimensions simultaneously. The dynamics are then more complex—to my knowledge no useful analytic approximation for the three dimensional case has been found—but qualitatively they follow the straight-forward expectation based on extrapolation from the one-dimensional case. In our lab, this configuration is implemented with a quadrupole magnetic field, which gives a magnetic field gradient in each dimension with all three components zeroed at the same position. Each of three orthogonal axes has laser beams with the proper circular polarization coming in and retro-reflections of those beams with opposite polarization coming back out (achieved with a quarter-wave plate before the retro mirror).

2.2 *Ytterbium*

While the theoretical results in this dissertation apply to a variety of atoms, the experiments were all performed using ytterbium-174 (^{174}Yb). Ytterbium has atomic number 70 and is found at the end of the lanthanide series on the periodic table. This means it has a completely filled $4f$ shell in addition to a filled valence $6s$ shell. This structure makes it similar to alkaline earth metals, such as calcium and strontium. It has seven stable isotopes, five bosons and two fermions. Each of the bosons has zero

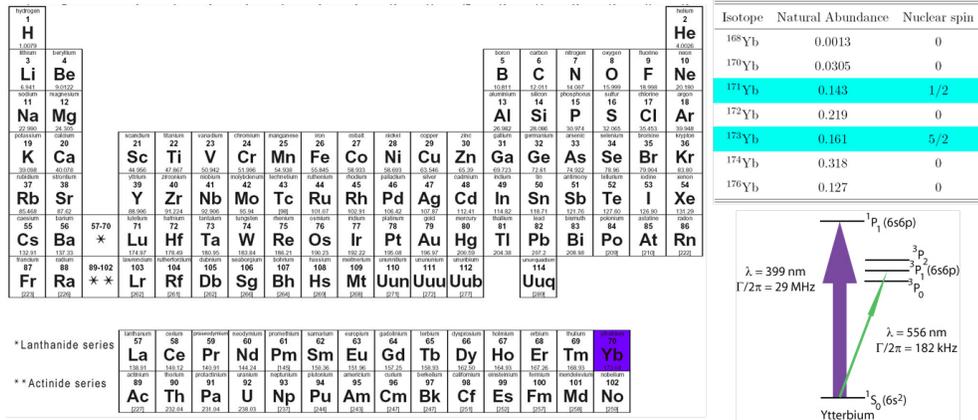


Figure 2.3: Ytterbium facts. Yb is found at the end of the lanthanide series on the periodic table. The top right table shows the stable isotopes of Yb along with their nuclear spin and natural abundance. Only ^{168}Yb has an abundance low enough to make laser cooling from an unenriched sample difficult. The fermions are indicated by cyan highlighting. The level structure relevant to laser cooling Yb is shown in the lower right.

nuclear spin.

The ground state is a 1S_0 state. Since the electrons have neither orbital nor spin angular momentum, it has no ground-state hyperfine structure. There are two transitions of interest for laser cooling. The $^1S_0 \rightarrow ^1P_1$ transition is a strong electric dipole transition with a width of 28 MHz and a wavelength of 398.9 nm (violet/near UV). The $^1S_0 \rightarrow ^3P_1$ transition, with wavelength 555.8 nm (green), is an intercombination transition. This refers to an electric dipole transition that violates a selection rule, but nonetheless occurs because the excited state actually contains a small admixture of a state that does not violate the selection rule. For the intercombination line of Yb, spin-orbit coupling mixes the 3P_1 and 1P_1 states. Mathematically, one sees that inserting the spin-orbit coupling perturbation term into the dipole matrix element gives a nonzero amplitude. Since Yb has a large nucleus, the spin-orbit coupling is

fairly strong, giving a transition width of 180 kHz. The other 3P states, 3P_0 and 3P_2 are metastable with lifetimes in excess of 10 s. While these are of interest to many experiments, their strong inelastic scattering rates make them infeasible for Bose-Einstein condensation, so we will ignore them for now.

For laser cooling, these two very different transitions fit nicely into the two requirements of Zeeman slowing and magneto-optical trapping. The strong, violet transition is used for Zeeman slowing since the fast scattering rate and relatively large per photon momentum give large acceleration. Our Yb slowers have interaction regions no more than 30 cm long. The weak, green transition allows for cold MOTs. We have achieved final temperatures as low as 12 μK for MOTs running on this transition, about 3 times the Doppler limit of 4.5 μK from $\Gamma = 2\pi \cdot 180$ kHz. In contrast, a MOT run on the violet transition would have a Doppler temperature around 700 μK .

Early laser cooling and trapping experiments were done exclusively with alkali metal atoms. Since these remain the dominant species in cold atoms experiments, we briefly review the differences in cooling alkalis versus Yb for readers who may be more familiar with alkali atoms.

The violet transition is five to ten times stronger than the D2 lines used for slowing and cooling alkalis. It is also shorter wavelength (i.e., higher photon momentum). While this gives a stronger scattering force, Yb's large mass means the acceleration is not as much greater as the force might suggest. Frequently the heavier alkalis, rubidium and cesium, are loaded from a room temperature vapor rather than requiring an oven and Zeeman slower. Ytterbium's low vapor pressure precludes this option.

The narrowness of the green transition gives our green Yb MOTs Doppler temperatures more than an order of magnitude lower than those of alkalis operated on the D2 line. However, the alkalis with resolved excited state hyperfine structure (Na, Rb, and Cs) partially make up for their high Doppler limits through sub-Doppler cooling. For bosonic Yb, this is not an option due to the lack of any hyperfine structure. Fermionic Yb can in principle benefit from sub-Doppler cooling as well. While hints

of such cooling in a green MOT have been reported[65], the temperatures achieved were still well above the Doppler limit and so not evidence of useful sub-Doppler cooling. Recently, sub-Doppler cooling was observed in a violet MOT[55]. The final temperature was $200 - 400 \mu\text{K}$. While this is not competitive with the green MOT for final temperature, it could be a way to increase transfer from an initial violet MOT into a green MOT.

While the green MOT is desirable for its low temperatures, the weakness of the transition also reduces the capture velocity of the MOT. Capture velocity is the largest speed an atom can have and still be slowed quickly enough to be caught in the MOT. We have found substantial improvement of the MOT capture by modulating the frequency of the MOT light, artificially broadening the laser to several megahertz, compared to our detuning of -8 MHz . Generally, one uses light relatively far detuned from resonance for MOT loading, to increase the capture velocity. Then, before loading into a conservative trap, the light level is reduced as the detuning is brought close to resonance, known as compressing the MOT. The far-detuned MOT cannot be as cold or as dense as the near-resonance MOT, so compression is key to good loading into conservative traps. In our case, the modulation is ramped to zero during the final MOT compression phase, allowing us to take full advantage of the low Doppler temperature for loading into our conservative trap. The end of our compression phase is generally at a detuning of -2Γ to $-\Gamma/2$ and an intensity well below the saturation intensity of $140 \mu\text{W}/\text{cm}^2$. For comparison, our MOT beams have an intensity around $4 \text{ mW}/\text{cm}^2$ for the loading phase. These parameters allow us to load about 5 million atoms into our conservative trap. The importance of the conservative trap is for evaporative cooling, which we now turn to.

2.3 Evaporative Cooling

While temperatures near $10 \mu\text{K}$ are quite cold, for our systems they are still about one and a half to two orders of magnitude above the condensation transition temperature.

From the MOT, we transfer our gas into a conservative potential and then cool it further via evaporative cooling.

The concept of evaporative cooling is straight-forward. A thermal distribution of particle energies in a gas has a long, high-energy tail. The small number of particles with energies in this tail have a per particle energy far higher than the average for the gas. Evaporative cooling consists in providing the particles in the high-energy tail a path out of the sample. The concept was first described[44] and implemented[66] in magnetically trapped spin-polarized hydrogen.

As the name implies, this is essentially the process at work in cooling a cup of coffee. The highest energy molecules in the coffee are the ones with sufficient energy to break the surface tension and be freed as vapor. As such, these evaporating molecules remove more energy than the average energy per particle. Once the liquid re-thermalizes (i.e., distributes energies into an equilibrium distribution via collisions), it will be at a lower temperature due to the lower energy per particle. Luckily, in coffee the particle energies able to leave the coffee remain constant. As the coffee cools and fewer particles have such high energies, the cooling rate drops.

For cold atom experiments, unlike coffee, we want to boost the speed of the cooling process rather than keep it slow. To do this, we progressively reduce the energy required to escape from the trap. This keeps the loss process (the “evaporation”) moving at a fairly constant rate, rather than allowing it to stall out as in cooling coffee. Figure 2.4 depicts this as a sequence of repeated discrete steps. First the energy required for a particle to leave the trap is reduced. Next, collisions re-establish the thermal distribution, possibly kicking some extra atoms out in the process. Then, we repeat. Realistic experiments generally reduce the energy needed to escape (i.e., the trap depth) continuously. This process of trap depth reduction is known as forced evaporative cooling. For this process to work efficiently, it is key that the particles continue to interact strongly enough to keep the energy distribution nearly thermal. The density and two-particle interaction strength set the time scale for re-

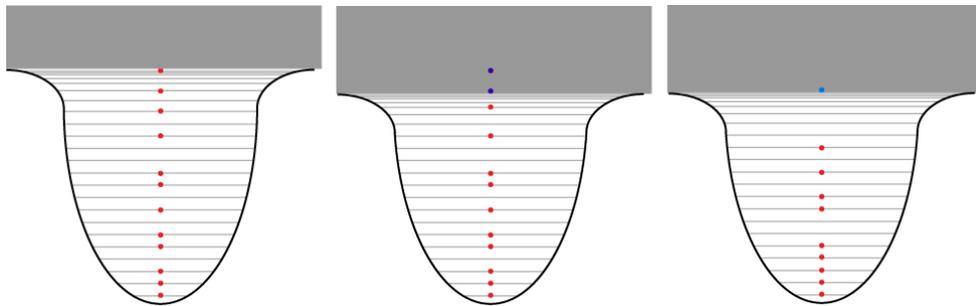


Figure 2.4: Forced evaporative cooling. The process of forced evaporative cooling can be visualized as having two steps. Beginning with a thermal distribution as shown on the left, the trap depth is reduced, as shown in the middle. This causes the highest energy atoms (purple here) to be no longer trapped. After these have gone away, collisions in the trap re-establish a thermal distribution, as shown on the right. In this process, some extra atoms may be lost (blue here). The trap is now ready for another reduction in depth. Real experiments generally have a continuous reduction of trap depth, done slowly enough as to keep the gas nearly thermalized.

thermalization and with it the rate at which the trap depth may be decreased without badly lowering efficiency.

The first gaseous Bose-Einstein condensates were produced in rubidium[2] and sodium[25] using forced evaporative cooling in magnetic traps. The specific path out of the sample used in these experiments is known as an rf-knife. In a magnetic trap, atoms in a low-field seeking hyperfine state are trapped near a minimum of the magnetic field³. The technique uses the larger Zeeman shift seen by higher energy atoms, due to their ability to move further from the center of the trap. By applying an rf field with frequency set to only address the most energetic atoms, their spin may be flipped to an untrapped state, releasing them from the trap. The frequency is steadily lowered until the gas is cooled to the BEC state. An advantage of this technique as compared to the other main technique used today is that the trap shape does not change with time. Thus, density increases steadily as temperature decreases, keeping the thermalization rate constant, or even causing it to increase, as temperature is reduced.

The other technique alluded to above is evaporation in an optical dipole trap (ODT). Since Yb has no ground-state electron angular momentum, magnetic trapping is not an option. An ODT takes advantage of the polarizability of atoms to trap them using light tuned to a frequency below an atomic resonance. For such light, referred to as red detuned, the atomic energy is shifted by $-\alpha E^2$, where $\alpha > 0$ is the polarizability and E is the amplitude of the light's oscillating electric field. Heuristically, one can imagine the valence electron as a classical oscillator. For frequencies below resonance, the electron can follow the direction of the electric field as a function of time. Thus,

³A straight-forward way to create a magnetic field minimum in a static configuration is to create a point with zero field. However, it is key to BEC in magnetic traps that the magnetic field zero not be accessible to the atoms. In the absence of magnetic field they may easily change hyperfine state to an untrapped state and be lost from the trap. Real magnetic traps use a blue detuned laser beam to make the zero crossing inaccessible, move the field zero in a quick orbit creating a time-average potential, or a zero in two components plus a curved bias field in the third component.

the atom has an induced electric dipole moment that is automatically anti-aligned to the electric field, reducing the atom's energy.

Our experiments load an ODT with ^{174}Yb from a cold, compressed MOT. The forcing process for optically trapped atoms involves simply reducing the intensity $\sim E^2$ of the trapping beam, thereby reducing the depth of the potential. Optical trapping allows for tighter confinement than magnetic traps. It also allows for trapping of any hyperfine state or atoms with no magnetic moment, unlike magnetic trapping. However, the reduction of intensity causes the trap to lower in both depth and confinement. This creates the potential for a forced evaporation ramp to stall at low temperatures as the thermalization rate can decrease under lowered confinement. With careful tailoring of the rate of intensity reduction, this problem can generally be avoided.

More specifically, we use a 1064 nm wavelength fiber laser to create our trap. Our trap consists of two focused beams—generated by splitting the single laser beam out of the fiber—crossed at an angle of about 20° in the center of the main chamber. The two beams have orthogonal linear polarizations to inhibit scattering of photons between the beams. For loading Yb, we have about 10 W in each beam (we found a slightly uneven splitting of the laser power between the two beams to give a small boost in loaded number). For interferometry we use an evaporation ramp that lowers the power in two exponential stages. There is a hold between the stages, with temperature around $1\ \mu\text{K}$, to allow any center of mass motion initiated by the first ramp to damp out. Such motion may be caused by the focal points of the two beams moving in response to temperature changes in optical elements as the power is ramped down. The second, slower ramp ends at a per-beam power of $\approx 30\ \text{mW}$, producing a pure BEC of up to 300,000 atoms. We now turn, finally, to discussing the BEC state.

2.4 Bose-Einstein Condensation

A Bose-Einstein condensate is a state of matter in which a macroscopic number of particles (i.e., a fraction large compared to $1/\sqrt{N}$) occupy a single quantum state. We begin by considering a non-interacting gas, as interactions introduce technical complications unnecessary to understanding the concept of BEC.

As the name implies, the bosonic nature of the particles is critical to the existence of this state of matter. We will quickly review some key statistical results as they relate to the BEC transition. The reader seeking more detail or thorough derivations should consult any standard statistical mechanics text, as I am doing right now[77, 83].

The expected occupation $\langle n_\epsilon \rangle$ of a state with energy ϵ in a system with chemical potential μ is $\langle n_\epsilon \rangle = (\exp[\beta(\epsilon - \mu)] - 1)^{-1}$, where $\beta = 1/(k_B T)$, T is the temperature, and k_B is Boltzmann's constant. For the simple picture of macroscopic occupation of a state given above, we see that the expected occupation blows up if $\mu = \epsilon$. Since actual divergence or negative values should not occur for a physical expected occupation, the boson's chemical potential should be below the energy of any state. Macroscopic occupation is accomplished as $\mu \rightarrow \epsilon_0$, where ϵ_0 is the energy of the ground state of the potential holding the system.

As a concrete example, consider a large, homogeneous box with volume V . Using periodic boundary conditions, it may be shown that the density of states is proportional to $\sqrt{\epsilon}$, because the energy is simply $p^2/(2m)$ in a flat potential. The density of states is a good approximation to the discrete energy levels so long as no level has large occupation. It also gives an incorrect result of zero states at zero energy. For a room temperature gas, or even one far below room temperature, these deviations are negligible for calculating the state of the system.

For a given μ , we define the fugacity $z \equiv e^{\beta\mu}$. Treating the states as continuous and integrating we find the expected number of particles to be

$$\langle N \rangle = \frac{2\pi V (2mk_B T)^{3/2}}{h^3} \int dx \frac{x^{1/2}}{z^{-1}e^x - 1} = \frac{V}{\lambda_T^3} g_{3/2}(z), \quad (2.2)$$

where $\lambda_T = h/\sqrt{2\pi mk_B T}$ is called the thermal de Broglie wavelength and $g_{3/2}(z) = \sum_{j=1}^{\infty} z^j/j^{3/2}$ —known as the Bose-Einstein function of order 3/2—simply comes from carrying out the integration. Making use of the number density n , this may be rewritten as

$$\langle n \rangle \lambda_T^3 = g_{3/2}(z) . \quad (2.3)$$

Since $0 < z < 1$, the right hand side of (2.3) is bounded to be no larger than $g_{3/2}(1) = \zeta(3/2) = 2.612\dots$. Then, for a given density, there is a temperature below which (2.2) cannot be correct. Since the equation ignores the ground-state occupation, the conclusion is that at this temperature a macroscopic fraction of particles begin to occupy the ground state. This is the critical temperature T_C for the BEC transition.

The quantity on the left side of (2.3) is referred to as the phase-space density. The phase-space density can be understood more clearly if we rewrite it as $n\lambda_T^3 = (\lambda_T/\ell)^3$, where ℓ is the inter-particle spacing. When this quantity is much less than one, as is true for $T \gg T_C$, λ_T essentially sets the scale for correlations in the gas. Quantum interference effects can only meaningfully impact the state of the system to the extent that correlated particles overlap. The ratio λ_T/ℓ is proportional to correlation length divided by inter-particle spacing, the key comparison for determining the relative importance of quantum interference. Thus, the phase-space density can be seen as a measure of the importance of quantum effects to the state of the system. As the transition is approached, $\lambda_T/\ell \rightarrow 1$, allowing many-body correlations to grow (i.e., if particle 1 is strongly correlated with particle 2 and 2 is strongly correlated with 3, then 1 and 3 will also be correlated). At this point the correlation length diverges, being no longer controlled by λ_T . This physical space picture of growing correlations leading to macroscopic occupation of the ground state is depicted in the left column of figure 2.5.

The term “phase-space density” can be clarified by rewriting it in yet another

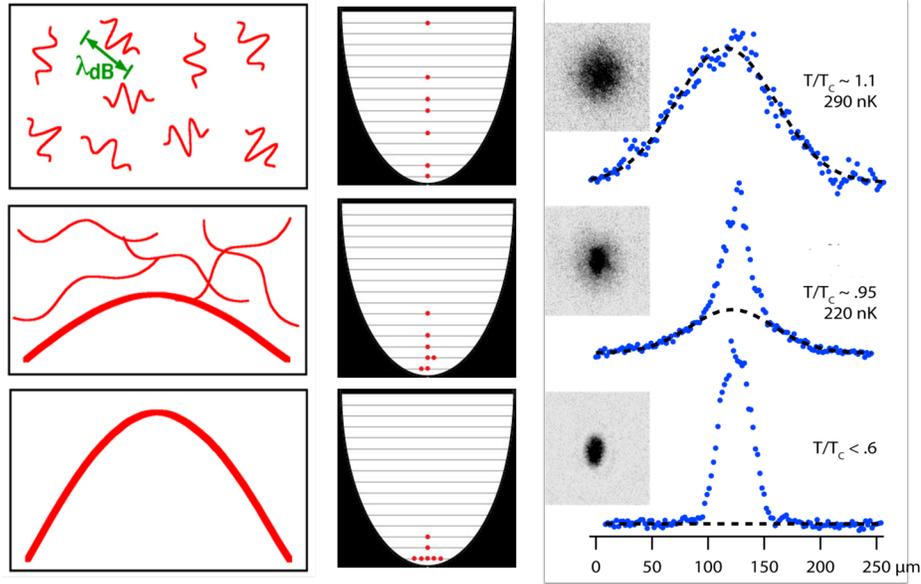


Figure 2.5: The condensation transition in position, phase, and momentum spaces. The three columns show the condensation phase transition in three different representations. In each column, the top picture shows a gas just above T_C , the middle picture a gas just below T_C and the bottom picture a gas well below T_C . The left column depicts the growth of spatial correlations as the phase-space density grows. The red curves represent interference between particles in the gas. The thick red curve shows the interference of many particles in the condensate as it grows. The middle column shows phase space as energy levels for a gas trapped in a harmonic oscillator. Each red dot indicates a single particle in a particular energy level. The right column shows time of flight absorption images as insets and plots of a slice through the center as blue points. Time of flight imaging essentially displays the momentum-space density of the gas in the trap. The dashed curves show a Gaussian fit from the edges of the distribution, which gives the fraction of the cloud effectively described by Maxwell-Boltzmann (classical) statistics.

form:

$$n\lambda_{\text{T}}^3 = N \left(\frac{V(mk_{\text{B}}T/(2\pi))^{3/2}}{\hbar^3} \right)^{-1} = N \left(\frac{V_{\text{s}}V_{\text{p}}}{\hbar^3} \right)^{-1} = \frac{N}{N_{\text{states}}}, \quad (2.4)$$

where V_{p} is the effective volume in momentum space, and N_{states} is the number of states accessible to a particle in the system at the given temperature. The second equality shows that the phase-space density can literally be seen as the density in phase space (i.e., particle number divided by the accessible volume of phase space). Since \hbar^3 is the smallest meaningful volume element in phase space—as seen from the indeterminacy relation $\Delta x \Delta p \geq \hbar$ —the accessible volume over \hbar^3 is effectively just counting accessible states. A phase space density approaching 1 means that the granularity of phase-space is becoming important. Again, phase-space density is measuring the importance of quantum effects.

This idea is shown in the middle column of figure 2.5. Here we consider instead a gas in a harmonic trap, since the available states then correspond to the evenly spaced energy eigenstates of the harmonic oscillator⁴. A red dot indicates a single particle occupying a particular energy level. As seen in this progression, increasing phase space density means fewer empty states available and, as T_{C} is approached, more levels with multiple occupancies. While the exact value of the phase-space density required for the phase transition is not the same in a harmonic trap as in a flat box potential, it turns out that if you use the density at the center of the trap (where density is highest) to calculate the phase space density, then the criterion becomes the same. This is an example of the “local density approximation.” Essentially, the transition begins when any part of the sample shows a phase space density above 2.6.

The final depiction in figure 2.5 is a momentum-space representation. The inset black and white pictures are time of flight absorption images from our lab. Such a picture is made by turning off the trap and allowing the gas to expand into the surrounding vacuum. After some time, light resonant with the strong violet transition

⁴We ignore the degeneracy of excited energy levels in a three-dimensional harmonic oscillator, such as our cold atom traps, as they do not qualitatively change the argument.

in Yb is shone at the expanded cloud. A CCD camera takes a picture of the shadow cast by the cloud. From these pictures we can determine the density by the fraction of resonant light absorbed. The position of an atom after expanding for a time t is $x_0 + v_0t$, where x_0 and v_0 are the position and velocity in the trap at the time it is turned off. For a sufficiently long time t , v_0t can be much larger than the variance of x_0 for the sample. This means position in the image actually corresponds to velocity, or momentum, in the trap. Thus, the pictures actually show the momentum-space density in the trap just before turn off⁵. The blue dots show the density in a cross-section of the image through the center. The dashed curves are Gaussian fits to the edges of the distribution. This essentially shows the “thermal” component of the gas, that is, the atoms not in the BEC. Notice also that representative temperatures are given for the top and middle pictures, which place $T_C \approx 250$ nK. Since the evaporative cooling process involves changing the confining potential as the temperature drops, the values of T_C implied by the two pictures do not agree exactly.

The key aspect of the BEC itself, a macroscopic number of atoms in the same state, is a great advantage for precision interferometry. With a large number of atoms in the same quantum state, the effect of an applied classical potential will be the same on each atom. We can, in effect, have a 100,000 atom condensate perform 100,000 identical single-atom experiments. This gives great advantages in signal-to-noise compared to a thermal atom source where 100,000 similar but not identical atomic experiments are performed in parallel. The coherence of the BEC source leads to long coherence times in the interferometer.

The downside of a BEC for precision interferometry is its higher density. We considered a noninteracting BEC for simplicity above. However, real BECs almost always have inter-atomic interactions. While the BEC transition still leads to macroscopic

⁵In a BEC, the effects of interactions are sufficiently strong that the expansion is dominated by the conversion of interaction energy to kinetic energy rather than the actual in-trap momentum. Nonetheless, there is a clear visual signature of BEC in the absorption images due to its anisotropic expansion. A non-interacting BEC would give a very similar picture, just smaller.

occupation in this case, the state so occupied changes with the number of atoms in the BEC. As described in Chapter 9, the lowest-energy state can be determined from a Hamiltonian that takes the atomic interactions into account. So, while the atoms are all in the same state, knowing the state is more complex than simply finding the single-particle ground state of the potential. More importantly, the interactions shift the atomic energy enough to impact the delicate measurements in a precision interferometry experiment. Chapter 9 is devoted to understanding and subtracting these effects from the results of a BEC interferometer.

Chapter 3

ATOM INTERFEROMETRY**3.1 Interference with Light**

One of the classic experiments in interference is the Young’s double slit experiment. A pair of thin apertures are illuminated with coherent light—either from a modern coherent source such as a laser or, as in the original experiment, from monochromatic light passed through a single thin aperture upstream, as depicted in figure 3.1. On a screen downstream from the two slits, a pattern of alternating light and dark bands appears. If you had no experience with waves (i.e., no knowledge of interference), you might be surprised to see more than two bright areas produced by only two apertures. Indeed, the phenomena of interference are definitional of wave behavior as compared to particle behavior.

In figure 3.1, the waves are represented by wavefronts. These green curves represent points which have identical phase at a given instant, for instance the crests of the waves. This picture is intuitive if you have ever watched ripples on the surface of a lake or river. We expect waves to spread with time in an orderly train. The ripples on a lake, or the ripples of electromagnetic field that we perceive as light, propagate according to wave equations. The electric field \vec{E} of some light propagating through space will follow the equation

$$\nabla^2 E_j - \frac{1}{c^2} \frac{\partial^2 E_j}{\partial t^2} = 0 . \quad (3.1)$$

The structure of this equation gives the diffraction and interference phenomena we associate with wave motion.

An alternative picture looks at the paths a particular wavefront could have taken

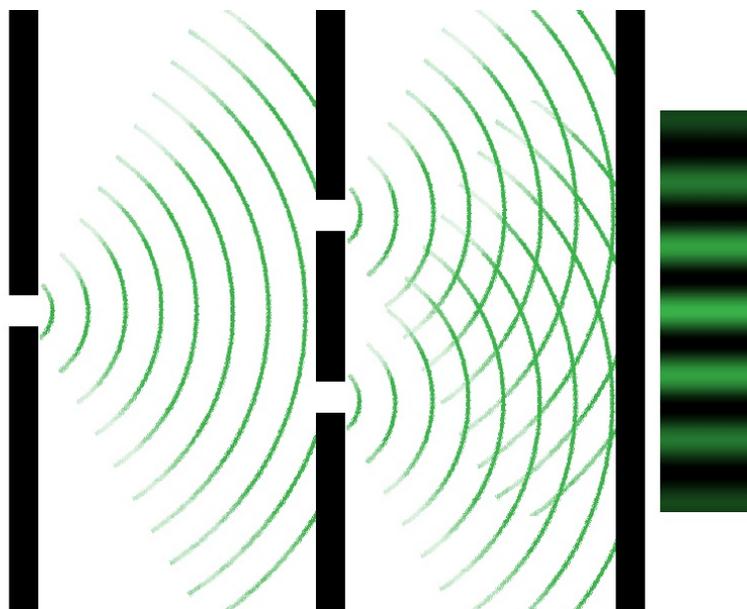


Figure 3.1: Young's double slit. Monochromatic green light incident on the single slit at the left, diffracts to produce coherent wave fronts, shown as green curves spreading as they move off to the right. The wavefronts represent crests (maxima of the electric field) at a particular moment in time. When this light reaches the pair of slits in the middle, most of it is absorbed, but some diffracts through each of the two slits, again spreading as it moves away from a slit. At the screen on the right, where two crests coincide, a bright spot appears. When a crest meets a trough (halfway between two crests), a dark spot appears. An interference pattern as would be seen on the screen is exhibited just to the right of the screen.

to arrive at the final observation point and adds them together. The idea is depicted in figure 3.2. When the paths to a point all have the same phase, they add constructively, giving a bright spot on the screen. In the figure, only two paths are depicted for the two observation points—one passing through each slit. To calculate the entire pattern, including the drop in overall intensity as you move away from the center of the screen, the many paths available through the finite extent of the slits would all need to be added together. While potentially less intuitive than the image of ripples spreading as they move, this picture nicely connects the wave theory of light to ray optics. For instance, the principle rays for finding the image point using a lens are defined so that where they intersect, all paths from the object point to the image point have the same phase. In fact, a convex lens is designed precisely so that the rays following straight paths are delayed in time by the thicker glass in the middle. The shape of the lens compensates shorter physical path length with longer delay in the glass, allowing extra phase to accumulate, so the longer paths can interfere constructively at the image point. This effective path length, in terms of phase accumulation is what is often referred to as optical path length.

While this picture gives a different sort of intuition, it in fact follows the same mathematics. The process of adding up all of the paths to determine the amplitude at the final point becomes, mathematically, a process for calculating the Green's function for the wave equation (3.1).

In most interferometers designed for making precise measurements, the number of paths light can possibly travel to reach the end point is intentionally limited to two. In this way, one knows exactly what phase shift is introduced on the arm that is being used for measuring. Because of this, the picture in terms of paths and optical path length is convenient for analyzing the workings of various interferometers. We will describe several useful interferometers, particularly those with analogs in atom interferometry.

In a Mach-Zehnder interferometer, collimated light falls on a beam splitter, which

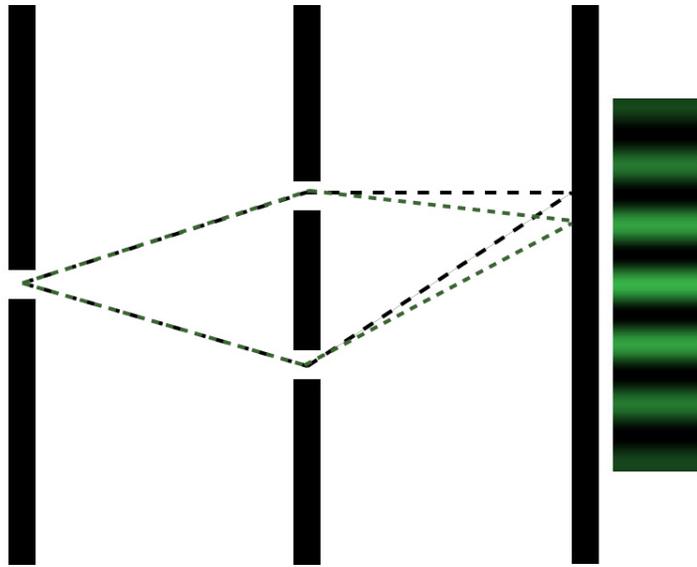


Figure 3.2: Path based calculation of Young's double slit. For two points on the screen, the paths passing through each of the two slits are shown. If the phases of the two paths are the same, the waves will interfere constructively, giving a bright spot. If the phases differ by π , they will interfere destructively, giving a dark spot.

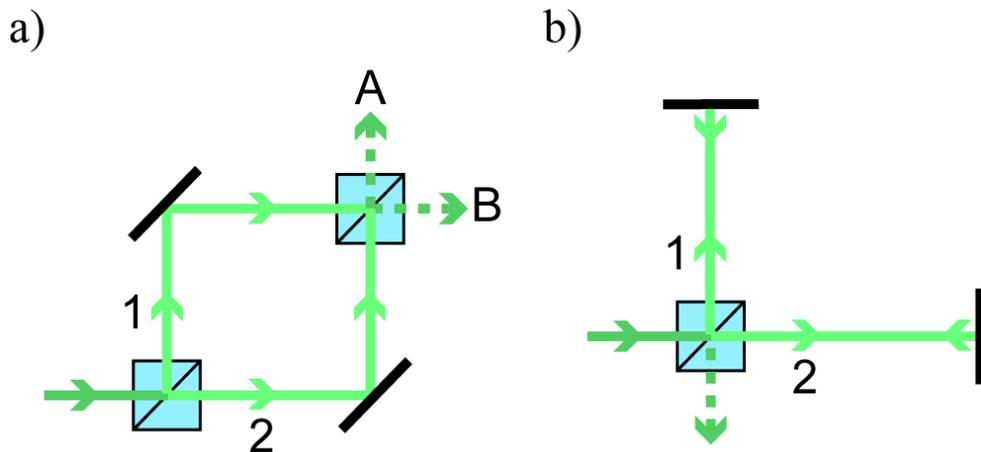


Figure 3.3: Mach-Zehnder and Michelson interferometers. Figure a) shows the design of a Mach-Zehnder interferometer. The input light at lower left is split into two paths by the beam splitter. Mirrors send the two paths into the second beam splitter. The brightness of the output at ports A and B is determined by the relative phase accumulated between paths 1 and 2. Figure b) shows a Michelson interferometer. The input light from the bottom left is again split into two paths. However, in the Michelson geometry the two paths are returned to the same beam splitter by mirrors. The brightness at the output port depends on the relative phase accumulated by the two paths during their round trips. The second output port (necessary for conservation of energy) is the input port.

coherently splits the electric field of the incoming light into two components. As depicted in figure 3.3 a), one component reflects upward, while the other passes through continuing to the right. Mirrors then direct the two components toward a second beam splitter. This beam splitter has two output ports, A and B. Suppose the initial beam splitter perfectly split the incoming intensity 50-50, giving electric field magnitudes of $E_0/\sqrt{2}$ for each path. If the final beam splitter is also a perfect 50-50 splitter, then the field at port A will be $(E_0/2)(1 + e^{\Delta\theta_{12A}})$, where $\Delta\theta_{12A}$ is the phase difference between paths 1 and 2 on leaving port A. This gives an intensity $I_A = E_0^2 \cos^2(\Delta\theta_{12A}/2)$. By conservation of energy, the intensity at port B must be $I_B = E_0^2 \sin^2(\Delta\theta_{12A}/2)$ ¹.

Measurements of phase shifts along the paths may, therefore, be made by observing the output intensities at A and B. As a particularly simple example, suppose you wish to know the index of refraction of a certain glass. It might be measured roughly by observing the refraction of light passing through it. To refine the value, a piece of the glass with known thickness could be placed into path 1 of a Mach-Zehnder interferometer. The change in output intensities caused by inserting the piece of glass would tell the change in optical path length caused by the glass, up to addition of integer multiples of π . If the index of refraction were known well enough from refraction studies to determine the number of integer multiples of π , then the interferometric measurement would increase the precision of the value substantially. Alternatively, a wedge of the glass with thickness growing from zero to some known value could be slowly slid into the path and the number of intensity oscillations used to count the number of π shifts at the known thickness.

This simple experiment points out two key aspects of interferometry. First, the $2\pi N$ ambiguity of phase measurements allows for extremely precise measurement,

¹This implies that $\Delta\theta_{12A} - \Delta\theta_{12B} = \pm\pi$. This must always be the case to conserve energy, though the calculation of the phase shifts of the individual beams due to the beam splitter will depend upon the design of the beam splitter and polarization of the incoming light.

if the ambiguity can be removed. The intensity ratio at the two output ports can be measured well with standard photodiode detectors. A one percent precise measurement of this phase can become a one part in 10^4 precise measurement by simply adding in the extra π phase shifts known from a one percent precise measurement looking at angle of refraction. This idea of “bootstrapping” through interferometry is key to success in fundamental measurements.

Second, for accurate measurements it is best to compare two experiments in which only one aspect of one path has been varied. With the glass example this seems relatively easy. However, simple seeming experiments often hide subtle difficulties. For example, reflection at the surfaces of the glass may decrease the intensity in path 1 and so spoil the equality of intensities from the two paths. This would shift the intensity ratio at the two output ports in a way having nothing to do with phase. Reflections would constitute a systematic shift to the measured value, which does not relate to the effect you actually wished to measure. Additionally, the extent to which the well-measured optical path length difference may be converted into an index of refraction will be limited both by knowledge of the glass’ thickness as well as how close to perpendicular the surface of the glass is oriented to the incoming light. This is a systematic shift that may be mitigated through careful alignment or measured by slowly rotating the block, knowing that the optical path length should be minimized at perpendicular orientation. Turning precise measurements into accurate measurements requires long, careful consideration of systematic effects and, ideally, additional measurements to allow their unambiguous removal.

Another useful interferometer is the Michelson interferometer. As shown in figure 3.3 b), collimated light is shone into one port of a beam splitter. The split light travels down arms 1 and 2, is reflected back down them to the same port of the beam splitter from which it came, and then comes out either the same port it came into or out the fourth, otherwise unused, port of the beam splitter. Generally, the end mirror for one of the arms will be movable, allowing the optical path length difference between the

arms to be easily varied.

Suppose the two arms have exactly equal optical path lengths in arriving at the output port of the beam splitter². Then, no matter what wavelength of light is shone into the interferometer, the output port will be bright. Now, suppose the movable mirror is moved by a distance x . The phases of light from the two arms will now differ at the output port by $2\pi x/\lambda$. The phase of the interference at the output port will depend on the wavelength of the light used. For example, if white light with wavelengths spanning from 400 – 700 nm was shone in, an x larger than a couple hundred nanometers would lead to totally different phases across the spectrum. So, while slowly moving the mirror might allow one oscillation from bright to dark to be observed, subsequent oscillations would be washed out by the broad spectrum of the light. Filtering the white down to just green light 550 – 560 nm would allow observation of roughly 30 oscillations before they became completely washed out.

For a given source of light, the distance x at which the contrast of the oscillations has dropped by a factor of $1/e$ is called the coherence length of the source. Its value is generally $\sim \lambda^2/\sigma_\lambda$, where σ_λ is the standard deviation of the wavelength distribution in the spectrum. For lasers used in atomic physics experiments, the coherence length can be more than a kilometer. For this reason, the change from less coherent sources to lasers had a tremendous impact on optical interferometry. In addition to coherence, lasers also generally bring far greater brightness, increasing their appeal still further. For a Michelson interferometer with a mirror that can be smoothly moved over a long distance, this connection between spectrum and coherence length can be inverted to measure spectra. Such a device is called a Fourier transform spectrometer. As a simple example, light from a gas emitting two closely spaced wavelengths, λ and $\lambda + \Delta\lambda$, will have a coherence with a beat note equal to the inverse of the wavelength difference. Thus, fringes will wash out and then return to contrast as the mirror

²Again, depending upon the characteristics of the beam splitter, this condition might or might not be the same as saying the two arms have the same physical length.

moves over a distance $\lambda^2/\Delta\lambda$. The question of source coherence will be returned to soon in the context of atom interferometry.

3.2 Interference with Atoms

As first postulated by de Broglie, matter also has a wave-like character, with wavelength determined by the de Broglie relation $\lambda = h/p$. So, interference effects may be seen with matter waves just as with electromagnetic waves. That we don't see such things on a daily basis is part of the difficulty in understanding how classical mechanics emerges from quantum mechanics. Two aspects of matter waves make observation of interference effects difficult: Their wavelengths are generally much shorter than those of visible light, and they interact with one another and their environment far more readily than light does.

Interferometry with matter waves began not with atoms but with electrons. The first electron interferometer[64], demonstrated in 1953, used the atomic lattice of thin-film crystals as diffractive elements. Because of the short wavelength of electrons in an electron beam, directly creating elements analogous to beam splitters and mirrors was out of the question. Instead, many-order diffraction from three thin films was used. Apertures were added to eliminate all but the desired paths through the interferometer. Twenty years later, a neutron interferometer was demonstrated[87]. This used a single crystal of silicon etched to have three "blades" at which Bragg diffraction occurred, simulating the effect of beam splitters and mirrors. This produced a Mach-Zehnder interferometer for neutrons. This demonstration included inserting material into one arm and observing modulation of the populations at the two output ports, exactly analogous to our description above of the index of refraction experiment with light. The analog of "bright" for matter waves is high density, or high count rate if particles are being detected one at a time.

Atoms also have the possibility for internal state interferometry. In fact, it may be argued that Ramsey spectroscopy is an atom interferometer[22]. In a Ramsey spec-

troscopy experiment, the paths 1 and 2 are not different spatial paths, but rather, different internal states. The initial $\pi/2$ pulse is like a beam splitter, setting up a coherent superposition between two internal states the same way that a beam splitter sets up a coherent superposition between two spatial states. After some free evolution, the second $\pi/2$ pulse acts as another beam splitter, now with inputs from both the 1 and 2 paths. Whether the atom comes out in state 1 or 2 depends on the phase evolution between the two beam splitters, just as in an optical Mach-Zehnder interferometer. These connections are important, as they allow ideas to flow more easily between spectroscopy and interferometry. Indeed, we will return to the conceptual connections in chapter 6, when we introduce a momentum-space Ramsey interferometer. For the purposes of this thesis, when we speak of an atom interferometer, we will be referring to a matter wave, not internal state, interferometer.

Despite another nearly twenty year gap between the first neutron interferometer and the first atom interferometers, 1991 saw demonstrations of four different atomic matter wave interferometers: A simple double-slit experiment using material slits[16], a Mach-Zehnder interferometer using material diffraction gratings[50], a Mach-Zehnder interferometer using stimulated Raman laser pulses as beam-splitters and mirrors[49], and a Sagnac interferometer made by rotating a Ramsey spectroscopy experiment. Material slits were made by cutting slits into thin gold foils, forming the exact same apparatus as described for the Young's double slit experiment above. Material gratings are the natural extension of thin membranes (silicon nitride in this case) with many precision etched slits in them; these are just transmission diffraction gratings but with shorter spatial periods than necessary for optical transmission gratings.

The use of light pulses, as in the third atom interferometer, is a new facet of atom interferometry. Though, in a certain sense it is the most natural approach: Just as matter is used to split, reflect, and recombine light for optical experiments, light may be used to split, reflect, and recombine matter waves for atomic experiments.

The third interferometer uses two counter propagating laser beams, with a frequency difference equal to the ^{23}Na hyperfine splitting, to drive stimulated Raman transitions. Since such a transition involves absorbing a photon from one beam and stimulatedly emitting into the other, the atom picks up a net momentum kick of twice the single photon momentum in the course of the Raman process. A $\pi/2$ pulse puts each atom in an equal superposition of the two hyperfine states, but it also entangles the momentum state with the internal state. Alternatively, this pulse may be thought of as a 50-50 beam splitter that happens to tag the two paths by internal state (roughly analogous to a polarizing beam splitter for light). It turns out that analogs of these Raman processes may be carried out for transitions in a single internal state. That is, counter-propagating laser beams may be used to drive two-photon transitions that leave the internal state unchanged but give the two-photon momentum kick. Such pulses can be tailored to act as beam splitters of any splitting ratio, including fully reflective mirrors. These sort of pulses will be considered in significant detail in section [5.2](#).

Rather than giving a detailed account of all that is possible with atom interferometry, we refer the reader to the excellent review of Cronin et al[22]. This contains an exhaustive list of various styles of atom interferometer, as well as a table of available tools for implementing atomic analogs of many optical components. Instead, we return briefly to see how ideas mentioned in discussing light interference appear in atom interference.

For the Young's double-slit experiment, we considered two different descriptions. Each has a direct carry over to atom interference. The quantity of interest for atoms is the wave function ψ , which must be squared to find the atomic probability density at a point, just as the electric field must be squared to find the light intensity at a point. The wave function may be propagated as a spreading wave just like light waves. The wave equation for the electric field is here replaced with the Schrödinger

equation

$$\nabla^2\psi - \frac{1}{i\hbar/(2m)} \frac{\partial\psi}{\partial t} = 0 , \quad (3.2)$$

for the wave function of a particle with mass m . In fact, for spinless particles the analogy can be made even stronger, as the Schrödinger equation is just a non-relativistic limit of the Klein-Gordon equation

$$\nabla^2\psi - \frac{1}{c^2} \frac{\partial^2\psi}{\partial t^2} = m^2\psi . \quad (3.3)$$

The other description made use of optical paths to build up the amplitude at a specific point. This optical path description becomes the path integral description for matter waves. Here rather than the simple $2\pi x/\lambda$ for light, the phase of the matter wave evolves according to the action (i.e., the time integral of the Lagrangian) of the particle. As with light, the two descriptions are equivalent, with the path integral representing a particular way to build the Green's function for the Schrödinger operator[91].

