Dynamic Properties of Dilute Bose-Einstein Condensates

by

Dallin S. Durfee

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 1999

© Massachusetts Institute of Technology 1999. All rights reserved.

Author .....
Department of Physics
August 30, 1999

Certified by .....
Wolfgang Ketterle
Professor of Physics
Thesis Supervisor

Accepted by .....
Thomas J. Greytak
Associate Department
Head for Education

Science
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
WITHDRAWN
SEP 21 1999
FROM
MIT LIBRARIES
LIBRARIES

Dynamic Properties of Dilute Bose-Einstein Condensates

by

Dallin S. Durfee

Submitted to the Department of Physics
on August 30, 1999, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis, a new apparatus for the study of Bose-Einstein condensation is described, and the first two experiments performed with the new device are discussed. The new instrument was constructed for the creation of dilute gas sodium Bose-Einstein condensates, and features an optical quality quartz cell, a high-flux spin-flip Zeeman slower, a tightly confining magnetic trap, and a high-resolution imaging system. The theory, design, and construction of each component is discussed, including a detailed explanation of non-destructive dispersive imaging. Bose-Einstein condensation was first achieved in the new apparatus in January of this year. Bose condensates consisting of 10 to 25 million atoms can be produced in this apparatus at a rate of two condensates per minute.

The first two experiments performed with the new instrument probed the dynamic properties of dilute Bose condensates, allowing comparisons to be made with long standing theories of weakly-interacting degenerate Bose fluids. The first experiment was the study of "surface wave" excitations of Bose condensates. Standing and rotating quadrupole and octopole excitations were driven with a novel scanned optical dipole potential, a new tool which allows us to generate arbitrary two-dimensional perturbations to the trapping potential which confines the atoms.

The second experiment studied the transition from dissipationless to dissipative flow in a Bose condensate. This study, performed by "stirring" the condensate with a focused laser, provided the first experimental evidence for the existence of a critical velocity for dissipation in dilute gas Bose condensates. This experiment is discussed in the context of earlier studies of the critical velocity of superfluid liquid helium.

Thesis Supervisor: Wolfgang Ketterle
Title: Professor of Physics

Acknowledgements

There are many people who deserve my gratitude for their role in the completion of my thesis work, and it would be impossible to mention them all in a few short paragraphs. There are several people in particular, however, who I wish to recognize specifically. First of all there are the many people which I have collaborated with and learned from while working in the sodium BEC lab at MIT. Most notable is Wolfgang Ketterle, the leader of the group. I have often been amazed at his knowledge and insight into all corners of atomic physics and all aspects of experimental technique — as evidenced by the phenomenal success of his lab. Furthermore, Wolfgang has a gift for being able to explain complex ideas in an intuitive way, and he possesses an incredible amount of enthusiasm for his work. I should also thank Wolfgang for his efforts to keep the lab funded (as I should thank the funding agencies as well), making possible the unique experiments which I have been fortunate to take part in. While writing this thesis, Wolfgang's comments have been invaluable, and the time he dedicated to reading and criticizing my writing has had a large impact on the clarity of the text which I have written.

I should next thank the people who, with me, designed and built the new BEC experiment — not only for their hard work and scientific knowledge, but also for their friendship during the long struggle to not only *make* the machine, but to make it *work*! Chris Kuklewicz was the first to join me in the early efforts on the new experiment. Chris has been a good friend from the start, and his skills in the lab and in front of a computer have helped to make the experiment reliable

and easy to control. Chandra Raman and Roberto Onofrio joined the effort soon after, bringing a large amount of experience with them. Chandra and I were soon bonded by a common enemy — the temperamental ring dye laser, and Roberto, with his patience, good nature, and sharp sense of humor has been the one who was always able to pull the group together. Michael Köhl and Johnny Vogels both spent a year working on the new apparatus and both made large contributions. Michael's work with the scanning AOMs was particularly important for the experiments discussed in this thesis. And Jeff Gore, our undergraduate UROP, was both productive and pleasant to work with.