Finally, let us return to the question of coherence length. For an optical interferometer, the coherence length of the source light is set by the distribution of wavelengths emitted by the source. White light has a coherence length on the scale of a micron, while nearly monochromatic light such as that emitted by a laser can have a coherence length of a kilometer or more. Similar considerations follow for atom interferometry. Through the de Broglie relation, we see that wavelength distribution maps directly to momentum or velocity distribution. So, the coherence length relates to the temperature of the source. Techniques for cooling an atomic gas, such as those described in chapter 2, can be used to increase the coherence length of atom interferometers. The limit of cooling bosons into a BEC corresponds, in both brightness and coherence, to use of lasers in optical experiments. We turn now to the question of how BECs have, and potentially could, impact matter wave interferometry.

3.3 BECs in Interferometry

The first observation of interference in a BEC came from the Ketterle group in 1997[4]. This constituted the first direct observation of the long-range coherence that defines the BEC state. A magnetically trapped cloud was split in two by a blue-detuned light sheet before evaporation to the BEC had completed, thus producing two independent condensates at the end of evaporative cooling. The fact that a single interference pattern was formed across the entire cloud demonstrates the long-range first-order coherence (i.e., that a macroscopic wave function describes the condensate well) of each individual condensate, just as interference between two laser beams demonstrates the first-order coherence of the individual lasers.

The Phillips group at NIST made further measurements of the first-order coherence of a BEC, both in a trap and after some expansion time following trap turn-off[38]³. Each of these measurements matched the predictions of mean-field theory⁴. The in trap measurement gave the amplitude of the condensate wave function. The post-expansion measurement revealed the loss of coherence across the condensate due to the expansion process. The experiments made use of a momentum-space Ramsey interferometer: a pair of short diffraction pulses with a variable time between pulses. As the phase of the moving state evolved relative to the stationary state, the second pulse acted as a continuation of the first pulse or a cancellation of the first pulse (or anything in between depending on the exact phase). In addition to measuring the coherence of the BEC, such a measurement is also a rudimentary recoil frequency measurement, as will be discussed further in 6.1.2.

Having mapped the in-trap condensate wavefunction's amplitude, given its constant phase, and having seen the effects of phase curvature in the expanded condensates, the Phillips group next mapped the phase curvature of expanded condensates[94].

³These measurements complimented non-interferometric measurements of the coherence such as Bragg spectroscopy[96].

⁴Mean-field theory is described in section 9.2.

This was accomplished with a modified Mach-Zehnder interferometer. The final recombination pulse was fired at times slightly off from the point of perfect overlap. This created a fringe pattern whose spacing changed as they varied the spatial offset of the two condensate branches at the time of the recombination pulse. With this technique they measured both the evolution of phase curvature during condensate expansion and the extra velocity kick given to the branches as they separate. These effects are relevant to the experiments described in this dissertation, as will be discussed in 9.5.

From these early experiments demonstrating properties of the BEC state, interferometry has moved on to using condensates as sources for external measurements. We highlight a few of the key results in the field.

The Ketterle group continued to explore interferometers using double-well geometries. They refined their methods, eventually switching to an optical dipole trap for coherently splitting a single condensate into two independent parts with a double-well trap⁵[92]. This produced deterministic phase differences, allowing phase shifts to be applied to one part and then measured by observing the resultant phase shift of the fringes.

Finally, the Ketterle group moved to magnetic traps on an atom chip (i.e., a small flat device with wires deposited on the surface). After some initial difficulties[93], they found a repeatable trap splitting process using radio frequency dressing of a single-well trap to create the double-well potential[48]. This splitting process produced demonstrable number squeezing due to the repulsive atom-atom interactions. In this same work they demonstrated a confined interferometer on the chip with a Sagnac-like geometry.

The Sackett group at UVA demonstrated a confined Michelson interferometer[32]. Their experiments have been carried out in a magnetic waveguide—essentially just

⁵The double-well was produced by simply driving the ODT acousto-optic modulator with two radio frequencies. This turned the single-well ODT into two as the two frequencies give slightly different deflections and therefore create two foci very close to one another.

a magnetic trap with low trap frequencies, one of which is particularly low and so treated as the axial direction in the waveguide. They have demonstrated several interferometer geometries in their confined waveguide, e.g. a Sagnac interferometer[13]. Magnetic waveguide interferometers tend to be limited by fluctuations and deviations from perfect harmonicity in the magnetic field induced potential.

A large collaboration, spearheaded by Ernst Rasel from Hannover, hopes to perform BEC interferometry in space. In the near term, they have demonstrated long experimental times on the order of one second in microgravity using a 100 m drop tower[75]. The use of “delta kicked cooling,” that is, turning off the trap and then turning it back on after one quarter of the trap frequency to reduce momentum width, has helped signal quality for such long interferometer times.

The Kasevich group recently demonstrated interferometric contrast with a momentum splitting of $102\hbar k$ [19]. While $\approx 18\%$ contrast was maintained, technical noise destroyed any interferometric information, i.e., there was a visibility of zero. With their 10 m high interferometry chamber, this could lead to dramatic arm separations and extremely high precisions in the future, if the technical noise can be rooted out.

Significant advances in the use of BECs for accelerometry and gravimetry have been made by the group of John Close at ANU. Utilizing a Mach-Zehnder geometry and large-momentum transfer (LMT) beam splitters—a combination of Bragg pulses and Bloch oscillations—they explored the role of source coherence in an atom interferometer for gravimetry[26]. They found markedly better fringe visibility with a BEC source compared to a thermal source with the same longitudinal momentum width.

Additionally, they demonstrated a similar interferometer in an optical waveguide (one beam from their crossed ODT) in which they achieved measurable fringe visibility even at momentum splittings of $80\hbar k$ [67, 68]. The optical waveguide provides a much smoother and more repeatable potential than the magnetic waveguides used by other groups. Using this trap, they made further study of the role of source coherence,

comparing a BEC to various thermal sources with nearly identical longitudinal momentum distributions. They again find a noticeable drop-off in fringe visibility for the thermal sources, as well as loss of coherence on much shorter time scales. However, the data do not seem convincing that this is strictly related to coherence rather than simply to the differences in transverse kinetic energy in the various sources⁶. For a dual species $^{85}\text{Rb} - ^{87}\text{Rb}$ BEC interferometer in the optical waveguide, they find a substantial phase shift due to interspecies interactions. Interaction shifts will likely be a strong effect in eventual high-precision experiments in such a waveguide, as the transverse confinement tends to keep atomic density relatively high.

A recent result from the Close group demonstrates a strong increase in fringe visibility and coherence time by use of bright solitons, as compared to non-interacting BECs, or those with weak interactions (attractive or repulsive)[69]. While the use of solitons may in the future be a boon to precision measurement, far more study is needed. However, this work does bring us full circle in the consideration of BEC interferometry. Initial BEC interferometers were used to study the properties of the BEC state. Now, potentially more exotic many-body states can be studied similarly and perhaps evolve into tools for ever finer precision measurements.

⁶While this could conceivably be considered to be a part of “coherence,” the transverse momentum spread is not generally a concern for coherence in optical interferometry, nor does it relate directly to the questions of long-range, first-order coherence discussed earlier.

Chapter 4

THE α AND THE ω_{REC}

This chapter details the larger context in which recoil measurements should be seen. Recoil measurements use a mass measurement to determine the fine structure constant. The particular importance of the fine structure constant is detailed. Afterward, the history of recoil measurements is sketched. Finally, the importance of recoil measurements in testing fundamental physics is explained.

4.1 The Fine Structure Constant

The fine structure constant, denoted α , is defined (in SI units):

$$\alpha = \frac{e^2}{4\pi\epsilon_0\hbar c}, \quad (4.1)$$

where e is the charge on the electron, ϵ_0 is the permittivity of free space, and c is the speed of light. The CODATA recommended value is $\alpha = 0.0072973525698(24)$, though it is more often written in terms of its inverse, $\alpha^{-1} = 137.035999074(44)$.

As the name suggests, this combination of constants got its own name in the course of trying to understand the fine structure splittings in the spectrum of hydrogen. Specifically, in 1916, Arnold Sommerfeld introduced a theory of elliptical orbits in the Bohr theory of the atom and calculated relativistic corrections to the energies of such orbits. In the course of these studies he defined the fine structure constant, which he called a “new abbreviation.” [76] Sommerfeld interpreted this new constant as the ratio of the speed of an electron in the first Bohr orbit to the speed of light.

From these humble beginnings, the fine structure constant has grown to immense importance in physics. It occurs in many fields of physics, including atomic, condensed matter, and high energy particle physics.

In atomic physics, it controls the hierarchy of energy splittings observed in the hydrogen atom[34]. The Bohr energy of the $n = 1$ level in hydrogen is $\frac{1}{2}\alpha^2 m_e c^2$. The fine structure, arising from a kinematic relativistic correction and the spin-orbit coupling, has size two powers of α smaller $\alpha^4 m_e c^2$. Yet another power smaller, $\alpha^5 m_e c^2$, is the Lamb shift, which arises from quantization of the electric field.

In condensed matter physics, many effects are related to the fine structure constant. In the quantum Hall effect, a strong (generally multiple Tesla) magnetic field is applied perpendicular to a two-dimensional electron gas (e.g., in a MOSFET). The system exhibits quantized conductance with values equal to $2n\alpha/(\mu_0 c)$, where n is an integer and μ_0 is the magnetic permeability of free space[54]. Since μ_0 and c are defined quantities in SI units, a measurement of the Hall resistance in SI units leads directly to a value for the fine structure constant.

In quantum field theory, the fine structure constant is the coupling constant for electromagnetic interactions. As such, any perturbative expansion in quantum electrodynamics (QED) is an expansion in powers of α . Desire to understand the Lamb shift, mentioned above, gave much of the impetus for developing consistent calculational tools in QED, which lead to the 1965 Nobel prize for Tomonaga, Schwinger, and Feynman. The gyromagnetic ratio for the electron is a fundamental physical property connected to α through QED, about which we will say more soon.

The fine structure constant has also taken on a quasi-mythical role for many physicists. As a pure number with no units, its value seems to be a fundamental structure of the universe. The values of other fundamental constants appearing in the standard model, for instance particle masses, are referenced to some human-devised unit; at best their ratios could have similar fundamental character. Some point out its clear significance as a combination of the fundamental constants relevant to relativity (c), quantum mechanics (\hbar), and electromagnetism (e)[34]. Wolfgang Pauli carried such a deep obsession with understanding the meaning of this number that his psychoanalytic sessions with Carl Jung often centered around it. He believed

that understanding the reason it takes the value it does should be the impetus for the next great revolution in physics[78]. In a similar vein, Richard Feynman commented that “all good theoretical physicists put this number $[\alpha^{-1}]$ up on their wall and worry about it.[30]”

Lest the reader be swept along in these mystical feelings that somehow this number represents a deep truth about existence, we should consider the following. In QED, whose great success gives α much of its prestige, the fine structure constant is, in fact, not constant. Rather, vacuum polarization corrections show that its value is actually a function of momentum transfer[80]. The α defined above is simply the coupling constant for QED in the limit of zero momentum transfer (i.e., $\alpha(q^2 = 0)$), which makes it sound less profound. One could even argue that the only reason we ascribe such great significance to α is that our daily lives are lived at such low energies that we can't perceive its non-constant nature.

Whether or not the next great revolution in physics reveals α to be of deep significance, we can see that since our daily lives are ruled by electromagnetic interactions at low energy, it is certainly important to physics as a human pursuit. The fact that it appears so often in three major fields of physics means we have opportunities to measure it to extremely high precision using vastly different experimental and theoretical tools. The extent to which these different measurements yield the same value constitutes a strong test of the consistency of physics over many domains.

4.2 History of Recoil Measurements

Atomic recoil measurements look to measure the ratio of an atomic mass to Planck's constant h/m_X where X will represent some particular atom. Such measurements are interesting for several reasons. Most simply, combined with an independent measurement of Planck's constant, an extremely active area of research at present, it gives the mass of the atom itself. High precision mass measurements (either m or h/m) can potentially be useful as inputs for other measurements, particularly gravitational

measurements (e.g., measuring g with atomic Bloch oscillations[84]).

Second, measures of h/m for two different atomic species can be compared to ratios of mass for the same pair determined by ion trap measurements. This allows a useful cross-check of systematics for each type of measurement. This is particularly powerful since atomic recoil measurements and ion balance measurements have largely unrelated systematics.

However, the most important motivation is measurement of the fine structure constant. Such measurement makes use of the formula

$$\alpha^2 = \frac{2R_\infty}{c} \frac{h}{m_e} = \frac{2R_\infty}{c} \frac{h}{m_X} \frac{m_X}{m_e}, \quad (4.2)$$

where $R_\infty = m_e e^4 / (8\epsilon_0^2 h^3 c)$ is the Rydberg constant with infinitely massive nucleus, and m_e is the mass of an electron. The formula was first suggested in 1978 by Wöger¹. As mentioned above, m_X can stand for the mass of any atom but in fact any test mass at all may be substituted. The first measurements of this kind found the ratio for a neutron h/m_n by using Bragg reflection from pure crystals of silicon in 1986[56]. This measurement was refined several times reaching a precision of 7.3×10^{-8} in 1999[58], which remains the highest precision measurement for a neutron[71]. These measurements depend critically on values of the lattice spacing in pure silicon crystals. However, the possibility that the spacing at the crystal surface—the actual spacing relevant to the Bragg reflections used in these measurements—could be different from the bulk value makes these measurements unreliable at their level of statistical precision.

The first precise measurement of h/m for an atom was made by the Chu group in 1993 on cesium (Cs) atoms[99], with a precision of 10^{-7} but no claim as to accuracy. This measurement used a laser cooled cloud that was subsequently velocity selected

¹The original publication by Wöger is not available to me, so I didn't feel it appropriate to list it as a regular citation. The paper of Krüger et al.[57] gives a citation to Wöger W., In *Probleme bei der Darstellung elektrischer Einheiten*, PTB-Bericht PTB-E- 12, Braunschweig, Physikalisch-Technische Bundesanstalt, 1979, 149-163.

using a Doppler sensitive Raman transition. This experiment made use of atom interferometry. The diffractive elements were produced by two counter-propagating laser beams driving Raman transitions. It used two pairs of $\pi/2$ pulses to generate a pair of interferometers moving in opposite directions. This geometry is referred to as a Ramsey-Bordé interferometer[88]. The signal was atomic population in each of the two hyperfine levels connected by the Raman beams, as a function of detuning between the beams for the second pulse pair. This signal oscillates with a fringe spacing of $1/T$ where T is the evolution time between the two pulse pairs. The recoil frequency is determined by comparing the positions of fringes for the two interferometers, whose frequency differences will differ by $4\omega_{\text{rec}}$. The experiment improved on this by adding extra π pulses of alternating orientation to increase this separation to $60\hbar k_{\text{rec}}$. This experiment was subsequently improved to achieve a precision of 5 ppb and an accuracy of 15 ppb[100].

A second group, led by Biraben at CNRS, announced a result for rubidium (Rb) in 2004 with a precision of 4×10^{-7} [9]. This experiment was subsequently refined to a precision of 9 ppb and an accuracy of 13 ppb. This experiment is quite similar to that of the Chu group. The key difference is that instead of alternating π pulses for acceleration, the Biraben group uses Bloch oscillations. These consist of loading the interferometer into a “moving standing wave” and then accelerating this standing wave. In practice this is achieved with two counter propagating laser beams with a frequency difference held fixed during an adiabatic intensity turn on followed by a linear increase of the frequency difference to provide the acceleration. This technique can be more easily scaled to very high momentum, achieving 1000 recoil momenta in their most recent paper[11], compared to a maximum of 120 reported with alternating π pulses. It also keeps the atoms in the same hyperfine state throughout the acceleration, simplifying the analysis of systematic effects.

A less substantial difference was the use of Raman π pulses rather than pairs of $\pi/2$ pulses. This was changed in later experiments, which used the Chu group’s

plan of $\pi/2$ pulse pair, acceleration, $\pi/2$ pulse pair. This allowed an improvement of accuracy to 1.3 ppb[11]. Since the formula for obtaining α from h/m contains a square root, the accuracy in their reported value for α is half as large. This brings the α measurement within a factor of three in accuracy compared to the most accurate current measurement using the electron $g - 2$, described in the next section.

Due to their similarity in design, these two sets of experiments share many characteristics. Both obtain the recoil frequency by measuring Doppler sensitive transitions in atoms that have been given a known relative momentum. That is, they essentially compare speed and momentum to determine mass: $v = p/m = k\hbar/m$, with k the atom's wavenumber. Since speed scales linearly with momentum, the precision of these measurements improves proportionally with N , the number of two-photon recoil kicks separating the two interferometers. Increasing T narrows their fringe patterns, thus improving the precision if the fringe peak position is found comparably well.

An interferometer that measured energy rather than speed was demonstrated in 2002 in the Pritchard group[36]. This new style of measurement was referred to as a contrast interferometer (CI). As contrast interferometry is the subject of much of this thesis, we will wait for a detailed description until chapter 5. One of the key points to distinguish it from the experiments discussed above is that it actually measured the recoil frequency directly. The recoil frequency ω_{rec} is the frequency associated with the kinetic energy of an atom:

$$E = \frac{p^2}{2m} \implies \omega_{\text{rec}} = \frac{\hbar k^2}{2m} . \quad (4.3)$$

Since the frequency is quadratic in p , the precision scales as N^2 . This has the potential to rapidly improve precision of such a measurement, as will be discussed in detail in ??.

Another technique which measures energy rather than speed is the grating echo interferometer. This technique has achieved precision of 37 ppb but with accuracy only at the level of 6 ppm[7]. Though this is an energy measurement, the scheme

cannot support additional momentum kicks to accelerate the arms, and so it cannot take advantage of the N^2 scaling available to the contrast interferometer. It is a temporal Talbot-Lau interferometer, using a pair of Kapitza-Dirac pulses separated by a time T as gratings. Just as a spatial Talbot-Lau interferometer obtains a matter-wave density image of the second grating at the Talbot distance, the grating echo interferometer sees such an image at the Talbot time, and determines that time precisely by Bragg reflecting light from the imaged grating.

As mentioned above, the current most accurate recoil measurement is that of the Biraben group. As of 2012, their future plans included moving to a BEC source to combat wave front curvature, their largest systematic uncertainty, and changing the interferometer geometry to allow acceleration within the Ramsey-Bordé interferometers. The work of the Chu group has passed to the hands of the Müller group. One important innovation is the insertion of accelerations between the Ramsey-Bordé arms, allowing them to take advantage of an N^2 scaling. Recently, they have demonstrated a frequency standard locked to the recoil frequency as a possible future definition of the kilogram[59]. This lock had a precision of 1.8 ppb and an accuracy of 4 ppb.

4.3 $g - 2$ and Tests of QED

The most precise determination of α comes from measurement of the anomalous part, $g - 2$, of the electron gyromagnetic ratio g [39]. The anomalous part is so named because it represents the deviation from the prediction of the Dirac equation. This equation takes into account quantum mechanics and special relativity with regard to the electron but fails to account for the quantum mechanics of the electromagnetic field. Its prediction for the magnetic moment of the electron is $\vec{\mu}_e = -2\mu_B\vec{S}/\hbar$ where \vec{S} is the electron spin and $\mu_B = e\hbar/(2m_e)$ is the Bohr magneton. More generally, the magnetic moment is written $\vec{\mu}_e = -g\mu_B\vec{S}/\hbar$

In Dirac theory the energy ΔE for an electron spin flip in a certain magnetic field

\vec{B} is identical to the cyclotron frequency $f_c = e|\vec{B}|/(2\pi m_e)$ times Planck's constant:

$$\Delta E = -\Delta\vec{\mu} \cdot \vec{B} = g_{\text{Dirac}} \frac{e\hbar|\vec{B}|}{2m_e} = \hbar \frac{e|\vec{B}|}{m_e} = hf_c . \quad (4.4)$$

The measurement uses this fact to effectively subtract out the 2 from g allowing a direct frequency measurement of $g - 2 \sim 10^{-3}$, thus saving three orders of magnitude in line splitting. This is accomplished by noting that with the equivalence in (4.4), the $n_c = 1, m_s = -1/2$ and $n_c = 0, m_s = 1/2$ states are degenerate. Therefore, the non-zero transition frequency f_a between the states corresponds exactly to $(g-2)f_c/2$, giving $g - 2 = 2f_a/f_c^2$.

The actual experiments are done by trapping a single electron in a cylindrical Penning trap. The electron's motion in the trap is cooled to sub-Kelvin temperatures. The cyclotron and spin-flip excitations are successively driven and then detected. Transitions are detected using quantum non-demolition techniques. A single electron can be trapped and interrogated for days at a time.

As mentioned in section 4.1, perturbative expansions in quantum electrodynamics are expansions in α . The calculation of $g - 2$, representing small corrections due to the quantum nature of the electromagnetic field, is just such a calculation. Measuring $g - 2$ and then inverting the QED calculation gives a value for α . However, here we run into a dilemma: Do we trust the theory of QED enough to believe a value of α determined in this fashion? QED has been a remarkably successful theory, however the $g - 2$ calculation is the most complex one undertaken in QED. The theory has, therefore, not been validated to the level of precision needed to reliably deduce α from $g - 2$ to the same accuracy as the $g - 2$ measurement.

Here we find the far stronger motivation for a recoil measurement than the simple dream of knowing α to one more decimal place. Even Pauli would likely admit that to be of dubious value to human understanding, unless that extra decimal place helped

²Small corrections to this formula due to weak axial confinement, relativistic effects, and the particular electromagnetic mode structure of the cavity must be accounted for to reach sub-ppb accuracy.

to test the revolution in physics he hoped for. The importance of measuring α to ever higher accuracy is that it allows us to test theories containing it as a parameter to ever higher accuracy. In comparing $g - 2$ to an independent value of α , we find a means to test QED to unprecedented levels, and even to test other parts of the standard model at lower energies than possible through other means. To see just how, we will briefly outline the $g - 2$ calculation.

The actual mechanics of the calculation involve use of Feynman diagrams. While it is often easy to speak of such diagrams as though they somehow represent physical processes with “virtual” particles, it should be borne in mind that they are really just calculational tools that simplify the use of perturbation theory in quantum field theories.

Figure 4.1 shows the Feynman diagrams for the lowest order contributions to the calculation of g . The zero-order diagram on the left represents the Dirac equation for a free electron in a classical electromagnetic field. The diagram on the right represents the correction from second-order perturbation theory for the quantum nature of the electromagnetic field. The details of the calculation for the second-order diagram fill about eight pages in an introductory quantum field theory textbook[80]. However, the complexity of the calculations grows like the factorial of the order in perturbation theory because both the number of diagrams and the complexity of calculating each individual diagram grows rapidly. The full analytic calculation of the sixth-order contribution was a Herculean task spanning nearly thirty years[70, 60]. The eighth-order corrections are completely out of reach of analytic computation and must be evaluated on super computers using complex computer codes.

The current best experimental value for $g - 2$ necessitates a calculation at tenth order in perturbation theory. At this level, contributions from loops involving muons and hadrons also become important. Figure 4.2 shows the importance of various perturbative contributions to the calculation of $g - 2$. Hadronic loops cannot be calculated directly due to the nonperturbative nature of QCD at low energies. Instead,

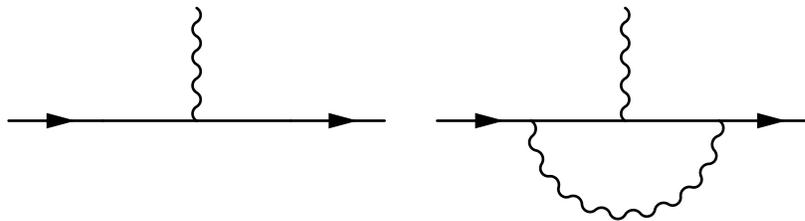


Figure 4.1: Low order contributions to g . These diagrams represent the zero- and one-loop contributions to the electron gyromagnetic ratio. The diagram on the left shows a single vertex connecting the horizontal electron line to the wavy photon line. The photon line having an external end point represents it having a real source, in this case a real magnetic field. This diagram gives the Dirac equation value $g = 2$. All subsequent diagrams add perturbative corrections. The diagram on the right gives the correction, α/π , from second order perturbation theory. (Second order because the expansion is technically in e or $\sqrt{\alpha}$, i.e., one factor of the charge for each copy of the interaction Hamiltonian.)

they are derived from low energy electron to hadron scattering data and structural properties of quantum field theories.

These experiments then test the standard model in several ways. The actual contributions of hadron and muon loops are testing the physics of these theories in a low-energy regime no other experiment can. In addition, the tenth-order (i.e., five-loop QED) calculation has recently been completed[5]. The value of the tenth-order coefficient, $A^{10} = 9.2 \pm 0.6$, was, in fact, nearly three times larger than the error bound indicated in figure 4.2, which at the time was considered quite conservative. The largeness of this coefficient may be significant. The perturbative expansion in QED is not a strictly converging expansion but rather an asymptotic expansion akin to a saddle point expansion. Thus, at some order it is expected to breakdown. Generally, the breakdown of an asymptotic expansion is signaled by a sudden rise in the size of the coefficients. Thus, understanding and testing the tenth-order contribution may be especially important to the understanding of QED.

Finally, testing the QED contributions to the electron $g - 2$ is indirectly important for searches for beyond the standard model physics. Due to the muon's much larger mass ($m_\mu/m_e \approx 200$), the muon $g - 2$ is quite sensitive to potential new physics at energy scales similar to those studied with the Large Hadron Collider. The current 3.6σ discrepancy between theory and experiment for the muon $g - 2$ has raised a number of interesting questions about new physics and our understanding of hadronic physics[35]. The electron $g - 2$ tests of QED theory essentially remove QED as a potential source of the discrepancy, allowing focus on more interesting possibilities.

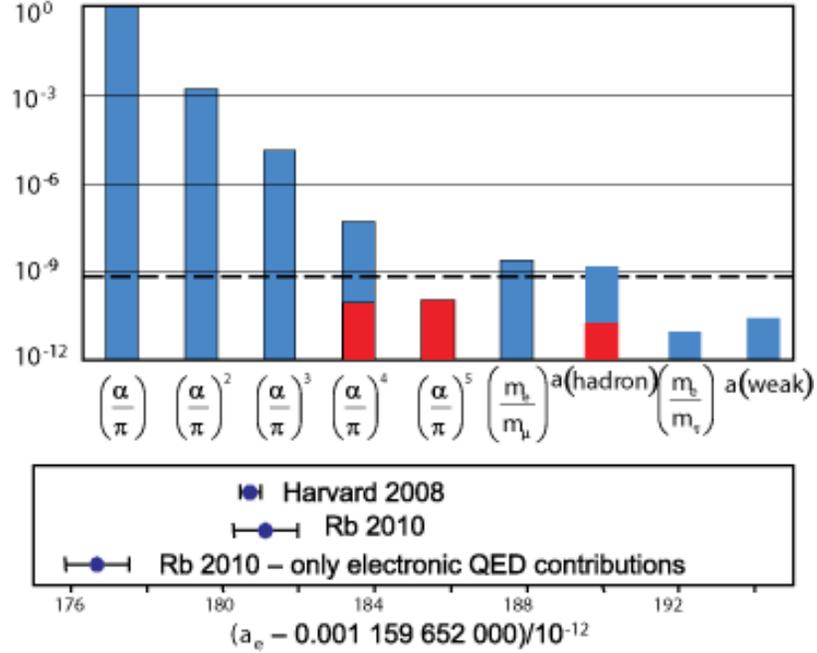


Figure 4.2: Contributions to $g - 2$ by loop content. This figure is adapted from [11]. The bars show the relative size of contributions to $g - 2$ from diagrams with loops of the described content. The bars labeled $(\alpha/\pi)^n$ indicate diagrams containing n loops with no components other than photon and electron lines. The m_e/m_μ , $a(\text{hadron})$, and m_e/m_τ bars show the contribution from two-loop diagrams containing muons, hadrons, and taus, respectively, rather than electrons. The bar labeled $a(\text{weak})$ represents one-loop diagrams involving the W^\pm or Z^0 bosons. The red bars indicate the relative uncertainties of the values. The dashed line indicates the precision of the $g - 2$ derived from the value of h/m_{Rb} . The three points plotted below show the measured $a_e = (g - 2)/2$ from Harvard[39], and the a_e derived from h/m_{Rb} with and without the muon and hadron contributions.

Part II

CONTRAST INTERFEROMETRY

Chapter 5

THE LIGHT PULSE CONTRAST INTERFEROMETER: THEORY

This chapter describes the contrast interferometer, the focus of this thesis. Both the geometry and the readout mechanism differ from those of more standard atom interferometers. As standing wave laser pulses are key elements of the contrast interferometer, the theory of such pulses is described in detail.

5.1 Geometry and Readout Mechanism

The first contrast interferometer was described by Gupta et al[36], though it bears some resemblance to earlier techniques[14]. The contrast interferometer is a three arm interferometer. Figure 5.1 shows the design schematically. We will focus on a contrast interferometer using a BEC source though, as discussed in section 6.6, use of a BEC is not necessary for contrast interferometry.

5.1.1 Basics

With a BEC source, the first step after turning off the trap is to allow the condensate to expand to reduce density. After some fixed expansion time, a standing-wave laser pulse, the splitting pulse, diffracts the condensate. The details of the optical standing wave gratings will be discussed in the next section. For the present discussion, it is sufficient to think of them as causing stimulated two-photon transitions that change the momentum of an atom in the electronic ground state by $\pm 2\hbar k_{\text{rec}}$ without causing a real electronic excitation. The splitting pulse will generally be tailored to populate three momentum states—0 and $\pm 2\hbar k_{\text{rec}}$ in the laboratory frame—though other states

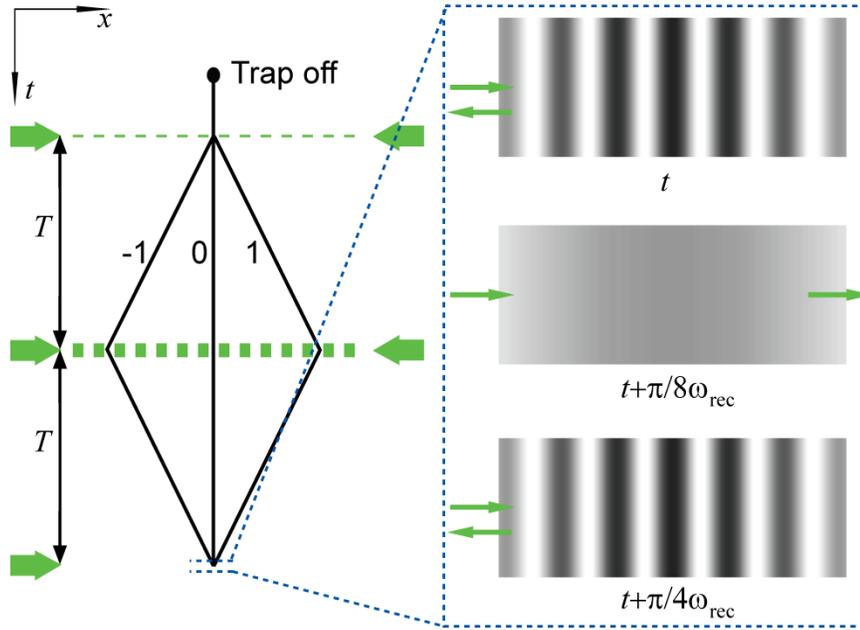


Figure 5.1: Schematic of a contrast interferometer. The left gives a space-time schematic of a contrast interferometer’s three arms. The arms are split at time 0 by a laser pulse, and at time T a second laser pulse turns the moving arms around. At the time $2T$ after the initial splitting pulse, the three arms again overlap. On the right is a representation of how the atomic density grating formed at this time rises and falls in contrast with a period of $2\pi/(8\omega_{\text{rec}})$. The readout light is reflected by a high contrast grating but passes through a zero contrast grating.

may have small populations as well. We will call the time of the splitting pulse $t = 0$.

Over a time period T , the ± 1 arms—so-called for their orders of diffraction or equivalently their momenta in units of $2\hbar k_{\text{rec}}$ —propagate freely away from the 0 arm. At time $t = T$, a second standing-wave laser pulse, the mirror pulse, reverses the momenta of the ± 1 arms. This pulse may be shaped so as to reverse these momenta without affecting the 0 arm.

Finally, the ± 1 arms are allowed to return to overlap with the 0 arm. This perfect

overlap occurs at $t = 2T$. This point of perfect overlap, called the closing time of the interferometer, is where the signal is generated.

First, consider the atomic density near the closing time if only the 0 and 1 arms existed. The wave function for a single atom would be

$$\psi(x) = A_0(\vec{x}, t)e^{-i\theta_0(t)} + A_1(\vec{x} - (2T - t)\vec{v}, t)e^{-i(2\vec{k}_{\text{rec}}\cdot\vec{x} + 4\omega_{\text{rec}}t + \theta_1(t))} ,$$

where $A_i(\vec{x}, t)$ describes the spatial extent of the wave packet for arm i , $\theta_i(t)$ collects any accumulated phases other than those related to the wavenumber and frequency, and $\vec{v} = -2\hbar\vec{k}_{\text{rec}}/m$. (The minus sign for arm 1 is due to this occurring after the mirror pulse.)

In a frame of reference moving at velocity $\vec{v}/2$, this wave function generates a standing wave of atomic density for times near $2T$. Thus, in the lab frame this corresponds to a sinusoidal density modulation moving at velocity $\vec{v}/2$. Due to the similarity of this density modulation to the structure of a diffraction grating, we will refer to it as a density grating. A similar consideration for an interferometer with only arms 0 and -1 concludes that this pair would produce a similar density grating but this time moving at velocity $-\vec{v}$ in the lab frame.

Heuristically, one may picture the full, three-arm interferometer as containing both of these density gratings moving past one another. At some time their density peaks will align, making a single, stronger density grating. At a slightly later time one grating's density maxima will coincide with the other's density minima, leading to a uniform density distribution. We describe these two times as having density gratings with high contrast and zero contrast, respectively.

More quantitatively, we can consider an atom's wave function in the three arm interferometer:

$$\psi(x) = A_0e^{-i\theta_0(t)} + A_1e^{-i(2\vec{k}_{\text{rec}}\cdot\vec{x} + 4\omega_{\text{rec}}t + \theta_1(t))} + A_{-1}e^{-i(-2\vec{k}_{\text{rec}}\cdot\vec{x} + 4\omega_{\text{rec}}t + \theta_{-1}(t))} . \quad (5.1)$$

The arguments for the amplitudes have been suppressed, but those of A_{-1} are identical to those of A_1 but with the sign of \vec{v} changed. For times near $2T$ we may essentially

ignore the spatial dependence of the A_i 's, since the extent of a condensate is much greater than the spacing between density peaks in one grating of $2\pi/k_{\text{rec}} = \lambda_{\text{laser}}/2$. The density then takes the form

$$\begin{aligned} |\psi|^2 = & |A_0|^2 + |A_1|^2 + |A_{-1}|^2 \\ & + 2A_0A_1 [\cos(2k_{\text{rec}}x + 4\omega_{\text{rec}}t + \theta_1 - \theta_0) + \cos(2k_{\text{rec}}x - 4\omega_{\text{rec}}t - \theta_{-1} + \theta_0)] \\ & + 2A_1^2 \cos(4k_{\text{rec}}x + \theta_1 - \theta_{-1}) \quad , \end{aligned} \quad (5.2)$$

where we've used the fact that the splitting pulse generally enforces a reflection symmetry making $A_1 = A_{-1}$. The second line of (5.2) may be rewritten using trig identities as

$$2A_0A_1 \cos(2k_{\text{rec}}x) \cos\left(4\omega_{\text{rec}}t + \frac{\theta_1 + \theta_{-1}}{2} - \theta_0\right). \quad (5.3)$$

The final signal to be measured is the time dependent contrast described above. Equation (5.3) gives an explicit form for this time dependence. However, this contrast has a spatial scale of $\lambda_{\text{laser}}/2$, which generally means that it is too small-scale to be reliably imaged. So, the readout must rely on a different technique.

The exactness of the spatial scale suggests the proper technique. Namely, since the density grating has spacing $\lambda_{\text{laser}}/2$, it makes an ideal Bragg reflector for the light that made the initial diffraction pulses. Further, the contrast of the grating translates directly into how good a Bragg reflector it makes. Therefore, if a traveling wave of light with the same wavelength as the diffraction beams is shone from one side some of the light will be back-reflected along the incoming path, and the magnitude of the reflected amplitude will exactly track the contrast of the density grating. This is represented on the right side of figure 5.1 by the green arrows, which show reflection for the high contrast grating and no reflection for the uniform density distribution.

This signal seems promising, however there are unknown time dependencies hiding in the θ_i 's. These contain contributions from any of a number of possible sources but we may approximate that they do not depend on the amount of free propagation T ¹.

¹That this is a reasonable approximation will be justified in section 5.1.2.

In this case, the slope of a phase at $t = 2T$, which we will call θ_{2T} , versus T curve should be simply $8\omega_{\text{rec}}$. The complete data-taking procedure is to find θ_{2T} at two or more values of T , extract the slope, and divide by 8.

5.1.2 Advantages of Contrast Interferometry

The readout procedure for the contrast interferometer is one key advantage, as it gives an entire interference fringe, in fact many fringes, in a single run of the interferometer. An example is shown in figure 5.2. In a traditional number-counting interferometer, external perturbations that shift the phase from run to run lead to loss of visibility of the signal. For a contrast interferometer, such perturbations have no effect on visibility but rather cause easily measured changes of the signal phase from one run to the next. This makes it possible to compare these phase changes to other parameters (e.g., exact diffraction pulse intensity profile) so that the source of the perturbations may be found. Other properties of the signal, such as coherence length, can also be assessed on a run by run basis. This could potentially be used as a diagnostic for certain systematic effects in the future.

A second key advantage arises from the three-arm geometry. This leads to the signal depending on the combination of phases indicated in (5.3):

$$\frac{\theta_1 + \theta_{-1}}{2} - \theta_0 . \quad (5.4)$$

Let $\delta E(\vec{x})$ denote the energy shift an atom would feel at a particular point in space. The position of the j arm of the interferometer at a given time is $\vec{x}_j = \vec{x}_0 + j(\Delta\vec{x})$. Since $\theta_j = \hbar^{-1} \int dt \delta E(\vec{x}_j)$, equation (5.4) can be rewritten as

$$\frac{1}{\hbar} \int_0^{2T} dt \frac{\delta E(\vec{x}_1) + \delta E(\vec{x}_{-1})}{2} - \delta E(\vec{x}_0) . \quad (5.5)$$

In a CI experiment, the arm separation $\Delta\vec{x}$ will be much smaller than the distance from the center of the interferometer to the walls of the vacuum chamber, which is the closest potential source of external fields. Therefore, it should be possible to expand

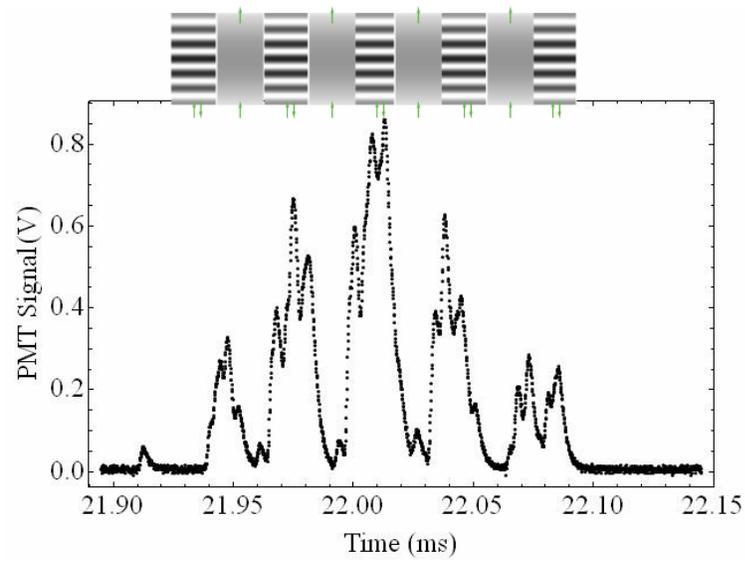


Figure 5.2: Example contrast signal. The figure shows the contrast signal from a single run of the contrast interferometer (PMT signal in volts versus time since the splitting pulse in milliseconds). Above the signal are cartoons of the atomic density, taken from figure 5.1, for the peaks and troughs of the signal.

δE around the point \vec{x}_0 in equation 5.5. Choosing coordinates such that $\Delta\vec{x}$ points in the x direction, the expansion reduces it to

$$-\frac{1}{2\hbar} \int_0^{2T} dt \left. \frac{\partial^2 \delta E}{\partial x^2} \right|_{\vec{x}=\vec{x}_0} (\Delta\vec{x})^2 + \mathcal{O}((\Delta\vec{x})^4) . \quad (5.6)$$

The key point is that the combination of phases controlling the readout signal causes constant energy offsets and gradients of energy to cancel. For instance, a contrast interferometer is insensitive to gravitational acceleration, since g is the gradient of the gravitational potential energy, while a two-arm interferometer is sensitive to g . Only the curvature of the gravitational potential energy (i.e., gradients of g) can affect a CI.

A third advantage of a CI is that the readout process does not require externally applied gratings. Just as the final beam splitter in a Mach-Zehnder interferometer turns phase difference between the two arms into differing brightness at the two output ports, in traditional atom interferometer designs some external grating is used to turn the phase difference into a population difference between two states. Any small movement of the external grating changes the measured phase. For example, in a light pulse interferometer, vibrations of the mirrors can change the position of the standing-wave optical grating between the initial splitting pulse (i.e., phase imprinting) and the final recombination pulse (i.e., phase readout). Vibrations rapidly damp out visibility in such an interferometer.

For the contrast interferometer, the two density gratings serve to read out each other's phase. Therefore, vibrations become almost completely irrelevant for the signal². Similarly, accelerations do not affect the interferometer.

While it has also been claimed that the CI is insensitive to rotations[36], this is not strictly true. A rotation between splitting and mirror pulses can cause the axis

²Vibrations of high enough frequency to substantially move an optical grating during a single pulse can affect the relative strength of the grating. This can have an effect on phase shifts due to the optical gratings. In practice, standard optical setups for cold atoms experiments provide sufficient damping of such high frequency vibrations.

for momentum transfer to differ enough between pulses to change the recoil frequency measured. This effect is discussed in section 7.1.

5.2 Optical Standing-Wave Gratings

Before delving into the technical details of atom-light interactions, we will briefly discuss a heuristic picture of optical standing-wave gratings that provides useful intuition.

When an atom in an optical grating transitions from $\vec{p} = 0$ to $\vec{p} = 2\hbar k_{\text{rec}}$, we can think of this as a coherent two-photon transition: absorption from one beam plus stimulated emission into the other. Such a transition must conserve energy and momentum. Since the atom picks up kinetic energy in such a transition, there must be a frequency difference between the absorbed photon and the emitted photon.

In practice, this frequency difference can arise in two ways. One is for the light in each beam to have sufficiently broad bandwidth to support absorption and stimulated emission differing in frequency by $4\omega_{\text{rec}}$. In such a laser pulse, the bandwidth symmetrically supports the transitions $0 \rightarrow 2\hbar k_{\text{rec}}$ and $0 \rightarrow -2\hbar k_{\text{rec}}$. Large bandwidth is usually achieved with a short pulse (e.g., a Fourier transform limited pulse of length $< 1/(4\omega_{\text{rec}})$ allows the transitions listed above).

The other method is to introduce a relative frequency shift between two counter-propagating beams with narrow bandwidth. This method breaks reflection symmetry and so gives control over whether to allow $0 \rightarrow 2\hbar k_{\text{rec}}$ or $0 \rightarrow -2\hbar k_{\text{rec}}$. Narrow bandwidth ($\Delta\omega < \omega_{\text{rec}}$) requires relatively long pulses (length $> 1/\omega_{\text{rec}}$).

Standing-wave pulses also allow for transitions of more than two photons. A short, intense pulse can support transitions such as $0 \rightarrow 10\hbar k_{\text{rec}}$, though such transitions are unwanted for contrast interferometry. As the splitting pulse is generally a short, high bandwidth pulse, care must be taken to tailor it to suppress transitions to higher momenta.

The mirror pulse is, in this picture, a four photon transition. The mirror pulse is

a low-bandwidth pulse with no frequency shift between the two beams. Therefore, it can only support transitions that do not require the atom to change energy, that is, transitions of the form $p \rightarrow -p$.

5.2.1 *Atom-Light Interactions*

The interaction of atoms and light can be viewed at many levels of sophistication. As seen in Foot’s textbook[31], a purely classical approach treating an atom as a charged, damped harmonic oscillator interacting with electromagnetic waves captures surprisingly many facets of reality: We find both the dipole force and the scattering force and may derive the correct line shape as a function of frequency. Only as the intensity of the light reaches the saturation intensity, or when coherent excitations are considered, does a semi-classical description become necessary. As we will be concerned with both, we will develop a semi-classical model to deal with condensate splitting and mirror pulses. A fully quantum mechanical treatment becomes necessary only for describing the back-action of the final readout pulse, where the granular nature of the light as a collection of photons becomes important.

The books of Cohen-Tannoudji et al.[20, 21] give extremely clear and detailed accounts of both semi-classical and fully quantum mechanical descriptions of quantum electrodynamics applied to atoms. We will attempt to cover only the bare minimum, collecting arguments scattered through several chapters of [20].

To properly include electromagnetic fields into quantum mechanics, we need the vector potential \vec{A} as well as the scalar potential ϕ . We also must account for the fact that these have too many degrees of freedom—that they may be gauge transformed without any effect on the predicted physics. For a semi-classical treatment, we may simply choose a gauge and then stay consistently within it. For atomic physics, the Coulomb gauge, $\nabla \cdot \vec{A} \equiv 0$, is particularly easy to work with. This ease follows from the way it separates the static Coulomb fields that give an atom its unperturbed structure and the radiation fields that we send from outside to manipulate the atom.

Considering the implication of this gauge choice for Gauss' law, $\nabla \cdot \vec{E} = \rho/\epsilon_0$, we see that the vector potential has no contribution to the static electric field. In fact, the scalar potential follows the expectation based upon the present positions of any charged particles, as though they were static. This is mathematically equivalent to the vector potential having no longitudinal component (i.e., in momentum space the condition becomes $\vec{k} \cdot \vec{A} \equiv 0$).

The Hamiltonian is modified by replacing the physical momentum \vec{p} with the canonical momentum $\vec{p} - q\vec{A}$, where q is the charge of the electron. There is also a term directly coupling the spin to the magnetic field, but we will exclude it, as it will be irrelevant to the atom-light coupling. The Hamiltonian takes the form

$$\sum_j \left[\frac{(\hat{p}_j)^2}{2m} - \frac{q}{m} \hat{p}_j \cdot \vec{A}(\hat{x}_j) + \frac{q^2}{2m} \left(\vec{A}(\hat{x}_j) \right)^2 + q\phi(\hat{x}_j) \right], \quad (5.7)$$

where j indexes a sum over all the electrons of the atom. (We assume the nucleus to be static.) In cold atom experiments, we wish to use light to manipulate atoms without markedly changing their properties. As such, $q|\vec{A}|$ will be much smaller than $q|\phi|$, which by the Virial theorem implies it will likewise be much smaller than $\vec{p}^2/2m$. Thus, the first three terms of (5.7) have a hierarchy of scale making the second a perturbation to the first and the third negligibly small for our purposes. Grouping the first and last terms we have the Hamiltonian for the atom's free-space structure, with the second term acting as a perturbation.

For visible light incident on an atom, the wavelength of the radiation is dramatically longer than the size scale of the atom itself (separated by three to four orders of magnitude). Thus, to a good approximation, the vector potential may be replaced by its value at the atom's center of mass, \vec{x}_{CM} . Then, (5.7) has the form, before expanding the square, of $(\vec{p} - q\vec{A})^2/(2m)$ where \vec{A} is now simply a (spatially) constant offset to the momentum. It then seems natural to use the momentum translation

operator to remove the offset. That translation operator takes the form

$$\hat{T}(t) = \exp \left(-\frac{iq}{\hbar} \sum_j \hat{x}_j \cdot \vec{A}(\hat{x}_{\text{CM}}, t) \right),$$

where we've made the time dependence of the vector potential explicit because the time dependence affects the transformation of the Hamiltonian under this operator.

Using the relation

$$\hat{H}'(t) = \hat{T}(t)\hat{H}(t)\hat{T}^\dagger(t) + i\hbar \frac{d\hat{T}}{dt} \hat{T}^\dagger(t)$$

we find the transformed Hamiltonian to have the form

$$\sum_j \left[\frac{(\hat{p}_j)^2}{2m} + q\phi(\hat{x}_j) \right] - \hat{d} \cdot \vec{E}(\hat{x}_{\text{CM}}, t), \quad (5.8)$$

where $\hat{d} = q \sum_j \hat{x}_j$ is the electric dipole operator and $\vec{E} = \partial \vec{A} / \partial t$ is the externally applied electric field (i.e., excluding the strong Coulomb fields within the atom, which ϕ accounts for). Again, we find a Hamiltonian with a form that separates the free-space atomic structure terms, in brackets, from an externally applied perturbation.

A peculiar facet of the application of the dipole approximation to cold atoms bears mention. We assumed above that in dealing with visible light there is a strong separation of scales—a few Bohr radii, $a_0 \approx 0.05$ nm, versus optical wavelengths $\lambda \approx 500$ nm—allowing us to replace the vector potential with its value at the atom's center of mass. However, for a BEC, the center-of-mass wave function will generally extend over larger length scales than the optical wavelength. We find ourselves in the peculiar position of arguing that the relative wave function is localized on scales much smaller than λ while the center of mass wave function is delocalized over scales substantially larger than λ . This does work mathematically. However, the fact that one can equally-well write the wave function in terms of the position of each individual electron, in which case each electrons' wave function is delocalized over a many- λ wide region, but the dipole approximation holds correctly at each of those points, causes me some level of consternation.

5.2.2 Modeling Diffraction

With the electric dipole Hamiltonian in hand, we turn now to the practical question of calculating how near-resonant light can diffract an atom. We move to the interaction picture, separating out the interaction Hamiltonian $\hat{H}_{\text{dipole}} = -q\vec{E}(\hat{x}_{\text{CM}}) \cdot \sum_j \hat{x}_j$. Using light of a frequency that interacts predominantly with only one pair of atomic levels, we can truncate the Hilbert space by ignoring all other levels. We thus reduce to a two-state model of an atom coupled to a light field. This entails inserting complete sets of the eigenstates of the unperturbed atomic Hamiltonian, $|n\rangle$:

$$\begin{aligned} \vec{E} \cdot \hat{r} &= \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} |i\rangle \langle i| \vec{E} \cdot \hat{r} |j\rangle \langle j| = \sum_{i,j \in \{g,e\}} \vec{E} \cdot \vec{\mu}_{ij} |i\rangle \langle j| \\ &= \vec{E} \cdot \vec{\mu}_{ge} |e\rangle \langle g| + \vec{E}^* \cdot \vec{\mu}_{ge}^* |g\rangle \langle e| , \end{aligned} \quad (5.9)$$

where for electronic wave functions $\psi(\vec{r})$:

$$\vec{\mu}_{ij} = \int d^3r \psi_i^*(\vec{r}) \vec{r} \psi_j(\vec{r}) .$$

As we are reducing the Hilbert space for the electronic states, we should also mention the form of the unperturbed Hamiltonian with this reduction. As these are eigenstates of the unperturbed Hamiltonian, it takes the simple form $\hbar\omega_0 |e\rangle \langle e| + (\hat{p}_{\text{CM}})^2/(2M)$, where M is the mass of the atom and the second term implicitly includes the identity operator on the internal electronic states (just as the other terms implicitly include the identity operator on position states). We define the so-called Rabi frequency $\omega_{\text{R}} \equiv \vec{E}_0 \cdot \vec{\mu}_{ge}/\hbar$, which implicitly contains the center-of-mass position operator through the electric field amplitude \vec{E}_0 . With this definition we may write the Hamiltonian as

$$\hat{H}' = \frac{(\hat{p}_{\text{CM}})^2}{2M} + \hbar\omega_0 |e\rangle \langle e| + \frac{\hbar\omega_{\text{R}}}{2} e^{i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}} - i\omega t} |e\rangle \langle g| + \frac{\hbar\omega_{\text{R}}^*}{2} e^{-i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}} + i\omega t} |g\rangle \langle e| . \quad (5.10)$$

The electric field may be separated into the phase term $e^{i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}} - i\omega t}$ and the slowly varying amplitude $\vec{E}_0(\hat{x}_{\text{CM}}, t)$. The difference in energy between ground and excited states, $\hbar\omega_0 \sim 1$ eV, is a dramatically larger energy scale than the kinetic energy of

the atoms $\sim 10^{-10}$ eV, which will eventually be our concern. As such, we would like to separate the extremely fast electronic dynamics from the atomic center-of-mass dynamics. For the case of tuning near resonance, ω , the frequency of the light, and ω_0 will be of similar scale. It turns out there is a single unitary transformation that removes each of these high frequencies from the problem. Consider the operator

$$\hat{T}(\omega) \equiv |g\rangle \langle g| + e^{i\omega t} |e\rangle \langle e| .$$

Applying this to our full Hamiltonian, we transform it into

$$\hat{H}' = \frac{(\hat{p}_{\text{CM}})^2}{2M} - \hbar\delta |e\rangle \langle e| + \frac{\hbar\omega_{\text{R}}}{2} e^{i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} |e\rangle \langle g| + \frac{\hbar\omega_{\text{R}}^*}{2} e^{-i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} |g\rangle \langle e| . \quad (5.11)$$

This form for the Hamiltonian will be important to the analysis here and in chapter 8. It is interesting to note that if we instead apply the related operator $\hat{T}(\omega_0)$, we get a slightly different form for the Hamiltonian:

$$\hat{H} = \frac{(\hat{p}_{\text{CM}})^2}{2M} + \frac{\hbar\omega_{\text{R}}}{2} e^{i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} e^{-i\delta t} |e\rangle \langle g| + \frac{\hbar\omega_{\text{R}}^*}{2} e^{-i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} e^{i\delta t} |g\rangle \langle e| . \quad (5.12)$$

This essentially moves the detuning from an energy offset on the excited state to a modulating frequency for the coupling between ground and excited states. Such trading may be useful for simplifying particular calculations.

For describing diffraction, we generally consider light far-enough detuned from resonance that the excited state is never substantially populated. It would be nice to be able to eliminate the excited state from our analysis altogether. A common approach is to consider equation (5.11), treating the last two terms as perturbations: $V = (\hbar\omega_{\text{R}}/2)e^{i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} |e\rangle \langle g| + (\hbar\omega_{\text{R}}^*/2)e^{-i\vec{k}_{\text{rec}} \cdot \hat{x}_{\text{CM}}} |g\rangle \langle e|$. Since these don't give a ground state to ground state matrix element until second order in perturbation theory, the standard technique is to find the matrix element in second order perturbation theory and then replace these two terms with that matrix element. This gives

$$\hat{V}' = \hbar \frac{|\omega_{\text{R}}|^2}{4\delta} |g\rangle \langle g| . \quad (5.13)$$

This result is unchanged if we include the spatial dependence of the electric field ($e^{i\vec{k}_{\text{rec}} \cdot \vec{x}_{\text{CM}}}$ in the current case) in ω_{R} . The effect of this perturbation—a detuning and intensity dependent shift of energy for ground-state atoms—is often called the AC Stark effect.

One might question the idea that these terms may be treated as a perturbation, since we will be interested in cases with $\omega_{\text{R}} \gg \omega_{\text{rec}}$, where ω_{rec} is the natural energy scale for the kinetic term. There are two regimes in which this approximation is valid. The first is $\omega_{\text{rec}}, \delta \gg \omega_{\text{R}}$, in which case the last two terms do constitute a small perturbation. The other regime is $\delta \gg \omega_{\text{R}} \gg \omega_{\text{rec}}$. In this case we can instead treat the kinetic term as a negligible perturbation. If we drop the kinetic term, we can then treat the last two terms as a perturbation in this truncated Hamiltonian.

The particular case of interest to us involves the electric field formed by two counter-propagating beams traveling in the $\pm x$ directions and linearly polarized in the z direction. This gives an electric field $E_0 \hat{z}(e^{-ikx} + e^{ikx})$ which gives a final form

$$\hat{V}' = \hbar \frac{\omega_{\text{R}}(t)^2}{4\delta} (2 + e^{-i(2k)\hat{x}} + e^{i(2k)\hat{x}}) = \hbar \frac{\omega_{\text{R}}(t)^2}{2\delta} (1 + \cos(2k\hat{x})) . \quad (5.14)$$

Notice that we have dropped the “rec” subscript, which is implicit in k being the wave vector of the light.

The first form in (5.14) is useful in the first case ($\omega_{\text{rec}}, \delta \gg \omega_{\text{R}}$), known as the Bragg regime. The second form is useful in the second case ($\delta \gg \omega_{\text{R}} \gg \omega_{\text{rec}}$), known as the Raman-Nath regime. We treat the Bragg regime first.

The operator $e^{i(2k)x}$ is a translation operator in momentum space, taking $p \rightarrow p + \hbar(2k)$. Since the first term in parentheses is a constant, representing an overall offset in energy which cannot affect the dynamics, it may be discarded³. This leaves a Hamiltonian that may be easily written using the momentum-state basis:

$$\hat{H} = \int dp \left\{ \frac{p^2}{2m} |p\rangle \langle p| + \hbar \frac{\omega_{\text{R}}^2}{4\delta} (|p + 2\hbar k\rangle \langle p| + |p - 2\hbar k\rangle \langle p|) \right\} . \quad (5.15)$$

³For real beams which are not perfect plane-waves, this term would introduce a spatially dependent energy offset, similar to an optical dipole trap.

To simplify future expressions, we define the two-photon Rabi frequency $\omega_{\text{R}}^{(2)} \equiv \omega_{\text{R}}^2/(2\delta)$.

At the beginning of 5.2 we discussed a picture in terms of two-photon transitions where pulse bandwidth and conservation of energy determined which momentum states were accessible. In this picture, the Bragg regime corresponds to a strict enforcement of conservation of energy for atoms (i.e., low bandwidth pulse). To have any appreciable effect, the length τ of a pulse must satisfy $\omega_{\text{R}}^{(2)}\tau \sim 1$. In the Bragg regime, $\omega_{\text{R}}^{(2)} \ll \omega_{\text{rec}}$ implies that the pulse will be long enough that its Fourier-limited bandwidth is much smaller than the recoil frequency, and therefore is unable to support transitions that change the magnitude of atomic momentum.

This insight allows us to treat the particular case of two-photon transitions between $p = \hbar k$ and $p = -\hbar k$. Restricting to these two states, the Hamiltonian reduces to a two-by-two matrix:

$$H = \begin{pmatrix} \omega_{\text{rec}} & \frac{\omega_{\text{R}}^{(2)}}{2} \\ \frac{\omega_{\text{R}}^{(2)}}{2} & \omega_{\text{rec}} \end{pmatrix} \quad (5.16)$$

with eigenvalues $\lambda = \omega_{\text{rec}} \pm \omega_{\text{R}}^{(2)}$ and eigenvectors

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ \pm 1 \end{pmatrix}. \quad (5.17)$$

From this we can determine the time evolution operator $U(t_i, t_f)$. Since the forms for the eigenvalues and eigenvectors are independent of $\omega_{\text{R}}^{(2)}$, we can consider the problem with an arbitrary time-dependence to the laser intensity (as long as it remains in the Bragg regime):

$$\begin{aligned} U(t_i, t_f) &= \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} e^{\int dt \omega_{\text{rec}} + \omega_{\text{R}}^{(2)}/2} & 0 \\ 0 & e^{\int dt \omega_{\text{rec}} + \omega_{\text{R}}^{(2)}/2} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \\ &= e^{i\omega_{\text{rec}} t} \begin{pmatrix} \cos(\theta/2) & i \sin(\theta/2) \\ i \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}, \end{aligned} \quad (5.18)$$

where $\theta = \int_{t_i}^{t_f} dt \omega_R^{(2)}(t)$ is referred to as the area of the pulse (i.e., the area under an intensity vs. time plot, with appropriate constants to render it dimensionless). As for the standard Rabi problem, we see that a state beginning in an unperturbed eigenstate (e.g., $|\hbar k\rangle$) will oscillate in probability between the two unperturbed eigenstates with frequency $\omega_R^{(2)}$.

To what extent are we introducing errors by neglecting states other than $|\pm\hbar k\rangle$? Essentially, this is controlled by the extent to which those other states are mixed into the $|\pm\hbar k\rangle$ states by the interaction. Since the Bragg regime has the optical coupling as a small perturbation, we may answer this question straightforwardly in perturbation theory. If we begin in the $|\hbar k\rangle$ state, the most important neglected state is $|3\hbar k\rangle$, as all other states are suppressed by larger orders in perturbation theory (and all except $|-3\hbar k\rangle$ are also suppressed by larger energy differences in the denominator). We find the first-order result

$$\langle 3\hbar k | \hbar k \rangle = \frac{\omega_R^{(2)}}{16\omega_{\text{rec}}} . \quad (5.19)$$

This weak mixing gives another perspective on the idea of conservation of energy. For weak coupling, a long, low-bandwidth pulse is needed. It also means the energy differences between recoil states are quite large compared to the coupling. So those states described in the previous picture as not allowed by conservation of energy for low bandwidth pulses, are not mixed with the initial states because of the large energy denominators. This gives us a systematic means for improving the conservation of energy argument away from the limit of infinitely weak pulse with infinite length, where it is exact.

The mirror pulse in our experiments causes a transition between $|2\hbar k\rangle$ and $|-2\hbar k\rangle$. This four-photon transition can be understood through the same procedure as the two-photon transitions discussed above. Namely, for atoms beginning in one of these two energy-degenerate states, we may eliminate the middle state, $|0\hbar k\rangle$, with exactly the same steps as we used to eliminate the electronically excited states for the two-photon transition. We find a coupling strength of $(\omega_R^{(2)}/2)^2/(4\omega_{\text{rec}})$, leading us to define the

four-photon Rabi frequency:

$$\omega_{\text{R}}^{(4)} \equiv \frac{\omega_{\text{R}}^4}{2^3 \delta^2 (4\omega_{\text{rec}})} . \quad (5.20)$$

The analysis of the new effective two-state system follows identically to that carried out for the two-photon transition.

The middle expression of (5.14), which we applied to the Bragg case, is essentially a momentum-space picture, when the exponentials are seen as momentum translation operators. The third expression of (5.14) gives a physical-space picture of a sinusoidally varying potential energy. In the Raman-Nath regime, we have said that the kinetic energy may be ignored. Therefore, the Schrödinger equation takes on a particularly simple form:

$$i\hbar \frac{\partial \psi}{\partial t} = V(x)\psi = \hbar\omega_{\text{R}}^{(2)}(t) \cos(2kx) . \quad (5.21)$$

This may be integrated directly, giving

$$\psi(x, t) = \exp \left[i \left(\int_0^t dt' \omega_{\text{R}}^{(2)}(t') \right) \cos(2kx) \right] \psi(x, 0) . \quad (5.22)$$

For a zero-momentum initial state this reduces to

$$e^{i\theta \cos(2kx)} = \sum_{n=-\infty}^{\infty} i^n J_n(\theta) e^{i(2nk)x} , \quad (5.23)$$

where θ is again the area of the pulse. The equality in (5.23) may be derived straight forwardly from the generating function, $\exp(x(t - 1/t))$, for the Bessel functions of the first kind[3]. In the Raman-Nath regime, the potential writes a particular phase structure onto the atoms. The decomposition into plane-waves with Bessel function amplitudes allows us to easily see the long-time dynamics generated by this phase imprinting.

As with the picture of the Bragg regime as a two-state transition enforced by energy conservation, the Raman-Nath regime is only exact for extreme pulses (i.e., arbitrarily high intensity, arbitrarily short pulses) and cannot be straight-forwardly

improved with the derivation given. Attempting to simply add the kinetic energy back in the position-space picture is problematic because the derivative representation needed in position-space makes this a singular perturbation. We will skip the derivation, as it is not terribly important to our present purposes, but by using the second expression from (5.14), the same form of plane waves with Bessel function coefficients may be derived. In this momentum-space picture the kinetic energy is not a singular perturbation, and so the Raman-Nath approximation may be systematically corrected. For our purposes it is only necessary to know that the corrections are suppressed as powers of $\omega_{\text{rec}}/\omega_{\text{R}}^{(2)}$, so that we may assess the validity of this approximation for real pulses used in experiments.

Chapter 6

THE LIGHT PULSE CONTRAST INTERFEROMETER: EXPERIMENTS

This chapter discusses, in detail, contrast interferometry experiments. The structure of the interferometry setup, various calibration procedures, and how to find a CI signal are detailed. Then two key innovations—decoupling the diffraction and readout paths and the ability to radically alter detuning thanks to the narrowness of the intercombination line—are discussed. Finally, new results are presented in the areas of precision interferometry and the study of many-body phenomena. It is worth noting that these experiments, performed with ^{174}Yb are the first reported interferometers using ytterbium atoms.

6.1 *Building the Interferometer*

Figure 6.1 shows absorption images of our contrast interferometry sequence. Key experimental parameters are introduced below in the text but collected for convenience in the figure caption. For diffraction pulses and readout light, we utilized light detuned $\lesssim 1000\Gamma$ from the $^{174}\text{Yb } ^1S_0 \rightarrow ^3P_1$ intercombination transition at 556 nm. This transition has a line width of $\Gamma = 2\pi \times 182$ kHz.

To get this light, we took the zero order from the Yb MOT acousto-optic modulator (AOM) and sent it, double passed, through a 200 MHz AOM. With our setup, this light would be detuned from the transition by ≈ -392 MHz before passing through the AOM. For the early experiments, the light would leave the second pass through the AOM ≈ -126 MHz detuned from the transition, where we use “ \approx ” because slightly different settings were used at various times. Actual control of the detuning

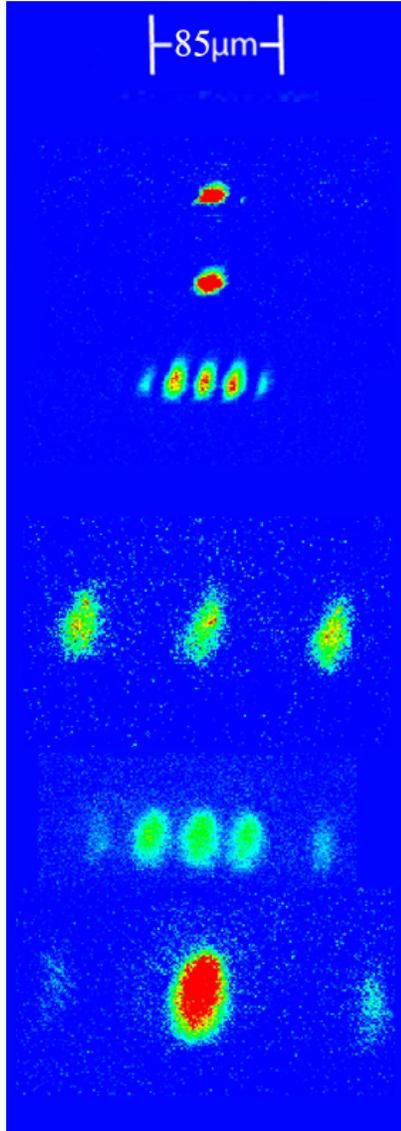


Figure 6.1: My CI. The sequence of absorption images shows the various stages of the contrast interferometer, which are shown schematically in figure 5.1. These depict a $T = 11$ ms interferometer with a 2 ms expansion prior to the splitting pulse. The splitting pulse was a $1.5 \mu\text{s}$ long square pulse. The mirror pulse was a Gaussian with temporal half-width $35 \mu\text{s}$. Though its effects are not detectable here, readout light was a weak $160 \mu\text{s}$ square pulse. Diffractive pulses used a detuning of $+700\Gamma$ from the $^1S_0 \rightarrow ^3P_1$ transition in ^{174}Yb , while the readout light was detuned by $+50\Gamma$ from the same transition.

was accurate to better than 50 kHz, as determined by absorption imaging with this light. The detunings were eventually varied quite a lot, as described in section 6.5.3.

From there the light was coupled into a single-mode fiber and sent to the interferometry breadboard described below. We achieved power out of the fiber of 40 mW, corresponding to a total efficiency including AOM and fiber coupling of 20 – 25%. We utilized a shutter in front of the fiber so that the AOM could be kept on at full power until about 3 ms before the diffraction pulses were applied. Keeping the AOM warmed up in this fashion seemed to improve the repeatability of the pulses.

6.1.1 Optical Setup

The Yb contrast interferometer described in this thesis was built in the existing Yb/Li mixture machine, described in Hansen et al.[42], and in extreme detail in two other theses[40, 52]. I was deeply involved in the mixture experiments[45, 41, 53, 42]. Particularly relevant to the current discussion, I was involved in achieving the first Yb BECs in the machine and in steadily improving Yb cooling techniques to produce BECs of 3.0×10^5 atoms, the largest ever reported.

With Yb BECs already generated in the machine, additions needed to be made to allow for standing-wave diffraction pulses and for detecting the back-reflection of the readout light. As the machine’s optical access was already spoken for by the mixture experiment, one of the MOT axes was modified to allow diffraction and readout along this axis after the MOT was extinguished. The mirror that coupled the MOT light into the chamber was replaced by a computer-controlled flipper mirror. Likewise, the retro-reflection mirror for this axis was replaced with an identical flipper mirror. Initially, the flipper mirror on the retro side was used to send the light to a quarter-wave plate and another mirror for the actual retro-reflection. This decision was made due to the lack of fine adjustment knobs on the flipper mount. Thus, alignment could be made on the rigidly mounted retro mirror. With time, space on the interferometry bread board became more valuable and this design was scrapped in favor of attaching

the flipper mount to a second mount with fine adjustment screws and placing the quarter-wave plate into the collar in contact with the mirror surface. With these designs, the mirrors flipped out of the way after the MOT was extinguished, during the evaporative cooling phase. This opened the axis for interferometry.

The interferometry bread board is shown in its final configuration in figure 6.2. The initial design included optics for shaping and directing the diffraction light as well as an optical path for detecting the back reflection signal. Later refinements appearing in the figure include greater polarization control of the diffraction light, a separate fiber for the readout light along with optics to shape and direct that light, and a port for detecting the intensity of the diffraction light. In the interests of brevity, we will simply discuss this depicted final state of the interferometry breadboard.

Before describing the board in detail, several generic design considerations are worth mentioning. Once the beam is expanded to its desired size at the atoms, it is important not to artificially distort the wavefronts by aperturing the beam. Aperturing can occur in several ways. An actual physical impediment is the general usage for aperture. For instance, using a one inch diameter lens clips any part of the beam that extends outside the area of the lens. Similarly, a one inch diameter mirror can only reflect light impinging on its surface, and so it makes a virtual aperture. Therefore, two inch diameter optics and careful alignment to the centers of all optics are import to avoid aperturing. The effective size of a mirror is reduced by the cosine of the angle the light's propagation direction makes to the mirror's normal vector, that is, the angle of incidence from geometric optics. Therefore, efforts were made in the layout of the optics to minimize the angle of incidence in each reflection. In considering sizes of optics, it is good to bear in mind that the beam waist is a radius, not a diameter, and that to avoid diffraction the beam should always have a clearance of several waists to either side of the beam's center.

The other general consideration is that somewhere in the setup there must be a beam splitter for separating the diffraction path or outgoing readout light from the

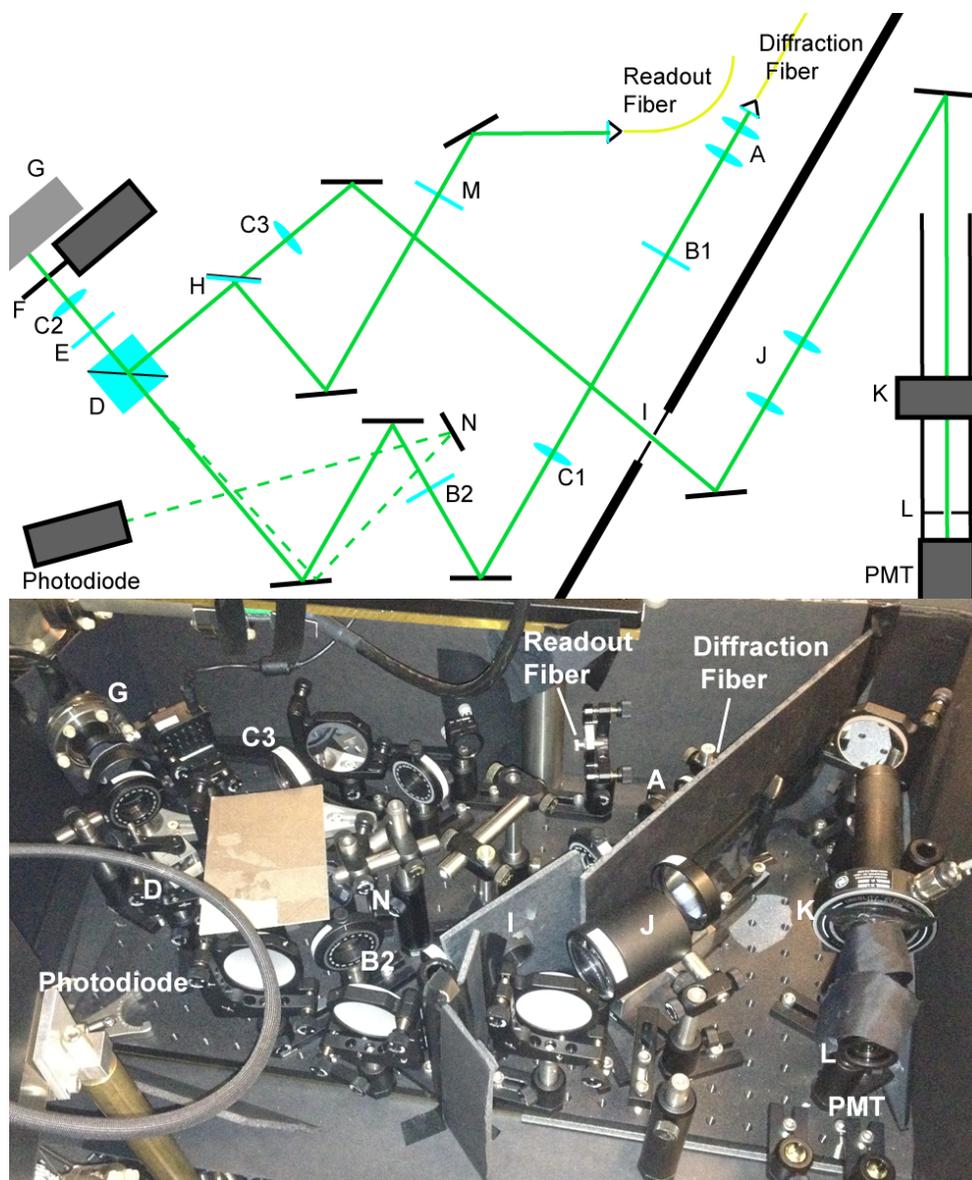


Figure 6.2: Interferometry breadboard. The top gives a schematic for the interferometry breadboard, discussed in detail in the text. The bottom is a photograph of the board, with several components labeled to aid in connecting the schematic to the photo. The photodiode and PMT have been removed before taking the photo but their positions are labeled.

back reflection signal. Any optics between this beam splitter and the atoms (including viewports) are potential sources of back reflections that act as backgrounds to the readout signal. Thus, it is advisable to have the diffraction beams and readout light enter the chamber at a nonzero angle of incidence to the viewports. Similarly, lenses should be tilted slightly off axis. Back reflections can be easily traced through the setup by turning up the beam powers and turning off the lights, so these adjustments may be made in a straight forward manner once the setup is built.

The diffraction light comes from a fiber at the top right corner of the breadboard. Lenses (labeled A in the figure) collimate the light. Since collection of the readout signal calls for a lens as close to the viewport as possible, the light is in fact not collimated through most of the setup. Rather, the next lens, C1, the light encounters is the first lens of a telescope completed by the lens near the viewport, C2. This is helpful as it keeps the beam smaller than the final waist (4 mm) through most of the setup, reducing the potential for aperturing. A quarter-wave plate, B1, between A and C1 linearizes the polarization coming from the fiber.

Next, the light reaches a set of three mirrors designed for steering the beam. While two would suffice for sufficient flexibility in steering the beam, three were used to minimize the angle of incidence. A half-wave plate, B2, between the first two mirrors sets the polarization to maximally transmit through the polarizing beam splitter D. The half-wave plate E is used to rotate the polarization for maximum diffraction efficiency. Due to birefringence in the viewports, the diffraction efficiency is a function of the polarization sent to the chamber. In this final configuration, the polarizing beam splitter also functions to separate the diffraction light and the back reflection signal.

The final piece before the chamber is a collimating lens, C2, the second half of the telescope that begins with C1. While the beam could be collimated anywhere, as mentioned above this lens also serves to image the atoms to maximize the back reflection light collected. Properly imaging the atoms also aids in finding the signal

initially. The flipper mirror, F, will be out of the path during the interferometer run. Finally, the light enters the viewport, G.

The optics on the far side of the chamber are much simpler. They consist essentially of a single lens and a retroreflection mirror at the focal point of the lens, forming a 1:1 telescope. The telescope is essential because the retro-reflection must be blocked during the readout step. Otherwise it would completely swamp the back reflection signal from the atoms. This requires a fast shutter, which is placed near the focus to minimize the shutter size and time needed to block the light completely. The decision to use a single lens was largely a concession to the confined space of the MOT breadboard onto which the retro optics had to fit. As will be discussed in section 6.1.2, this is strongly recommended against for future designs.

Next, we have the optics for imaging the back-reflection signal. The back reflection from the atoms will be collimated by the lens C2. The polarization of the readout light is set so the signal will reflect out of the polarizing beam splitter toward H, which is a 10% reflecting beam splitter. The 90% of the signal that passes through it is focused by C3, which together with C2 forms an imaging system for the atoms with magnification -1. The iris I is placed at the focal point of C3 to block unwanted back reflections from optics and stray light from elsewhere in the room. Next the light passes through a second telescope, J, making a complete imaging system (including C2 and C3) with a magnification of five. This makes high quality images of the back-reflection signal possible. A second iris, L, is placed near the focal point of J to further block stray light from reaching the photomultiplier tube (PMT), which is positioned at the image point of J. A fast mechanical shutter, K, is placed before the PMT to block the back reflection of the diffraction light during splitting and mirror pulses. This shutter is critical, as it blocks several milliwatts of light from reaching a sensor that can be overwhelmed by a nanowatt. Initially we used a Stanford Research Systems SR475 shutter. These shutters promised (and, in fact, delivered) extremely low vibrational noise, which seemed ideal since our shutter needed to sit on the interferometry board

itself. However, we found that a sufficient amount of light leaked through the shutter to provide a strong signal on the PMT. To avoid potentially degrading the PMT over thousands of exposures to such strong signals, we opted for the loud, but actually opaque, Uniblitz LS6 shutter.

Finally, we have the readout optics. While the previous descriptions of contrast interferometers used the same light for diffraction and readout, as discussed in section 6.4.2, we found great improvement from using a separate path for the readout light. Just to the left of the diffraction light's fiber at the top of the board is the readout fiber. The light is roughly collimated out of the fiber. The half-wave plate M rotates its polarization to maximally reflect at D. Before reaching D, 10% reflects at H. The other 90% is dumped into the black cap from an optical fiber to avoid creating excess background light in the readout path. After reflecting at D, it follows the same path to the atoms as the diffraction light. Since the light is roughly collimated before the beam splitter, the final lens will focus it onto the atoms. The reason for this focusing is discussed in 6.4.2.

Though not optical elements, per se, the large amount of black cardboard visible in the image is important enough to merit a comment. Protecting a PMT driven to the single-photon-counting regime can be quite difficult. We found life was dramatically improved in the lab once we had done a sufficiently thorough job as to allow data taking with some of the room lights turned on. This required three stages of protection. The PMT was unceremoniously crammed into a lens tube. Grooves had to be cut to accommodate the PMT's unfortunately square cross section, necessitating a great deal of black electrical tape to be wrapped around the tube to return its light-blocking integrity. This was screwed onto the adjustable aperture L connected to a long, black lens tube. The whole assembly was inside a small box constructed of black cardboard, electrical tape and various bits of black plastic sheeting to seal the edges. Admittance to this box was granted through the adjustable aperture I. In addition, the entire breadboard was surrounded by a second box made of cardboard,

electrical tape, black sheeting, and prayers. It is strongly recommended that future experiments give more thought than we gave to how the PMT will be boxed before constructing the board.

6.1.2 Alignment, Calibration, and the First Recoil Measurement

Once gross alignment was completed on the outgoing diffraction light (e.g., aligning it by eye to the centers of the viewports on both sides of the chamber) careful alignment proceeded by blowing atoms out of the optical dipole trap. With the retro-reflection shutter closed, a short, high-power pulse of diffraction light was fired at an evaporatively cooled cloud in the ODT. Pointing of the diffraction beam was iteratively adjusted to keep it pointed parallel to the desired diffraction axis while increasing the number of atoms kicked out of the trap.

As alignment progressed, the length of the pulse could be steadily reduced, from roughly 100 ms to less than 1 ms. At such lengths it becomes feasible to align to the atoms after some time of flight. This centers the beam closer to the region where diffraction will actually occur in experiments. The cloud can be fit after some additional time of flight and the x displacement of the cloud can be maximized to bring the center of the beam to the position of the atoms.

Once the outgoing beam is aligned, the retro beam was aligned by eye to counter-propagate with it. From there, Kapitza-Dirac diffraction (another name for diffraction by optical gratings in the Raman-Nath regime) was used to optimize the overlap of the brightest part of the beam with the atoms. Under the assumption that the outgoing beam had been properly aligned to have the atoms in the center of the beam, optimizing Kapitza-Dirac diffraction efficiency would also bring the center of the retro to the atoms. Since the two beams certainly coincided at the retro mirror, this would guarantee them to be counter-propagating.

This procedure is certainly not ideal. A far better way to align the retro is by maximizing the coupling efficiency of the retro light into the fiber from which the light

originally comes. Our initial attempts at this met with a great lack of success, as the coupling efficiency was inexplicably insensitive to adjustments of the retro-reflection mirror's orientation. After some time, a distant memory of a comment from Joel Hensley's thesis[43] suggested the answer. The 1:1 telescope formed by the lens and retro mirror served essentially as a point source at the mirror and a lens imaging the point source. On this view, the optics on the outgoing side of the chamber served to complete an imaging system for this point source with image point at the focal point of the fiber optic output coupler. Thus, the orientation of the retro could be changed, causing the overlap of the retro beam with the outgoing beam to change, without substantially changing the coupling efficiency into the fiber. This was finally understood at a point in the experiment where there was too little time left to us for data-taking to justify rebuilding the retro setup¹.

Once the diffraction beams were aligned, we calibrated the intensity at the atoms by looking at Kapitza-Dirac diffraction in the zero and first orders for increasing intensity. A sample plot is provided in figure 6.3. More careful calibrations of intensity at the atoms versus measured power in the beams will be discussed in section 8.3, where careful calibration is used to understand an important systematic effect.

The mirror pulses were initially planned to be in the Bragg regime. However, we found that pulses of sufficient length for the Bragg approximations to be valid tended to cut the center out of the cloud. Essentially, we were doing Bragg spectroscopy[96]. We thus moved to shorter pulses, picking Gaussians with half-width $\approx 35 \mu s$. This showed a good experimental compromise between reflecting essentially all parts of the moving arms while not diffracting new states out of the non-moving arm.

We also demonstrated superradiance and high-order diffraction. A set of pictures

¹This sounds bizarre, as building a telescope should be simple. However, the MOT breadboard was out of easy real estate. Rebuilding the retro setup correctly would have required changing paths for either the MOT or the Zeeman slower. Since we had an approaching deadline by which the machine had to be returned to the mixture experiment for good, tinkering with the laser cooling at this level seemed unlikely to be a net win.

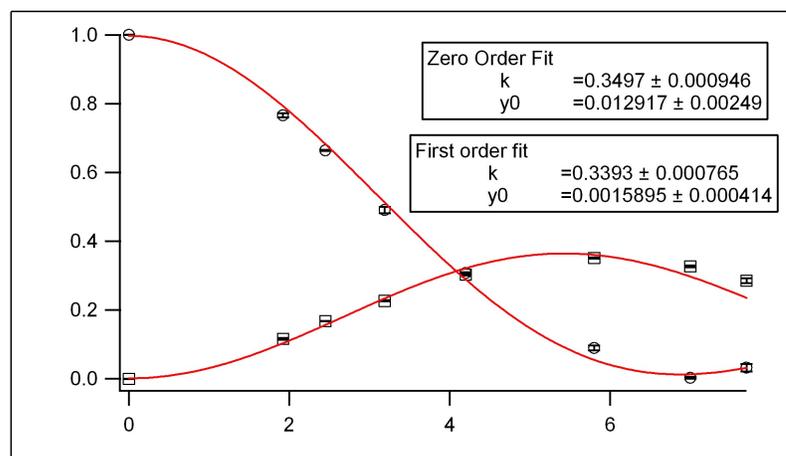


Figure 6.3: Calibrating with Kapitza-Dirac. The fraction of the population in the zero diffraction order and the average of the fraction in the \pm first order are plotted versus control voltage. The fits are to Bessel functions of zero and first order, respectively. The fitting function is of the form $y = J_\nu(kx) + y_0$, and the fit values are shown on the figure.