I also have great appreciation for the people who did early work in the Pritchard/Ketterle ultracold atoms group for laying the foundation upon which my thesis work was based, and I am greatly indebted to the people I worked with on the old BEC experiment for the many things which they taught me. Mark-Oliver Mewes and Michael Andrews, who were already well-experienced when I joined the Ketterle lab, introduced me to laser cooling and magnetic trapping. Dan Stamper-Kurn, who joined the lab the same year that I did, has been a constant source of interesting conversation (on all topics both regarding physics and not), and has often taken the time to share his experience and insight with the group of us working on the new experiment. I also had the opportunity of working with several very bright post-docs while working on the older BEC experiment. Klass-jan van Druten, Hans-Joachim Miesner, and Chris Townsend taught me a lot while I was a junior graduate student. And when I first began work on the new apparatus, Hans was the person who most often lent me advice or a helping hand. I have also benefited greatly from my interactions with Shin Inouye, Joern Stenger, and Ananth Chikkatur, who joined the work on the older experiment after my attention had been turned to the new one. And I cannot forget to thank Carol Costa for her assistance dealing with all of the administrative tasks associated with experimental work, and Peggy Berkovitz and the people at the physics department education office and at RLE headquarters.

The many professors who introduced me to physics at Brigham Young University are responsible for instilling in me the desire to scrutinize the inner workings of the universe, and for showing me the incredible predictive power of applied mathematics. Special thanks are in order for my undergraduate research advisor, J. Ward Moody, for giving me the knowledge and confidence necessary to pursue a graduate degree in physics. His leadership of the GoldHELOX project has had an immeasurable influenced on the students lucky enough to have worked on it.

A special thanks goes out to the people who encouraged and uplifted me outside of the halls of science. My parents, who encouraged my educational pursuits in every way possible as I grew up, and who gave me a peaceful and happy home to study in, were also a great, although long-distance, support during my time at MIT. And my brothers and sister, as well as my wife's family, have always been ready to lend a listening ear or a helping hand, and Georgianna has been a great help over the last few weeks. My friends from the Cambridge first ward and all of the "Balldrivers" past and present gave me the opportunity to get away from the darkened lab now and then, and to see what was going on in the world outside of building 26.

The greatest of all thanks goes to my wife Memorie. During the good times and the hard times at MIT, she always did everything that she could to encourage, love, and support me. Throughout my graduate studies, she has always been the one to rebuild my confidence during times of discouragement. In short, she is the best friend that I've ever had. I especially have to thank her for her help and support while I was writing my thesis. Despite being nine months pregnant at the time that my thesis was completed, she worked hard to make sure that I was undisturbed and free from other concerns so that I could focus on my thesis writing. And with her background in physics, she has served as the primary reader of my thesis, in addition to searching out and photocopying countless references and creating many beautiful figures.

Lastly, I would like to thank my soon-to-be first-born child for waiting until after my thesis defense to be born. Due to a great number of circumstances, my thesis defense ended up taking

place on the very date that the child was due. I am grateful that this child had enough compassion on me to not to be born on the date prescribed by the doctors, allowing me to finish my work as a graduate student before beginning my life as a father. Furthermore, the hope, excitement, anticipation, and the impending responsibility of becoming a father has been a great motivator during my thesis writing.

Contents

Chapter 1

Introduction and History

The realization of Bose-Einstein condensation (BEC) in dilute alkali gases in 1995 [1, 2, 3] was the fruition of many years of work by various experimental groups. This achievement has allowed the study of a new quantum fluid with properties very different from superfluid helium, and has allowed physicists to control atoms with quantum limited precision similar to the way in which the optical laser provides the ultimate control over photons. Much of the theory of dilute gas condensates stems from decades-old work. But since the experimental realization of alkali gas condensates, many new properties have emerged — due to the finite number of atoms, the existence of internal degrees of freedom, and the inhomogeneity in these new systems. Although many of the early experimental studies of alkali gas condensates parallel older studies of liquid helium (the study of elementary excitations and sound, for example), other studies have been quite different from work done in liquid helium (including the observation of the spatial interference of overlapping condensates, and the time-resolved measurement of the condensate's formation). There are several *key* aspects of superfluid helium, however, which are just beginning to be explored in alkali gas Bose condensates. Among these are the existence of surface waves and a critical velocity for dissipationless flow. Studies of these phenomena are the major scientific results reported in this thesis.