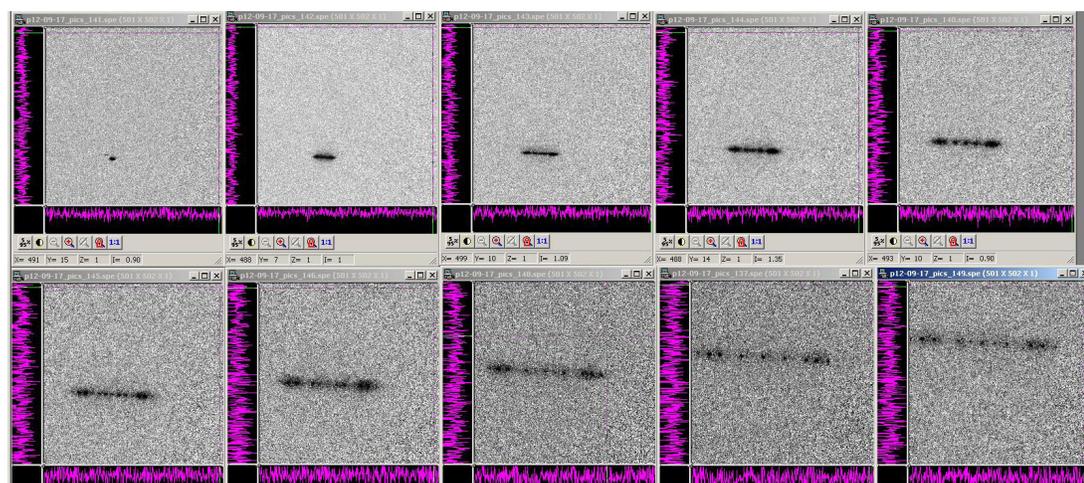


Figure 6.4: High-order diffraction. A set of time of flight absorption images for increasing time of flight shows the separation of many orders of diffraction after a short intense splitting pulse.

showing high-order diffraction are shown in figure 6.4. Superradiance was induced by firing a high power, long ($\approx 100\mu\text{s} \gg \Gamma^{-1} \approx 1\mu\text{s}$) diffraction pulse with the retro-reflection shutter closed. This allowed enough time for spontaneous scatter to seed the matter-wave and for matter-wave amplification to proceed. We generated diffraction out to seventh order using square pulses of width $2.3\mu\text{s}$. This pulse width is small enough to fall into the Raman-Nath regime ($\tau \ll 1/(4N^2\omega_{\text{rec}})$) for low diffraction order N , but is well outside it for $N = 7$. For taking data, shorter square pulses with width $1.5\mu\text{s}$ were used. The decision to use shorter pulses was made to simplify the diffraction phase analysis, as discussed in 8.4.

The last important result from before we achieved a working contrast interferometer was our first measurement of the ^{174}Yb recoil frequency. This measurement was useful in three regards. First, though it had only 1% precision, it provided a number from which to bootstrap to higher precision using the contrast interferometer. Second, it established our ability to do traditional number-counting interferometry in our setup. This would be important later on for demonstrating the vibration insensitivity of the contrast interferometer. Finally, it gave us a value for the coherence time, which was also relevant for the contrast interferometer.

The measurement was done using a momentum-space Ramsey interferometer. The technique involves a pair of identical Kapitza-Dirac pulses separated by a variable time T . The pulses are designed to have a relatively small amplitude to populate the first diffraction order and negligible amplitudes to higher orders. For $T = 0$, the atoms essentially experience a single, longer Kapitza-Dirac pulse. Therefore, a much larger population is found in the first order than would be found for a single pulse. For $T = \pi/(4\omega_{\text{rec}})$, the phase of the diffracted atoms has gone π rad out of phase, and so the second pulse will reverse the effect of the first, leaving no population in the first order. As seen in figure 6.5, we observed these oscillations for two full cycles. The plot gives a value $\omega_{\text{rec}} = 2\pi \times 3.72 \pm 0.06$ kHz. In addition to a value for the recoil frequency, this measurement also gave a value for the coherence time of the

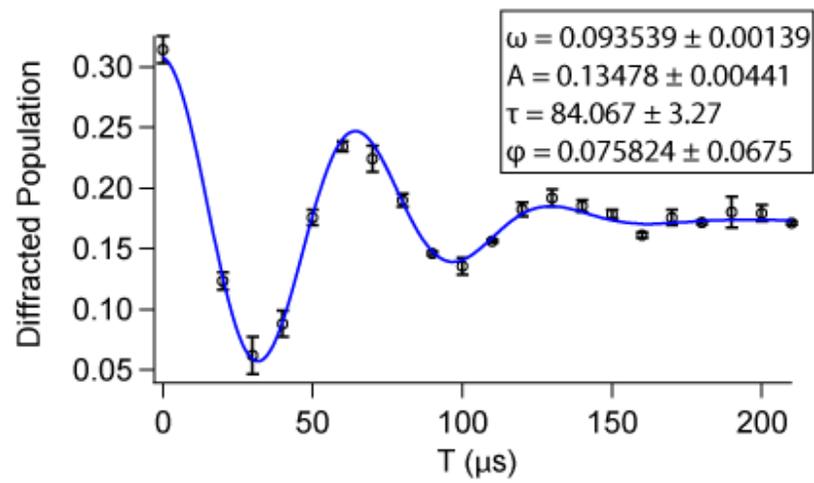


Figure 6.5: Momentum-space Ramsey interferometer. The fraction of the atoms diffracted after the second pulse in the Ramsey sequence is plotted versus the time between the pulses. The fit parameters from the functional form $A \exp(-t/\tau) \sin(\omega t + \phi)$ are given on the plot. The fine structure constant implied by this measurement is $\alpha^{-1} = 136.9 \pm 1.0$.

BEC source. Later on, we would develop more careful evaporation ramps achieving lower trap frequencies before trap turn off, thus increasing the coherence time.

6.2 Finding the Signal

Once the splitting and mirror pulses were calibrated we could cause the atoms' wave functions to follow the prescription of the contrast interferometer. All that remained was to observe the back-reflected readout light. In the course of looking for the signal, short intense readout pulses centered on $t = 2T$ were used, as we simply wanted to see a signal. Later, the intensity was reduced as much as possible to avoid decohering the density gratings. Generally, a pulse of length $150 - 200 \mu\text{s}$ centered on $t = 2T$ was used for taking data. Pulse length and power were periodically re-optimized to give the best signals with the least disruption of the clouds visible in absorption images after the shot.

For a CI using diffraction light as readout light as well, stray back-reflections are a formidable background. A substantial amount of effort was spent in iteratively attempting to block and attenuate these various tiny spots of light, which though tiny could completely overwhelm a PMT driven by enough voltage to see the CI signal. Viewports, lenses, and beam splitters were all guilty. The proximity of these specular reflections to the actual signal light made blocking them while not occluding the signal itself particularly bothersome. In addition to specular reflections from transparent optics, diffuse scatter of the readout light at particular places in the setup could be focused into the imaging path for the CI signal, making an essentially irreducible background.

A key technique for dealing with these is to turn up the readout light (at this time this was just the diffraction light) to maximum intensity with the back-reflection shutter closed. Then, the various unwanted back reflections may be viewed on a white card. The fact that 556 nm is near the peak of human eye response is quite helpful here. Further progress can be made by turning off the room lights and allowing your

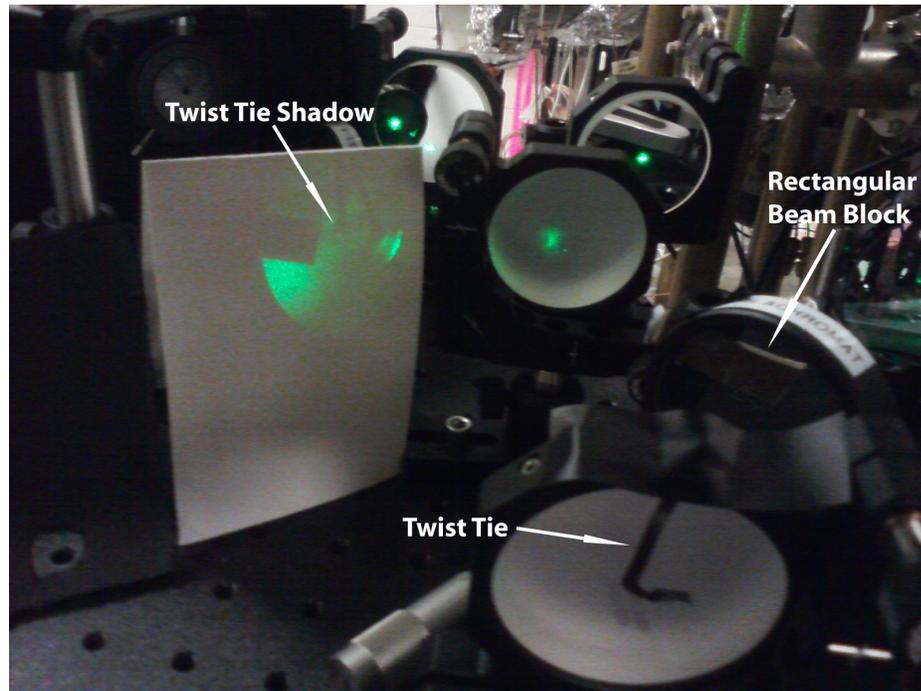


Figure 6.6: Stray light backgrounds. The image shows a small portion of the interferometry breadboard. For this image, the diffraction light was turned up to maximum intensity and the back-reflection shutter on the far side of the chamber closed. This allows back reflections to be observed with the naked eye. The white card on the right shows the light headed for the readout detector. Most of the diffuse glow will be blocked by irises downstream and so is not of concern. A specular back reflection that would be overwhelmingly bright on this card was blocked with the twist tie bent over the mirror indicated at the bottom. The shadow of the twist tie and the Fresnel spot from the small round block made at the end of the twist tie are indicated. The shadow of a more substantial beam block further upstream may be seen and the block itself is indicated on a lens.

eyes to adjust. A poignant example of the woe back reflections caused may be seen in figure 6.6. In the foreground you can see a twist-tie taped to a mirror. This device was designed to provide the least occlusion possible of the CI signal will providing a block for a reflection that fell extremely close to the signal axis. The end of the twist-tie in front of the mirror was shaped into a small ball. The white card on the left shows the rectangular shadow of a larger beam block (visible on the lens in front of the twist-tied mirror) and the shadow of the twist-tie. The green spot right in the center of the twist-tie shadow is the Fresnel spot formed by diffraction of the unwanted reflection around the improvised beam block.

Since the signal path was designed to image the atoms, we first performed absorption imaging of BECs along this axis by placing a camera² with its CCD chip at the position we intended for the PMT to take in the future. The atoms were illuminated by the diffraction light, with the retro-reflected light acting to back light the atoms for absorption imaging. While some prior concerns were expressed about the fact that the light passed through the atoms twice, no adverse effects were observed in the images. Having the absorption imaging setup allowed us to optimize the lens positions for signal collection. More importantly, it allowed for beam blocking and aperturing around the telescope focal points to be guided by knowledge of where the image of the atoms was within the beam. Blocking too close to the absorption shadow would harm the image quality, warning that this would also block some of the signal light coming from the atoms.

Once the unwanted lights were blocked to such a level that we could hope to see the CI signal above them, we looked for the CI signal on the camera. This required running the camera with unnervingly high electron multiplier gain, though the camera came out of the ordeal unscathed. Having only one camera, we frequently changed from absorption imaging to signal hunting and back again, to be sure we still had

²Specifically, an Andor iXon3 885.

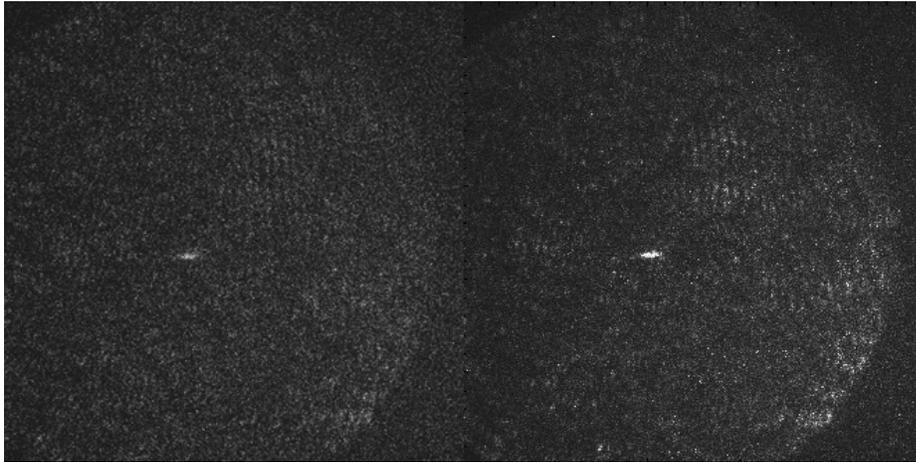


Figure 6.7: Images of CI signals. On the left is an example image of the CI signal. On the right is an average of five such images to bring out the signal shape more clearly.

atoms despite having no signal on the camera. Figure 6.7 shows an example of the CI signal as imaged by the camera, along with an average of several signal images. It is interesting to note that the shape of the signal more closely matches the shape of the original BEC before expansion than the shape of the BEC at the time of the readout, most likely reflecting that the faster expansion in the vertical direction does not substantially increase the coherence length of the condensate despite increasing the total size by more than a factor of ten.

Once the signal is located on the camera, irises at the telescopes' foci can be closed down until they begin to reduce the signal strength. After these were adjusted to satisfaction, the camera was replaced with a PMT and time-resolved CI signals were captured.

6.3 Analysis and Small T Results

The CI signal combines modulation at $8\omega_{\text{rec}}$, which carries the desired phase information, and an envelope related to the condensate coherence and decoherence caused by

the readout light. Any fit to this data will be somewhat tricky due to the necessity of doing nonlinear curve fitting. With nonlinear fitting, it is easy to become trapped in a local minimum which is not the global minimum or have the optimization routine run some parameter out to absurdly high or low values. As a particularly perverse example, fitting a sine wave with a Gaussian envelope to data with only two peaks can lead to a fit with a low frequency sine wave and an enormous Gaussian (amplitude ≈ 100 times the amplitude of the data) that when multiplied together yield a two-peaked structure remarkably similar to the data.

Our early data had short coherence times, likely due to excessively strong readout light. As such, much effort was expended trying to fit the signal envelope to obtain better phase information from the underlying sine wave. However, the best results were found by selecting the highest quality portion of the data and fitting that to a sine wave. While this is an effective technique, automating the decision process for the highest quality portion of the data would be quite difficult. Instead, the analysis software uses a proxy that has been found to agree well with human judgment as to the best quality portion of the data.

The procedure utilized in the analysis software involves picking a number of different windows from the data. It starts by finding the point of maximum signal. The actual maximum in a file is often a noise spike on top of actual signal. However, it is extremely unlikely for noise in a low quality portion of the signal to spike high enough to win out over the strength of the signal in its best quality regions. The routine then defines windows extending in either direction from the maximum point by different numbers of oscillation cycles. For example, $\{1, 3/2, 2, 5/2, 3, 7/2, 4\}$, for a total of 49 different windows. The data in each window are fit to $A [\sin(8\omega_{\text{rec}}(t - 2T) + \phi) + 1]$, with A and ϕ as fitting parameters and ω_{rec} as a previously determined value of the recoil frequency. At the end, the fit values from the window with the smallest uncertainty in ϕ are used.

Notice that we fit the signal to a simple sine wave with twice the frequency ex-

pected from (5.3). The density modulation described by this equation will determine the reflected electric field amplitude, however, the PMT detects the intensity, which is the square of the amplitude. Squaring the time dependence of (5.3) we arrive at an expected signal proportional to

$$\cos^2 \left(4\omega_{\text{rec}}t + \frac{\theta_1 + \theta_{-1}}{2} - \theta_0 \right) = 1 + \cos(8\omega_{\text{rec}}t + \theta_1 + \theta_{-1} - 2\theta_0) . \quad (6.1)$$

The decision to fit to a sine wave rather than a squared sine wave came from a desire to minimize the “nonlinearity” of the fitting function, though in all likelihood the two would work equivalently. However, we care about the phase of the atoms themselves, so the phase form on the left side of (6.1) is more directly relevant. This is also the phase that models for various systematic effects calculate. Thus, after extracting ϕ , we divide by 2 and bear in mind that the 2π ambiguity in the fit corresponds to a π ambiguity in the final value³.

The first recoil result from the CI used values for T of 0.99, 1.00, 1.01, and 1.02 ms. By taking such small steps in T , we could be sure that we were not missing multiple changes by π in the phase reported for each value of T . The π ambiguity in phase requires several small adjustments during the data analysis. For a given T , all measurements are shifted by a multiple of π to bring them within $\pi/2$ of the initial mean value of ϕ . Next, corrections for systematic effects are applied. Particularly important are the diffraction phase shifts, described in chapter 8.

To obtain a value for the recoil frequency from the data sets at T values listed above, we must add multiples of π to all of the data sets except the smallest T ($T = 0.99$ ms) set. To determine the correct number of π 's, we calculate

$$\frac{\langle \phi(0.99 \text{ ms}) \rangle + 2(T - 0.99 \text{ ms}) \cdot 4\omega_{\text{rec}} - \langle \phi(T) \rangle}{\pi} ,$$

³To divide by 2 or not amounts essentially to an aesthetic choice. It has the satisfaction of showing a smaller standard deviation in a set of identical measurements, but this of course does not affect the relative uncertainty of the measurement. The only important thing is that analysis remains consistent with the decision made.

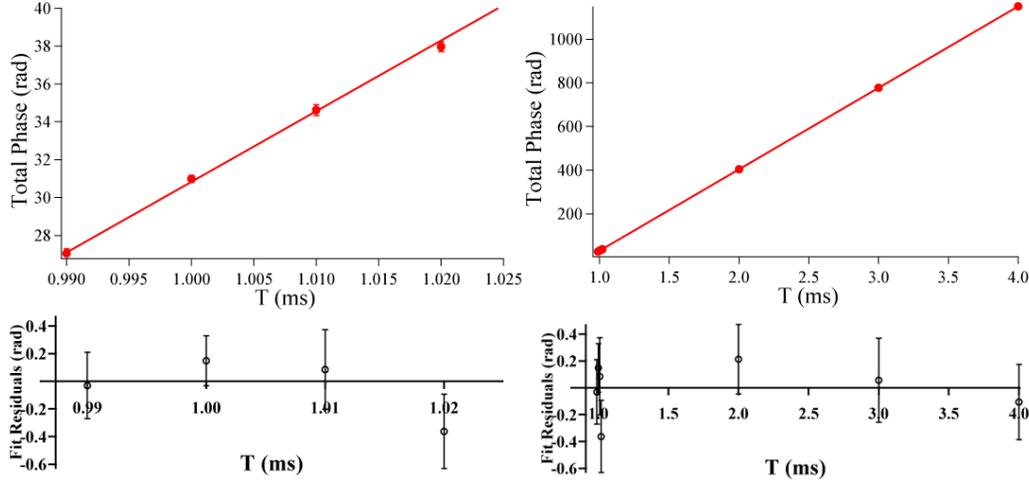


Figure 6.8: Small T Results. The left side of the figure shows data points and a linear fit for $T \in \{0.99, 1.00, 1.01, 1.02\}$ ms. The right side shows the same data with the addition of points for $T \in \{2.00, 3.00, 4.00\}$ ms. Fit residuals for the linear fits are plotted beneath the main data plots.

using the value for ω_{rec} determined in the momentum-space Ramsey experiment, and then round the result to the nearest integer. This integer result multiplied by π is then added to each value in the $\phi(T)$ set. Once all of the data sets are adjusted, we fit a straight line to $\phi(T)$ vs. T . The slope of the line is $8\omega_{\text{rec}}$. From the fit uncertainty for the slope from the least squares fit, we determine the uncertainty in our recoil frequency. The linear fit and the fit residuals for this data set are shown on the left side of figure 6.8. The result is $\omega_{\text{rec}} = 2\pi \times 3.62 \pm 0.09$ kHz.

We may then use this value of the recoil frequency to extend to larger values of T . Adding data sets at $T = 2, 3,$ and 4 ms gives the data and residuals shown on the right side of figure 6.8. The result from these data is $\omega_{\text{rec}} = 2\pi \times 3.7122 \pm 0.0007$ kHz. These data were taken and analyzed before the diffraction phase systematic was fully understood, and so they did not benefit from the large reductions in uncertainty that this gave. This result corresponds to a relative precision of 190 ppm. It gives a value

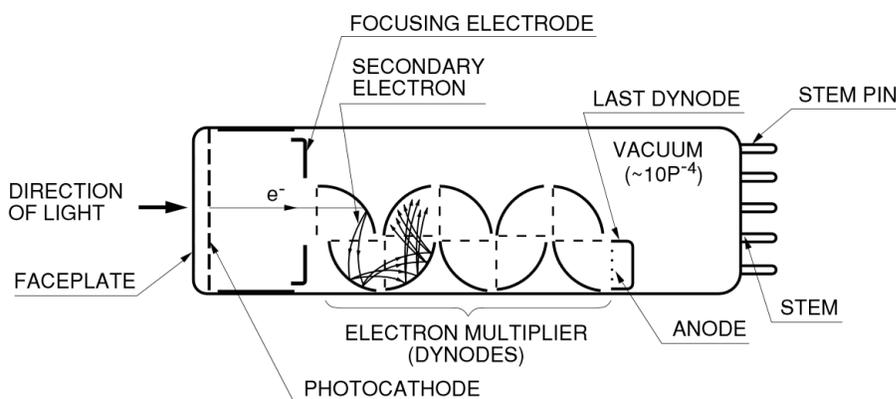


Figure 6.9: This schematic of a photomultiplier tube shows a single photon striking the photocathode. The emitted electron is directed to the dynode chain where large potential differences between successive dynodes cause the number of electrons to be multiplied on collision with each plate. The single photon is thus converted into a measurable current spike. Figure from *Photomultiplier Tubes: Basics and Applications*, available from the Hamamatsu website http://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE-Chapter2.pdf

of $\alpha^{-1} = 137.053 \pm 0.013$, which differs from the CODATA value by 1.3 standard deviations. With hind sight, it is likely that there was a small systematic shift due to diffraction phase, the magnitude of which was not appreciated at the time.

6.4 Improving the Readout

The initial readout system was plagued by several problems leading to poor signal-to-noise ratios. These problems were largely overcome by upgrading the PMT and amplifier and installing a separate fiber and optical path for the readout light.

6.4.1 *Better Equipment*

To better appreciate the nature of the improvements to the signal detection equipment, we'll briefly review the workings of a photomultiplier tube. The actual detector in a PMT is a photocathode that emits an electron when hit by a photon of sufficiently high frequency. This process is an example of the classic photoelectric effect, of which Einstein's explanation served to help begin the understanding of quantum mechanics. The photocathode material has some work function, which determines the minimum frequency a photon must have to dislodge an electron. The rest of the PMT consists of a series of plates (called dynodes) at progressively higher potentials. The initially emitted electron is accelerated toward the first of these plates. When it strikes this plate, having much greater energy than when initially emitted, it dislodges many electrons. These freed electrons are then accelerated toward the second plate, where each electron dislodges many new electrons on impact. This process continues through to a final collector, the anode. Thus, a single photon at the photocathode is converted into a measurable current pulse at the anode.

Since the PMT gives a certain current response for a single photon, the output current, not voltage, is linear in light intensity. As such, we want an amplifier that gives a linear conversion from current to voltage, voltage signals being more robust. Such an amplifier is referred to as a transimpedance amplifier. In a very crude sense, a large resistor is a transimpedance amplifier, since it will convert a small current signal into a large voltage signal. However, the large resistance can distort the current signal, and the small current would not support a robust voltage output. Ideally, a transimpedance amplifier should combine a very small input impedance with a large output impedance. A number of considerations go into the design of such amplifiers to balance bandwidth and amplification.

The initial data were taken with a Thorlabs PMM02 photomultiplier module. These general purpose PMTs come with a built-in transimpedance amplifier. While

this is a fine general purpose light detector, far more sensitive than standard photodiodes, it fails to live up to the full potential of a PMT. Our two devices never achieved the amplifications deemed possible by the specs. Perhaps equally important, the built-in transimpedance amplifier response had a sharp roll off above 30 kHz. Since our signal oscillates at $8\omega_{\text{rec}} = 2\pi \cdot 30 \text{ kHz}$, this almost certainly introduced distortion of signals.

To improve the system we changed to a Hamamatsu H10721-210 PMT. This is just the raw PMT and so requires a separate amplifier (as well as a bit of time to build voltage control electronics for the internal high-voltage). This PMT easily worked to specification, despite occasional mistreatment⁴. At maximum safe control voltage, and even just a touch below max, the PMT is sensitive enough to pick up individual photons. We chose this particular PMT with an eye to the future. Its quantum efficiency at 556 nm is only about 11%. However, it jumps to 30% for 399 nm light. As we have good reason to believe that readout of the CI signal using 399 nm in the future will substantially improve the experiment, we opted for a PMT that would increase those gains further.

We paired this with the Stanford Research Systems SR570 low noise current preamp. This preamp has a variety of settings allowing for bandwidth up to 1 MHz. We found optimal results with a low-pass setting of 300 kHz. The amplifier has outstanding levels of amplification with imperceptible electronic noise. Together, these made an impressive improvement in our signal quality.

⁴I hesitate to say this, lest it embolden future grad students. However, on at least two separate occasions, the experiment was run with the PMT protecting shutter open during the diffraction pulses. While this drops unconscionable amounts of light onto the detector, the ill effects in terms of increased noise were only temporary. Undoubtedly this would lead to sustained poor performance if such abuse happened very often, and so should be avoided at all costs.

6.4.2 *Decoupling Diffraction and Readout*

As discussed in 6.2, back reflections from optics form a background to the CI signal which is difficult to reduce. Installing separate fibers and a new optical path for delivering the readout light to the atoms solved this in several ways. Even with the second fiber, we continued to use a single AOM for both light sources. The readout light was split from the diffraction light and sent to its own fiber. This fiber had a shutter just in front of it, as did the diffraction fiber.

The most important improvement was the major reduction of readout light required. The diffraction beams had a waist of 3 to 4 mm over the life of the experiment. However, the BEC at the time of the readout has dimensions of, at the most, 100 μm . This suggests that $> 99\%$ of the light is of no value for making the CI signal, since it doesn't hit the atoms at all. This light may, however, contribute to both the diffuse background and to specular reflections. Switching to separate delivery of the readout light allows it to be mode-matched to the BEC. The total power available for making unwanted backgrounds is immediately reduced by a factor of 100. Future experiments will need to move to roughly twice the waist for diffraction beams, while hopefully using BECs with smaller momentum spread and thus smaller physical size at time of readout. Thus, the wins from using separated readout may be even more substantial.

The mode-matching allowed by separated readout would also be essential to any future studies of back action or decoherence caused by the readout light. Quantifying the amount of light actually reaching the atoms would be necessary and becomes far easier if all of the delivered light reaches the atoms. By using absorption imaging along the diffraction/readout axis, the readout light could be adjusted to overlap with the BEC or even to address specific regions of the BEC. This last possibility could be used to understand how coherence contributes to the signal in a manner complimentary to taking direct images of the CI signal itself, as in figure 6.7.

Separate readout light also allowed use of a polarizing beam splitter to connect the readout path and the diffraction light. With the same light for diffraction and readout, the CI signal has the same polarization as everything else. Thus, a partially reflecting beam splitter had to be used. The reflection ratio had to be chosen to trade off fraction of transmitted power for diffraction beams against fraction of signal reflected to the detector. With the PBS, each may be nearly 100%. The readout light and the reflected signal still have to be combined on a nonpolarizing beam splitter, but since there is vastly more readout light available than is actually needed, we could throw away 90% of the readout light to allow collecting 90% of the signal.

6.5 Large T Results

Over the course of the year the interferometer was taking data, a number of incremental improvements were made. An initially poor choice for the retro optics, which included a mirror dangling in space from the edge of a third breadboard, was rectified. Various tweaks were made to improve the working of the mirror pulse, mostly related to getting the width just right. Changes were made to the evaporation ramp, including improved intensity feedback control for the ODT laser, allowing for stable running with lower final trap frequencies, and a 0.5 – 1 s hold at low trap depth to damp out excitations of the cloud just before condensation. The end result was a pair of measurements with 7 ppm precision using different values for T .

6.5.1 Vibrational Immunity

A key advantage of contrast interferometry over traditional atom interferometers is the absence of an external grating for reading out the phase, discussed in 5.1.2. For an external grating readout interferometer, the position of the external grating relative to the initial splitting grating must be carefully stabilized. In interferometers using physical gratings, this may be done through optical interferometry on the physical gratings with some feedback control[50]. For optical gratings, this entails stabilizing

the positions of the optics responsible for the standing wave, since sub-wavelength scale changes in optical path length can markedly change the positions of nodes and antinodes in the optical potential[49]. If the two beams interfering to make the standing wave come from different sources, those sources must also be phase locked[43].

The contrast interferometer avoids these issues by beating two atomic density gratings against one another for its readout. So, in principle, the CI should be immune to vibrations. This was conceptually demonstrated in the original CI paper[36] by changing the phase between the two interfering laser beams in a deterministic manner. However, a CI actually returning a signal over interferometer times long enough for vibrations to be a concern was lacking.

To demonstrate the vibrational immunity we compared the coherence time for a CI and for a momentum-space Ramsey interferometer. The interferometry setup had no active vibration control. Being built on an optics table and adding brass plates to the breadboards that were on legs above the optics table were the only passive measures to combat vibration. The Ramsey experiment followed the procedure for the CI exactly, with the one change that the final readout was a standing-wave pulse identical to the splitting pulse. The experiment differed from that described in 6.1.2 in that a mirror pulse was inserted into the sequence so that our coherence time was not limited by the spatial extent of the BEC. So, as for the CI, the Ramsey readout pulse occurred at the time of perfect overlap for the interferometer arms.

Figure 6.10 shows the loss of fringe visibility in the Ramsey interferometer with increasing T , as compared to the constant visibility for the CI signal over a far larger span. The visibility of the fringe comes from fitting a sine wave to the data with form $A \sin(\omega t + \phi) + y_0$. The visibility is then simply $2A/(A + y_0)$. This is to be compared to the interferometer contrast, defined as $y_{\max} - y_{\min}/(y_{\max} + y_{\min})$. The essential difference is that the contrast measures the level of coherence in the atomic interference itself, while the visibility measures the level of coherence relative to the

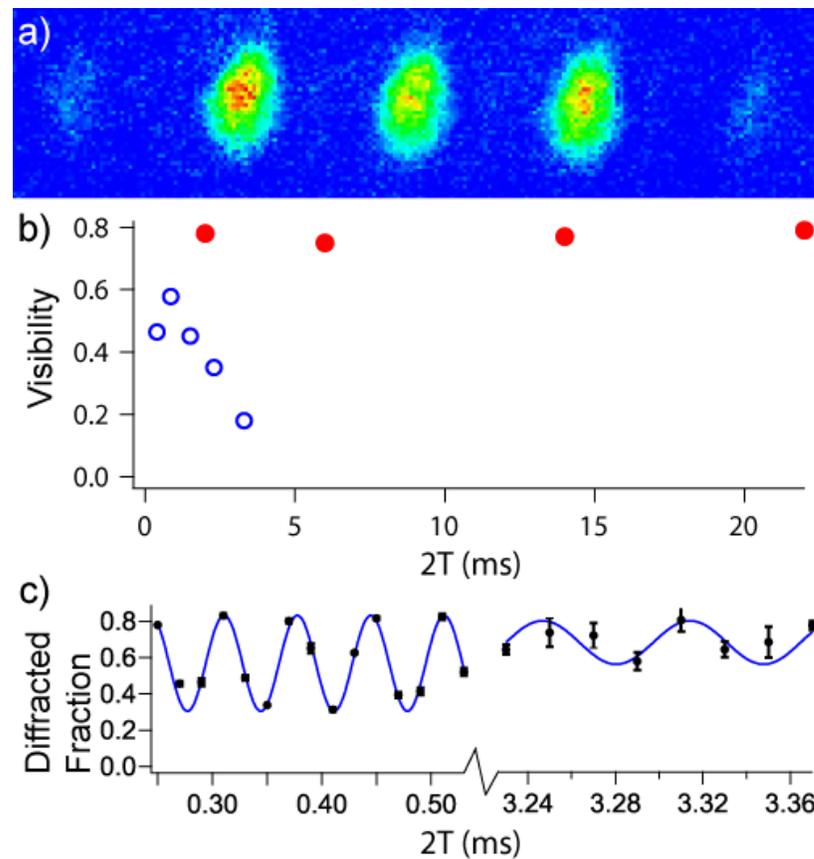


Figure 6.10: Vibration immunity. In a) we see an absorption image of the separation of the arms of the CI at the time of the mirror pulse for a $T = 11$ ms interferometer. In b) the fringe visibility is plotted as a function of T for the CI (red disks) and the Ramsey interferometer (blue circles). In c) two sample signals from the Ramsey interferometer are shown to demonstrate the loss of visibility.

readout process. For instance, a single point in the right-hand plot of c) could include shots ranging from 0.2 to 0.8, which would give a contrast of 0.6. However, the fact that these all occurred for the same readout settings means that while the atoms maintain their coherence (otherwise the value would be essentially 0.5 for all shots) the interferometer gratings are not maintaining their phase relationship from shot to shot. A good summary would be contrast indicates coherence while visibility indicates a useful measurement with that coherence.

We see the visibility in the Ramsey experiment begin a quick decline for $T > 1$ ms. This suggests that vibrations of frequency $\lesssim 1$ kHz afflict the setup. The visibility for the CI is unaffected by this noise, as seen from the fact that it is unchanged out to $T = 11$ ms. We also took CI data where we intentionally coupled vibrations into the table by running the Webtrol booster pump. This high pressure pump for cooling the MOT and Helmholtz coils coupled sufficient vibrational noise into the optics table as to make the Li spectroscopy, located on the same table, unlockable. There was no discernible difference in the data quality from the CI with or without the booster pump running.

6.5.2 *Below 10 ppm and Scaling to Better Precision*

Two different data sets were taken to the ≤ 10 ppm level. One (set A) included 500 data points at each of $T = 1$ ms and $T = 7$ ms. The other (set B) used 150 data points at each of $T = 3$ ms and $T = 11$ ms.

Each data set was fitted using a holistic fitting routine. This is similar to that described in section 6.3. However, rather than using the window which gives the best value of the phase uncertainty for a single shot, the entire data set is fit with different windows and then the window that gives the lowest standard deviation for the data set is used. This gives a better metric for goodness of a window, since many windows give small phase uncertainty for any given trace. It also addresses a potential concern that if there are some nonlinear frequency changes across the time of the signal

(e.g., from interaction induced phase curvature) that sampling from different parts of the window for different runs could introduce systematic shifts and/or increased uncertainty in the sample. The diffraction phase systematic described in chapter 8 is strong, and highly susceptible to laser intensity noise. An example in that chapter shows a factor of three reduction in sample standard deviation upon application of the diffraction phase correction. Therefore, we first calculated the shift for each shot and then applied those corrections before determining which window gave the best standard deviation of the data set.

One other technique was used to give a modest boost to the precision of the measurement. The mean and standard deviation are statistical estimators that are extremely susceptible to skewing by outliers. Estimators with less susceptibility to outliers are referred to as robust statistics. Two of the simplest are trimmed mean and trimmed standard deviation. Essentially, these consist of the mean and standard deviation calculated after removing some percent of the lowest and highest values from the data set. In this way, extreme outliers are removed and the central trends of the data set are preserved. Data sets without extreme outliers are essentially unchanged by this procedure, as long as the percent removed is not too great⁵.

Data set A was analyzed with and without a 10% outlier trimming (i.e., the lowest 5% and highest 5% of data points removed). Without trimming, the result is $\omega_{\text{rec}} = 2\pi \times 3712.110 \pm .035$ Hz. With trimming, the result is $\omega_{\text{rec}} = 2\pi \times 3712.134 \pm 0.028$ Hz. As the results agree with one another to within error bars, it seems reasonable to conclude that the trimming procedure does not introduce a systematic shift. This should be rechecked in future high precision data sets, as some physical systematic effect could cause a significantly skewed distribution, though this seems unlikely. The relative uncertainty in the trimmed data is 7 ppm.

Data set B was analyzed similarly. Without trimming, the result is $\omega_{\text{rec}} = 2\pi \times$

⁵The extreme limit of this removal process would give the median as the 49.999% trimmed mean.

$3713.190 \pm .042$ Hz. With trimming, the result is $\omega_{\text{rec}} = 2\pi \times 3713.156 \pm 0.038$ Hz, giving a relative uncertainty of 10 ppm. Again we see agreement between the two values to within error bars.

Unfortunately, discussion of accuracy is not possible for these data sets. As discussed in chapter 9, trap uncertainties limit the ability to correct for interaction shifts at a level commensurate with the precision demonstrated in these data sets.

6.5.3 *One Detuning, Two Detuning, Red Detuning, Blue Detuning*

A technical advantage of Yb over alkali atoms is the ease of working near the narrow intercombination line. To achieve detunings of 500 linewidths or more from resonance for an alkali atom requires shifting frequencies by multiple gigahertz. With the 556 nm transition for Yb, this reduces to about 100 MHz. Therefore, with a single AOM, the diffraction light may be set to 500Γ of red detuning or blue detuning or set to resonance. This flexibility may be enormously helpful in understanding and correcting for the index of refraction shift in future experiments.

To test for any effect from detuning, we took data with several different combinations of detuning for diffraction light and for readout light. Diffraction light was generally $\sim 700\Gamma$ detuned, either to the red or the blue of the transition, while readout light was generally $\sim 50\Gamma$ detuned. We tried all four possible combinations of red detuning and blue detuning for readout and diffraction lights. We found essentially no difference in data quality or phase accumulation between these. At the precision we worked to, the index of refraction shift to the recoil frequency, described in 7.3.4, would have been too small to notice.

One motivation for the change from red to blue detuning in the diffraction pulses was the avoidance of molecular photo-association resonances. The positions of these resonances are well-known[97], and so avoiding them directly is fairly easy. However, there was some concern that they might indirectly influence the interferometer by contributing to the coupling of the diffraction light to the atoms. Since the strength of

the photoassociation line is density dependent, this would induce a density dependence into the diffraction phase shifts described in chapter 8. An alternative description would be in terms of an optical Feshbach resonance[29]. While we found no measurable difference between far red and far blue detuned diffraction beams, it is recommended that future experiments stick to blue detunings. There is seemingly no down side, while in principle these photoassociation/optical Feshbach effects could be visible at higher precision with red detuning.

6.6 To BEC, or Not to BEC

The discussion thus far has assumed that the atomic source for the CI is a Bose-Einstein condensate. A contrast interferometer using a cloud at a temperature above the condensation temperature—a “thermal” cloud—is an interesting prospect for two broad reasons. First, a sizable body of work must be applied to the BEC to correct for interaction shifts due to the condensate’s high density, as described in chapter 9. A thermal cloud has much smaller density and as such far less severe interaction shifts.

Second, the ability to do CI with a thermal cloud would allow study of the effect of inter-particle coherence in a CI, which could be enlarged to a general study of many-body physics using CI. The envelope of the CI signal is essentially the two-point correlation function of the cloud along the axis of the interferometer. Thus, phase transitions with signatures in the two-point function could be studied. A particular example would be to measure critical exponents for the condensation transition[28]. Additionally, the effect of inter-particle coherence on rates of decoherence due to controlled external perturbations could be studied. The decoherence rate for uncorrelated atoms could be determined using a thermal cloud CI, and then decoherence could be studied at various temperatures crossing the transition down to pure BECs.

To demonstrate this ability, we performed CI with essentially the same parameters as for taking high precision data. The evaporation ramp was modified by stopping

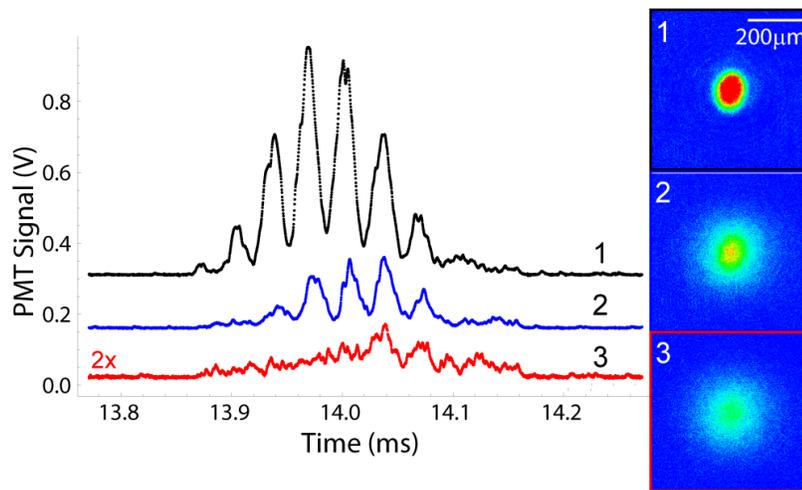


Figure 6.11: CI across the BEC transition. The figure shows CI signals from three clouds. Each trace is an average over twenty signals with the same evaporation ramp. Trace 1 is from a pure BEC. Trace 2 is from a cloud just below the transition temperature with $T/T_C \approx 0.9$. Trace 3 is from a thermal cloud just above the transition temperature but below the recoil temperature (i.e., the diffraction orders are well-resolved in absorption images with long time of flight). The time of flight absorption images on the right show the momentum space distributions of the clouds. A small condensate is visible in the center of image 2.

and holding the trap depth just before the BEC begins to form. For various hold times at this trap depth the cloud will continue to cool and so cross the BEC phase transition. The goal was to study a variety of T/T_C values at fixed trap frequencies. Hold times were 0, 100, 200, 300, 400, and 600 ms. Additionally, data was taken with no hold at a slightly higher trap depth and after a long hold at a slightly lower trap depth. A second modification was to use a longer, lower power readout pulse. The pulse was lengthened to 300 μs to allow for the possibility of seeing many peaks in a situation with long coherence time. An attempt was made to move to 400 μs , but this seemed to artificially decohere the density gratings.

Figure 6.11 shows some results of this effort. The three traces show the averaged signals from twenty runs with a pure BEC, a cloud with $T/T_C \approx 1$, and a thermal cloud just above T_C . The emergence of a small condensate, visible in absorption image two, seems to be accompanied by a qualitative change in the signal, namely the appearance of three clearly defined central peaks. The results of this study, necessarily quite preliminary due to lack of time on the machine, are quite promising for the use of CI to study many-body physics in the future.

One potentially interesting twist would be to use a two arm configuration. That is, use a Bragg pulse to populate only the 0 and $2\hbar k_{\text{rec}}$ states. Then, simply look at the amplitude of the back reflection as a function of time. This signal would have no recoil frequency modulation, and so would provide cleaner data for the coherence length. Such an asymmetric splitting pulse would require two separate sources for the two diffraction beams. While this introduces some small extra experimental effort, it removes the time-scale of closing the back-reflection shutter as a concern. Thus, the reflection signal could be observed immediately after the Bragg pulse (i.e., during separation rather than after a mirror pulse and recombination). Studies of grating coherence at these times might also give insight into whether some coherence is lost due to collisions between the interferometer arms during the splitting process.

6.7 Recommendations for Future Experiments

Describing how we did things is important, since what we did worked. However, we learned a number of lessons along the way and frequently wished to have time to start anew. Some subset of these insights are collected here, as advice for future experiments.

- Two-lens telescope for the diffraction light retro. As explained in 6.1.2, using a single lens with the retro mirror at its focus makes it impossible to use optical fiber return coupling for aligning the retro diffraction beam. This is such a

powerful technique that it'd be foolish to give it up.

- Image with outgoing diffraction light. For high-precision interferometry, it will be helpful to know where the atoms are in the diffraction beams. This will also be necessary if experiments to test the fully quantum mechanical treatment of wavefront curvature are to be carried out. A telescope must be constructed on the retro side of the interferometry axis for the shutter to work. As mentioned above, having a complete two-lens telescope for the diffraction light retro is key. If such a telescope is to be constructed anyway, it is not much more work to add a magnetic stage with a mirror after the second lens such that a mirror may be inserted before the retro mirror to send the light to a camera. This should also simplify the initial alignment process as the atoms can simply be imaged with the light and then the light position adjusted to center it on the atoms at the proper time of flight. Admittedly, this does require the telescope to be adjusted such that the first lens has the atoms at its focus, but in the long run the extra information on where the atoms are in the beam should be worth the trouble. Also of interest, this would allow imaging of the readout light to place it on the atoms and adjust its focus to mode-match the BEC at the time of readout. This could be important to experiments designed to understand the dynamics—including quantum back-action and decoherence—of the readout process.
- Avoid the Hall of Mirrors. While it is important to keep angles of incidence small at all mirrors, the interferometry paths should be planned so as to accomplish this goal without increasing the number of optics beyond the bare minimum. The three steering mirrors, dubbed the “Hall of Mirrors” lead to ambiguity in adjusting alignments, and occasional confusion as to which mirror had most recently been adjusted. Additionally, keeping optics clean is essential for avoiding odd diffraction effects from dust and other contaminants. Keeping the optics to the bare minimum reduces the number of places for contaminants

to affect the beam.