1.1 A Short History of BEC

The phenomena of Bose-Einstein condensation was first described by Albert Einstein in 1925 [4]. While working on a generalization of Satyendra Nath Bose's theory of the quantum statistics of photons, Einstein realized that an ideal Bose gas, cooled below a critical temperature T_c, would have a macroscopic population of atoms in a single quantum state (the ground state). Stated this way, Einstein's prediction is not surprising; it makes sense that a large ground state occupation would occur if an ideal Bose gas were cooled to a temperature lower than the first excited state of the system. What *is* surprising, however, is the fact that (1) the temperature at which this occurs can be many times *higher* than the energy of the first excited state, and (2) the number of atoms in the ground state, as a function of temperature, was predicted to have a discontinuous derivative at the critical temperature T_c. Until the identification of superfluid helium as a Bose condensed system, however, there were many doubts as to whether this "ideal gas" phenomena would occur in real systems with interparticle interactions (the ideal gas Bose-Einstein condensate is discussed in many textbooks [5, 6], and will not be elaborated upon here).

1.1.1 Early Studies of BEC

Bose-Einstein condensation was first experimentally observed in superfluid liquid helium. In 1911 (just three years after producing the world's first sample of liquefied helium), Kamerlingh Onnes discovered that the density of liquid helium, as a function of temperature, had a sharp maximum near 2.2 Kelvin [7]. In 1926 Onnes and Dana measured another strange maxima in the specific heat near this same temperature [8] (this temperature later became known as the λ-point, due to the λ-shaped specific heat curve of liquid helium). In 1928, Wolfke and Keesom expressed the opinion that these measurements were the result of a phase transition in liquid helium (the higher temperature phase being labeled He I, and the lower temperature phase He II) [9, 10]. Other

strange properties of liquid helium associated with this phase transition were later discovered. For example, a sudden transition to extremely low viscosity [11, 12, 13] and extremely high heat conductivity [14] was seen to occur as the liquid was cooled through the λ-point. And an extremely high thermomolecular pressure in He II (illustrated dramatically by the "fountain effect" [15]), and the adsorption of mobile "Rollin" films, many atomic layers thick, by any surface in contact with He II were also reported.

In 1938 Fritz London, noting the similarity of the λ-shaped heat capacity curve in liquid helium to the behavior of the heat capacity of an ideal gas near the BEC transition, proposed that superfluidity *could* be a manifestation of Bose-Einstein condensation [16]. This proposal, and its later verification, constitutes a landmark in modern physics. For the very first time a system had been discovered which was described by a macroscopic wavefunction, and to this date superfluid helium continues to be the subject of a considerable amount of research (for a more exhaustive treatment of superfluid helium, I refer the reader to some of the many books written on the subject [17, 18, 19]).

Since the connection was made between BEC and superfluid helium, several similar systems have been discovered. In 1956 a theory was presented in which pairs of fermionic electrons formed composite bosons in superconductors, linking superconductivity to a phenomena very similar to Bose-Einstein condensation [20]. More recently, fermionic ^{3}He has been shown to undergo a similar "BCS" transition [21, 22], and excitons in semiconductors have been made to Bose condense [23, 24]. More "exotic" Bose condensates have also been theorized, including pion and kaon condensates in neutron stars [25, 26, 27], and the condensation of Higgs bosons responsible for the mass of elementary particles [28].

1.1.2 BEC in Dilute Gases

The handful of Bose condensed systems mentioned in the paragraphs above have been studied in great detail over the years. But there has long been a desire for BEC in a long-lived, weakly interacting system which could be easily described from first principles. Long before dilute gas condensates were realized, the properties of weakly-interacting degenerate Bose gases had been studied theoretically in an attempt to gain insight into the more complicated, more strongly interacting superfluid helium. But until recently, no physical system had been realized which could provide a direct experimental test of these theories.

1.1.2.1 Development of the theory of weakly interacting bosons

Due to the complicated interactions between atoms in liquid helium, a full "first principles" description of the liquid helium phenomena discovered in the 1920's and 30's was not immediately available. Instead, two competing phenomenological theories emerged. The two-fluid model of Tisza and London considered liquid helium to be composed of two interpenetrating components, each with its own density and its own independent (to first approximation) flow field [29]. The so-called "normal" fluid was assumed to obey classical laws of fluid dynamics, while the "superfluid" component approximated an ideal gas Bose-Einstein condensate.

Landau's quantum hydrodynamic theory took a different view of superfluidity. Landau's theory was based on the quantization of classical hydrodynamics, de-emphasizing the quantum statistics of individual atoms and instead focusing on excitations in the fluid.