- **Boxes Within Boxes.** Plan the layout of the interferometry optics such that the optics from the aperture on the first signal telescope to the PMT itself can all be boxed up, with a second box surrounding the entire setup, including this inner box. Also, robust boxes which are easy to open and then close light-tight again will make things go much more quickly.
- **Interlock Protection for PMT.** The PMT must be run at maximum gain settings to see the CI signal well. Thus, any leakage of room light or inadvertent shining of lasers into the PMT can be very hazardous to the equipment. The home-built voltage gain controls could be fairly straight-forwardly retrofitted with an interlock system that would drop the control voltage to zero if a signal output above a certain threshold were detected. This is not perfect as it still requires the PMT being connected to a working transimpedance amplifier for it to be effectively protected, but it might go a long way to safe guarding this critical piece of equipment.

Part III
SYSTEMATICS

Chapter 7

SYSTEMATIC SHIFTS IN A CONTRAST INTERFEROMETER

This chapter discusses a wide variety of physical effects that could systematically shift the output of the contrast interferometer. The size of such effects and how they scale with T will be derived. Also, we will consider the effect on an interferometer where the moving arms are accelerated to $p = \pm 2N\hbar k$ for free propagation. Such acceleration may be accomplished in several different ways[74, 19]. Since such acceleration will be key to future experiments with sub part per billion precision, we will also consider scaling with N . We intentionally avoid describing the techniques for acceleration, simply noting that they involve two- or multi-photon transitions akin to those made in the mirror pulse. An interferometer with $N > 1$ requires four stages of acceleration: speed up after the splitting pulse, slow down before the mirror pulse, speed up again after the mirror pulse, and final slow down before read out.

Two larger/more complex systematic shifts will be described separately in chapters 8 and 9.

7.1 Reference Frame Effects

Reference frame effects refers to any systematic shifts due to relative movement of the diffraction beams and the atoms. While general relativity shows that gravity should properly be classified as a reference frame effect, we opt to use the simpler Newtonian picture of an external force field, which is more than sufficiently accurate for our needs.

The standing optical wave defines a reference frame. In a traditional interfer-

ometer, the fact that an external grating reads out a phase relative to itself makes unintended relative motion between grating and atoms a serious concern. Such issues were discussed in section 6.5.1 where we demonstrated the vibrational immunity of the contrast interferometer. Similarly to the vibration results, the CI is generically less sensitive to reference frame effects than traditional interferometers because the external reference frame is not directly involved in the final phase measurement.

Since the CI relies on momentum kicks delivered by the optical standing wave, a change in its orientation between momentum kicks will change the magnitude of the total momentum kick, thus skewing the results. To take an extreme example, if the apparatus rotated $\pi/2$ between the splitting pulse and the mirror pulse, the momentum kick from the mirror pulse would be orthogonal to the original momentum kick. In fact, this interferometer would not close. If a signal could be obtained (e.g., for extremely short T), the kinetic energy would be too large by a factor of two for the second half of the interferometer. We will assume that the effects are only small perturbations of the momentum kicks. However, for an $N = 100$ interferometer, tiny imperfections in the momentum kicks could add up to unacceptable inaccuracies.

If the atoms are moving along the standing wave axis, they will see photons from one diffraction beam as having less momentum but the other having correspondingly more. Thus, initial velocity or any acceleration along the standing wave axis will not affect the net momentum kick obtained in a two-photon transition¹. Such motion may affect the efficiency of diffraction pulses by making the two-photon transition slightly off-resonant but this does not affect the momentum kick. Similarly, motion perpendicular to the axis can change the intensity at the atoms, thereby changing the diffraction pulse efficiency but not affecting the momentum kick².

¹Equivalently, we may think of the standing wave as having a physical spacing between intensity maxima of $\lambda/2$. This physical spacing can be length contracted by motion along the axis, but the relative size of these corrections for an $N = 1000$ interferometer (atomic speeds of ≈ 8 m/s) would be 10^{-15} .

²It has been argued that moving within the beam affects the wavefront curvature effects described

Rotations, as suggested in our extreme example above, may alter the momentum kick. Specifically, a rotation that changes the direction of the standing wave axis relative to the inertial reference frame of the freely falling atoms by an angle $\delta\theta$ will in turn change the direction of the momentum kick by the same amount.

As a simple case, consider an $N = 1$ CI. Defining the initial standing wave axis to be the x direction, the two momentum kicks for the $+1$ arm would be $(2\hbar k, 0, 0)$ and $(-4\hbar k[1 - \delta\theta^2/2], -4\hbar k\delta\theta, 0)$, where for simplicity we've defined axes where the axis of rotation is the z axis and assumed the validity of the small-angle approximation. This gives kinetic energy for the two free propagation times of $4\omega_{\text{rec}}$ and $4\omega_{\text{rec}}(1 + 2\delta\theta^2)$, respectively. So, how large might this effect be?

Cold atoms experiments are generally bolted to a heavy table (certainly the case for all experiments described in this dissertation). As such, they are unlikely to rotate relative to the Earth, leaving the rotation of the Earth itself as the major concern. The Earth rotates at a rate of $\Omega_{\oplus} = 2\pi/(8.64 \times 10^4 \text{ s})$. Consider the worst-case scenario of a vertical interferometer at the equator, which will have its axis rotate at the same rate. For an experiment with $T = 10 \text{ ms}$, we find $\delta\theta = 7.3 \times 10^{-5} \text{ rad}$, which would result in a 5 ppb shift. So, we see an effect large enough to be taken seriously, particular since moving to vertical geometry has as a principal bonus the ability to increase T by factors of 5 or more over the $\sim 10 \text{ ms}$ that is feasible in horizontal geometries. More generally, for a vertical interferometer run at latitude ϕ , we have

$$\frac{\delta\omega_{\text{rec}}}{\omega_{\text{rec}}} = \left(\frac{2\pi T}{8.64 \times 10^4 \text{ s}} \right)^2 \cos^2 \phi .$$

Seattle's latitude wins a factor of two, but even the $T = 10 \text{ ms}$ interferometer would have a 2.5 ppb shift. Of course, latitude and the length of the day being well-known and unchanging³ makes this an easy shift to correct for.

in 7.3.1. We consider this question there.

³Technically, each of these changes on geological time scales due to tectonic plate movement and tidal interactions with the moon, respectively. If time to complete a data set should ever approach geological scales, this dissertation would be obsolete.

For a large N experiment, rotation during the sequence of acceleration pulses could also affect the momentum during free propagation. For an acceleration process wherein each individual momentum kick is small compared to $2\hbar kN$, the actual momentum at the end may be found by treating the process as continuous. Such a continuous process may be modeled by adding the individual, infinitesimal momentum vectors. Geometrically, this gives an arc of a circle. Since the total length of the arc must be $2\hbar kN$ and the angle subtended by the arc must be the difference in angle between the first and last momentum kicks, $\Delta\theta$, we must have an arc from a circle of radius $2\hbar kN/\Delta\theta$. Using this construction, we find a net momentum kick of $2\hbar kN(1 - \Delta\theta^2/6, \Delta\theta/2, 0)$.

Defining $\Omega'_\oplus \equiv \Omega_\oplus \cos \phi$, using a particular latitude of interest, we can write down the momentum during the two free propagation periods for an interferometer where each of the four acceleration phases lasts a time t_{acc} . The momentum for the first time comes directly from the above formula:

$$\vec{p}_1 = 2\hbar kN \left(1 - \frac{(\Omega'_\oplus t_{\text{acc}})^2}{6}, \frac{\Omega'_\oplus t_{\text{acc}}}{2}, 0 \right). \quad (7.1)$$

The momentum at the second time requires applying a rotation to the above formula, yielding

$$\vec{p}_2 = -2\hbar kN \left(1 - \frac{7}{6}(\Omega'_\oplus t_{\text{acc}})^2 - (\Omega'_\oplus T)^2 - 2\Omega'^2_\oplus T t_{\text{acc}}, -\Omega'_\oplus(2T + 3t_{\text{acc}}/2) \right). \quad (7.2)$$

Combining (7.1) and (7.2) we find the total effect to be

$$\frac{\delta\omega_{\text{rec}}}{\omega_{\text{rec}}} = \Omega'^2_\oplus \left(T^2 + T t_{\text{acc}} + \frac{1}{12} t_{\text{acc}}^2 \right). \quad (7.3)$$

For orientations other than vertical, the result of (7.3) holds but with a different effective rotation rate Ω'_\oplus . For an interferometer oriented along a line of longitude we have $\Omega'_\oplus = \Omega_\oplus \sin \phi$, while for one oriented along a line of latitude $\Omega'_\oplus = \Omega_\oplus$. Generically, $\Omega'_\oplus = \Omega_\oplus |\hat{r} \times \hat{x}|$, where \hat{r} is the unit vector in the direction of Earth's rotation axis and \hat{x} is the unit vector in the direction of the interferometer's axis.

Accurate subtraction of this effect should be achievable with a simple compass and suitable correction for the difference between magnetic and rotational north.

7.2 External Field Effects

In this section we discuss systematic shifts due to interactions of the atoms with external fields not deliberately imposed on the atoms. These include the gravitational and magnetic fields of the earth, stray electric fields (either static or radiative), and stray magnetic fields from sources other than the earth. The combination of Yb as source atom with the symmetry of the contrast interferometer renders all of these shifts quite small.

7.2.1 Gravitational Fields

As mentioned in 7.1, Newtonian gravity viewed as a potential energy is more than sufficiently accurate for our purposes. General relativistic effects need not be considered.

The potential energy due to gravity is given by $m\Phi_g$ where Φ_g is the gravitational potential. For an interferometer oriented along the x axis, the first relevant effect is due to $\partial^2\Phi_g/\partial x^2$. It should be noted that, like the electrostatic potential, the gravitational potential in free space satisfies Laplace's equation: $\nabla^2\Phi_g = 0$. Thus, in principle, the interferometer could be run in each of three orthogonal directions. The sum would then exactly cancel gravitational effects (out to fourth derivatives, which will be negligible). This is, however, infeasible in practice.

It is standard to discuss gravitational effects in terms of the acceleration due to gravity $\vec{g} = -\nabla\Phi_g$. In terms of g , the first relevant effect is due to gradients of g . Gradients of g are used in geodesy for discovery of oil and mineral deposits. As such, commercial gravity (g) gradiometers are available with extremely good accuracy and precision. We will return to this point once the magnitude of the effect is established.

We calculate the energy difference responsible for the differential phase evolution of an interferometer oriented along the x axis:

$$\frac{E_{-1}(t) + E_1(t)}{2} - E_0(t) \approx m \frac{1}{2} \frac{\partial^2 \Phi_g}{\partial x^2} \delta^2(t) = -\frac{m}{2} \frac{\partial g_x}{\partial x} \delta^2(t) \quad (7.4)$$

where $\delta(t)$ is the separation between arm 0 and either of the ± 1 arms. For an interferometer with a deceleration/mirror/acceleration phase of time span $2t_{\text{acc}}$ near maximum separation, we calculate the total phase difference $\Delta\theta_{\text{grav}}$:

$$\begin{aligned} \Delta\theta_{\text{grav}} &= -\frac{m}{2\hbar} \frac{\partial g_x}{\partial x} \int dt \delta^2(t) \\ &= -\frac{m}{2\hbar} \frac{\partial g_x}{\partial x} \left[2 \int_0^T dt \{ (2Nv_r t)^2 \} + 2t_{\text{acc}} (2Nv_r T)^2 \right] \\ &= -8\omega_{\text{rec}} N^2 T \frac{\partial g_x}{\partial x} T \left(\frac{1}{3} T + t_{\text{acc}} \right), \end{aligned}$$

giving

$$\frac{\Delta\theta_{\text{grav}}}{\theta_{\text{rec}}} = -\frac{\partial g_x}{\partial x} \left(\frac{1}{3} T^2 + t_{\text{acc}} T \right). \quad (7.5)$$

As an example, consider the energy difference for an interferometer outside of a spherical mass M , with axis oriented along the radius vector of the spherical mass:

$$\begin{aligned} \frac{E_{-1} + E_1}{2} - E_0 &= -\frac{GMm}{2} \left[(R + d(t) + \delta(t))^{-1} + (R + d(t) - \delta(t))^{-1} - 2(R + d(t))^{-1} \right] \\ &= -GMm \left[\frac{\delta^2(t)}{(R + d(t))^3 - \delta^2(t)(R + d(t))} \right] \\ &= -\frac{GMm\delta^2(t)}{(R + d(t))^3} \left(1 - \frac{\delta^2(t)}{(R + d(t))^2} \right) \\ &= -\frac{GMm\delta^2(t)}{R^3} + \mathcal{O}(d/R, (\delta/R)^2), \end{aligned}$$

where R is the distance from the center of the massive body to the initial position of the interferometer before the splitting and $d(t)$ is the displacement of arm 0 from its initial position. The expression has the form we would anticipate based on equation (7.4), but the order symbol allows us to see the magnitude of error we make by considering only gradients. For the Earth's gravitational field we have $\partial g_r / \partial r =$

$2GM_E/R_E^3 = 3.1 \times 10^{-6} \text{ s}^{-2}$, where $M_E = 5.97 \times 10^{24} \text{ kg}$ and $R_E = 6.37 \times 10^6 \text{ m}$. For $T = 10 \text{ ms}$ and $t_{\text{acc}} = 1 \text{ ms}$ we find a relative phase shift of -1.2×10^{-10} . From Laplace's equation and an approximate cylindrical symmetry of the local field around the radius vector, we determine that this shift will have the opposite sign and half the size for a horizontally oriented interferometer.

Alternatively, we might ask: For the given T and t_{acc} how accurately must we measure the gradient in g to subtract the phase shift down to the 0.1 ppb level? We see from above that the answer will be $\approx 10^{-6} \text{ s}^{-2}$. Commercial gravity gradiometers, such as those built by Gravitec, can measure gradients of size $5 \times 10^{-9} \text{ s}^{-2}$ or better.

Since the gradient due to the spherical Earth is barely at the level needing correction, we might wonder if anything in the lab could cause gravity gradients sufficient to concern us. Keeping the gradient in g fixed and simply rescaling M and R simultaneously, we find that a mass one meter from the atoms would have to be over 22,000 kg to cause a gradient comparable to that due to earth. One can estimate that Mt. Rainier has a mass $\sim 10^{13} \text{ kg}$ (base area $\approx 5 \text{ km}^2$ and height $\approx 4 \text{ km}$ with density $\approx 3 \text{ g/cm}^3$) and is a distance $\sim 10^5 \text{ m}$ away and so creates a completely negligible gravity gradient. Tidal changes in Puget Sound constitute a change of water mass of approximately $5 \text{ km}^3(1 \text{ g/cm}^3) = 5 \times 10^{12} \text{ kg}$ [63]. If the entire mass were located 10 km from the atoms (roughly the distance to the closest part of Puget Sound) the gradient due to the changed water mass is roughly 10^{-4} times that due to the spherical earth. So, we can safely neglect any time-varying gradient. For the accuracy required to subtract this effect, the gradient calculated with the spherical Earth model should be correct.

7.2.2 *Electromagnetic Fields*

External (quasi-static) electric fields can affect the interferometer through the Stark shift. The energy of an atom in a static electric field is $E_{\text{Stark}} = p\mathcal{E}^2$, where p is the atomic polarizability and \mathcal{E} is the magnitude of the electric field. The polarizability

of ytterbium is $\approx 2.3 \times 10^{-39} \text{ C}^2 \cdot \text{m}^2/\text{J}$ [10]. The curvature of the energy along the interferometer axis, z , gives

$$\Delta E_{\text{Stark}} = p \left[\left(\frac{\partial \mathcal{E}}{\partial z} \right)^2 + \mathcal{E} \frac{\partial^2 \mathcal{E}}{\partial z^2} \right] (\Delta z)^2 . \quad (7.6)$$

Integrating this over the trajectory of the moving arms gives a result identical to (7.5) with $-dg_x/dx$ replaced by p/m times the bracketed expression in (7.6).

Treating the closest possible charge (accumulated charge on the vacuum viewports, $d = 5 \text{ cm}$ from the atoms) as a point source, we find

$$\begin{aligned} \frac{\Delta \theta_{\text{EM}}}{\theta_{\text{rec}}} &= 10 \frac{p}{m} \frac{k^2 q^2}{d^6} \left(\frac{1}{3} T^2 + t_{\text{acc}} T \right) \\ &= 4 \times 10^{14} \text{ s}^{-2} \left(\frac{q}{1 \text{ C}} \right)^2 \left(\frac{1}{3} T^2 + t_{\text{acc}} T \right) . \end{aligned} \quad (7.7)$$

For a $T = 10 \text{ ms}$ experiment with $t_{\text{acc}} = 1 \text{ ms}$, the charge must be no larger than 0.076 nC to keep the systematic shift to 0.1 ppb . A charge this large could easily be detected with an electrometer⁴. The Stark shift due to blackbody radiation can be considered as a quasi-static effect as well. The variation of the black body field in a vacuum chamber is far too small to have a noticeable impact on the CI phase[89].

Static magnetic fields couple to the spin of an atom. Since the Yb ground state has spin zero, magnetic fields cannot affect it. This, in fact, is one of the key advantages of Yb for precision measurement.

7.3 Diffraction Beam Effects

For the vast majority of interaction effects between single atoms and light, the semi-classical approximation—treating the electromagnetic field as a classical potential for

⁴For example, the Standard Imaging SuperMAX Electrometer can measure charges as small as 1 fC . In practice, 0.076 nC is sufficiently large that confetti thrown at the viewport would stick due to polarization of the paper. In fact, a charge as small as 0.025 nC should cause confetti to stick

quantum mechanical atoms—gives quite accurate results ⁵. For a contrast interferometer, the splitting and acceleration pulses all involve interactions with strong laser pulses for which the quantum nature of light is essentially irrelevant. Though, as discussed earlier, talking about photons may be a convenient description of these interactions for a single atom, the semi-classical approximation is more than sufficient for the level of precision required of our experiment.

In this section we discuss the possible systematic shifts due to interactions between the atoms and the diffraction beams. We omit the most important of these effects, the diffraction phase. This effect will be treated separately in chapter 8.

7.3.1 Beam Geometry

We already used the semi-classical theory to explain diffraction processes in 5.2. In this section we address the effects of real beam geometry as compared to the simple plane wave used in that discussion. We will first look at how this geometry modifies the momentum transferred to an atom during a laser pulse. Then we will see how it affects the dispersive energy shifts of the spatially separated branches of the wave function.

For simplicity, we consider a laser beam propagating in the z -direction plane polarized such that its only non-zero electric field component is E_x . We then postulate a solution to the wave equation of the form $E_x = \psi(\vec{x}, t)e^{ikz - i\omega t}$ where the function ψ changes slowly in space and time. Specifically, we assume that

$$\frac{\partial^2 \psi}{\partial z^2} \ll k \frac{\partial \psi}{\partial z} \quad \text{and} \quad \frac{\partial^2 \psi}{\partial x^2}, \frac{\partial^2 \psi}{\partial y^2} \gg \frac{\partial^2 \psi}{\partial z^2}. \quad (7.8)$$

The first assumption corresponds to the envelope function changing only on scales much larger than the wavelength of the light while the second corresponds to focusing with small numerical aperture. Both of these assumptions are valid to high precision

⁵Even the photoelectric effect, originally cited as evidence for the existence of quanta of light, can be explained with semi-classical theory.

for our diffraction beams. A final assumption is of a stationary envelope (i.e., $\frac{\partial \psi}{\partial t} = 0$) Plugging this ansatz into the wave equation and dropping the term $\frac{\partial^2 \psi}{\partial z^2}$, in accord with the above assumptions, yields a two-dimensional, free-particle Schrödinger equation for ψ with z serving the role of the time coordinate:

$$\frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2} = -i(2k) \frac{\partial \psi}{\partial z}, \quad (7.9)$$

where $k = \omega/c$. Analogously, the Schrödinger equation in the context of single-particle quantum mechanics may be viewed as an approximation to a relativistic equation (e.g., the Klein-Gordon equation for ^{174}Yb) where the wave function changes slowly in the ct direction of space-time⁶. The choice of k both simplifies the equation and gives the plane wave term the wave vector corresponding to that determined experimentally, since we measure the laser frequency directly and then use the speed of light to find the wave vector. We choose the plane with uniform phase to be $z = 0$. As we use a single-mode optical fiber to purify the laser mode, a Gaussian is the appropriate shape for ψ in this plane. We may use the scaling solutions, described fully in chapter 9 for a Gaussian wave packet to determine the form of ψ for all z . For now, the reader is invited to check that the form given below indeed solves the equation; the full solution technique is explained in chapter 9. With $\psi_0(x, y, z = 0) = E_0 \exp(-(x^2 + y^2)/w_0^2)$, we find a solution of the form $\psi = E_0 \exp(if(t) - (x^2 + y^2)/(w_0^2 \alpha^2) - \dot{\alpha}(x^2 + y^2)/(2\alpha))/\alpha$

$$\frac{df}{dz} = -\frac{1}{2k} \frac{1}{\psi_0 \alpha^2} (2) \left. \frac{\partial^2 \psi_0}{\partial x^2} \right|_{(x,y)=(0,0)} = \frac{2}{kw_0^2} \alpha^{-2} \quad \text{and} \quad (7.10)$$

$$\frac{d^2 \alpha}{dz^2} = \frac{1}{2k^2} \frac{1}{\alpha^3} \frac{\partial^2}{\partial x^2} \left[\frac{1}{\psi_0} \frac{\partial^2 \psi_0}{\partial x^2} \right] \Big|_{(x,y)=(0,0)} = \frac{4}{k^2 w_0^4} \alpha^{-3} \quad (7.11)$$

⁶Linear media may be included in this formalism as terms analogous to potential energies in the quantum mechanical Schrödinger equation. Nonlinear media can also be included leading to nonlinear Schrödinger equations. In particular, a uniform medium with third-order nonlinearity leads to an equation identical to the two-dimensional Gross-Pitaevski equation, leading to many fruitful analogies between the propagation of strong laser pulses and the dynamics of Bose-Einstein condensates.

with solutions

$$\alpha = \left(1 + \left(\frac{z}{z_R}\right)^2\right)^{\frac{1}{2}} \quad \text{and} \quad f = \text{atan}\left(\frac{z}{z_R}\right) \quad (7.12)$$

where $z_R = kw_0^2/2$ is the so-called Rayleigh range. The electric field $\vec{E}(x, y, z)$ at all points in space is then given by

$$E_0 \hat{x} \frac{w_0}{w(z)} \exp\left(i(kz - \omega t) - \frac{x^2 + y^2}{w(z)^2} + i\frac{z}{z_R} \frac{x^2 + y^2}{w(z)^2} - i\text{atan}\left(\frac{z}{z_R}\right)\right), \quad (7.13)$$

where

$$w(z) = w_0 \left(1 + \left(\frac{z}{z_R}\right)^2\right)^{1/2}.$$

While cumbersome when written out in such detail, each term has a simple physical interpretation. The first term in the exponential is that expected for a simple plane wave. The second term in the exponent and the similar z dependence of the amplitude ensure conservation of energy. That is, as the beam cross-section shrinks or grows the power flowing into a thin spatial slice (perpendicular to the beam axis) equals the power flowing out the other side. The third term in the exponential is known as the wavefront curvature term. This accounts for the necessary curvature of the phase fronts to allow focusing of a beam coming in from $z = -\infty$ and subsequent defocusing for $z > 0$. The final term in the exponential is the Gouy phase. The wavefront curvature slightly increases the local longitudinal momentum density away from the beam axis. To conserve net momentum flowing through a thin spatial slice perpendicular to the beam axis, the longitudinal momentum must be reduced throughout the slice. This is the role of the Gouy phase. So, we see the Gaussian beam takes exactly the form needed to allow for focusing while conserving energy and momentum.

With the solution in hand, we can now check our approximations for consistency. We have characteristic length scales for change in the x , y , and z directions of w_0 , w_0 , and z_R , respectively. Thus, neglecting $\partial^2\psi/\partial z^2$ compared to $\partial^2\psi/\partial x^2$, $\partial^2\psi/\partial y^2$, and $k\partial\psi/\partial z$ corresponds to neglecting terms of order $(w_0/z_R)^2$, $(w_0/z_R)^2$, and $1/(kz_R)$.

These orders all equal roughly $(\lambda/w_0)^2$ where λ is the wavelength corresponding to k . For our experiment, with $\lambda = 556$ nm and $w_0 = 4$ mm, these corrections are smaller than those we account for by a factor of 10^{-8} . As will be seen below, these corrections would be far smaller than our desired precision of 10^{-9} or 10^{-10} .

As discussed in 5.2, when an atom makes a momentum-state transition induced by the standing wave of light, the momentum is quantized in units of $2\hbar\vec{k}_{\text{rec}}$, called the recoil momentum, where \vec{k}_{rec} is the local wave vector of the laser beam at the atom's position. This is determined by the gradient of the phase term in the exponential, which gives

$$\frac{2x}{w(z)^2}\hat{x} + \frac{2y}{w(z)^2}\hat{y} + \left[k + \frac{x^2 + y^2}{w(z)^2} \left(\frac{1}{z_{\text{R}}} - \frac{z^2}{z_{\text{R}}^3} \left(1 + \left(\frac{z}{z_{\text{R}}} \right)^2 \right)^{-1} \right) - \frac{1}{z_{\text{R}}} \left(1 + \left(\frac{z}{z_{\text{R}}} \right)^2 \right)^{-1} \right] \hat{z}.$$

By imaging using the diffraction light, x_0 and y_0 can be determined to within 20 μm . Using a Shack-Hartmann wavefront analyzer, z_0 may be found to within 1 cm. These should be compared to $w_0 = 4$ mm and $z_{\text{R}} = 90$ m, respectively.

First, consider a horizontal interferometer geometry. Let y be the vertical direction. Then, the x and z coordinates will be constant. The uncertainty in the x component of \vec{k}_{rec} leads to an uncertainty in the final recoil frequency measurement. Since $z/z_{\text{R}} \ll 1$ we may approximate $\delta k_x/k = \delta x \delta z/z_{\text{R}}^2 = 2.5 \times 10^{-11}$, giving a relative uncertainty in the recoil frequency of $(2N)^2 \cdot 6.3 \times 10^{-22} = 2.5 \times 10^{-21} N^2$. The uncertainty due to the y component could be calculated analogously by using $y = 3$ mm in place of δx . This gives a relative uncertainty in the recoil frequency of $4 \times 10^{-17} N^2$, still markedly below our desired precision for an $N = 100$ experiment.

The relative shift in the z component of \vec{k}_{rec} may be approximated by again using $z/z_{\text{R}} \ll 1$ to find

$$\frac{\Delta k_z}{k} = \frac{1}{k z_{\text{R}}} \left[\left(\frac{\rho}{w_0} \right)^2 - 1 \right] = \frac{\lambda^2}{2\pi^2 w_0^2} \left[\left(\frac{\rho}{w_0} \right)^2 - 1 \right] = 9.8 \times 10^{-10} \left[\left(\frac{\rho}{w_0} \right)^2 - 1 \right]$$

where $\rho = \sqrt{x^2 + y^2}$ is the distance from the beam axis.

A realistic sequence begins at $y = 3/4w_0$, finishing at $y = -3/4w_0$. For our parameters of $w_0 = 4$ mm this gives a total time of 35 ms. With an expansion time of 5 ms and $T = 15$ ms the initial momentum is given at a point $y = (3/4)w_0 - (3/2)w_0(5/35)^2 = .72w_0$ and the momentum reversal occurs at a point $y = 3/4w_0 - 3/2w_0(20/35)^2 = .26w_0$. The relative shift in frequency is twice the relative shift in wavenumber. Thus, we see a frequency shift for the first time T of -9.4×10^{-10} . The momentum reversal gives a net momentum $k_2 = 2N(k + \Delta k_1) - 4N(k + \Delta k_2) = -2N(k + 2\Delta k_2 - \Delta k_1)$ for Δk_1 and Δk_2 the wavenumber shifts at the initial acceleration point and the turn around point. The relative shift in recoil frequency for the second time T is, therefore, $2 \cdot 9.8 \times 10^{-10}(2 \cdot .068 - .52 - 1) = -2.7 \times 10^{-9}$, giving an overall shift of -1.8×10^{-9} . Using the more modest $T = 7$ ms gives an overall shift of -1.3×10^{-9} .

The preceding discussion makes the vertical geometry straight forward to analyze. As seen above, the x and y components of the momentum are irrelevant at our level of precision. As discussed below, it may be advantageous to work with $\rho = w_0$. Allowing for z to change by 10 cm in the course of such a run gives a relative frequency shift of $7 \times 10^{-15} N^2$.

For a vertically oriented interferometer, x and y may be held essentially fixed during the run. Placing the interferometer at a point with $\rho = w_0$ causes a near-perfect cancellation of Δk_z between the wavefront curvature and Gouy phase terms. The residual Δk_z is smaller than those discussed for a horizontal geometry by a factor of $(z/z_R)^2$. The relative uncertainty would then be dominated by the uncertainty in ρ :

$$\frac{\delta k_z}{k} = \frac{\lambda^2}{\pi^2 w_0^2} \frac{\delta \rho}{w_0} = 2.0 \times 10^{-9} \frac{\delta \rho}{w_0} . \quad (7.14)$$

While equation 7.14 suggests that in a vertical interferometer geometry, the wavefront curvature and Guoy phase shifts may be made to cancel by careful placement of the atoms in the diffraction beams, there is a potential issue from beyond the

semi-classical approximation that ought to be mentioned. If we consider first-order Bragg diffraction as a two-photon transition, momentum is conserved by matching the atomic momentum change, Δp_{atom} , to the momentum change of the light field due to a stimulated absorption/emission pair $\Delta p_{\text{light}} = \hbar(k_1 - k_2)$, where k_1, k_2 are the momenta of the absorbed and emitted photons, respectively.

The uncertainty in Δp_{atom} should be $\sqrt{2}$ times larger than the uncertainty in k_1 and k_2 . The precision of the most recent recoil measurement of Biraben's group ($\delta\omega_{\text{rec}} = 1.3 \times 10^{-9}$ with 1000 photon recoils) implies $\delta k_1 < (1/\sqrt{2})1.3 \times 10^{-9}\sqrt{1000} \approx 3 \times 10^{-8}k_1$. If this absorption event localizes the photon somewhere within the atomic wave packet, typically of order $10^2/k_1$ in size, this would lead to a gross violation of the Heisenberg indeterminacy principle for the absorbed and emitted photons (i.e., $\Delta x \Delta p \approx 3 \times 10^{-6}\hbar \ll \hbar/2$).

While one might argue that the stimulated absorption/emission pair technically only constrains the uncertainty in the momentum difference between the two photons, this seems a bit of a cheat. One could imagine a single stimulated absorption event for an atom with a long-lived metastable state (e.g., the Yb 3P_0 state) allowing a similarly high precision measurement of the momentum transfer due to a single photon being absorbed. The question of the fundamental quantum indeterminacy of such a measurement cannot be settled with current experiments in the way that the indeterminacy of Δp_{atom} is settled by precision recoil measurements.

A potential solution of this conundrum is to suggest that the photons interact with the atom as a mode of the electromagnetic field, roughly analogous to a single-particle wave function for a massive particle. This essentially amounts to saying the atom generates a coupling between the modes but does not localize a photon to the size of the atomic wave packet. Under this interpretation, the precise position of the atomic wave packet within the mode would not affect the transferred momentum. It would still affect the strength of interaction of the mode and the atom by changing the overlap integral between the mode function and the atomic wave function. Thus,

we recover the well-established variation of Rabi frequency within a Gaussian beam without violating the Heisenberg indeterminacy relation in the course of a two-photon transition. The uncertainty in the atomic momentum change would then be related to the uncertainty of the momentum for the EM field mode, which is set by the Rayleigh range for k_z in a Gaussian field mode.

This question could be settled experimentally by running a vertical orientation interferometer with the atomic cloud placed at different points in the beam and looking for variation of the measured recoil frequency with position in the beam.

7.3.2 Counter-propagation

Since the momentum transferred to an atom is $\hbar(\vec{k}_1 - \vec{k}_2)$, if the two laser beams are not perfectly counter-propagating, the z direction momentum transfer will be $2\hbar k \cos \theta$, where θ is the half-angle between the two beams. From the small angle expansion, we see that the total momentum transfer is $2\hbar k(1 - \theta^2/2)$ for small misalignments from perfect counter-propagation. Since $\delta\omega_{\text{rec}}/\omega_{\text{rec}} = 2\delta k/k$, we have

$$\frac{\delta\omega_{\text{rec}}}{\omega_{\text{rec}}} = (\delta\theta)^2 . \quad (7.15)$$

Shrinking this uncertainty to the level required for a 0.7 ppb measurement (i.e., $\delta\theta < 2.6 \times 10^{-5}$ rad) may be accomplished for a retro-reflected beam or for two independent beams by maximizing the coupling of one beam into the output fiber for the other beam. As discussed in 6.7, this technique can go awry if the light is focused at the retro-reflection mirror. Avoiding this problem makes fiber coupling an excellent tool for aligning the two beams to 10^{-5} rad precision.

Similar considerations of accidental imaging will arise in aligning two independent beams. The uncertainty in θ should be $\sqrt{2}$ as large for this case because in retro-reflection, you know that the two beams are perfectly overlapped at the retro mirror and so the uncertainty in θ is related only to the coupling efficiency at one fiber. For independent beams the misalignment at each fiber should be uncorrelated, and so the

uncertainties will add in quadrature. This constitutes an argument (admittedly, not a particularly strong one) in favor of horizontal geometry, as retro-reflection is not a good option in vertical interferometer geometries.

7.3.3 AC Stark Effect

Another concern related to the geometry of the diffraction beams is a differential AC Stark energy shift. The $w_0/w(z)$ factor, and the $-\rho^2/w(z)^2$ in the exponent, in equation (7.13) gives a varied intensity along the diffraction axis, which leads to a different AC Stark shift for the three arms of the interferometer. For a two-level atom, the AC Stark shift is $\Omega_1\Omega_2/(2\Delta)$, where Ω_1, Ω_2 are the Rabi frequencies for the two laser beams and Δ is the detuning from resonance. Our diffraction beams are sufficiently far from resonance for any but the $^1S_0 \rightarrow ^3P_1$ transition, that the two-level atom approximation is valid. As the beams are essentially mode-matched, we may treat the variation of $\Omega_1\Omega_2$ along the z axis as

$$\Omega_1(z)\Omega_2(z) \propto E(z)^2 = E(0)^2 \left(\frac{w_0}{w(z)} \right)^2 e^{-2\frac{\rho^2}{w(z)^2}}. \quad (7.16)$$

For each $2\hbar k$ acceleration, an integrated Stark shift divided by \hbar that is $\sim \pi$ is needed. This may be seen, for instance, in the Bragg regime calculation in 5.2. Thus, we have an integrated phase evolution per acceleration ramp of πN . The exact multiple will depend upon the technique used for acceleration; for now we will use π as an example.

For a single acceleration ramp, the fractional difference in AC Stark shift will be roughly

$$\begin{aligned} \frac{1}{E(z_0)^2} \Delta z_{01}^2 \left. \frac{\partial^2}{\partial z^2} E^2 \right|_{z=z_0} &= \left(\frac{\Delta z_{01}}{z_R} \right)^2 \frac{2\rho^2 - w_0^2}{w(z_0)^2} \left[w_0^2 \frac{4z_0^2}{z_R^2} \left(\frac{2\rho^2}{w(z_0)^4} - \frac{2}{w(z_0)^2} \right) + 2 \right] \\ &\approx 2 \left(\frac{\Delta z_{01}}{z_R} \right)^2 \frac{2\rho^2 - w_0^2}{w(z_0)^2}, \end{aligned}$$

where Δz_{01} is the distance between arms 0 and 1. We've used the fact that $z_0 \ll z_R$ for any conceivable experiment to obtain the more wieldy form on the second line. For

our benchmark experiment ($N = 100, T = 10$ ms) $\Delta z_{01} = 8$ mm. So, at most, we face an effect of order $\pi(100)(10^{-8}) \approx 3 \times 10^{-6}$ rad. Since the benchmark experiment calls for an accuracy of order 10^{-3} rad, this effect can likely be neglected. For completeness, we give an upper bound formula for the relative phase shift due to the two acceleration sequences assuming peak separation throughout:

$$\begin{aligned} \frac{\Delta\theta_{\text{ACS}}}{\theta_{\text{rec}}} &\leq 4(\pi N) \left(\frac{2Nv_{\text{rec}}T}{z_{\text{R}}} \right)^2 \frac{2\rho^2 - w_0^2}{w(z_0)^2} \frac{1}{8\omega_{\text{rec}}N^2T} \\ &= 4\pi N \frac{T\hbar}{mz_{\text{R}}^2} \frac{2\rho^2 - w_0^2}{w(z_0)^2} \\ &= 5 \times 10^{-15} N \left(\frac{T}{10 \text{ ms}} \right) \left(\frac{4 \text{ mm}}{w_0} \right)^4 \frac{2\rho^2 - w_0^2}{w(z_0)^2}. \end{aligned} \quad (7.17)$$

Again, it should be borne in mind that different acceleration techniques may increase this by a small multiplicative factor (no greater than 4).

A far more serious effect is the momentum-state dependent Stark shift of atoms in a periodic potential, which leads to the diffraction phase. This is discussed in detail in chapter 8.

7.3.4 Index of Refraction

While the preceding three effects were due to controllable parameters of the diffraction beams alone, the index of refraction effect is due to the interactions of the atoms with the diffraction beams. However, as this is also amenable to a semi-classical analysis, we include the discussion here.

The question of the exact momentum of light in a dispersive medium stood as a vexing question for many decades[6]. One school of thought, following Minkowski, arrives at $p_{\text{photon}} = n\hbar k_{\text{vac}}$ where $k_{\text{vac}} = \omega/c$ is the vacuum wave vector. The other school, following Abraham, arrives at $p_{\text{photon}} = \hbar k_{\text{vac}}/n$. The Minkowski formula follows from using the de Broglie formula $p = h/\lambda$ with the wavelength in the medium, λ/n , while the Abraham formula follows by assuming uniform motion of the center of mass-energy for a pulse of light passing through a block of medium.

An experiment in 2005 by Campbell et al.[15], showed definitively that for the case of an atom in a dilute gas, the Minkowski formula is correct. Two extreme cases help motivate the final formulas. For a single atom being diffracted out of a large cloud, the entire cloud acts as a medium, so a two-photon transition gives the single atom a momentum kick $2n\hbar k$, where we return to the practice of using k for the vacuum wave vector. Outside the cloud the photons have momentum $\hbar k$ and so the net change of momentum for the light in this process is $-2\hbar k$. This implies that the medium, the cloud itself, must recoil with momentum $2(1-n)\hbar k$ to conserve momentum, i.e., each atom gets a momentum kick $2(1-n)\hbar k/N_{\text{at}}$, where N_{at} is the total number of atoms in the cloud.

For acceleration pulses or the mirror pulse, if the three arms are physically separated and the efficiencies are 100%, conservation of momentum dictates that the transferred momentum is $2\hbar k$ per two-photon transition. This may also be understood as the net result of an atom receiving a $2n\hbar k$ momentum kick in a two-photon transition plus a total momentum kick of $N \cdot 2(1-n)\hbar k/N = 2(1-n)\hbar k$ for its role as medium in the N transitions required for full transition efficiency.

For the initial splitting pulse in the interferometer, the above analysis may be repeated to give the defect of the transferred momentum. In the +1 arm, each atom receives a momentum kick $2n\hbar k$ plus by serving as a medium for transitions into the +1 and -1 arms, it receives momentum kicks $2(1-n)\hbar k f_{+1}$ and $-2(1-n)\hbar k f_{-1}$, respectively, where $f_{\pm 1}$ is the fraction of atoms diffracted into the ± 1 state. For a symmetric splitting pulse, these two contributions cancel, leaving a net momentum of $2n\hbar k$. By identical reasoning, the -1 arm has a momentum $-2n\hbar k$. For a short splitting pulse these cancellations should be essentially exact. For longer Bragg-type splitting pulses, the motion of the arms relative to one another during the pulse might cause small alterations to these results.

Assuming the ± 1 arms have initial momenta $\pm 2n\hbar k$ and that subsequent acceleration and mirror pulses transfer in exact multiples of $2\hbar k$, the recoil energy during the

outgoing period is $4\omega_{\text{rec}}(N + \delta n)^2$, where we use the notation $\delta n \equiv n - 1$. However, since the mirror pulse is not a perfect mirror but rather a transfer in multiples of $2\hbar k$, the recoil energy during the returning period is $4\omega_{\text{rec}}(N - \delta n)^2$. For example, in the experiments described in this thesis, the mirror pulse changes the momentum of the +1 arm from $2n\hbar k = 2\hbar k(1 + \delta n)$ to $2n\hbar k - 4\hbar k = -2\hbar k(1 - \delta n)$. This works in the interferometer's favor by canceling the cross terms from the $(N \pm \delta n)^2$ expressions when calculating the total recoil phase:

$$\theta_{\text{rec}} = 4\omega_{\text{rec}}(N + \delta n)^2 T + 4\omega_{\text{rec}}(N - \delta n)^2 T = 8\omega_{\text{rec}} T [N^2 + (\delta n)^2]. \quad (7.18)$$

Additionally, we can account for the small index of refraction shifts due to imperfect diffraction in the acceleration and mirror pulses. For simplicity, we will assume that each two-photon transition has the same efficiency e , which is close to 1, and that atoms not diffracted in a given transition are too far away to act as media for all subsequent transitions. This gives a momentum per transition of $2n\hbar k - 2\delta n\hbar ke = 2\hbar k(1 + \delta n(1 - e))$. Since δn is proportional to the atomic density, we arrive at a geometric series for the outgoing momentum

$$p_{\text{out}} = 2\hbar k \left(1 + \delta n_0 + \sum_{i=2}^N [1 + \delta n_0(1 - e)e^{i-2}] \right) = 2\hbar k (N + \delta n_0(2 - e^{N-1})),$$

where n_0 is the initial index of refraction. The returning momentum will be

$$\begin{aligned} p_{\text{return}} &= p_{\text{out}} - 2\hbar k \sum_{i=N+1}^{3N} [1 + \delta n_0(1 - e)e^{i-2}] \\ &= p_{\text{out}} - 2\hbar k (2N + \delta n_0(e^{N-1} - e^{3N-1})) \\ &= -2\hbar k (N + \delta n_0(e^{3N-1} - 2)). \end{aligned}$$

Combining with the outgoing momentum, this gives a recoil phase of

$$\begin{aligned} \theta_{\text{rec}} &= 8\omega_{\text{rec}} T \left[N^2 + N\delta n_0 (e^{3N-1} - e^{N-1}) \right. \\ &\quad \left. + (\delta n_0)^2 \left(4 + 2(e^{3N-1} - e^{N-1}) + \frac{1}{2}(e^{6N-2} + e^{2N-2}) \right) \right], \end{aligned}$$

implying the relative error is

$$\frac{\delta\theta_{\text{index}}}{\theta_{\text{rec}}} = \left(\frac{\delta n_0}{N}\right) (e^{3N-1} - e^{N-1}), \quad (7.19)$$

where we dropped the terms of order $(\delta n_0)^2$ because, as will be shown next, $\delta n_0 \lesssim 10^{-6}$ for our current experiments, and likely will be even lower for future sub-ppb-level experiments.