In 1947 Bogoliubov solved, from first principles, the ground state and fundamental excitations of a weakly interacting Bose gas [30]. Bogoliubov's assumptions in this study were not strictly valid for liquid helium (the temperature was assumed to be zero Kelvin, the "depletion" of the ground state was assumed to be small, and interactions were limited to mean-field interactions,

binary collisions, or hard sphere interactions). Nevertheless, this work showed that (A) Bose condensation can occur in the presence of interactions, and (B) the low-lying excitations of the ground state in this system are phonons. Bogoliubov's theory, then, contained the essential elements of both Landau's and Tisza's theories, suggesting that the two were complementary, not opposing.

The study of theoretical systems with more complicated interactions soon yielded fruitful results (including Penrose and Onsager's arguments for the existence of BEC in systems with more general types of interactions and, in particular, in liquid helium [31], and several works which lowered the restrictions of Bogoliubov's assumptions [32, 33, 34]). The theoretical study of weakly interacting degenerate Bose systems, however, continued. Early work done included studies of the ground and lowest excited states of dilute Bose systems [35, 36], the fugacity of such systems [37], the nature of the phase transition [38, 39], and sound propagation at finite temperature [40]. It wasn't until recently, however, that a physical system in which the interactions could well be described by a mean-field term was realized, allowing these theories to be directly tested.

1.1.2.2 Experimental realization of dilute gas Bose condensates

Since most atoms tend to form solids at the temperatures and densities required for BEC, one could easily be led to believe that dilute gas condensates are simply not possible. For a long time liquid helium (being the *only* known substance which did not become solid at absolute zero) seemed to be unique. In 1959, Bose-Einstein condensation in weakly-interacting spin-polarized hydrogen was proposed by Hecht [41], who predicted that the formation of molecules in atomic hydrogen could be suppressed all the way down to absolute zero by the use of strong magnetic fields. Hecht considered the nature of interparticle interactions in such a system, and concluded that they could possibly be described well by a hard-sphere model. By 1976 calculations had been done which confirmed that hydrogen would, indeed, remain a gas under these circumstances [42], and the experimental pursuit of BEC in spin-polarized hydrogen was begun.

The goal of BEC in hydrogen was eventually achieved in 1998 [43] (see [44] and [45] for further information about BEC in hydrogen). In the mean time, however, laser cooling and trapping techniques had been developed [46]. By combining these methods with the evaporative cooling technique developed for spin-polarized hydrogen, Bose-Einstein condensation in dilute alkali gases had been achieved in 1995 [1, 2, 3]. Since this time, many experiments have been performed probing the static and dynamic properties of weakly interacting Bose condensates, and dilute gas BEC has been produced by twenty experimental groups (see [47] for a review of many studies done and methods used in dilute gas BEC experiments).

1.2 Characteristics of Alkali Gas Bose Condensates

There are several properties which make alkali gas condensates of great interest. First of all, Bose condensates constitute a pure quantum state of matter. As such, they can be described by a macroscopic wave function with a global phase. The existence of this phase has been directly measured by interfering two condensates (see Chapter 8), and can be considered to be a result of the stimulated process by which condensates form [48]. The existence of a macroscopic wave function is responsible for the many strange properties of superfluid helium, and it is to be expected that similar properties should also exist in dilute gas Bose condensates. Alkali gas condensates, however, differ from superfluid helium in many ways. For example, as mentioned above, the alkali gas condensates which have been studied consist of a relatively dilute cloud of weakly interacting atoms, and are therefore highly compressible. The weak interatomic interactions can be described accurately by a mean-field, allowing comparisons to be made with the results of mean-field theories developed since the 1940's.

Another property of dilute gas condensates is the finite number (much less than Avogadro's number) of atoms involved. This results in a system with slightly different properties than a gas in

the thermodynamic limit [49]. In addition, dilute gas condensates are held in an inhomogeneous potential, resulting in a non-uniform condensate density in space. The inhomogeneous potential has the interesting result of spatially "separating" the condensate (which settles into a small "cold" cloud at the center of the trap) from the non-condensate atoms (which, having higher energies, are localized to a larger region of the trap).