Now, we need an expression for δn . For a classical ideal gas, the dielectric function is

$$\frac{\epsilon}{\epsilon_0} = 1 + \alpha_{1S_0}\rho, \quad (7.20)$$

where

$$\alpha_{1S_0} = -\frac{\mu^2}{\epsilon_0\hbar} \frac{1}{\delta + i\Gamma/2} = -\frac{\mu^2}{\epsilon_0\hbar} \frac{\delta - i\Gamma/2}{\delta^2 + \Gamma^2/4} \quad (7.21)$$

is the Yb ground-state polarizability and ρ is the atomic density[72]. There are modifications to this result based on bosonic statistics, specifically related to the deviation of the two-point correlation function from ρ^2 [72]. For a pure BEC, these corrections vanish, due to the triviality of the two-point function⁷. Additionally, for a condensate high density corrections are possible[85]. In the case of the diffraction pulses, the detuning is sufficient to render these corrections negligible. However, they could become important in a detailed analysis of the readout dynamics.

Since the index of refraction is just the square root of the real part of the dielectric function, we arrive at an expression

$$\delta n = \sqrt{1 - \rho \frac{\mu^2}{\epsilon_0\hbar} \frac{\delta}{\delta^2 + \Gamma^2/4}} - 1 \approx -\frac{1}{2}\rho \frac{\mu^2}{\epsilon_0\hbar} \frac{\delta}{\delta^2 + \Gamma^2/4} = -\rho \frac{3\lambda^3}{8\pi^2} \frac{2\delta}{\Gamma} \frac{1}{1 + \left(\frac{2\delta}{\Gamma}\right)^2}, \quad (7.22)$$

where we used $\mu^2/(\epsilon_0\hbar) = 3\Gamma\lambda^3/(8\pi^2)$ in the final equality[31]. Combining (7.19) and (7.22) we find our final form for the index shift:

$$\frac{\delta\theta_{\text{index}}}{\theta_{\text{rec}}} = -\rho \frac{3\lambda^3}{8\pi^2} \frac{2\delta}{\Gamma} \frac{1}{1 + \left(\frac{2\delta}{\Gamma}\right)^2} \frac{e^{3N-1} - e^{N-1}}{N}. \quad (7.23)$$

⁷For a thermal cloud interferometer, these corrections would need to be considered more carefully.

For the 556 nm transition in Yb with large detuning, we find an approximate value

$$0.0065 \left(\frac{\rho}{10^{12} \text{ cm}^{-3}} \right) \left(\frac{2\delta}{\Gamma} \right)^{-1} \frac{e^{N-1} - e^{3N-1}}{N} . \quad (7.24)$$

With a short expansion time (2 ms) the splitting pulse is fired at a density of $1.0 - 1.5 \times 10^{12} \text{ cm}^{-3}$. Together with our detuning of $\delta = 700\Gamma$, this gives a recoil shift less than 4 ppm. For a future ppb-level experiment, a density at least 10 times lower is recommended for dealing with atomic interaction shifts. Coupled with a larger detuning and a large N , this systematic should be reducible to the ppb level. Additionally, detuning and pulse efficiency can be measured extremely well. The uncertainty in the density is the only limit in the ability to subtract these effects. Finally, though it was recommended earlier that blue detuning be used for diffraction pulses, a pair of data sets with equal magnitude detuning, one red and one blue, would have their index shifts cancel exactly, assuming constant density. Thus, the ability to cancel index shifts exactly would depend not on the ability to measure density but simply to have a repeatable density.

Chapter 8

DIFFRACTION PHASES

Diffraction phases is a term used to describe differential phase shifts on interfering states induced by the diffraction process. Such phases have previously been measured for diffraction from physical absorption gratings[79]. They have been studied theoretically[12] and experimentally[102] for optical gratings in and near the Raman-Nath regime. In this chapter, we develop numerical tools for accurately calculating such phase shifts in any regime of pulse length or strength and show a dramatic reduction of uncertainty when applying these corrections to a contrast interferometer.

8.1 Basic Framework

We will build up a means for calculating the populations in various momentum states, and their phases, after a diffraction pulse. We start with the simplest case of a single, two-state atom with a plane wave, $\vec{p} = 0$ wave function. Additional complications, such as wave functions with nonzero momentum width, will be added to the model subsequently.

In section 5.2 we developed several different forms for the Hamiltonian describing a two-level atom interacting with an optical standing wave. We first developed equation (5.10). To solve problems numerically, we must project this onto a finite set of states. Since we begin with the state $|\vec{p} = 0, g\rangle$, the complete set of accessible states is $\{\dots, |2\hbar k, g\rangle, |\hbar k, e\rangle, |0\hbar k, g\rangle, |-\hbar k, e\rangle, |-2\hbar k, g\rangle \dots\}$. In practice, the Hilbert space may be truncated to include only a manageable number of states without introducing errors large enough to affect final results. For the purposes of this discussion, we will

truncate to five states¹. With this truncation, (5.10) has the matrix representation

$$H = \begin{pmatrix} 4\hbar\omega_{\text{rec}} & \frac{\hbar\omega_{\text{R}}}{2}e^{i\omega t} & 0 & 0 & 0 \\ \frac{\hbar\omega_{\text{R}}}{2}e^{-i\omega t} & \hbar\omega_0 + \hbar\omega_{\text{rec}} & \frac{\hbar\omega_{\text{R}}}{2}e^{-i\omega t} & 0 & 0 \\ 0 & \frac{\hbar\omega_{\text{R}}}{2}e^{i\omega t} & 0 & \frac{\hbar\omega_{\text{R}}}{2}e^{i\omega t} & 0 \\ 0 & 0 & \frac{\hbar\omega_{\text{R}}}{2}e^{-i\omega t} & \hbar\omega_0 + \hbar\omega_{\text{rec}} & \frac{\hbar\omega_{\text{R}}}{2}e^{-i\omega t} \\ 0 & 0 & 0 & \frac{\hbar\omega_{\text{R}}}{2}e^{i\omega t} & 4\hbar\omega_{\text{rec}} \end{pmatrix}. \quad (8.1)$$

For numerical computation, this Hamiltonian suffers from a severe separation of scales. For typical experiments, $\omega \approx \omega_0 \sim 10^{15}\text{Hz}$ while $\omega_{\text{rec}} \sim 10^4\text{Hz}$. The high scale may be dealt with by passing to the interaction picture. Briefly, this entails separating the Hamiltonian into a non-interacting piece, H_0 , and an interaction potential V , so $H = H_0 + V$. The time evolution due to H_0 is then removed by changing to interaction picture states:

$$|\psi, t\rangle_{\text{I}} = e^{iH_0 t} |\psi, t\rangle_{\text{S}}, \quad (8.2)$$

where $|\psi, t\rangle_{\text{I}}$ and $|\psi, t\rangle_{\text{S}}$ are the interaction and Schrödinger picture states, respectively. Time evolution of the interaction picture state is governed by the interaction picture Hamiltonian, V_{I} , given by

$$V_{\text{I}}(t) = e^{iH_0 t} V e^{-iH_0 t}. \quad (8.3)$$

Physically, it is tempting to treat the diagonal as the non-interacting Hamiltonian. However, this choice continues to connect the large and small frequency scales through electronically excited states (e.g., with energy $\hbar(\omega_{\text{rec}} + \omega_0)$). With this in mind, a more fruitful choice might be $H_0 = \hbar\omega_0 |e\rangle\langle e|$, thereby removing the high scale from the diagonal. This yields a Hamiltonian that may be easily simulated, however, this Hamiltonian remains time dependent.

¹In practice, we find 17 states to be necessary for sufficient accuracy in modeling our $2\hbar k \leftrightarrow -2\hbar k$ mirror pulses.

Numerical computations are dramatically sped up by the use of a time-independent Hamiltonian. For our Hamiltonian, the choice $H_0 = \hbar\omega |e\rangle \langle e|$ gives

$$V_I = \begin{pmatrix} 4\hbar\omega_{\text{rec}} & \frac{\hbar\omega_R}{2} & 0 & 0 & 0 \\ \frac{\hbar\omega_R}{2} & \hbar\omega_{\text{rec}} - \hbar\Delta & \frac{\hbar\omega_R}{2} & 0 & 0 \\ 0 & \frac{\hbar\omega_R}{2} & 0 & \frac{\hbar\omega_R}{2} & 0 \\ 0 & 0 & \frac{\hbar\omega_R}{2} & \hbar\omega_{\text{rec}} - \hbar\Delta & \frac{\hbar\omega_R}{2} \\ 0 & 0 & 0 & \frac{\hbar\omega_R}{2} & 4\hbar\omega_{\text{rec}} \end{pmatrix}. \quad (8.4)$$

Equation (8.4) is just the projection of (5.11) onto our five-state basis. In addition to dramatically improving numerical efficiency, this choice also shows a nice connection to the dressed-atom picture for atom-light interactions. Since $\hbar\omega$ may be viewed as the single photon energy, the choice made above is mathematically equivalent to considering the non-interacting systems as “ground state atom plus N photons” and “excited state atom plus $N - 1$ photons,” then subtracting off the energy of N photons from the total energy.

The time evolution operator, $U(t', t)$, for fixed laser intensity can then be calculated as $U(t', t) = \exp(-iV_I(t' - t)/\hbar)$. For typical diffraction pulses with time-varying intensities, the time evolution may be calculated using small time steps approximated as having constant intensity. The SciPy extension to the Python scripting language makes such calculations quick by virtue of its ability to call the LINPACK and BLAS highly optimized linear algebra libraries.

Real experiments are not conducted with perfect plane wave atomic wave functions, nor with ideal laser standing waves. We now consider some refinements which allow the theory to connect more closely with experiments.

8.2 Refinements

8.2.1 Non-Zero Momentum Width

A BEC in a harmonic trap has a non-zero momentum width in accordance with the Heisenberg indeterminacy relation and its non-zero spatial extent. On release from the trap this momentum width grows as atomic interaction energy is converted into kinetic energy.

This issue may be treated by simply calculating the time evolution for each plane-wave in the momentum-space wave function of the condensate. In practice, an extremely coarse-grained approximation using five or seven plane-wave states is sufficient for finding the expected population diffracted by a Bragg pulse for a BEC.

The Hamiltonian (8.4) assumes that the atomic plane wave has initial momentum a multiple of $\hbar k$. The kinetic energy terms in V_I can be written as $(n\hbar k)^2/(2m) = n^2\hbar\omega_{\text{rec}}$ for integer n . Other plane-wave states may be accommodated by adding a parameter δ which gives the fractional offset from the closest integer multiple of the recoil momentum. For example, a BEC released from a trap may have a momentum spread from $-0.3\hbar k$ to $0.3\hbar k$. For the plane wave with $p = -0.3\hbar k$, we use $\delta = p/(\hbar k) = -0.3$. This transforms the kinetic terms of V_I to $(n + \delta)^2\hbar\omega_{\text{rec}}$ but otherwise has no effect on the evolution operator.

This construction also lends itself to the simulation of Bloch oscillations (BO). In BO, a time varying frequency difference between the two counter-propagating laser beams makes an effective standing wave in a frame of reference moving with respect to the lab frame. In this picture, there is a standing wave in an accelerating reference frame. By tuning the depth of the effective periodic potential appropriately, the atoms may be kept in the lowest band, effectively being accelerated along with the “standing” wave in the lab reference frame. The δ parameter may be continuously varied to simulate this accelerated frame picture of the BO. In this way, V_I is viewed as being calculated in the standing wave’s rest frame, thus causing the atomic kinetic

energy in a given state to vary with time.

8.2.2 Imbalanced Laser Beams

Another assumption of the theory in section 8.1 is that the two counter-propagating laser beams have identical intensity, manifested in their having identical Rabi frequencies ω_R . In practice this can be difficult to achieve. In the retro-reflected geometry used for these experiments, inevitable losses in the retro-reflection optics (small absorptions in mirrors, reflections from lens faces, etc.) guarantee an imbalance in power between the outgoing and reflected beams. Careful mode-matching between the two beams is important for dealing with wavefront curvature and relative alignment of the beam wave vectors, so in practice power imbalance is equivalent to intensity imbalance at the atoms.

To model such imbalances, we replace ω_R with two Rabi frequencies ω_1 and ω_2 for the two beams. This modifies the form of V_I to have off-diagonal terms alternating between ω_1 and ω_2 :

$$V_I = \begin{pmatrix} 4\hbar\omega_{\text{rec}} & \frac{\hbar\omega_2}{2} & 0 & 0 & 0 \\ \frac{\hbar\omega_2}{2} & \hbar\omega_{\text{rec}} - \hbar\Delta & \frac{\hbar\omega_1}{2} & 0 & 0 \\ 0 & \frac{\hbar\omega_1}{2} & 0 & \frac{\hbar\omega_2}{2} & 0 \\ 0 & 0 & \frac{\hbar\omega_2}{2} & \hbar\omega_{\text{rec}} - \hbar\Delta & \frac{\hbar\omega_1}{2} \\ 0 & 0 & 0 & \frac{\hbar\omega_1}{2} & 4\hbar\omega_{\text{rec}} \end{pmatrix}. \quad (8.5)$$

8.3 Calibration and Diffraction Phase Calculation

To accurately calculate the phase differences due to a mirror pulse, the parameters in the evolution operator must be fixed. Several are straightforwardly fixed. The recoil frequency ω_{rec} is measured sufficiently well by low precision interferometry before reaching the level of precision where diffraction phase correction becomes important. The detuning Δ is simply read off of the AOM driving signal after calibrating the

resonance frequency. The resonant frequency is found by absorption imaging with a variety of frequencies to map out the Lorentzian line shape. The momentum-space wave function, which determines the relevant values of δ , is measured with time of flight images of the interferometer sequence (allowing direct comparison of the momentum width to the recoil momentum).

It is reasonable to assumed a fixed ratio ω_1/ω_2 from shot to shot. The chief difficulty lies in finding ω_1 for a given laser pulse. A photodiode trace recorded for each experimental run characterizes the total laser power as a function of time for each pulse. However, the relevant parameter is the intensity at the atoms as a function of time. While one could in principle precisely image the position of the atoms in the laser beam and try to connect the total power to the intensity using this information, such a tactic makes difficult to quantify assumptions.

Instead, we calibrate the intensity vs. photodiode signal by producing a sequence of experimental runs with increasing mirror pulse intensity. In each run the pulse width is kept roughly fixed. Absorption images of the resulting momentum states allow us to determine the effect of each pulse on atoms in the $\pm 2\hbar k_{\text{rec}}$ momentum states. This produces a set of photodiode traces and associated fraction of $\pm 2\hbar k_{\text{rec}}$ atoms diffracted by the pulse.

Next, we calculate the expected fraction diffracted by each pulse by guessing a fixed calibration from intensity to ω_1^2 . This process is repeated for a variety of calibration values. Then, the expected fractions are compared to the actual diffracted fraction for each pulse. The calibration values giving the closest fit to the data are determined by calculating the summed squared error (SSE) and taking the calibration that minimizes this. Explicitly,

$$\text{SSE}(C) = \sum_{\text{traces}} [f(\text{trace}, C) - f_{\text{data}}(\text{trace})]^2,$$

where $f(\text{trace}, C)$ is the diffracted fraction calculated for a given photodiode trace and multiplier, C , from photodiode voltage to Ω_1^2 in MHz^2 . The actual diffracted fraction

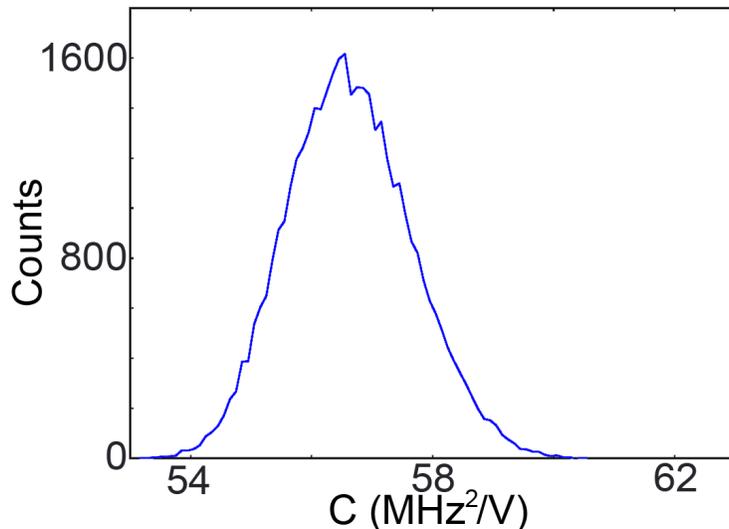


Figure 8.1: Intensity calibration. The histogram gives the number of trials in which the given value of the calibration C gave the smallest error. This particular calibration curve, for the $T = 11$ ms mirror pulse, shows the results from 40,000 trials. The value extracted from this plot is $C = 56.4 \pm 0.8 \text{ MHz}^2\text{/V}$.

from the data is $f_{\text{data}}(\text{trace})$ for the given trace.

To find a robust value for C and quantify the uncertainty, we randomly sample half of the traces from the calibration data set and find the C which best fits this subset. This process is repeated many times and a histogram is constructed for the best values of C over a large number of samplings. Figure 8.1 gives an example of such a histogram, this one for calibrating the $T = 11$ ms data. With a value of C in hand, the diffraction phases may be calculated for each of the data points.

8.4 Results for Current Experiments

For each shot in a data set, the power profile of the mirror pulse is recorded on a fast photodiode (200 MHz bandwidth, ν -Focus model 1801). An example trace is shown in

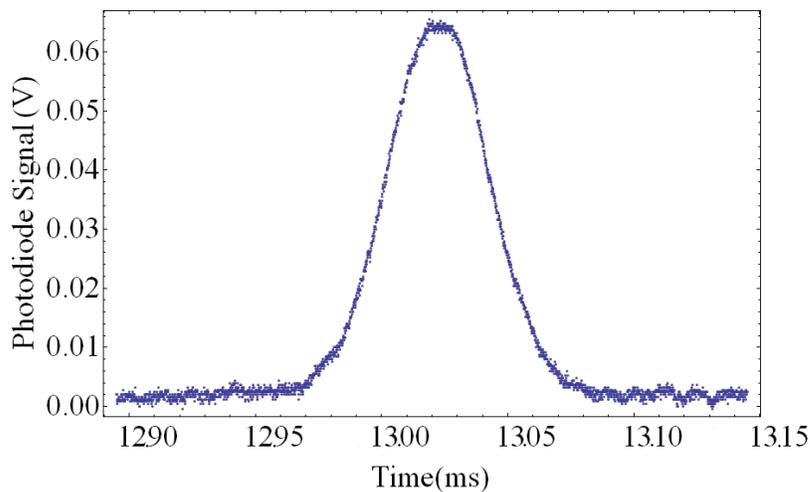


Figure 8.2: Mirror pulse. The trace shows a typical mirror pulse, recorded on a fast photodiode. This particular file is from a $T = 11$ ms data set. The time on the x axis includes a 2 ms offset for the initial expansion of the condensate.

figure 8.2. The photodiode looked at a weak reflection from the beam splitter face. For some time a Thorlabs PMT (Model PMM02) was used to record the pulses, looking at a reflection from a lens face on the retro-reflection side of the chamber. The change to a fast photodiode was made because of the slowness (30 kHz bandwidth) of the Thorlabs PMT. However, the greater sensitivity of the PMT had given better signal to noise. It also consistently had a background voltage of zero, while the photodiode background could vary by a few percent of the peak pulse amplitude from shot to shot. For future experiments, a diffraction light monitor using a high-bandwidth PMT should be incorporated at the design phase. Using one identical to the one detecting the readout light would reduce concerns about differences in timing.

The value of C from the calibration data is used to calculate the diffraction phase for each shot. After calculating the diffraction phase we subtract it from the phase of the contrast signal, $\phi(2T)$. As described in 6.3, we extract $\phi(2T)$ by fitting a sine wave to the signal and finding its phase at time $2T$. Figure 8.3 shows $T = 1$ ms data

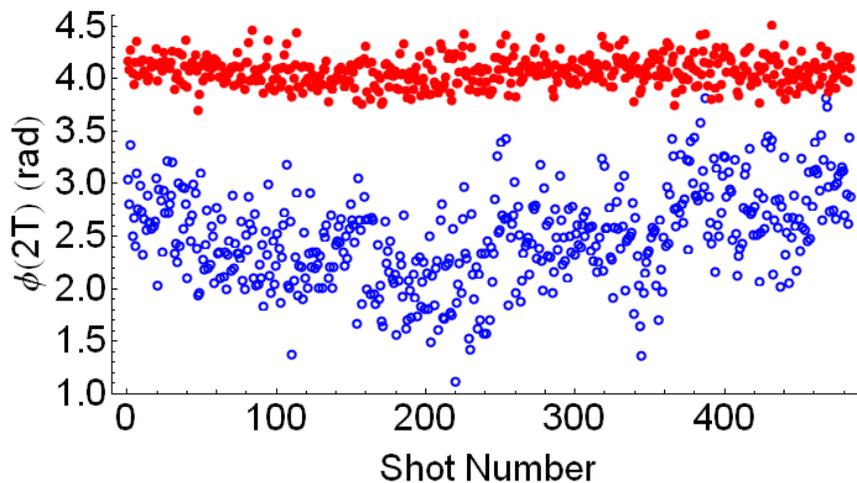


Figure 8.3: Diffraction phase correction. A single large data set is shown as a function of shot number. The open blue circles show the extracted phases before diffraction phase correction. The red filled disks show the same phases after shot-by-shot correction for diffraction phase.

before and after the diffraction phase correction. In addition to a uniform noise width, we also see a marked drift in phase over time. This drift is eliminated in the corrected data, showing that it arose from drifts in laser intensity. (The data were taken over the course of 8 hours.) We also note that the noise width is reduced by the correction. Overall, the standard deviation of the data set drops from 0.44 rad to 0.14 rad. The ability to correct for phase shifts induced by both laser drift and random laser noise is a powerful confirmation of our model. The diffraction phase correction procedure was so useful that it was incorporated into a holistic fitting procedure, described in section [6.5.2](#).

Chapter 9

ATOMIC INTERACTIONS IN A BEC INTERFEROMETER

In this chapter, we will consider the most difficult systematic effect for the Yb BEC contrast interferometer: Atomic interactions. We begin by discussing atomic interactions in the ultra-cold regime. This leads to a discussion of mean-field theory, an extremely accurate technique for describing interactions in Yb BECs. With the technique in place, we develop two special tools for solving the mean field equations. From these technical tools, we extract both qualitative understanding and quantitative predictions.

9.1 Ultracold Interactions

Interaction potentials are often most easily understood in terms of scattering events. For a scattering event, two particles are prepared at infinity where their interactions are negligible. They then propagate toward one another, collide, and some products are again detected at infinity. Such a scenario has the advantage that we needn't consider the complications of many-particle systems to understand the two-body interaction potential. Many-body systems may then be built up by considering the various particles to interact via that same potential.

The theory of quantum scattering has a vast literature. Two particularly good references are the text of Gottfried and Yan[33] for general scattering theory and the Varenna notes by Dalibard[24] for discussion specific to cold atomic physics. Considering the partial-wave expansion of a scattering event, we may see that interactions in the ultra-cold regime have a particularly simple form.

The amplitude for scattering between two plane waves can be decomposed into a

sum of contributions due to different relative angular momentum states. For example, a direct head-on collision has relative angular momentum $\ell = 0$, referred to as *s*-wave, by analogy to the $\ell = 0$ orbitals of hydrogen. This decomposition is possible because, in the absence of spin, the interaction Hamiltonian between two particles will be rotationally symmetric. Therefore, any collision must conserve angular momentum.

The ultra-cold regime simplifies this picture by essentially disallowing any but *s*-wave interactions. In the center-of-mass frame for the incoming particles, the collision problem reduces to an effective single particle scattering from a fixed potential. The radial Schrödinger equation for this effective particle has the form

$$\frac{\hbar^2}{2\mu} \left(-\frac{1}{r^2} \frac{d}{dr} \left[r^2 \frac{d}{dr} R_\ell(r) \right] \right) + \left(\frac{\hbar^2 \ell(\ell + 1)}{2\mu r^2} + V(r) \right) = E_{n\ell} R_\ell(r) , \quad (9.1)$$

where μ is the reduced mass of the pair, $V(r)$ is the interaction potential, R_ℓ is the radial wave function, and $E_{n\ell}$ is the energy of the incoming particles in the center of mass frame. The interaction potential is modified by the addition of the centrifugal barrier term, except for $\ell = 0$. For a potential localized within a radius b , if the relative momentum of the colliding particles is less than $\hbar^2 \cdot 2/(2\mu b^2)$, then the potential lies in the classically disallowed region for any partial wave with $\ell > 0$. Thus, the interaction in the *p*-wave ($\ell = 1$) channel is exponentially suppressed by the probability to tunnel through the centrifugal barrier (as are all higher ℓ channels). For the range of the interatomic potentials encountered in most atomic physics experiments¹, the energy scale for the centrifugal barrier corresponds to temperatures of order hundreds of microkelvin[27]. Thus, for experiments occurring at temperatures $< 1 \mu\text{K}$, the interaction is purely *s*-wave to an extremely high precision.

Treating the interactions as purely *s*-wave simplifies the analysis. The condition of having low enough kinetic energy to ignore $\ell > 0$ is equivalent (up to a factor of π) to $\lambda > b$. As the wavelength becomes much larger than the range of the interaction,

¹Strictly speaking, the potentials encountered in atomic physics are not completely localized. However, since they drop off with distance faster than $1/r^3$, a specific centrifugal barrier energy may still be defined.

the shape of the interaction potential becomes less important. For a case as extreme as that of atoms at $T < 1 \mu\text{K}$, to a very good approximation the interaction potential may be replaced by a simple delta function².

A single number characterizes the strength of such an interaction. We will use the s -wave scattering length, a_s . The scattering length arises out of the effective range expansion[101]. For our purposes, it may be thought of simply as the radius corresponding to the scattering cross-section, if the atoms are considered as solid spheres. Thus, it correlates with the strength of atomic interactions, not the physical range of the interaction potential. By matching this cross-section to that calculated with the Born approximation for a potential $\eta\delta(\vec{r})$, we find the value of the coefficient η to be $2\pi\hbar^2 a_s/\mu = 4\pi\hbar^2 a_s/m$, where we used the assumption of identical particles to replace the reduced mass.

With the interaction potential characterized in this low-temperature limit, we may now describe the collective interaction effects in the BEC.

9.2 Mean-Field Theory

To describe the interactions between atoms in the condensate, we make use of the idea that all atoms are in the same (ground) state. This allows us to consider a single atom to be experiencing interactions with all of the other, identical atoms. The single state in which they all find themselves may then be solved for self-consistently. The atom is said to interact with the average distribution of atomic density in the system, hence the name “mean-field.” Mean-field theory is well described in several review articles[23] and texts[81]. We will more closely follow the treatment of Leggett[61], which stresses the structure of the many-body wave function.

The mean-field treatment reduces the often intractable problem of finding a many-

²This approximation is actually part of a controlled approximation scheme known as the gradient expansion. The next term has the form $\alpha\nabla^2\delta(\vec{r})$. Discussing the connections between this expansion and effective field theory would take us too far afield from our goal of describing interactions in a BEC. The interested reader should enjoy the excellent review by LePage[62].

body wave function into the simpler problem of solving for a single-particle wave function. This approach leads to a nonlinear version of the Schrödinger equation. Nonlinear equations are more complex both to understand and to solve numerically than linear equations of comparable size. However, the truly bewildering size of the Hilbert space for a many-body wave function means that the nonlinear Schrödinger equation can, in fact, be both more comprehensible and more amenable to simulation than the linear Schrödinger equation for a large number of atoms. In fact, direct numerical solution of the Schrödinger equation for more than a few particles in three dimensions is beyond the reach of computers.

With the picture of all atoms in the same state, and the requirement that bosonic wave functions be symmetrized, we find a many-body wave function of the form

$$\Psi(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N) = \prod_{j=1}^N \phi(\vec{x}_j) , \quad (9.2)$$

where $\phi(\vec{x})$ is a single-particle wave function. Since ϕ should be the ground state, we may use the variational method to find the function. However, to do this we need a Hamiltonian including the interaction potential.

For a many-body system, it is easier to write the Hamiltonian in terms of creation and annihilation operators. Define $\hat{\psi}(\vec{x})$ and $\hat{\psi}^\dagger(\vec{x})$ as the operators which annihilate or create a particle at point \vec{x} , respectively. The Hamiltonian can be written as

$$\begin{aligned} \hat{H} &= \int d^3x \hat{\psi}^\dagger(\vec{x}) \left[-\frac{\hbar^2}{2m} \nabla^2 + U(\vec{x}) \right] \hat{\psi}(\vec{x}) \\ &\quad + \frac{1}{2} \int d^3x d^3x' \hat{\psi}^\dagger(\vec{x}') \hat{\psi}^\dagger(\vec{x}) V(\vec{x} - \vec{x}') \hat{\psi}(\vec{x}) \hat{\psi}(\vec{x}') \\ &= \int d^3x \left\{ \hat{\psi}^\dagger(\vec{x}) \left[-\frac{\hbar^2}{2m} \nabla^2 + U(\vec{x}) \right] \hat{\psi}(\vec{x}) + \frac{2\pi\hbar^2 a_s}{m} \hat{\psi}^\dagger(\vec{x}) \hat{\psi}^\dagger(\vec{x}) \hat{\psi}(\vec{x}) \hat{\psi}(\vec{x}) \right\} . \quad (9.3) \end{aligned}$$

We may then use this equation to find the ground state subject to the ansatz of equation (9.2) through a variational calculation. To keep $\phi(\vec{x})$ normalized, we introduce a Lagrange multiplier, $-N\mu$, in anticipation of the direct connection that emerges between the multiplier and the chemical potential μ . Taking the variational derivative,

with respect to ϕ^* and then dividing through by N gives

$$\mu\phi = -\frac{\hbar^2}{2m}\nabla^2\phi + U(\vec{x})\phi + (N-1)g|\phi|^2\phi, \quad (9.4)$$

where $g \equiv 4\pi\hbar^2 a_s/m$. Equation (9.4) is referred to as the Gross-Pitaevskii equation³.

Similarly, we may plug the mean-field ansatz into the time-dependent Schrödinger equation. For this purpose it may be instructive to rewrite the state in bra-ket notation:

$$|\Psi\rangle = \bigotimes_{j=1}^N |\phi\rangle_j. \quad (9.5)$$

This clarifies the action of the annihilation operators:

$$\begin{aligned} \hat{\psi}(\vec{x})|\Psi\rangle &= \phi(\vec{x})|0\rangle \otimes \left(\bigotimes_{j=2}^N |\phi\rangle_j\right) + \phi(\vec{x})|\phi\rangle_1 \otimes |0\rangle \otimes \left(\bigotimes_{j=3}^N |\phi\rangle_j\right) + \dots \\ &\quad + \phi(\vec{x}) \left(\bigotimes_{j=1}^{N-1} |\phi\rangle_j\right) \otimes |0\rangle, \end{aligned}$$

which leads to

$$\hat{\psi}^\dagger(\vec{x})\hat{\mathcal{O}}\hat{\psi}(\vec{x})|\Psi\rangle = \phi^*(\vec{x})\mathcal{O}_x[\phi(\vec{x})]|\Psi\rangle, \quad (9.6)$$

where \mathcal{O}_x is the \vec{x} -space representation of the operator $\hat{\mathcal{O}}$. With this result, it is easy to apply the Hamiltonian operator to the state $|\Psi\rangle$ in writing down the Schrödinger equation for this state. Finally, multiplying by the bra $\langle\Psi|$ and dividing by N leads to

$$i\hbar\frac{\partial}{\partial t}\phi = -\frac{\hbar^2}{2m}\nabla^2\phi + U(\vec{x})\phi + (N-1)g|\phi|^2\phi, \quad (9.7)$$

known as the time-dependent Gross-Pitaevskii equation. This is a particularly simple nonlinear Schrödinger equation. Such equations arise in the description of laser pulses in nonlinear media or the propagation of certain types of water waves, in addition to the description of BECs.

³Technically, the equation referred to as the Gross-Pitaevskii equation is missing the factor of $(N-1)$ on the interaction term. This comes about by defining the condensate order parameter, which is just $\sqrt{N}\phi(\vec{x})$, and then finding the equation for the ground-state order parameter (recalling that $N-1 = N$ to good accuracy in the mean-field regime).

Equations (9.4) and (9.7) allow us to determine both the ground state of a BEC in a trap and all of its subsequent dynamics upon release from the trap. In the following sections, we develop a set of tools for solving those problems. Solutions to nonlinear partial differential equations are notoriously difficult to come by. We will employ a mix of analytic approximations and numerical algorithms to both understand the physics and arrive at quantitative predictions for interaction effects in a BEC interferometer. While we will eventually apply these techniques specifically to a ^{174}Yb contrast interferometer, it is worth noting that they are applicable to any atom interferometer using a BEC as source.

Before turning to the development of these solution techniques, we return briefly to the discussion at the beginning of this section. We have dramatically reduced the Hilbert space by picking the ansatz in equation (9.2). We have thus surely lost some aspects of the many-body state. In this case, what is lost is correlations. Namely, the n -point density correlation function for this mean-field state will always have the form $|\phi(\vec{x}_1)|^2 \dots |\phi(\vec{x}_n)|^2$. This correlation function is factorisable and so cannot exhibit quantum mechanical fluctuations. That is to say, the reduced correlation function is just

$$\frac{\langle \hat{n}(\vec{x}_1) \hat{n}(\vec{x}_2) \dots \hat{n}(\vec{x}_n) \rangle}{\langle \hat{n}(\vec{x}_1) \rangle \langle \hat{n}(\vec{x}_2) \rangle \dots \langle \hat{n}(\vec{x}_n) \rangle} = 1 . \quad (9.8)$$

For instance, considering the two-point function shows that the density at any point has zero variance: $\langle \hat{n}^2(\vec{x}) \rangle - (\langle \hat{n}(\vec{x}) \rangle)^2 = 0$. This again connects to the idea that an atom interacts with the mean-field: The field has no fluctuations, so the mean is all that the atom can possibly interact with. While this is a heavy loss of quantum mechanical information, it turns out that mean-field theory is incredibly precise for calculating the interaction energy in a BEC with ≥ 1000 atoms.

We can estimate the regime of validity for the mean-field approximation with a simple heuristic. Assuming Gaussian fluctuations, one expects that density fluctuations should be of order \sqrt{N} at any point, whereas the total energy should be of order N . This would make us expect that mean-field theory, much like averaging

methods in classical statistical mechanics, should have a relative precision of order $N^{-1/2}$. We can expect an added benefit for calculations of time integrated quantities such as the accumulated phase, because a single experimental run covers much longer times than the time scale for density fluctuations (a time scale that should, naïvely be related to the chemical potential). Thus, there is, in fact, some averaging out of those fluctuations over the course of a run, so we likely do much better than $N^{-1/2}$ precision. For $N = 1000$, this suggests that mean-field theory is correct at roughly the 3% level, even before accounting for the time-averaging gain.

The break down of mean-field theory is due to the emergence of strong interatomic correlations. Since the mode in which condensation occurs can be described exactly by a state of the form in equation (9.5), correlations can only emerge from the presence of noncondensed atoms. Thus, another way to describe the breakdown point for mean-field theory is the place where a sizable fraction of atoms are uncondensed. Small system size is one way, as described above, for these correlations to become important for the system dynamics.

The other is through strong interactions in a large system. This breakdown is discussed in detail by Castin and Dum[18]. Essentially, the fraction of atoms still uncondensed at zero temperature scales with $\sqrt{na_s^3}$. In fact, they show that, for large N , mean-field theory can be viewed as an expansion in $\sqrt{na_s^3}$, and so mean-field theory is valid for $\sqrt{na_s^3} \ll 1$. The first correction to mean-field theory in this expansion is Bogoliubov-de Gennes theory⁴. For all of the work described in this thesis, the Bogoliubov corrections are irrelevant since, even at the highest, in-trap densities of ^{174}Yb encountered in our experiments, $\sqrt{na_s^3} \approx 0.03$.

With the nature and applicability of mean-field theory established, we now turn to solution techniques of particular value for BEC interferometry.

⁴In line with the discussion of mean-field theory as lacking non-trivial correlations, the Bogoliubov-de Gennes theory represents the addition of non-trivial two-particle correlations to the description.

9.3 Scaling Solutions

9.3.1 Deriving the Scaling Solutions

Many of the results of this section were published in Jamison et al.[47]. In this section we describe an approximate solution to the GPE which is accurate for both small and large interaction strength.

Our scaling solution builds from a previous result valid in the Thomas-Fermi (TF) approximation, derived by Castin and Dum [17]. The TF approximation describes condensates where the interaction energy is much larger than the kinetic energy, by simply ignoring the kinetic energy. The time-independent GPE then reduces to

$$|\phi_{\text{TF}}(\vec{x})|^2 = \mu - U(\vec{x}) . \quad (9.9)$$

This turns out to be an extremely good approximation to the condensate wave function for cold atom BECs.

While the TF approximation may be quite accurate in calculating ground-state wave functions, removing the kinetic energy term from the GPE (or the linear Schrödinger equation) removes all interesting dynamics. In real space, potentials and the nonlinear term can cause space-dependent phase evolution, but they do not allow different points in space to communicate with one another. Only the kinetic term can do this⁵. Thus, a straight-forward extension of the TF approximation to problems in dynamics fails.

Said differently, the kinetic energy term is the one responsible for diffraction. Therefore, the TF approximation ignores the wave-like character of the particle and so should be considered as a semi-classical approximation. To derive a “dynamical Thomas-Fermi” approximation, Castin and Dum explicitly considered the semi-classical nature of the TF approximation. They used an analogy to the expansion of

⁵This can be seen especially clearly by discretizing space and writing the GPE as a set of ordinary differential equations: $\dot{\phi}_j = U_j \phi_j + g|\phi_j|^2 \phi_j - (\phi_{j+1} + \phi_{j-1} - 2\phi_j)/4$, where the index j indicates a specific point in space.

a classical interacting gas to find a self-similar scaling solution to the GPE. We briefly recount their approach, as it gives a satisfying intuition to the problem.

Consider a classical gas in which a particular particle feels a force $\vec{F} = -\nabla[U(\vec{x}) + g\rho(\vec{x})]$, where ρ is the mass density, in analogy to the mean-field Hamiltonian discussed above. The equilibrium ($\vec{F} = 0$) density distribution will match the TF result in equation (9.9). It is thus reasonable to look for inspiration in the dynamics of this system.

For a harmonic potential (even one with time dependence), this system has an exact solution. The solution corresponds to a dynamical rescaling of the initial density distribution:

$$\rho(x_1, x_2, x_3, t) = \frac{\rho\left(\frac{x_1}{\lambda_1}, \frac{x_2}{\lambda_2}, \frac{x_3}{\lambda_3}, 0\right)}{\lambda_1 \lambda_2 \lambda_3},$$

where the λ_j 's are time-dependent scale factors. In addition to simple expansions and contractions, such solutions describe phenomena such as quadrupole oscillations induced by sudden trap changes.

To make scaling solutions work for the quantum case, the local kinetic energy induced by the scaling must be accounted for in the phase. The velocity of a particle would be $v_i = \dot{\lambda}_i r_i$. Recalling that the wave vector in the phase for a plane wave connects to the velocity as $v_i = \hbar k_i / m$, by analogy we expect a phase factor of the form $C \sum_i \dot{\lambda}_i x_i^2$ for some prefactor C that needs to be computed. Thus, we arrive at an ansatz that treats the initial kinetic energy as negligible but correctly accounts for the kinetic energy induced by interactions and changes in U .

Using an alternate derivation we can find scaling solutions that are accurate near the origin (i.e., the center of the BEC) for any interaction strength. In the TF regime ($g \rightarrow \infty$), these solutions reduce to those of Castin and Dum and are exact. In addition, these scaling solutions reproduce the exact solutions for the noninteracting ($g = 0$) gas as well.

The region near the origin is where most signal comes from in an interferometer

since this is the area of highest density as well as greatest phase coherence. For precision interferometry experiments, bringing the mean-field shift as low as possible without compromising the advantages of a BEC is essential, so the residual interaction effects will likely be well outside the regime where TF is applicable. As will be seen below, interaction effects for such BECs can still be large enough to spoil the results of precision interferometers. However, a BEC released from a trap with enough atoms to give strong signals will likely start in or near the TF regime. Therefore, solutions that are good for all interaction strengths in a region near the cloud center are ideal theoretical tools for interferometry. As seen below, scaling solutions naturally give just such tools.

Suppose the initial state has a known form $\phi_0(\vec{x})$. Now, consider an ansatz of the form suggested above by the semi-classical arguments:

$$\phi(\vec{x}, t) = \frac{\phi_0\left(\frac{x_1}{\lambda_1(t)}, \frac{x_2}{\lambda_2(t)}, \frac{x_3}{\lambda_3(t)}\right)}{(\lambda_1(t)\lambda_2(t)\lambda_3(t))^{\frac{1}{2}}} e^{-i\theta(\vec{x}, t)}, \quad (9.10)$$

where λ_i , θ , and ϕ_0 are all real valued. We will not initially assume the form of θ suggested above, but it will in fact follow from the analysis.

We will use the notation $y_i \equiv x_i/\lambda_i$ to simplify the appearance of equations. This gives $\partial/\partial x_i \rightarrow \lambda_i^{-1}\partial/\partial y_i$. Additionally, this means there is some time dependence in $\phi_0(\vec{y})$ due to the λ_i 's.

To proceed, we substitute this ansatz into (9.7). Since all functions are real-valued, we can separate the result into two separate equations for the real and imaginary parts. The equation for the imaginary part gives:

$$\sum_{j=1}^3 \left[\frac{\phi_0}{2} \left(\frac{\dot{\lambda}_j}{\lambda_j} + \frac{\hbar}{m\lambda_j^2} \frac{\partial^2 \theta}{\partial y_j^2} \right) + \frac{\partial \phi_0}{\partial y_j} \left(\frac{\dot{\lambda}_j y_j}{\lambda_j} + \frac{\hbar}{m\lambda_j^2} \frac{\partial \theta}{\partial y_j} \right) \right] = 0. \quad (9.11)$$

Setting each of the two expressions in parentheses equal to zero we find the following condition on θ :

$$\theta(\vec{y}, t) = f(t) - \frac{m}{2\hbar} \sum_{j=1}^3 \dot{\lambda}_j \lambda_j y_j^2, \quad (9.12)$$

which matches the expectation from the semi-classical arguments above.

Using this form for θ , the equation from the real part yields

$$f - \frac{m}{2\hbar} \sum_{j=1}^3 \lambda_j \ddot{\lambda}_j y_j^2 = \frac{g\phi_0^2}{\hbar\lambda_1\lambda_2\lambda_3} - \frac{\hbar}{2m} \sum_{j=1}^3 \frac{1}{\lambda_j^2 \phi_0} \frac{\partial^2 \phi_0}{\partial y_j^2} + \frac{m}{2\hbar} \sum_{j=1}^3 \Omega_j^2 \lambda_j^2 y_j^2, \quad (9.13)$$

where $\Omega_j/(2\pi)$ is the trapping frequency in the j direction. (Recall that we are not requiring the Ω_j 's to be constant.) If (9.13) can be satisfied exactly for a given ϕ_0 , the ansatz gives us an exact solution to the GPE. We will see that the solutions are exact at the two extremes of interaction strength.

In the TF limit, a condensate initially in a harmonic trap gives

$$\phi_0^2(\vec{y}) = \frac{1}{g} \left(\mu - \frac{1}{2} m \sum_{j=1}^3 \Omega_j^2 y_j^2 \right), \quad (g \rightarrow \infty). \quad (9.14)$$

The second term on the righthand side of (9.13) becomes negligible. In this limit, (9.13) can be solved exactly, recovering the original result of Castin and Dum:

$$\ddot{\lambda}_j = \frac{\Omega_j^2(0)}{\lambda_j \lambda_1 \lambda_2 \lambda_3} - \Omega_j^2(t) \lambda_j.$$

As discussed above, this result may be interpreted as the scaling dynamics of a classical gas driven by pressure due to interactions. From this perspective the second term, neglected in the TF limit, can be viewed as the first quantum correction: an extra ‘‘quantum pressure’’ due to wave packet dispersion.