The techniques used to make and study alkali gas condensates also allow for the existence of multicomponent condensates—either by mixing two different atoms or isotopes in a single cloud, or by putting a single species into multiple internal states. The tendency of the two components to intermix or phase separate, as well as the tendency for atoms in the condensate to change their internal state through collisions with other atoms, can be studied under various conditions. Multi-component condensates have been studied in Boulder by trapping atoms in two different hyperfine states [50]. Multicomponent "spinor" condensates have been studied here at MIT by putting the condensate into a superposition of m_F spin states [51].

1.3 Outline of the Thesis

This thesis can be roughly divided into four sections. In the section following this introduction, I present a detailed description of a new apparatus which we have constructed for dilute sodium gas BEC studies. Work on the new apparatus began about two and a half years ago, and its completion was marked in January of this year with the production of its first Bose condensate. Since I was the student who initially began work on the new apparatus, and since this is the first thesis written which contains data produced by it, I will both document the relevant parameters of the apparatus, as they relate to experiments which have or will be done on it, and record some of the knowledge gained while building and optimizing it. The third section of this thesis contains a discussion of the first two experiments performed on the new apparatus. The first experiment studied collective

excitations with angular momentum, using time-averaged optical potentials to generate rotating surface waves on the condensate. I will briefly describe a second ongoing experiment, which was begun while I was writing my thesis. This study has measured the temperature rise in a condensate when it was "stirred" by a focused off-resonant laser, producing evidence for a critical velocity for dissipative flow. The last section of this thesis contains several appendices which augment the discussions found in the rest of the thesis.

1.4 My Work at MIT

The work described in this thesis only represents the second half of my research done at MIT. Before beginning the work described in this thesis, I participated in many experiments performed with our older BEC apparatus. I started working with Wolfgang Ketterle's group just six months before we produced the first sodium Bose condensate, and I consider myself fortunate to have been part of the many ground breaking experiments which took place in the first years after this event. These experiments have been reported on extensively in journal articles and in other students dissertations, and will only be briefly mentioned here.

In 1995 we observed our first evidence for the existence of a sodium gas Bose condensate [2, 52]. Soon afterwards we implemented a new type of trap for BEC studies using only DC magnetic fields (the "cloverleaf" trap [53]). The flexibility of this new trap facilitated the study of some of the condensate's collective oscillation modes [54]. The frequencies of these modes were compared to those described by mean-field theory and found to be in good agreement. In the spring of 1996, we used dispersive imaging methods to view condensates in-situ and non-destructively [55]. This method was later improved upon (see Chapter 6), and was used in many future experiments to record real-time "movies" of condensate dynamics [56].

The second half of 1996 was an extremely exciting time in the Ketterle lab. By exposing

condensates to a variety of RF magnetic fields, we demonstrated Rabi oscillations of the trapped condensate and studied non-adiabatic Landau-Zener transitions into untrapped m_F states. These studies led to the creation of an "output coupler" which would allow an arbitrary fraction of the condensate atoms to be released from the trap [57]. Then, soon after this work was finished, we did an experiment in which we were able to observe the spatial interference pattern of two overlapping Bose condensates [58]. This experiment furnished the first direct proof of long-range coherence in a Bose condensate, drawing analogies between optical lasers and BEC. These two experiments constitute the first demonstration of an atom laser. In the course of these experiments we tried several methods to increase the time-of-flight of the expanding condensates (in order to increase the spatial period of the observed fringe pattern). One such attempt resulted in a sub-recoil atomic fountain [59].

Early in 1997, non-destructive imaging and the use of optical dipole forces allowed us to study the propagation of sound pulses in a Bose condensate [60, 61]. This work included the first measurement of the speed of sound in a Bose-Einstein condensate, allowing a direct comparison with Bogoliubov theory. Non-destructive imaging also made possible the direct observation of the formation of a Bose condensate. In this experiment, a cloud of atoms near the BEC transition temperature was "shock-cooled," and the formation of a condensate was observed as the cloud relaxed to equilibrium (see [48] and Fig. 6.5). This experiment revealed evidence of Bose stimulation in the formation of a condensate, showing an additional similarity of BEC and optical lasers.

Some of my major contributions to these early experiments include the implementation of non-destructive phase-contrast imaging (discussed in Chapter 6 and Appendix E) and rapid sequence imaging (see Section 6.2.2), and the construction of the cloverleaf trap (with fellow graduate student Dan Stamper-Kurn).