In the opposite limit of a noninteracting BEC (or a single atom) initially in a harmonic trap, ϕ_0 is Gaussian, with the well-known form

$$\phi_0 = (\pi^3 a_1 a_2 a_3)^{-1/4} \exp \left(\frac{1}{2} \sum_{j=1}^3 \left(\frac{x_j}{a_j} \right)^2 \right),$$

where $a_j = \sqrt{\hbar/(m\omega_j)}$. The exact solution to the dynamics after switching off the harmonic potential (i.e., a Gaussian wave packet expanding in free-space) can be found in standard quantum mechanics texts[90]. Adding in the possibility of non-zero, time-dependent trap frequencies leads to the result:

$$\ddot{\lambda}_j = \Omega_j(0)^2 \frac{1}{\lambda_j^3} - \Omega_j^2(t) \lambda_j. \quad (9.15)$$

The key, mathematically, in both cases is that the right hand side of (9.13) has no terms of order higher than y^2 .

From the exactness in the limits of both arbitrarily strong and arbitrarily weak interactions, we expect (9.10) subject to (9.12) and (9.13) to give good results for very small or very large g . To extend this solution to all g values, we enforce the condition on order in y by treating the above equations as an expansion in y . Thus we find solutions valid for all interaction strengths in a region near the center of the condensate.

In addition to the importance of the center of the BEC to interferometry signals discussed above, there is a second reason to desire these centrally valid solutions. Since one primary goal in a precision measurement is to suppress the phase shift due to mean-field effects, setting an upper bound on that shift is useful. Accurately knowing the mean-field phase shift at the center of the condensate, where it will be greatest, sets a tight upper bound.

We expand the right hand side of (9.13) in a power series in y and keep only terms up to order y^2 . Equating the coefficients for the y^0 terms and the y_j^2 terms yields the following set of equations for f and the λ 's:

$$\dot{f} = \frac{g}{\hbar} \frac{\alpha_0}{\lambda_1 \lambda_2 \lambda_3} - \frac{\hbar}{2m} \sum_{k=1}^3 \frac{\alpha_k}{\lambda_k^2} \quad (9.16)$$

$$\alpha_0 \equiv (\phi_0(0))^2 \quad , \quad \alpha_k \equiv \frac{1}{\phi_0(0)} \left. \frac{\partial^2 \phi_0}{\partial y_k^2} \right|_{\vec{y}=0}$$

$$\ddot{\lambda}_j = \frac{g}{m} \frac{-\beta_{0j}}{\lambda_j \lambda_1 \lambda_2 \lambda_3} + \frac{\hbar^2}{2m^2} \frac{1}{\lambda_j} \sum_{k=1}^3 \frac{\beta_{kj}}{\lambda_k^2} \quad (9.17)$$

$$\beta_{0j} \equiv \left. \frac{\partial^2 \phi_0^2}{\partial y_j^2} \right|_{\vec{y}=0} \quad , \quad \beta_{kj} \equiv \left. \frac{\partial^2}{\partial y_j^2} \left[\frac{1}{\phi_0} \frac{\partial^2 \phi_0}{\partial y_k^2} \right] \right|_{\vec{y}=0}$$

where the α 's and β 's are constants calculated from the initial state.

9.3.2 Numerical Simulations

As the intermediate interaction regime has no exact solutions, we tested our approximation by comparison to full numerical simulations of the three-dimensional GPE.

These simulations were coded in MATLAB[®]. Time evolution was performed in momentum space using fourth-order, adaptive Runge-Kutta via the Matlab 'ode45' function. Implementation of spatial potentials and the interaction term required Fourier transforming from momentum space to physical space, computing these terms, then Fourier transforming the computed terms to momentum space for time evolution. Using momentum-space time evolution rather than finite-differences speeds up the calculation of the second derivative term enough to make this seemingly baroque sequence the most efficient way to simulate the GPE in three dimensions⁶. As is true of so many numerical challenges, the absurdly good scaling of fast Fourier transforms (FFTs) makes this algorithm fast. As such, these calculations were always done on grids with number of grid points equal to a power of 2 to maximize the speed of the FFTs.

Initial in-trap states were found by imaginary-time evolution to find the lowest energy steady state. The algorithm was essentially identical to that of the real time evolution described above. However, imaginary-time evolution does not conserve probability. So, after a small imaginary-time step, the condensate wave function was renormalized before taking the next imaginary-time step.

All experiments were simulated on a variety of grids to check numerical convergence. Both size of box and fineness of grid sampling were varied. Use of momentum-space time evolution lead to the use of periodic boundary conditions. Thus, sufficient sampling as well as avoiding interaction with periodic images of the condensate were

⁶Atomic physicists frequently make use of split step algorithms for simulating the GPE. Though I did not write a separate simulation code using such methods to test which is faster, my understanding is that momentum-space time evolution using efficient fourth-order evolution algorithms is essentially as fast but more flexible than split-step methods.

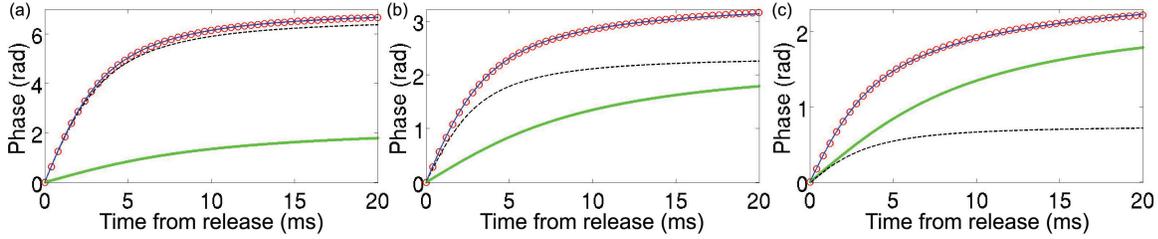


Figure 9.1: Accuracy of scaling solutions at center of BEC. We plot the phase at the center of the condensate as a function of expansion time for (a) high density, (b) intermediate density, and (c) low density. The red circles show results from full numerical simulations, while the solid blue curves show our scaling solutions using the initial condensate wave function calculated in the trap. For reference, the TF scaling solutions are plotted as dashed black curves and the noninteracting solutions are plotted as thick green curves. Neither the TF nor the noninteracting solution is sufficient for high-accuracy measurements at any of these densities.

concerns. Both problems were easily avoided with grids that fit comfortably on commodity personal computers.

To test the scaling laws, the evolution of a ^{174}Yb ($a_s = 5.6$ nm) BEC was simulated. The condensate was formed in a harmonic trap with frequencies $(\Omega_x, \Omega_y, \Omega_z) = 2\pi \times (50, 50, 20)$ Hz. Then, it was allowed to expand for 20 ms, roughly the relevant time-scale for the contrast interferometer.

Figure 9.1 shows the phase at the center of the condensate as a function of time for condensates with $N_{\text{at}} = 10^4$, 10^3 , and 10^2 (corresponding to high, intermediate, and low interactions). For comparison, the results of the TF scaling laws and the noninteracting solution are plotted along with the numerical data and the central scaling laws derived here. In the early stages of the expansion with $N_{\text{at}} = 10^4$, the TF result agrees well with the full numerical solution. This shows the initial density (peak density in trap of $6.2 \times 10^{13} \text{ cm}^{-3}$) falls in the TF regime. However, we see that

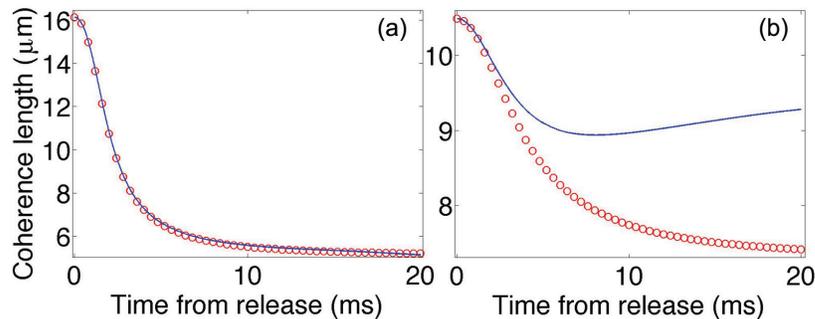


Figure 9.2: Scaling solutions for coherence length. The coherence length is plotted versus expansion time for (a) high density and (b) intermediate density. The red circles show results from full numerical simulations, while the solid blue curves show our scaling solutions using the initial condensate wave function calculated in the trap. The coherence length is not well approximated with the scaling solutions for intermediate densities where the $\mathcal{O}(y^2)$ expansion breaks down.

the TF result begins to diverge from the numerical solution after around 5ms, growing to a substantial deviation at the end of the expansion. This deviation is important for precision measurements. While a condensate may begin an experiment in the TF regime, its subsequent expansion lowers the density, eventually spoiling the accuracy of the TF approximation. For condensates deeper in the TF regime (such as those used in our experiments) this feature may become even more pronounced due to the faster expansion.

It should be noted that this deviation in the phase is more pronounced than the deviation in density profile. The TF scaling laws are shown[17] to agree with the time of flight density distribution for less strongly interacting condensates over even longer times of expansion.

For $N_{\text{at}} = 10^2$ (peak density in trap of $6.5 \times 10^{12} \text{ cm}^{-3}$) the low density makes the TF approximation inaccurate throughout. However, the central scaling solutions derived above maintain good validity. A common suggestion for precision BEC inter-

ferometry is to adiabatically lower the in-trap density before beginning an experiment. The failure of both the TF and the noninteracting solutions to describe the $N_{\text{at}} = 10^2$ results highlights the need to be able to analyze interaction effects, even for seemingly low initial densities.

The middle case of $N_{\text{at}} = 10^3$ (peak density in trap of $2.4 \times 10^{13} \text{ cm}^{-3}$) shows marked departure from both the TF and noninteracting results. The continued agreement with the central scaling solution in the intermediate regime highlights the robustness of this technique.

To further discuss the extent of validity of our central scaling solutions, we now address more global properties. The density and phase profiles across the condensate generally show good agreement between numerical simulation and central scaling solutions for all three interaction regimes. However, these properties individually are not usually of direct experimental importance for atom interferometry. Instead, we consider here the more relevant parameter of coherence length l_c which depends sensitively on both density and phase profiles. For a condensate wave function $\phi(\vec{x})$ the coherence length is defined by:

$$\frac{\int d^3x \phi(\vec{x})^* \phi(\vec{x} \pm (l_c/2)\hat{x}_3)}{\int d^3x \phi(\vec{x})^* \phi(\vec{x})} = \frac{1}{e}. \quad (9.18)$$

For a two-arm interferometer this is directly related to the $1/e$ coherence time of the signal through the relative velocity of the two arms at recombination [38]. There is a similar relationship, up to some numerical factors related to the splitting, for a three-arm interferometer .

As figure 9.2 shows, the scaling solutions reproduce this global property well for high density ($N_{\text{at}} = 10^4$) but show clear (25% in fig 9.2b) deviations for intermediate densities ($N_{\text{at}} = 10^3$). These deviations grow to 50% for the $N_{\text{at}} = 10^2$ case and then decrease sharply to below 5% for $N_{\text{at}} = 10^1$. Given that coherence length is a global quantity, it is unsurprising that our central scaling solutions are inadequate to calculate it, since the central scaling solutions are only accurate near the origin.

In the small interaction case, the deviations fall in line with the relative size of the y^2 and y^4 terms in the small y expansion of the right-hand side of (9.13) found using first-order perturbation theory. This suggests, again unsurprisingly, that the failure to accurately reproduce the coherence length in intermediate densities directly follows from truncating the expansion at $\mathcal{O}(y^2)$. We will see in section 9.6 that central scaling solutions accurately reproduce phase offsets even for the difficult case of a long experiment that moves from deep in the TF regime to weak interactions. Thus, extension of the scaling technique beyond $\mathcal{O}(y^2)$ is unnecessary for adequately accounting for systematic shifts in precision BEC interferometers.

9.4 Slowing-Varying Envelope Approximation

BEC interferometers typically contain widely disparate scales. The laser pulses used to manipulate condensates have durations ranging from hundreds of nanoseconds to tens of microseconds[37], while the entire experiment can last for tens or hundreds of milliseconds. Accurate simulation with such a separation of scales is computationally intensive—in some cases it may be prohibitively intensive. A similar separation of scales in distances arises, comparing the healing length of the condensate, which is of order 100 nm to the separation between interferometer arms, which is of order 100 μm for the experiments reported in this thesis and could grow to order 1 cm for future sub-ppb measurements. To separate these scales computationally, we can decompose the condensate wave function into quasi-independent pieces for each arm. The decomposition technique, known as the slowly-varying envelope approximation (SVEA), is well-known in the context of optics⁷. Its application to BECs was pioneered by Trippenbach et al[98].

The SVEA leverages one of the key experimental advantages of a BEC: its narrow momentum spread. The momentum-space wave function of a condensate typically

⁷The use closest in application to the present case comes in modeling wavelength division multiplexing systems. See the book by Agarwal[1].

has a width well below the laser recoil momentum. Thus, when the condensate wave function is split using a light grating, the momentum-space wave function consists of a series of clearly separated peaks. The SVEA may be applied to any splitting method that creates clearly separated momentum states, and so is applicable to many possible experiments.

We generalize the standard SVEA, finding certain terms dropped in the usual description of the technique can be important at the level of accuracy needed for modeling precision experiments. We also find a description of this modified SVEA as an expansion whose small parameter gives us an estimate of the accuracy of the method.

To start, we postulate a form for the condensate wave function:

$$\phi(\vec{x}, t) = \sum_j \phi_j \left(\vec{x} - \frac{\hbar \vec{k}_j}{m} t, t \right) e^{i \vec{k}_j \cdot \vec{x} - i \omega_j t} \quad (9.19)$$

where the \vec{k}_j 's are the relevant wave vectors and the ω_j 's the corresponding frequencies. We will consider $\vec{k}_j = 2j \vec{k}_{\text{rec}}$ where j is an integer and \vec{k}_{rec} is the laser wave vector. This form corresponds naturally to the language we use to describe an interferometer in terms of arms. After inserting this ansatz we reorganize the GPE into the following suggestive form:

$$\begin{aligned} & \sum_{j=-\infty}^{\infty} i \hbar \frac{\partial \phi_j}{\partial t} e^{i \vec{k}_j \cdot \vec{x} - i \omega_j t} \\ &= \sum_{j=-\infty}^{\infty} \left[-\frac{\hbar^2}{2m} \vec{\nabla}_{\xi_j}^2 \phi_j + g \sum_{l_1, l_2, l_3=-\infty}^{\infty} \phi_{l_1}^* \phi_{l_2} \phi_{l_3} \delta_{j, -l_1 + l_2 + l_3} e^{-i(-\omega_j - \omega_{l_1} + \omega_{l_2} + \omega_{l_3})t} \right] e^{i \vec{k}_j \cdot \vec{x} - i \omega_j t} \quad , \end{aligned} \quad (9.20)$$

where coordinates $\vec{\xi}_j = \vec{x} - \hbar \vec{k}_j t / m$ were chosen for the ϕ 's and $\delta_{a,b}$ is the Kronecker delta function. The choice of coordinates cancels a term of form $\vec{k}_j \cdot \vec{\nabla} \phi_j$, and the usual non-relativistic dispersion relation $\omega = \hbar k^2 / (2m)$ is used to cancel another pair of terms.

Consider the Fourier transform of this equation. If all ϕ_j 's (“envelopes”) have bounded support centered at zero with diameters smaller than one third⁸ of the minimum of $|k_i - k_j|$, then this equation separates exactly into an infinite set of equations of the form

$$i\hbar \frac{\partial \phi_j}{\partial t} = -\frac{\hbar^2}{2m} \vec{\nabla}_{\xi_j}^2 \phi_j + g \sum_{l_2+l_3-l_1=j} \phi_{l_1}^* \phi_{l_2} \phi_{l_3} e^{-i(\omega_{l_2}+\omega_{l_3}-\omega_{l_1}-\omega_j)t} . \quad (9.21)$$

This form differs slightly from that found in the literature [1]. Standard derivations of the SVEA refer to collecting “phase matched” terms into separate equations. In optics, where $\omega_j \propto k_j$, spatial phase matching is equivalent to temporal phase matching. For matter waves, where $\omega_j \propto k_j^2$, the two conditions may be considered separately. Equation (9.21) retains pieces with $\vec{k}_{l_2} + \vec{k}_{l_3} - \vec{k}_{l_1} - \vec{k}_j = 0$ but $\omega_{l_2} + \omega_{l_3} - \omega_{l_1} - \omega_j \neq 0$ (i.e., it enforces spatial but not temporal phase matching). Such pieces can have large enough effects to be important in simulating precision experiments, specifically during the fast dynamics of diffraction pulses.

Given initial conditions satisfying the above criterion, this set of equations is equivalent to the full GPE. To simplify, we select a subset of the envelopes to keep, setting the rest to zero. This truncation affects the envelopes we keep by dropping terms from the nonlinear piece. To understand the limits of this approximation, we want to know how large an error this induces.

Since we only set $\phi_m \equiv 0$ if $\phi_m(t=0) = 0$, for short times (i.e., times in which the modulus of ϕ_m remains much smaller than the modulus of the retained envelopes) we

⁸For a linear equation one expects the condition to be one rather than one third, which is equivalent to the $\phi_j \exp(i\vec{k}_j \cdot \vec{x})$ terms being orthogonal. The stronger condition is necessary to keep momentum separation in the nonlinear term. This can be extended to higher nonlinearities, with a fifth order nonlinearity necessitating the stronger “one fifth” condition on the support of the ϕ_j 's and so on. In fact, these conditions for nonlinear equations are only sufficient initially, as momentum-space mixing can slowly broaden the momentum width of the envelope functions. For our purposes, the BEC gives far better than the factor of one-third in separation, and so there is no concern about mixing the envelopes.

can drop the kinetic term to find

$$i \frac{\partial \phi_m}{\partial t} = \frac{g}{\hbar} \sum_{l_2+l_3-l_1=m} \phi_{l_1}^* \phi_{l_2} \phi_{l_3} e^{-i4\omega_{\text{rec}}t(l_2^2+l_3^2-l_1^2)} .$$

For time scales short enough that the other envelopes do not appreciably move or expand, this can be integrated. The initially unpopulated terms would oscillate with frequency at least $4\omega_{\text{rec}}$. By considering the average growth of ϕ_m , we see that the amplitude will be smaller than that of the initially populated states by the ratio of the mean-field energy to the $4\hbar\omega_{\text{rec}}$. This result gives us a good error bound, since envelopes expanding and moving relative to one another will decrease the right hand side. It also suggests a way to systematically improve the accuracy of the approximation by retaining one or more of the initially unpopulated terms.

To verify these uncertainties, a condensate wave function with an initial superposition of three momentum states (0 and $\pm 2\hbar k_{\text{rec}}$), all populated with equal density, was simulated. These simulations considered Na condensates ($a_s = 2.9$ nm) with $N_{\text{at}} = 10^4$ and trap frequencies $(\Omega_x, \Omega_y, \Omega_z) = 2\pi \times (50, 50, 20)$ Hz. The results from SVEA simulations were compared to results from full simulations of the GPE. Figure 9.4 shows the fractional error in the phase accumulated by the zero momentum branch during separation. The percent error in simulations with both spatial and temporal phase matching fluctuates with a frequency $2 \times (4\omega_{\text{rec}})$. This matches the expectation for error due to neglecting a term $\phi_0^* \phi_1 \phi_{-1} e^{-i(4\omega_{\text{rec}})t(1^2+1^2)}$. The sharp decline in fluctuations just after 0.2 ms corresponds to complete separation of the condensate branches.

The SVEA approach can be used to simulate the effects of external potentials as well. Consider the potentials created by laser standing waves used for splitting and acceleration the arms of the interferometer. Using the light-shift potential formalism, these pulses may be described by a potential of the form

$$V(\vec{x}, t) = A(\vec{x}, t) \left| e^{i\vec{k}_{\text{rec}} \cdot \vec{x}} + e^{-i\vec{k}_{\text{rec}} \cdot \vec{x}} \right|^2 \quad (9.22)$$

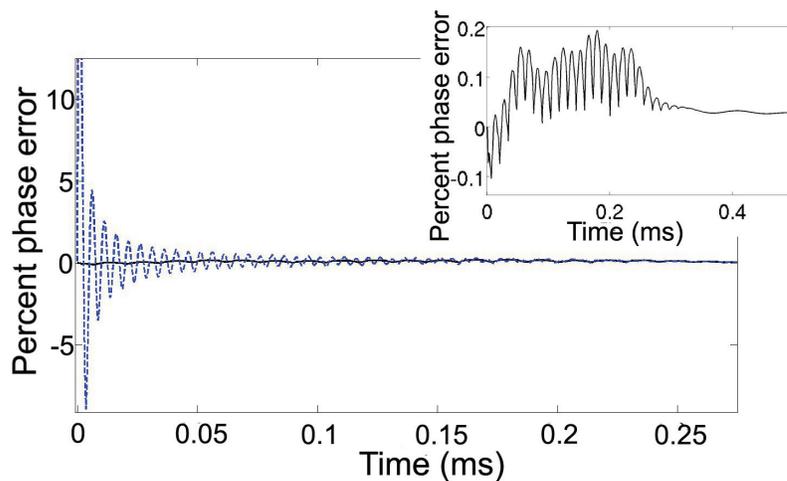


Figure 9.3: Errors in SVEA from temporal phase matching. The phase at the center of the zero-momentum branch of the condensate in SVEA simulations is compared to the result of full simulations of the GPE. The blue dashed curve shows the SVEA with both spatial and temporal phase matching. The black solid curve shows the modified SVEA with only spatial phase matching. The initial peak of the dashed curve rises to 26%, but has been cropped from the image to make other details visible. The inset shows the solid curve on a scale where details are visible.

where A is the amplitude of the potential, which may depend on time and space. The use of relative detunings may create potentials moving relative to the lab frame. In the rest frame of such a potential, the $e^{\pm 2i\vec{k}_{\text{rec}}\cdot\vec{x}}$ pieces connect branches of the condensate wave function with momenta differing by $\pm 2\hbar\vec{k}_{\text{rec}}$. Since dynamics during a laser pulse occur at the time-scale of the recoil frequency, removing the $e^{-i\omega_j t}$ phase is no longer acceptable. Equation (9.21) is modified by the reappearance of the term $(\hbar k_{\text{rec}})^2 \phi_j / (2m)$, which was canceled by the time dependence of $e^{-i\omega_j t}$ in the original derivation. Thus, for modeling laser interactions in the SVEA it is key that temporal phase-matching *not* be enforced.

Accurate simulation of the physics during laser interactions necessitates retaining a number of initially unpopulated condensate branches. For both Raman-Nath and Bragg regime pulses it is sufficient to consider two extra accessible branches on each side of the range you expect to populate (for better than percent-level accuracy of all final wave functions). For instance, a Bragg pulse that takes $2\hbar k_{\text{rec}}$ to $-2\hbar k_{\text{rec}}$ will also require keeping track of the 0 momentum branch, two more states above $2\hbar k_{\text{rec}}$ and two more states below $-2\hbar k_{\text{rec}}$. However, once a laser interaction is complete, the branches that are no longer populated can be removed from the simulation, keeping the number of states tracked from growing during simulation of an experiment with many light gratings.

9.5 Modeling a Complete Experiment: Analytic Results

Since the equations for the envelope functions in the SVEA look essentially like a set of Gross-Pitaevskii equations, the scaling solutions may be combined with the SVEA to reduce the description of a full experiment to a small set of ordinary differential equations. While these equations cannot be solved analytically, they represent a tremendous simplification in terms of numerical simulation. In section 9.6 we will describe the results of such simulations. First, we consider what analytic analysis of these various approximations reveals about the physics of atom-atom interactions in

a contrast interferometer. Readers interested only in intuitive pictures of the physics, not their derivations, should skip ahead to section 9.5.3.

Using the scaling solutions, it is straight-forward to follow the time evolution of the condensate between its release from the trap and the initial splitting pulse. This pulse is, by design, too short for any meaningful dynamics to occur during the pulse. Thus, we may treat it as occurring instantaneously for calculations of the condensate dynamics.

9.5.1 Modified Kinematics and Scaling Behavior

Let the initial condensate wave function in the trap be $\phi_0(x_1, x_2, x_3)$. Further, suppose that the scale factors, $\lambda_1, \lambda_2, \lambda_3$, have been calculated up to the time of the splitting pulse, t_s . The instant before the splitting pulse the condensate wave function is

$$\phi(x_1, x_2, x_3, t_s) = \frac{\phi_0\left(\frac{x_1}{\lambda_1(t_s)}, \frac{x_2}{\lambda_2(t_s)}, \frac{x_3}{\lambda_3(t_s)}\right) e^{i\theta(x_1, x_2, x_3, t_s)}}{(\lambda_1(t_s)\lambda_2(t_s)\lambda_3(t_s))^{1/2}},$$

where

$$\theta(x_1, x_2, x_3, t_s) = \int_0^{t_s} dt \dot{f}(t) - \frac{m}{2\hbar} \sum_{n=1}^3 \frac{\dot{\lambda}_n(t_s)}{\lambda_n(t_s)} x_n^2$$

with

$$\dot{f}(t) = \frac{g}{\hbar} \frac{\phi_0^2(0, 0, 0)}{\lambda_1(t)\lambda_2(t)\lambda_3(t)} - \frac{\hbar}{2m} \sum_{n=1}^3 \left[\frac{1}{\lambda_n^2(t)} \left(\phi_0^{-1} \frac{\partial^2 \phi_0}{\partial x_n^2} \right) \Big|_{\vec{x}=\vec{0}} \right], \quad (9.23)$$

as dictated by the scaling solutions. Then, a splitting pulse that instantaneously diffracts the condensate into three momentum states would give a wave function of the form

$$\phi(x_1, x_2, x_3, t_s) (\alpha_1 e^{ikx} + \alpha_0 - \alpha_{-1} e^{-ikx}),$$

where $k = 2k_{\text{rec}}$. For a pulse that split the momentum states equally, we would have $\alpha_1 = i/\sqrt{3}$, $\alpha_0 = 1/\sqrt{3}$, and $\alpha_{-1} = -i/\sqrt{3}$. This form may now be reduced, using the slowly-varying envelope approximation (SVEA) to a set of coupled equations for the wave functions of the separate momentum states.

We need not confine ourselves to only three states. In fact, depending upon the parameters of the initial splitting, higher diffraction orders may be essential to the proper understanding of interaction shifts. Generically, we begin the SVEA analysis at time $t = t_s$ with initial states $\psi_n(\vec{x}, t_s) = \alpha_n \phi(\vec{x}, t_s)$.

The SVEA equations are (arguments of functions will be suppressed for brevity)

$$i\hbar \frac{\partial \psi_n}{\partial t} = -\frac{\hbar}{2m} \nabla^2 \psi_n + g \sum_{l+m-j=n} \psi_j^* \psi_l \psi_m e^{-i(l^2+m^2-j^2-n^2)\omega}, \quad (9.24)$$

where $\omega = 4\omega_{\text{rec}}$. For instance, if we truncate to include only three states, the interaction term for the ψ_0 equation becomes

$$g (|\psi_0|^2 \psi_0 + 2|\psi_1|^2 \psi_0 + 2|\psi_{-1}|^2 \psi_0 + 2\phi_0^* \phi_1 \phi_{-1} e^{-2\omega t i}). \quad (9.25)$$

If we intend to simulate only the propagation dynamics of the condensate, using separate techniques to simulate the effects of diffractive laser pulses, then we can discard any term for which $l^2 + m^2 - j^2 - n^2 \neq 0$. Thus, from equation (9.25), the final term would be dropped. The discarded terms amount to four-wave mixing terms, which are suppressed at the densities for which precision interferometry is feasible. For instance, the final term of (9.25) corresponds to atoms scattering into or out of the 0 arm by scattering from the atomic density grating made by other arms (e.g., an atom from the -1 arm scattering off of the grating made by the 0 and $+1$ arms into the 0 arm). These processes can add coherently to the effects of laser diffraction and so cannot be ignored in simulations of laser pulses whose temporal length is $\gtrsim 1/\omega$.

Applying these insights to equation (9.24), we obtain a simplified SVEA:

$$i\hbar \frac{\partial \psi_n}{\partial t} = -\frac{\hbar}{2m} \nabla^2 \psi_n + g |\psi_n|^2 \psi_n + 2g \sum_{m \neq n} |\psi_m|^2 \psi_n. \quad (9.26)$$

The equation for each branch has the structure of the GPE with an extra potential generated by the densities of the other momentum states. Therefore, the scaling solutions may be applied to each branch's wave function. Recalling that the scaling solutions are derived as an expansion around the center of the condensate, the densities of the branches need be known only to second order in their arguments.

For the single condensate case, reflection symmetry is a good approximation, and so terms linear in x don't appear. However, as the branches of a split condensate move past each other, the potentials due to inter-branch interaction will have non-negligible linear terms. These linear potentials can be absorbed by allowing the wave vector k_n to have a slow time dependence. Specifically, the left side of equation 9.26 contributes a term $-\hbar\dot{k}_n \cdot \vec{x}\psi_n$, where the dot represents a time derivative. Since we are interested in the center of branch n , at a time t , we set $x - v_n(t - t_s) = 0$. This makes the argument of $\nabla|\psi_m|^2$ equal to $(v_n - v_m)(t - t_s)$. Thus we have

$$\dot{k}_n = -\frac{2g}{\hbar} \sum_m \nabla|\psi_m(y)|^2|_{y=(v_n-v_m)(t-t_s)}. \quad (9.27)$$

This may be straight-forwardly integrated. Note that once all branches have completely separated, the final change in wavenumber is

$$\Delta k_n = -\frac{2g}{\hbar} \sum_m \frac{1}{v_n - v_m} \int_0^\infty d\xi \nabla_\xi |\psi_m(\xi)|^2 = \frac{2g}{\hbar} \sum_m \frac{|\psi_m(\vec{0})|^2}{v_n - v_m}. \quad (9.28)$$

By treating the velocities as constants, we introduce an error in Δk of order $\Delta k/k$. Note that—as may be easily seen in the case of only three branches—the Δk given by equation 9.28 agrees with the value obtained by converting interaction energy into kinetic energy and calculating the resulting wavenumber.

With equations 9.26 and 9.27 we can now determine the important corrections to the single-branch scaling solutions. First, the time varying k_n will lead to a time variation in $\omega_n = \hbar k_n^2/(2m)$. Second, the time-varying k_n will affect the kinematics of the interferometer causing it to close at a later time than anticipated. This may be accounted for by pre-calculating the closing time—the time of perfect overlap for the three arms—and reading out the phase at that time. Alternatively, data taken at the naive closing time, $2T$, may be corrected by taking into account a small shift due to the phase curvature, which will be discussed below.

To fully account for other effects, we must finish combining the scaling solutions with the SVEA. Each branch will feel an effective potential due to each of the other

branches of the condensate. This leads to corrections both to the central phase function $f_n(t)$, but also to the evolution equations for the scale factors. For a given branch, n , we refer to the scale factors as λ_{n1} , λ_{n2} , and λ_{n3} . The scaling form for each branch gives us

$$\psi_n(\vec{\xi}) = \alpha_n \frac{\phi_0 \left(\frac{\xi_{n1}}{\lambda_{n1}}, \frac{\xi_{n2}}{\lambda_{n2}}, \frac{\xi_{n3}}{\lambda_{n3}} \right)}{(\lambda_{n1}\lambda_{n2}\lambda_{n3})^{1/2}} e^{i\theta(\xi_{n1}, \xi_{n2}, \xi_{n3}, t)}, \quad (9.29)$$

where $\xi_{nj} = x_j - v_{nj}t$.

To facilitate using the SVEA, we write equation 9.23 in a more compact form:

$$\dot{f}(t) = \frac{g}{\hbar} \left| \phi(\vec{0}) \right|^2 - \frac{\hbar}{2m} \sum_{n=1}^3 \left(\phi^{-1} \frac{\partial^2 \phi}{\partial x_n^2} \right) \Big|_{\vec{x}=\vec{0}}. \quad (9.30)$$

In the SVEA, this becomes

$$\begin{aligned} \dot{f}_n(t) &= \frac{g}{\hbar} \left| \psi_n(\vec{0}) \right|^2 - \frac{\hbar}{2m} \sum_{j=1}^3 \left(\psi_n^{-1} \frac{\partial^2 \psi_n}{\partial \xi_j^2} \right) \Big|_{\vec{\xi}_n=\vec{0}} + \frac{2g}{\hbar} \sum_{m \neq n} \left| \psi_m(\vec{\xi}_{mn}) \right|^2 \\ &= \frac{g}{\hbar} \frac{\alpha_{n0}}{\lambda_{n1}\lambda_{n2}\lambda_{n3}} - \frac{\hbar}{2m} \sum_{j=1}^3 \frac{\alpha_{nj}}{\lambda_{nj}^2(t)} + \frac{2g}{\hbar} \sum_{m \neq n} \frac{\alpha_{n0m}}{\lambda_{m1}\lambda_{m2}\lambda_{m3}} \end{aligned}$$

where $\vec{\xi}_n = \vec{x} - \vec{v}_n(t - t_s)$ and $\vec{\xi}_{mn} = \vec{\xi}_m - \vec{\xi}_n = (\vec{v}_n - \vec{v}_m)(t - t_s)$. The coefficients are

$$\alpha_{n0} = \phi_0^2(\vec{0}), \quad \alpha_{nj} = \left(\phi^{-1} \frac{\partial^2 \phi}{\partial x_n^2} \right) \Big|_{\vec{x}=\vec{0}}, \quad \text{and} \quad \alpha_{n0m} = \phi_0^2(\vec{\xi}_{mn}) \quad (9.31)$$

Similarly, the inter-branch interactions add an extra term to the equation for the evolution of scale factors:

$$\begin{aligned} \ddot{\lambda}_{nj} &= \lambda_{nj} \frac{\partial^2}{\partial \xi_j^2} \left\{ \frac{g}{m} |\psi_n|^2 + \frac{\hbar}{2m} \psi_n^{-1} \nabla^2 \psi_n \right\} \Big|_{\vec{\xi}=\vec{0}} + \lambda_{nj} \frac{2g}{m} \sum_{m \neq n} \frac{\partial^2}{\partial \xi_j^2} |\psi_m|^2 \Big|_{\vec{\xi}=\vec{\xi}_{mn}} \\ &= \frac{g}{m} \frac{-\beta_{n0j}}{\lambda_{nj}\lambda_{n1}\lambda_{n2}\lambda_{n3}} + \frac{\hbar^2}{2m^2} \frac{1}{\lambda_{nj}} \sum_{k=1}^3 \frac{\beta_{nkj}}{\lambda_{nk}^2} + \frac{g}{m} \sum_{m \neq n} \frac{-\beta_{nm0j}}{\lambda_{mj}\lambda_{m1}\lambda_{m2}\lambda_{m3}}, \end{aligned}$$

where

$$\beta_{n0j} = \frac{\partial^2 \phi_0^2}{\partial x_j^2} \Big|_{\vec{x}=\vec{0}}, \quad \beta_{nkj} = \frac{\partial^2}{\partial x_j^2} \left[\frac{1}{\phi_0} \frac{\partial^2 \phi_0}{\partial x_k^2} \right] \Big|_{\vec{x}=\vec{0}}, \quad \text{and} \quad \beta_{nm0j} = \frac{\partial^2 \phi_0^2}{\partial x_j^2} \Big|_{\vec{x}=\vec{\xi}_{mn}}.$$

In addition to the standard phase evolution due to the interaction energy within a branch, an extra phase evolution due to the interaction energy between branches appears. Similarly, expansion in one branch is accelerated by the curvature of the density in other branches, when they overlap with it.

9.5.2 Effects of Phase Curvature

While the central phase shifts have an obvious effect on the overall phase measured by the interferometer, the phase curvature also affects the final signal. Failure to take data at exactly the closing time of the interferometer adds an extra complication to this analysis. For this analysis, we will assume that the wave functions for the different momentum states have been calculated with the techniques of the previous section.

Let the closing time of the interferometer be $t = 0$. We consider a case with motion only in the x_1 direction. To simplify, we ignore the phase curvature in other directions, though it may be added to the analysis straight-forwardly. The wave function for a contrast interferometer at times near $t = 0$ can then be written as

$$\begin{aligned} \psi(x, t) = & A_1(\vec{x} - \vec{v}t) \exp \left[i \left(kx_1 - f_1(t) - C_1(x_1 - vt)^2 \right) \right] \\ & + A_0(\vec{x}) \exp \left[i \left(-f_0(t) - C_0x_1^2 \right) \right] \\ & + A_{-1}(\vec{x} + \vec{v}t) \exp \left[i \left(-kx_1 - f_{-1}(t) - C_{-1}(x_1 + vt)^2 \right) \right], \end{aligned}$$

where A_j is the real amplitude of the j -momentum branch and f_j and C_j are the corresponding central phase and phase curvature. The signal from a contrast interferometer is a time-varying back-reflection whose amplitude is roughly proportional to the amplitude of the $\cos(kx)$ component of the atomic density. Computing $|\psi|^2$ and keeping only terms whose exponential contains $\pm ikx_1$ yields

$$\begin{aligned} & A_0(\vec{x}) e^{i(f_0(t) + C_0x_1^2)} \left[A_1(\vec{x} - \vec{v}t) \exp \left[i \left(kx_1 - f_1(t) - C_1(x_1 - vt)^2 \right) \right] \right. \\ & \left. + A_{-1}(\vec{x} + \vec{v}t) \exp \left[i \left(-kx_1 - f_{-1}(t) - C_{-1}(x_1 + vt)^2 \right) \right] \right] + \text{h.c.} \end{aligned} \tag{9.32}$$

The contrast interferometer has reflection symmetry, meaning that $A_1 = A_{-1}$, $C_1 = C_{-1}$, and $f_1 = f_{-1}$. For $x_1 < 0$ and $t < 0$, $A_1(x - vt) > A_1(x + vt)$. The reflection symmetry guarantees that the signal contribution for $x_1 > 0$ will equal the contribution from $x_1 < 0$. This simplifies the model for times relatively near to the closing time of the interferometer where $|A_1(x - vt) - A_1(x + vt)| \ll A_1(x + vt)$, allowing for terms proportional to the difference in the moving branch amplitudes to be dropped.

For a situation of nearly even splitting of the initial condensate, we have the further simplifications $A_1 \approx A_0$ and $C_1 \approx C_0$. Then, dropping terms proportional to $|A_1(x - vt) - A_1(x + vt)|$, equation 9.32 reduces to

$$4A_0(\vec{x})A_1(\vec{x} + \vec{v}t) \cos[(k + 2C_1vt)x_1] \cos[f_1(t) - f_0(t) + C_1(vt)^2]. \quad (9.33)$$

The first cosine describes the spatial density modulation leading to the optical Bragg reflection signal. The second cosine describes the time variation of this reflection, from which the recoil frequency is to be extracted. From the scaling solutions we have $C_1 = -(m/2\hbar)\dot{\lambda}_1/\lambda_1$, which makes the argument of the first term $k(1 - t\dot{\lambda}_1/\lambda_1)$, which for our experiments is close enough to k to have negligible effect on the reflection. The argument of the second term becomes $\omega t - \omega t^2\dot{\lambda}_1/(4\lambda_1)$.

To complete the analysis we must compute the closing time of the interferometer. This is essentially just kinematics. Let the time of the center of the mirror pulse be T (now defining the time of the initial splitting pulse as $t = 0$). For a condensate density of the form $A - Bx_1^2$ along the direction of motion, eq. 9.27 leads to

$$\dot{u}_1 = \frac{4g}{m}u_1tB(t)[\theta(vt - R(t)) + 2\theta(2vt - R(t))], u_1(0) = v \quad (9.34)$$

where u_1 is the time-dependent velocity of the branch with initial velocity v , θ is the Heaviside function, and $R(t) = A/B(t)$ is the radius of the condensate. The theta functions' use of vt rather than $\int dt' u_1(t')$ introduces errors of order $\Delta u_1/v \sim 10^{-2}$ in the endpoint of the integral, which make negligible changes to the kinematics. Define

t_{short} as the solution to $2vt = R(t)$ and t_{long} as the solution to $vt = R(t)$. For the time interval $[t_{\text{long}}, T]$, the velocity is fixed at $u_1(t_{\text{long}})$. At time T the velocity changes to $u_1(t_{\text{long}}) - 2v$, so the moving branches return to the center of the interferometer more slowly than they moved away. This pushes the closing time to be later than $2T$. From (9.28) we know that $u_1(t_{\text{long}}) - v \propto \rho$, where ρ is the condensate density at the time of the splitting pulse. For data taken at $2T$, the deviation Δt from the proper closing time is proportional to ρ .

9.5.3 Summary of the Physics

A key factor in the influence of interaction effects on the final interferometer signal is the splitting parameter x , defined such that the densities in the $+1$, 0 , and -1 arms are in the ratio $1 - x : 1 + 2x : 1 - x$. We describe four separate effects. First, there are intra-branch energy shifts. These shifts arise from the interaction energy of a single branch of an atom with the total atomic density in that branch of the interferometer. Recalling the combination of phases that generate the CI signal's phase, equation 5.3, the CI phase will be shifted by the difference between intra-branch energy for the moving versus non-moving branches. The phase shift is proportional to $-x\rho T$, where ρ is the atomic density just before the initial splitting and x is the splitting parameter.

There are similar inter-branch interactions. During the time the branches of the interferometer are overlapped in space, the $2\hbar k$ branch of an atom will interact with the total atomic density in the 0 branch of the BEC. This gives an energy shift analogous to that from intra-branch interactions. Similar shifts arise for all other pairs of momentum states.

Inter-branch interactions are twice as strong as intra-branch interactions. This effect appears in the SVEA from the fact that there are two different ways to get a

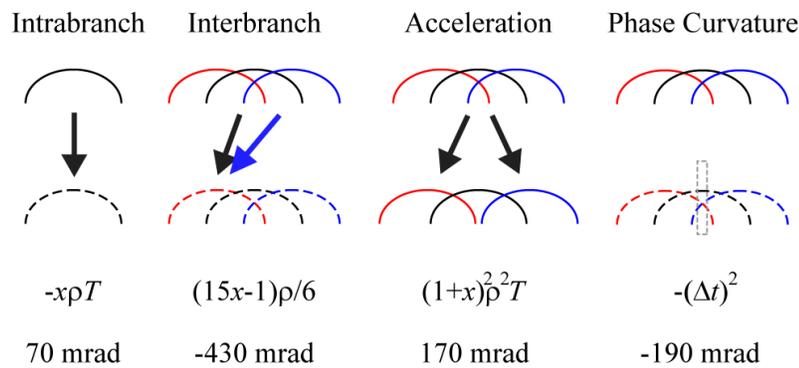


Figure 9.4: Summary of interaction effects. The four important interaction effects are illustrated. Below each illustration are the scaling of the effect with splitting parameter, x , density at time of splitting, ρ , the difference between $2T$ and the time of perfect overlap, Δt , and T . The branches are represented by red, black, and blue lumps. Solid lines show density profiles, and dashed lines show phase profiles. The gray box in the phase curvature illustration highlights that the center of the black $0\hbar k$ branch interferes with the wings of the $\pm 2\hbar k$ branches. Values given to illustrate the relative sizes are for the experiment described in the text.