Chapter 2

A New Apparatus for BEC Studies

Early in 1997 we began to develop plans for a second BEC apparatus in our lab. Due to the multitude of experiments which we wished to pursue, we realized that we were limited by the finite run time available on a single experiment. By building a second apparatus, we also had the chance to implement several design improvements, and to create an experiment which would complement the other experiment — limiting its usefulness for some studies, but simplifying experiments which would be difficult in our first apparatus. The first Bose-Einstein condensate was produced in the new experiment in January of this year (see Section 2.2). Since this time the apparatus has been used to study several features of the dynamics of dilute sodium Bose condensates (see Chapters 9 and 10). This chapter discusses some of the important design features of the new apparatus, and reviews some of the key optimization steps taken in the newly constructed apparatus.

2.1 Overview of the New Apparatus

The purpose of a BEC apparatus is to cool and compress atomic vapors until the thermal de Broglie wavelength of the atoms is on the order of the spacing between atoms. At this point, the individual

wavefunctions of the atoms overlap, and a transition to a single multiparticle wavefunction occurs. The process of cooling atoms to this transition temperature involves several steps, each of which is discussed in detail in the following chapters. To give a brief synopsis: First, an atomic sodium vapor is produced in an oven and formed into a beam (Section 4.1). The beam is slowed by a Zeeman slower (Section 4.2) which generates a high flux of atoms slow enough to be captured, compressed, and cooled by a magneto-optical trap (Section 4.3). After the MOT is loaded to capacity, a few milliseconds of polarization gradient cooling is applied to further cool the atoms (Section 4.3.2), and then the atoms are transfered into a purely magnetic trap (Chapter 5). The magnetic trap is then strengthened, compressing the atoms to higher densities, and the final stage of RF-induced evaporation is applied to the atoms, cooling them through the BEC transition (Section 5.6). The density of the atoms at the BEC transition temperature is close to the density in the atomic beam oven, about 10^{14} atoms per cm^3. The temperature, however, has been reduced by a factor of a billion, resulting in a phase space density increase by 13 orders of magnitude (laser cooling and evaporative cooling each providing roughly half of this increase). After the critical temperature has been reached, the phase space density of the atoms increases rapidly with very small changes in the temperature of the gas, quickly producing a nearly pure Bose-Einstein condensate.

2.1.1 New Features of the Apparatus

The cooling sequence described above is identical to the one used in our earlier apparatus, and indeed the two experiments have a lot in common. There are, however, many important differences. One feature of the new apparatus is a Zeeman slower designed to deliver a higher flux of slow atoms to the MOT. The new slower is somewhat longer than the one on our original experiment, allowing it to slow faster atoms (see Section 4.2.3). More importantly, unlike the increasing field slower used on our first apparatus, the spin-flip design of the new slower allows it to be placed very

close to the MOT, resulting in a ten-fold enhancement in the number of slow atoms which actually enter the MOT's trapping volume (Section 4.2.2).

Another feature of the new experiment is the use of an optical quality quartz cell. In our older experiment the atoms are trapped in the center of a large stainless steel vacuum chamber. As a result, optical access to the atoms is limited, and it is hard to get optics and magnetic coils very close to the atoms. The quartz cell on the new experiment, however, allows optics and coils to approach within less than an inch of the condensate, and allows almost unlimited optical access to the atoms. The quartz cell in our experiment has optically flat surfaces ($\lambda/10$), allowing for distortion free imaging. To achieve this flatness, we had the cell manufactured by a special process — the optical quality of typical glass cells is destroyed when the cell walls are fused together. The cell that we have employed, however, was constructed by diffusion bonding — a process performed at much lower temperatures, thereby preserving the flat surfaces of the quartz.

Using a quartz cell does have some drawbacks. Time-of-flight measurements are limited, for example, by the decreased distance that the atoms can fall before hitting a wall (with a maximum expansion time of 50 ms, the optical density of a large condensate does not reach unity, forcing us to detune our probe light for many of our time-of-flight images). And the limited vacuum conduction through the narrow cell has required greater care to be taken to insure a good vacuum near the trap. Furthermore, in a large steel chamber one always has the option of introducing new objects into the vacuum (such as evanescent-wave or magnetic mirrors, etc.). As a result, the steel chamber of our first experiment and the quartz cell on our new experiment complement each other well, giving us the ability to perform more types of experiments than could be done on a single type of apparatus.