ϕ_1 and a ϕ_0 from the second and third terms of the product in the nonlinear piece. Physically, this is a manifestation of the distinguishability of the two branches as final states in the quantum scattering problem⁹. The moving branches have a shift proportional to $\rho((1 + 2x) + 0.5(1 - x))/3$ (the 0.5 comes from the moving branches overlapping with one another half the time they overlap with the non-moving branch) while the non-moving branch has a shift proportional to $\rho(2)(1 - 3x)/3$. Thus, the overall effect scales as $\rho(5x/2 - 1/6)$.

The momentum of a moving arm from $t = 0$ to $t = T$ is increased by an amount $\hbar\Delta k \propto (1 + x)\rho$. By the same argument as given in section 7.3.4 for the index of refraction, only the term quadratic in Δk from the total kinetic energy survives. Thus, the total phase shift from this effect is proportional to $(\Delta k)^2 T \propto (1 + x)^2 \rho^2 T$.

A final, less obvious effect of interactions involves the phase curvature across the condensate. Due to the acceleration effect, the branches may not be perfectly overlapped at time $2T$. The actual time of perfect overlap is referred to as the closing time of the interferometer. In this case, the grating is formed by the interference between non-analogous parts of the different branches. In figure 9.4 the gray box highlights the fact that the center of the 0 momentum branch interferes with the wings of the $\pm 2\hbar k$ branches. The phase accumulated due to interactions in the BEC before splitting is curved like the density profile that generates it, as first demonstrated in [38]. Therefore, the phase is greatest in the center of a branch and decreases into the wings. This effect scales like $(\Delta t)^2$, where Δt is the difference in time between the proper closing time and the time when data is taken. This shift can be made negligibly small by taking data at $\Delta t = 0$ rather than at $2T$. Taking data at $\Delta t = 0$ spoils the cancellation of the $2\hbar^2 k \Delta k$ term in the acceleration shift. Thus, a trade-off must be made in deciding between $2T$ and the closing time. For all work described

⁹For bosons, the scattering cross section is twice as large for distinguishable final states as for indistinguishable because the cross-sections for the two permutations of particle and final state add. There is no such ambiguity for indistinguishable final states.

in this thesis, data was taken around $2T$.

9.6 Modeling a Complete Experiment: Numerical Results

In this section we use our theoretical tools to numerically calculate the interaction phase shifts in a contrast interferometer measurement of ω_{rec} . First, we consider comparisons of the SVEA plus scaling solutions to more complete numerical simulations.

Two full experiments were simulated. The first allows us to assess the validity of the SVEA decomposition as compared to full numerical solutions of the GPE. We find that the potentially dramatic reduction of computational cost allowed by the SVEA does not diminish the accuracy of the simulations. The second shows the efficacy of central scaling solutions combined with SVEA in several parameter regimes as compared to three-dimensional SVEA simulations. This involves simulating experiments that could not be adequately simulated on commodity PCs using the full GPE.

First, we simulated a short experiment using a ^{23}Na condensate with 1 ms of free expansion out of the trap ($(\Omega_x, \Omega_y, \Omega_z) = 2\pi \times (50, 50, 20)$ Hz) and $T = 0.2$ ms. This experiment is sufficiently short that full simulations of the GPE may be run in reasonable time. Figure 9.5 shows the signal from the full GPE simulation along with the signal from the SVEA simulation. The signals were generated by extracting the magnitude of the $2\hbar k_{\text{rec}}$ component of the total atomic density as a function of time. The phase at time $2T$ agrees within the granularity of the time steps in the simulated signal¹⁰.

Since the final phase of this signal is sensitive to slight differences between simulations in the free propagation, condensate separation/recombination, or laser interaction periods, the agreement seen in figure 9.5 shows the power of the SVEA to accurately model all periods of an experiment. The actual experiment considers the slope of a ϕ_{sig} vs. T plot in which many of these details are expected to cancel out.

¹⁰The small differences in signal envelope may arise from spurious gratings created by higher momentum states, which we remove from the SVEA simulations.

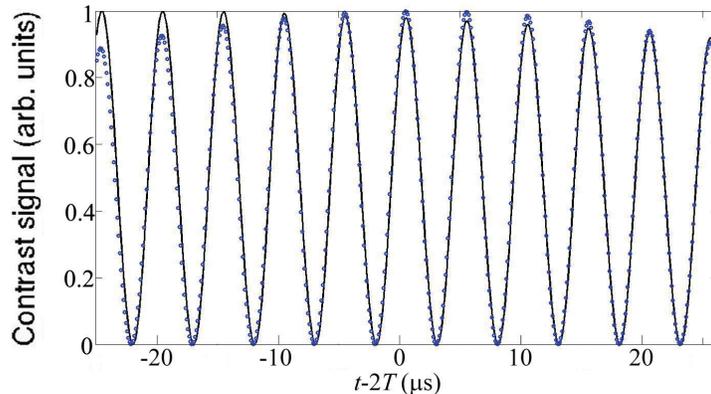


Figure 9.5: Simulated output of a contrast interferometry experiment. The solid (black) curve shows the output signal simulated using the full GPE. The circles (blue) show the output signal simulated using the SVEA.

Being able to accurately model these details can thus improve the confidence in these cancellations, allow modeling of laser intensity fluctuation effects, etc.

Several recent studies of atom-light interactions with varying levels of complexity [95, 51, 12, 73] can easily be adapted to describe the condensate splitting. In the SVEA plus scaling framework, the complexity of the laser interaction model may be changed as needed without substantially affecting the models for the rest of the interferometer. For the simulations described above we used the light-shift potential formalism described in section 9.4. For the following, we use a much simpler model ignoring phases accumulated during the laser interactions. These phases should not affect the final result of an experiment in principle. The effects of diffraction phases discussed in chapter 8 can be handled separately.

The second simulated experiment is a ^{174}Yb contrast interferometer. We used $N_{\text{at}} = 10^4$ and had the condensate initially confined in an isotropic trap with frequency $\Omega = 2\pi \times 200$ Hz. After a 10 ms period of free expansion, the BEC is split into three branches. The densities in arms -1 , 0 , and 1 are in the ratio $(1-x) : (1+2x) : (1-x)$,

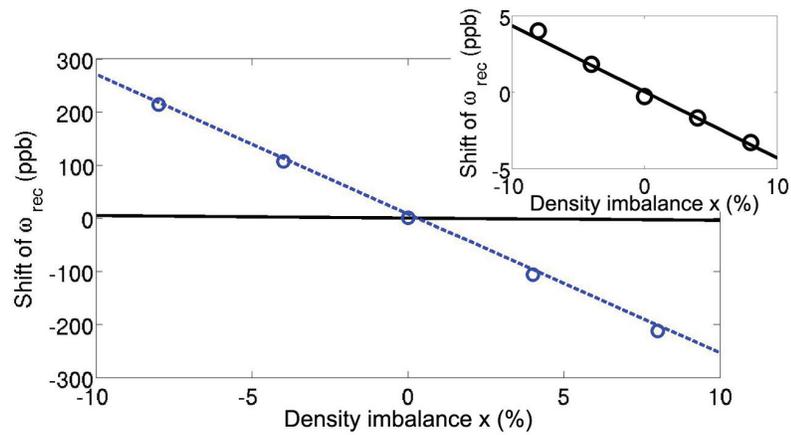


Figure 9.6: Mean-field shifts versus splitting parameters. The black solid line shows the scaling solution value for the shift in measured recoil frequency as a function of density imbalance for the experiment described in the text with $\Omega = 2\pi \times 200$ Hz. The blue dashed line shows the same but for an initial trapping frequency of $\Omega = 2\pi \times 5$ Hz. The circles show data points from SVEA simulations. The inset shows a closer view of the solid line. For imbalance levels of $\leq 2\%$ the mean-field shift is reduced to the 1 ppb level.

using the splitting parameter defined above. Notice that x is the difference in fraction of original density between the stationary branch and the accelerated branches (e.g., density splitting in the ratio 1:2:1 has $x = 0.25$). We make use of our scaling solutions from section 9.3 to quickly probe a large region of parameter space to find suitable conditions a ppb-level experiment.

To apply our scaling solutions, we break a full experiment into stages of free propagation and stages of laser interaction. These stages are simulated using the results of section 9.5.3.

For these simulations, we ignored the back-reaction—the change in shape of each branch due to its interactions with other branches. This approximation gives the first term of an expansion in a small parameter. We can derive its value for the case of release from a trap followed by some period of expansion before the splitting laser pulse. Let the z axis be the direction of laser propagation. The relative size of the neglected back-reaction effects is $[(\Omega_z/\omega_{\text{rec}})(\mu/(\hbar\Omega_z))^3]^{1/2}\lambda_x^{-1}\lambda_y^{-1}$ for a condensate in the TF regime and $[(\Omega_z/\omega_{\text{rec}})(\mu/(\hbar\Omega_z))^2]^{1/2}\lambda_x^{-1}\lambda_y^{-1}$ in the small interaction limit.

With this approximation the entire post-splitting pulse propagation is simulated with a single scaling solution and the inter-branch interaction phase may be added at the end as a perturbation. With the parameters described for this simulation, these make good approximations. Later, in simulating actual experiments, we work with ten times more atoms, higher trap frequencies, and a shorter expansion time. In such a case, the full solutions from section 9.5 must be solved numerically. However, future sub-ppb measurements will use several techniques to reduce density before the splitting pulse. In such a case, adding in inter-branch interactions and their resultant accelerations as small perturbations to a single scaling solution should be a good approximation, at least for choosing trap parameters in the design phase.

Returning to the simulation, we find the fractional shift of the measured recoil frequency due to atomic interactions as a function of x (see figure 9.6). To do so, we calculated the phase of the signal at $t = 2T$, $\phi(2T)$ for runs of $T = 2$ ms and $T = 5$ ms

and then found the slope of $\phi(2T)$ vs. T . The result from the SVEA plus central scaling solution agrees well with the results of three-dimensional SVEA simulations. For $|x| < 0.02$ the mean-field shift contributes at less than the ppb level. We have achieved control of the splitting parameter at the level of ± 0.01 , as will be seen in the data presented below.

While the simulated conditions were somewhat contrived, they achieve densities that should be feasible in the newly built interferometry machine. As such, we see that interactions can indeed be controlled to the ppb-level with our current level of control for the splitting pulse. However, it should also be said that accelerations to $N = 100$, which are the goal for the new machine, turn this ppb-level result into a sub-ppb result, more than adequate for longer-term goals of pushing to 0.1 ppb.

For a given atom number and available total time for a run, the mean-field shift is generally smaller for larger trap frequencies because the condensate expands much more rapidly after release from strong traps than weak traps. The rapid expansion quickly makes up for the higher initial density. However, the more rapid expansion is a direct result of increased momentum spread, potentially diminishing the advantage of a BEC's narrow momentum distribution.

An example considered in [47] used $N_{\text{at}} = 10^4$ at $\Omega = 2\pi \times 200$ Hz. There, the final momentum spread was still less than one tenth of the recoil momentum, while the mean-field shift was well below the ppb level. While this may sound like an appealing avenue, there is, in fact, a strong argument for keeping the momentum spread as low as possible. As seen in section 7.3.4, efficient acceleration and mirror pulses are key both to avoiding atom loss and to avoiding strong shifts due to index of refraction. The best course of action, considering all factors, is to reduce the final trap frequencies and density as far as possible, understanding that the interaction effects may still have to be carefully modeled and subtracted to achieve the desired precision.

Finally, we compare these models to data from our contrast interferometer. We

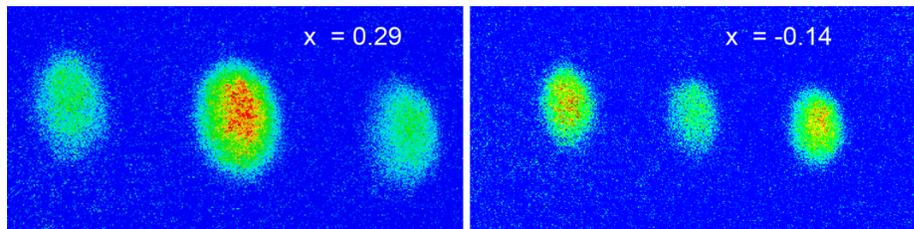


Figure 9.7: Different splittings. The density distributions produced by two different splitting parameters— $x = 0.29 \pm 0.01$ and $x = -0.14 \pm 0.01$ —are shown in these absorption images (average of three images, each). Mean atom number dropped from 1.3×10^5 to 1.0×10^5 between the two sets.

created a dataset specifically to maximize the effects of interactions by using $T = 11$ ms and only allowing the BEC to expand for 2 ms before the splitting pulse to magnify the interaction effects. We then compare two data sets differing only in their density splittings, one with $x = 0.29 \pm 0.01$ and another with $x = -0.14 \pm 0.01$. These splittings are displayed in figure 9.7. The interaction effects for these BECs, having $\approx 10^5$ atoms per run, can be calculated using the full SVEA plus scaling solutions formalism. While this requires numerically integrating a small number of coupled nonlinear ordinary differential equations, shifts may be calculated for each run of the interferometer at a rate of a couple seconds per run on a commodity PC.

Before applying interaction shift corrections, the phase difference between these two data sets was 0.70 ± 0.03 rad. After applying the corrections, the difference dropped to 0.02 ± 0.1 rad. The large error bar is due to uncertainty in trap geometry and turn-off. For future experiments in the new interferometry machine, trap parameters can be both better controlled and better measured, reducing the uncertainty. Together with longer expansion times, these enable correction of interaction effects at the < 5 mrad level. For the design goals of $N = 100$ and $T = 5$ ms, this is safely below the ppb level and can likely be pushed below the 0.1 ppb level in the long term.

We thus have a set of tools for understanding and calculating interaction effects in BEC interferometers. While we have applied them exclusively to contrast interferometers in the interest of brevity, it should be noted that they may also be applied to any other geometry of condensate interferometer. By adding in the terms of the scaling solutions for trapping potentials, these techniques may even be applied to confined atom interferometers such as those of Hughes and Close[13, 67].

Appendix A

**EFFECTS OF GRAVITATIONAL CONFINEMENT ON
BOSONIC ASYMMETRIC DARK MATTER IN STARS**

This appendix presents the paper *Effects of gravitational confinement on bosonic asymmetric dark matter in stars* published in (and copyright to) Physical Review D[46]. The paper corrects some mistakes in the literature on asymmetric dark matter specifically related to the formation of Bose-Einstein condensates of such matter. The essential insight added by this paper is that the gravitational potential in the center of a neutron star forms a harmonic trap for such dark matter, modifying the density of states. Additionally, the relativistic corrections to this potential are found to be quite substantial.

Effects of gravitational confinement on bosonic asymmetric dark matter in stars

Alan O. Jamison*

Department of Physics, University of Washington, Seattle, Washington 98195, USA
(Received 25 April 2013; published 5 August 2013)

Considering the existence of old neutron stars puts strong limits on the dark matter-nucleon cross section for bosonic asymmetric dark matter. Key to these bounds is formation of a Bose-Einstein condensate (BEC) of the asymmetric dark matter particles. We consider the effects of the host neutron star's gravitational field on the BEC transition. We find this substantially shifts the transition temperature and so strengthens the bounds on cross section. In particular, for the well-motivated mass range of 5–15 GeV, we improve previous bounds by an order of magnitude.

DOI: [10.1103/PhysRevD.88.035004](https://doi.org/10.1103/PhysRevD.88.035004)

PACS numbers: 95.35.+d, 14.80.-j, 26.60.Kp, 97.60.Jd

I. INTRODUCTION

While the existence of dark matter is at present well established, the nature of this dark matter is still a matter of considerable speculation [1]. Constraints from precision cosmology [2] and observations of the Bullet Cluster [3] give a picture of the generic phenomenology that dark matter models must satisfy. Direct detection experiments have also placed constraints on the interactions of the dark matter with baryons. Recently, attempts have been made to use astrophysical observations to further constrain the nature of the dark matter. However, such constraints tend to be strongly model dependent.

One particularly well-motivated class of dark matter models are known as asymmetric dark matter (ADM) models [4,5]. These models attempt to explain the similarity of dark matter and baryonic densities $\Omega_{\text{DM}} \approx 5\Omega_{\text{baryon}}$ by postulating a mechanism for transferring the baryonic asymmetry of the standard model sector to the dark matter sector and visa versa. Such models favor a mass of the dark matter particles in the range of 5–15 GeV [6]. The asymmetry in the dark sector implies that dark matter particles are stable and lack antiparticles with which to annihilate.

This stability allows for interesting bounds to be derived from the existence of old neutron stars, particularly for bosonic dark matter [7,8]. The stability of ADM allows a neutron star to slowly accumulate a cloud of ADM particles by scattering them in such a way as to reduce their speed below the neutron star's escape speed. Once captured in the neutron star's gravitational field, repeated scattering between the star and the ADM particles allows the ADM to thermalize with the star. Once a sufficient amount of dark matter has accumulated, the ADM cloud can become self-gravitating and collapse into a black hole. This black hole then consumes the host neutron star. Bosonic ADM can undergo a phase transition to a Bose-Einstein condensate (BEC) state. The BEC, having higher

density than a thermal gas cloud, considerably increases the strength of the bounds obtained by this method.

However, all analyses of this situation, to date, have treated the thermal ADM cloud as homogeneous. In this note, we consider the shortcomings of this approximation for a cloud of noninteracting bosons with a radius much smaller than that of the neutron star. Taking account of the effects of the confining gravitational potential improves the strength of the bounds for all masses up to 15 GeV. The improvement is particularly strong, a factor of an order of magnitude, in the well-motivated mass range of $\approx 5\text{--}15$ GeV.

II. EFFECTS OF CONFINEMENT

The well-known Newtonian theory of gravitation in a uniform density sphere suggests that the ADM cloud within the neutron star should experience a harmonic potential energy given by $V(r) = (2\pi G\rho m/3)r^2$. Calculations [9], subsequently verified by experiments with dilute alkali gases [10], show the density of states for a gas trapped in a harmonic potential is markedly different from that of a homogeneous gas. Specifically, for a gas confined in an isotropic harmonic potential the density of states, $g(\epsilon)$, at energy ϵ is given by $g(\epsilon) = \epsilon^2/(2(\hbar\omega)^3)$, where ω is the angular frequency of the harmonic oscillator potential. This may be compared to the homogeneous gas result $g(\epsilon) = \sqrt{\epsilon}(2\pi V)(2m/h^2)^{3/2}$, where m is the dark matter particle mass and V is the volume occupied by the gas.

For a fixed temperature T , the critical number N_c in the harmonic trap for which a BEC begins to form is $N_c = \zeta(3)(kT/\hbar\omega)^3$, with $\zeta(x)$ the Riemann zeta function and $\zeta(3) \approx 1.2$. Using the Newtonian potential above we find a critical number

$$N_{\text{sho}} = \zeta(3) \left(\frac{kT}{\hbar\sqrt{4\pi G\rho/3}} \right)^3.$$

Comparing this to the value N_{hom} , found using the homogeneous gas in a volume determined by gravitational binding in [7,8], we see the same temperature scaling but a larger prefactor:

*jamisona@uw.edu

$$\frac{N_{\text{hom}}}{N_{\text{sho}}} = \sqrt{\frac{6}{\pi}} \frac{\zeta(3/2)}{\zeta(3)} = 3.0.$$

Using the Newtonian potential for a relativistic system such as a neutron star, particularly near the core, misses important effects. We therefore derive the metric near the center of the star and from it derive an effective gravitational potential energy. To do so, we model the star as spherically symmetric with a metric of the form

$$ds^2 = -e^{2\Phi} dt^2 + e^{2\lambda} dr^2 + r^2 d\Omega,$$

$$e^{2\lambda} \equiv \left(1 - \frac{2Gm(r)}{r}\right)^{-1},$$

where $m(r)$ is the mass contained within a sphere of radial coordinate r . For the temperatures of interest, $T \approx 10^5$ – 10^7 K at the neutron star core, the thermal ADM cloud should be localized within a sphere of radius roughly 1 m around the center of the neutron star. Thus it is reasonable to treat the density as uniform in the region of interest. This gives us $m(r) = 4\pi\rho r^3/3$, making $g_{rr} = 1 - 8\pi r^2\rho/3$. Using $\rho = 1.5 \times 10^{15}$ g/cm³ shows that for $r \leq 1$ m, $|g_{rr} - 1| \leq 10^{-8}$. Thus, we can neglect the curvature of the spatial components of the metric and treat the radial coordinate as the proper length, simplifying the analysis. This allows us to treat the system via an effective Newtonian potential Φ . Using the Oppenheimer-Volkoff equation [11] we find

$$\frac{d\Phi}{dr} = \frac{Gm}{r^2} \left(1 + \frac{4\pi Pr^3}{m}\right) \left(1 - \frac{2Gm}{r}\right)^{-1}$$

$$= \frac{4\pi}{3} G\rho r \left(1 + 3\frac{P}{\rho}\right),$$

where P is the pressure at the star's core.

From this expression, we see that the Newtonian analysis can be carried over by simply replacing ρ with $\rho + 3P$.

The pressure depends sensitively on the neutron star equation of state. Recently, observations of x-ray emission from neutron stars have progressed to the point where good approximations for this equation of state can be had [12,13]. The core density used above corresponds to $\rho \approx 0.9$ GeV/fm³ with a pressure of $P \approx 0.3$ GeV/fm³. The relativistic analysis therefore gives a gravitational potential roughly twice as strong as the Newtonian analysis. This further decreases the critical number for the condensation transition. Our final result for the critical number is thus

$$N_{\text{sho}} = \zeta(3) \left(\frac{kT}{\hbar\sqrt{4\pi G(\rho + 3P)/3}} \right)^3.$$

For the mass range 5–15 GeV the constraints on dark matter-nucleon scattering are dominated by the critical number. As such, using the corrected expression above strengthens the bounds in this range by roughly a factor of 8 as compared to those from Ref. [8]. Further refinements to the neutron star equation of state are anticipated in the near future, based on discrepancies between complementary techniques for deriving it from observation [13,14] as well as the accumulations of further data. Bounds on dark matter cross sections may be easily updated in light of such new information, based on the above results.

ACKNOWLEDGMENTS

I thank A. Nelson and S. Reddy for helpful discussions and readings of the manuscript. This work was supported by a NIST Precision Measurement Grant and the Department of Energy under Grant No. DE-FG02-96ER40956.

-
- [1] G. Bertone, D. Hooper, and J. Silk, *Phys. Rep.* **405**, 279 (2005).
 - [2] X. Wang, M. Tegmark, and M. Zaldarriaga, *Phys. Rev. D* **65**, 123001 (2002).
 - [3] D. Clowe, M. Brada, A.H. Gonzalez, M. Markevitch, S.W. Randall, C. Jones, and D. Zaritsky, *Astrophys. J.* **648**, L109 (2006).
 - [4] S. Nussinov, *Phys. Lett.* **165B**, 55 (1985).
 - [5] D. B. Kaplan, *Phys. Rev. Lett.* **68**, 741 (1992).
 - [6] D. E. Kaplan, M. A. Luty, and K. M. Zurek, *Phys. Rev. D* **79**, 115016 (2009).
 - [7] C. Kouvaris and P. Tinyakov, *Phys. Rev. Lett.* **107**, 091301 (2011).
 - [8] S. D. McDermott, H.-B. Yu, and K. M. Zurek, *Phys. Rev. D* **85**, 023519 (2012).
 - [9] V. Bagnato, D. E. Pritchard, and D. Kleppner, *Phys. Rev. A* **35**, 4354 (1987).
 - [10] J. R. Ensher, D. S. Jin, M. R. Matthews, C. E. Wieman, and E. A. Cornell, *Phys. Rev. Lett.* **77**, 4984 (1996).
 - [11] S. L. Shapiro and S. A. Teukolsky, *Black Holes, White Dwarfs, and Neutron Stars* (Wiley Interscience, New York, 1983).
 - [12] F. Özel, G. Baym, and T. Güver, *Phys. Rev. D* **82**, 101301 (2010).
 - [13] A. W. Steiner, J. M. Lattimer, and E. F. Brown, *Astrophys. J.* **765**, L5 (2013).
 - [14] S. Guillot, M. Servillat, N. Webb, and R. Rutledge, [arXiv:1302.0023](https://arxiv.org/abs/1302.0023).

BIBLIOGRAPHY

- [1] G. Agarwal. *Nonlinear Fiber Optics*. Academic Press, Burlington, MA, 2006.
- [2] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269(5221):198–201, 1995.
- [3] Larry C. Andrews. *Special Functions for Engineers and Applied Mathematicians*. MacMillan Publishing Company, New York, NY, 1985.
- [4] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. Observation of Interference Between Two Bose Condensates. *Science*, 275(5300):637–641, 1997.
- [5] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Tenth-Order QED Contribution to the Electron $g-2$ and an Improved Value of the Fine Structure Constant. *Phys. Rev. Lett.*, 109:111807, Sep 2012.
- [6] Stephen M. Barnett and Rodney Loudon. The enigma of optical momentum in a medium. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1914):927–939, 2010.
- [7] B. Barrett, A. Carew, S. Beattie, and A. Kumarakrishnan. Measuring the atomic recoil frequency using a modified grating-echo atom interferometer. *Phys. Rev. A*, 87:033626, Mar 2013.
- [8] J.-L. Basdevant and J. Dalibard. *The Quantum Mechanics Solver, 2e*. Springer, 2006.
- [9] Rémy Battesti, Pierre Cladé, Saïda Guellati-Khélifa, Catherine Schwob, Benoît Grémaud, François Nez, Lucile Julien, and François Biraben. Bloch Oscillations of Ultracold Atoms: A Tool for a Metrological Determination of h/m_{Rb} . *Phys. Rev. Lett.*, 92:253001, Jun 2004.
- [10] K. Bely. Experimental constraints on the polarizabilities of the $6s^2 \ ^1S_0$ and $6s6p^3 \ ^1P_0^o$ states of Yb. *Phys. Rev. A*, 86:022521, Aug 2012.

- [11] Rym Bouchendira, Pierre Cladé, Saïda Guellati-Khélifa, François Nez, and François Biraben. New Determination of the Fine Structure Constant and Test of the Quantum Electrodynamics. *Phys. Rev. Lett.*, 106:080801, Feb 2011.
- [12] M. Büchner, R. Delhuille, A. Miffre, C. Robilliard, J. Vigué, and C. Champenois. Diffraction phases in atom interferometers. *Phys. Rev. A*, 68:013607, 2003.
- [13] J. H. T. Burke and C. A. Sackett. Scalable Bose-Einstein-condensate Sagnac interferometer in a linear trap. *Phys. Rev. A*, 80:061603, Dec 2009.
- [14] 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.
- [15] Gretchen K. Campbell, Aaron E. Leanhardt, Jongchul Mun, Micah Boyd, Erik W. Streed, Wolfgang Ketterle, and David E. Pritchard. Photon Recoil Momentum in Dispersive Media. *Phys. Rev. Lett.*, 94:170403, May 2005.
- [16] O. Carnal and J. Mlynek. Young’s double-slit experiment with atoms: A simple atom interferometer. *Phys. Rev. Lett.*, 66:2689–2692, May 1991.
- [17] Y. Castin and R. Dum. Bose-Einstein condensation in time dependent traps. *Phys. Rev. Lett.*, 77:5315–5319, 1996.
- [18] Y. Castin and R. Dum. Low-temperature Bose-Einstein condensates in time-dependent traps: Beyond the U(1) symmetry-breaking approach. *Phys. Rev. A*, 57:3008–3021, Apr 1998.
- [19] Sheng-wei Chiow, Tim Kovachy, Hui-Chun Chien, and Mark A. Kasevich. $102\hbar k$ Large Area Atom Interferometers. *Phys. Rev. Lett.*, 107:130403, Sep 2011.
- [20] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Photons and Atoms: Introduction to Quantum Electrodynamics*. John Wiley, New York, NY, 1989.
- [21] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-Photon Interactions: Basic Processes and Applications*. John Wiley, New York, NY, 1992.
- [22] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.*, 81:1051–1129, Jul 2009.

- [23] Franco Dalfovo, Stefano Giorgini, Lev P. Pitaevskii, and Sandro Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.*, 71:463–512, Apr 1999.
- [24] J. Dalibard. Collisional dynamics of ultra-cold atomic gases. In M. Inguscio, S. Stringari, and C. Wieman, editors, *Proceedings of the International School of Physics Enrico Fermi, Course CXL: Bose – Einstein condensation in gases, Varenna*. 1998.
- [25] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein Condensation in a Gas of Sodium Atoms. *Phys. Rev. Lett.*, 75:3969–3973, Nov 1995.
- [26] J. E. Debs, P. A. Altin, T. H. Barter, D. Döring, G. R. Dennis, G. McDonald, R. P. Anderson, J. D. Close, and N. P. Robins. Cold-atom gravimetry with a Bose-Einstein condensate. *Phys. Rev. A*, 84:033610, Sep 2011.
- [27] B. DeMarco, J. L. Bohn, J. P. Burke, M. Holland, and D. S. Jin. Measurement of p-Wave Threshold Law Using Evaporatively Cooled Fermionic Atoms. *Phys. Rev. Lett.*, 82:4208–4211, May 1999.
- [28] T. Donner, S. Ritter, T. Bourdel, A.ttl, M. Khl, and T. Esslinger. Critical Behavior of a Trapped Interacting Bose Gas. *Science*, 315(5818):1556–1558, 2007.
- [29] K. Enomoto, K. Kasa, M. Kitagawa, and Y. Takahashi. Optical Feshbach Resonance Using the Intercombination Transition. *Phys. Rev. Lett.*, 101:203201, Nov 2008.
- [30] Richard P. Feynman. *QED: The Strange Theory of Light and Matter*. Princeton University Press, Princeton, NJ, 1985.
- [31] C. J. Foot. *Atomic Physics*. Oxford University Press, Oxford, UK, 2005.
- [32] O. Garcia, B. Deissler, K. J. Hughes, J. M. Reeves, and C. A. Sackett. Bose-Einstein-condensate interferometer with macroscopic arm separation. *Phys. Rev. A*, 74:031601, Sep 2006.
- [33] K. Gottfried and T-M. Yan. *Quantum Mechanics: Fundamentals*. Springer, New York, NY, 2003.

- [34] David J. Griffiths. *Introduction to Quantum Mechanics*. Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [35] Particle Data Group. Review of Particle Physics. *Phys. Rev. D*, 86:010001, Jul 2012.
- [36] 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.
- [37] S. Gupta, A. E. Leanhardt, A. D. Cronin, and D. E. Pritchard. Coherent manipulation of atoms with standing light waves. *CR Acad. IV*, 2:479–495, 2001.
- [38] E. W. Hagley, L. Deng, M. Kozuma, M. Trippenbach, Y. B. Band, M. Edwards, M. Doery, P. S. Julienne, K. Helmerson, S. L. Rolston, and W. D. Phillips. Measurement of the Coherence of a Bose-Einstein Condensate. *Phys. Rev. Lett.*, 83:3112–3115, Oct 1999.
- [39] D. Hanneke, S. Fogwell, and G. Gabrielse. New Measurement of the Electron Magnetic Moment and the Fine Structure Constant. *Phys. Rev. Lett.*, 100:120801, Mar 2008.
- [40] Anders H. Hansen. *Interacting Quantum Gases of Lithium and Ytterbium*. PhD thesis, University of Washington, 2013.
- [41] Anders H. Hansen, Alexander Khramov, William H. Dowd, Alan O. Jamison, Vladyslav V. Ivanov, and Subhadeep Gupta. Quantum degenerate mixture of ytterbium and lithium atoms. *Phys. Rev. A*, 84:011606, Jul 2011.
- [42] Anders H. Hansen, Alexander Y. Khramov, William H. Dowd, Alan O. Jamison, Benjamin Plotkin-Swing, Richard J. Roy, and Subhadeep Gupta. Production of quantum-degenerate mixtures of ytterbium and lithium with controllable interspecies overlap. *Phys. Rev. A*, 87:013615, Jan 2013.
- [43] Joel Moses Hensley. *A Precision Measurement of the Fine Structure Constant*. PhD thesis, Stanford University, 2001.
- [44] Harald F. Hess. Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen. *Phys. Rev. B*, 34:3476–3479, Sep 1986.

- [45] Vladyslav V. Ivanov, Alexander Khramov, Anders H. Hansen, William H. Dowd, Frank Münchow, Alan O. Jamison, and Subhadeep Gupta. Sympathetic Cooling in an Optically Trapped Mixture of Alkali and Spin-Singlet Atoms. *Phys. Rev. Lett.*, 106:153201, Apr 2011.
- [46] Alan O. Jamison. Effects of gravitational confinement on bosonic asymmetric dark matter in stars. *Phys. Rev. D*, 88:035004, Aug 2013.
- [47] Alan O. Jamison, J. Nathan Kutz, and Subhadeep Gupta. Atomic interactions in precision interferometry using Bose-Einstein condensates. *Phys. Rev. A*, 84:043643, Oct 2011.
- [48] G.-B. Jo, Y. Shin, S. Will, T. A. Pasquini, M. Saba, W. Ketterle, D. E. Pritchard, M. Vengalattore, and M. Prentiss. Long Phase Coherence Time and Number Squeezing of Two Bose-Einstein Condensates on an Atom Chip. *Phys. Rev. Lett.*, 98:030407, Jan 2007.
- [49] Mark Kasevich and Steven Chu. Atomic interferometry using stimulated Raman transitions. *Phys. Rev. Lett.*, 67:181–184, Jul 1991.
- [50] David W. Keith, Christopher R. Ekstrom, Quentin A. Turchette, and David E. Pritchard. An interferometer for atoms. *Phys. Rev. Lett.*, 66:2693–2696, May 1991.
- [51] W. Ketterle and S. Inouye. Collective enhancement and suppression in dilute Bose-Einstein condensates. *arxiv:cond-mat/0101424 Cargese Summer School on Bose-Einstein condensates and atom lasers*, 2001.
- [52] Alexander Khramov. *Experiments in the Ultracold Lithium - Ytterbium System*. PhD thesis, University of Washington, 2013.
- [53] Alexander Y. Khramov, Anders H. Hansen, Alan O. Jamison, William H. Dowd, and Subhadeep Gupta. Dynamics of Feshbach molecules in an ultracold three-component mixture. *Phys. Rev. A*, 86:032705, Sep 2012.
- [54] K. v. Klitzing, G. Dorda, and M. Pepper. New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance. *Phys. Rev. Lett.*, 45:494–497, Aug 1980.
- [55] Nikita Kostylev, Eugene Ivanov, Michael E. Tobar, and John J. McFerran. Sub-Doppler cooling of ytterbium with the $^1S_0 - ^1P_1$ transition including ^{171}Yb ($I=1/2$). *arXiv:1405.1161*, 2014.

- [56] E. Krüger, W. Nistler, and W. Weirauch. Determination of h/m_n . *Metrologia*, 22(3):172, 1986.
- [57] E. Krüger, W. Nistler, and W. Weirauch. Determination of the fine-structure constant by a precise measurement of h/m_n . *Metrologia*, 32(2):117, 1995.
- [58] E. Krüger, W. Nistler, and W. Weirauch. Re-evaluation of a precise measurement of h/m_n . *Metrologia*, 36(2):147, 1999.
- [59] Shau-Yu Lan, Pei-Chen Kuan, Brian Estey, Damon English, Justin M. Brown, Michael A. Hohensee, and Holger Müller. A Clock Directly Linking Time to a Particle's Mass. *Science*, 339(6119):554–557, 2013.
- [60] S. Laporta and E. Remiddi. The analytical value of the electron $(g - 2)$ at order α^3 in QED. *Physics Letters B*, 379(1?4):283 – 291, 1996.
- [61] Anthony J. Leggett. Bose-Einstein condensation in the alkali gases: Some fundamental concepts. *Rev. Mod. Phys.*, 73:307–356, Apr 2001.
- [62] Peter Lepage. How to Renormalize the Schrödinger Equation. *arXiv:nucl-th/9706029*, 1997.
- [63] John H. Lincoln. The Puget Sound Model: Summary. http://exhibits.pacsci.org/puget_sound/PSSummary.html, (retrieved June 2014), 2000.
- [64] L. Marton, J. Arol Simpson, and J. A. Suddeth. Electron Beam Interferometer. *Phys. Rev.*, 90:490–491, May 1953.
- [65] R. Maruyama, R. H. Wynar, M. V. Romalis, A. Andalkar, M. D. Swallows, C. E. Pearson, and E. N. Fortson. Investigation of sub-Doppler cooling in an ytterbium magneto-optical trap. *Phys. Rev. A*, 68:011403, Jul 2003.
- [66] Naoto Masuhara, John M. Doyle, Jon C. Sandberg, Daniel Kleppner, Thomas J. Greytak, Harald F. Hess, and Greg P. Kochanski. Evaporative Cooling of Spin-Polarized Atomic Hydrogen. *Phys. Rev. Lett.*, 61:935–938, Aug 1988.
- [67] G. D. McDonald, H. Keal, P. A. Altin, J. E. Debs, S. Bennetts, C. C. N. Kuhn, K. S. Hardman, M. T. Johnsson, J. D. Close, and N. P. Robins. Optically guided linear Mach-Zehnder atom interferometer. *Phys. Rev. A*, 87:013632, Jan 2013.

- [68] 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, Nov 2013.
- [69] G. D. McDonald, C. C. N. Kuhn, K. S. Hardman, S. Bennetts, P. J. Everitt, P. A. Altin, J. E. Debs, J. D. Close, and N. P. Robins. Bright Solitonic Matter-Wave Interferometer. *Phys. Rev. Lett.*, 113:013002, Jul 2014.
- [70] J.A. Mignaco and E. Remiddi. Fourth-order vacuum polarization contribution to the sixth-order electron magnetic moment. *Il Nuovo Cimento A Series 10*, 60(4):519–529, 1969.
- [71] Peter J. Mohr, Barry N. Taylor, and David B. Newell. CODATA recommended values of the fundamental physical constants: 2010. *Rev. Mod. Phys.*, 84:1527–1605, Nov 2012.
- [72] Olivier Morice, Yvan Castin, and Jean Dalibard. Refractive index of a dilute Bose gas. *Phys. Rev. A*, 51:3896–3901, May 1995.
- [73] H. Müller, S-w. Chiow, and S. Chu. Atom-wave diffraction between the Raman-Nath and the Bragg regime: Effective Rabi frequency, losses, and phase shifts. *Phys. Rev. A*, 77:023609, 2008.
- [74] Holger Müller, Sheng-wei Chiow, Quan Long, Sven Herrmann, and Steven Chu. Atom Interferometry with up to 24-Photon-Momentum-Transfer Beam Splitters. *Phys. Rev. Lett.*, 100:180405, May 2008.
- [75] 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.
- [76] Abraham Pais. *Inward Bound: Of Matter and Forces in the Physical World*. Oxford University Press, Oxford, UK, 1986.
- [77] R. K. Pathria. *Statistical Mechanics, 2e*. Elsevier, Oxford, UK, 2004.

- [78] Wolfgang Pauli. *Writings on Physics and Philosophy*. Springer-Verlag, Berlin, 1994.
- [79] John D. Perreault and Alexander D. Cronin. Measurement of atomic diffraction phases induced by material gratings. *Phys. Rev. A*, 73:033610, Mar 2006.
- [80] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.
- [81] C. J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases, 2e*. Cambridge University Press, Cambridge, UK, 2008.
- [82] William D. Phillips and Harold Metcalf. Laser Deceleration of an Atomic Beam. *Phys. Rev. Lett.*, 48:596–599, Mar 1982.
- [83] Michael Plischke and Birger Berbersen. *Equilibrium Statistical Physics, 2e*. World Scientific, River Edge, NJ, 1994.
- [84] N. Poli, F.-Y. Wang, M. G. Tarallo, A. Alberti, M. Prevedelli, and G. M. Tino. Precision Measurement of Gravity with Cold Atoms in an Optical Lattice and Comparison with a Classical Gravimeter. *Phys. Rev. Lett.*, 106:038501, Jan 2011.
- [85] H. David Politzer. Light incident on a Bose-condensed gas. *Phys. Rev. A*, 43:6444–6446, Jun 1991.
- [86] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Phys. Rev. Lett.*, 59:2631–2634, Dec 1987.
- [87] H. Rauch, W. Treimer, and U. Bonse. Test of a single crystal neutron interferometer. *Physics Letters A*, 47(5):369 – 371, 1974.
- [88] F. Riehle, Th. Kisters, A. Witte, J. Helmcke, and Ch. J. Bordé. Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer. *Phys. Rev. Lett.*, 67:177–180, Jul 1991.
- [89] M. S. Safronova, S. G. Porsev, and Charles W. Clark. Ytterbium in Quantum Gases and Atomic Clocks: van der Waals Interactions and Blackbody Shifts. *Phys. Rev. Lett.*, 109:230802, Dec 2012.
- [90] J. J. Sakurai. *Modern Quantum Mechanics*. Addison Wesley, 1993.

- [91] L. S. Schulman. *Techniques and Applications of Path Integration*. Dover Publications, 2005.
- [92] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt. Atom Interferometry with Bose-Einstein Condensates in a Double-Well Potential. *Phys. Rev. Lett.*, 92:050405, Feb 2004.
- [93] Y. Shin, C. Sanner, G.-B. Jo, T. A. Pasquini, M. Saba, W. Ketterle, D. E. Pritchard, M. Vengalattore, and M. Prentiss. Interference of Bose-Einstein condensates split with an atom chip. *Phys. Rev. A*, 72:021604, Aug 2005.
- [94] J. E. Simsarian, J. Denschlag, Mark Edwards, Charles W. Clark, L. Deng, E. W. Hagley, K. Helmerson, S. L. Rolston, and W. D. Phillips. Imaging the Phase of an Evolving Bose-Einstein Condensate Wave Function. *Phys. Rev. Lett.*, 85:2040–2043, Sep 2000.
- [95] D. M. Stamper-Kurn and W. Ketterle. Spinor condensates and light scattering from Bose-Einstein condensates. *arxiv:cond-mat/0005001 Les Houches*, 72:Course 2, 2001.
- [96] J. Stenger, S. Inouye, A. P. Chikkatur, D. M. Stamper-Kurn, D. E. Pritchard, and W. Ketterle. Bragg Spectroscopy of a Bose-Einstein Condensate. *Phys. Rev. Lett.*, 82:4569–4573, Jun 1999.
- [97] Satoshi Tojo, Masaaki Kitagawa, Katsunari Enomoto, Yutaka Kato, Yosuke Takasu, Mitsutaka Kumakura, and Yoshiro Takahashi. High-Resolution Photoassociation Spectroscopy of Ultracold Ytterbium Atoms by Using the Intercombination Transition. *Phys. Rev. Lett.*, 96:153201, Apr 2006.
- [98] M. Trippenbach, Y. B. Band, M. Edwards, M. Doery, P. S. Julienne, E. W. Hagley, L. Deng, M. Kozuma, K. Helmerson, S. L. Rolston, and W. D. Phillips. Coherence properties of an atom laser. *J. Phys. B*, 33:47–54, 2000.
- [99] David S. Weiss, Brenton C. Young, and Steven Chu. Precision measurement of the photon recoil of an atom using atomic interferometry. *Phys. Rev. Lett.*, 70:2706–2709, May 1993.
- [100] Andreas Wicht, Joel M Hensley, Edina Sarajlic, and Steven Chu. A Preliminary Measurement of the Fine Structure Constant Based on Atom Interferometry. *Physica Scripta*, 2002(T102):82, 2002.
- [101] Samuel S. M. Wong. *Introductory Nuclear Physics, 2e*. Wiley, 1999.

- [102] Xuguang Yue, Yueyang Zhai, Zhongkai Wang, Hongwei Xiong, Xuzong Chen, and Xiaoji Zhou. Observation of diffraction phases in matter-wave scattering. *Phys. Rev. A*, 88:013603, Jul 2013.