The magnetic traps used on both the old and new apparatus are in many ways similar, both producing Ioffe-Pritchard type field patterns. However the new "modified Ioffe-Pritchard trap" has a mean trapping frequency which is more than twice that of the cloverleaf trap on the old

experiment. This is a direct result of the use of the quartz cell discussed above, allowing us to place the coils closer to the atoms (see Section 5.3). The new trap features optical access similar to that of the cloverleaf trap, but due to the geometry of the coils, the new trap allows for higher resolution imaging (see Fig. 5.4). Several additional improvements to the experiment have been explored which are not currently implemented, including a recirculating sodium oven (Appendix B) and the use of optical fibers to stabilize optical molasses (Section 4.3.2).

2.1.2 Why Sodium?

The decision to use sodium in our second BEC apparatus was based on several considerations. When selecting an atom to Bose condense, the only *absolutely* required property is that the atom be bosonic. This criterion is not difficult to satisfy since all stable elements, with the exception of beryllium, have at least one bosonic isotope. The choice of atom, then, is mainly determined by the practical concerns of the cooling and trapping techniques used.

Advantages of alkali atoms Magnetic trapping works best with atoms having a strong magnetic moment, and laser cooling favors atoms with strong transitions in the visible or infrared region (where commercial cw lasers powerful enough to saturate the transition are available). Given these requirements, the alkali atoms are especially good candidates for BEC. Their simple and well studied spectra make them the obvious choice for laser cooling, and their unpaired electron gives them a relatively strong magnetic moment.

Alkali atoms also have several convenient (though less critical) properties. Alkali metals have low melting points, making it possible to generate a dense atomic vapor or beam with a very simple low temperature oven (see Section 4.1). Although atoms and molecules with higher melting points have been magnetically trapped, they require a more complicated mechanism for vaporization [62]. In addition, the thermal vapor formed in an alkali oven consists almost entirely of free monatomic

atoms (a great simplification over experiments with hydrogen, for example, in which the hydrogen dimer must be split to produce an atomic gas). At this time, BEC has been achieved with all of the stable alkalis except cesium and potassium, as well as with hydrogen (by a combination of cryogenic and RF-evaporative cooling [43]).

Collisional properties of sodium Collisions among atoms are an important part of evaporative cooling (see Section 5.6 and Appendix C.3). In a gas with absolutely no interactions, RF evaporation would simply remove atoms from higher-energy trap states without increasing the population of lower-energy trap states. So in order for evaporation to work, interactions must be present to maintain thermal equilibrium, and the rate at which effective RF evaporation can occur is determined by the collision rate of the atoms. Collisions, however, can do more than just redistribute kinetic energy. When two atoms collide, it is possible for spin angular momentum to be exchanged as well. As a result, collisions can put atoms into internal states which are not confined by the magnetic trap, and can mediate the release of internal energy into kinetic energy, both effects leading to trap loss. Thus the ratio of "good" elastic collisions to "bad" inelastic collisions under a given set of conditions is a very important figure of merit when selecting an atom to Bose condense.

Figure 2.1 maps out some of the collisional properties of sodium. In this figure, the rate of elastic collisions, at the density required for Bose-Einstein condensation, is compared to the dominant inelastic collision rates as a function of temperature. The region where the elastic rate is much larger than the inelastic rate spans most of temperature range displayed. Within this region, the shaded "BEC window" represents the range where the lifetime of the sample exceeds 0.1 seconds and where the rate of elastic collisions is faster than 1 Hz. The fact that it covers several orders of magnitude in temperature and density allows studies of BEC in sodium over a large range of parameters.

One further thing to note is that unlike ^{87}Rb, sodium does not have a low enough spin relaxation

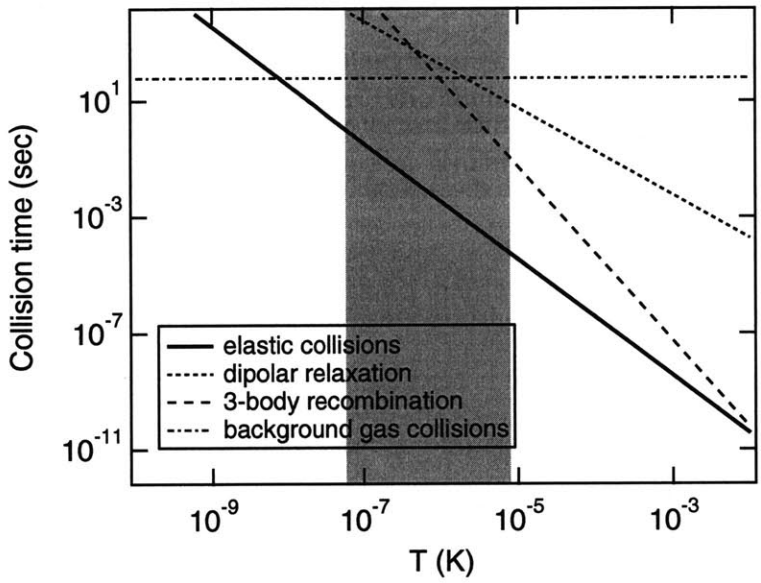

Figure 2.1: Elastic and inelastic collision times for sodium at the BEC transition. The mean collision time for several elastic and inelastic processes in a sodium gas are shown as a function of temperature at the critical density for Bose-Einstein condensation. The "BEC window," where the lifetime of the sample exceeds 0.1 seconds and the rate of elastic collisions is faster than 1 Hz is shaded. (This figure uses a scattering length of $50a_0$ and rate coefficients for two- and three-body inelastic collisions of 10^{-16} cm^3/s and 10^{-30} cm^6/s respectively). (Figure reprinted from reference [47])

rate to allow simultaneous magnetic trapping of two different hyperfine states (a useful property for multicomponent BEC studies). This property of ^{87}Rb is very unusual, and is not expected to occur in many atoms (in general, the collision between two hyperfine states involves many channels through which atoms can end up in an untrapped state — the amplitudes of these processes simply happen to cancel in ^{87}Rb). Work by members of our lab has demonstrated, however, that multiple m_F spin states of sodium *can* be held in an optical trap, opening the door to studies of BEC with a multi-dimensional order parameter [63, 51].

Additional properties of sodium An additional influencing factor in our choice of sodium was the fact that sodium is the atom used in our first BEC experiment, so the required experience and infrastructure were already in place. One *drawback* of sodium, for the time being, is that commercial solid state lasers are not available at the sodium wavelength. Unlike rubidium experiments, which typically utilize low-maintenance diode lasers, our sodium experiment relies on the use of maintenance-intensive dye lasers (although the MIT hydrogen BEC experiment reminds us that much more complicated laser systems are possible and practical). Working with visible light does, however, make it very easy to align optics and to diagnose and optimize a MOT. Another property of sodium is the existence of only one stable isotope. This can be seen as a disadvantage (potassium and lithium, for example, have both fermionic as well as bosonic isotopes, making it possible to study both Fermi and Bose gases in the same apparatus with a minimum of difficulty), or as an advantage (unwanted isotopes do not diminish the useful flux of atoms from the oven or contribute to the background gas pressure).

2.2 Realization of BEC in the New Apparatus

BEC was first realized in our new apparatus on January 28 of this year. As with earlier experiments, the first evidence seen for BEC was the existence of a bimodal, asymmetric velocity distribution measured in time-of-flight images (see Fig. 2.2). Our first condensates consisted of about 5 million atoms. By optimizing each cooling phase, we were eventually able to produce condensates of 10 to 25 million atoms, produced at a rate of just over two condensates per minute.

The most critical optimization in the experiment seems to be the polarization gradient cooling (see Section 4.3.2). This is optimized by first carefully zeroing magnetic fields by adjusting currents in a set of three Helmholtz coils, while measuring the magnetic field near the quartz cell with a Hall probe (once the correct currents are established, we have found that the field cancellation is stable, and daily adjustment of these currents is not necessary). Then the intensity of each of the counter-propagating molasses beams is carefully balanced. This step involves repeatedly loading the MOT for several seconds, and then switching to dark molasses for several seconds and watching the atoms diffuse away. By observing which way the atoms tend to move in the optical molasses, we can determine which beams are too weak or too strong. Fine readjustment of the beam balance typically has to be done several times during the course of a run. In addition to carefully adjusting beam balance, we have also found that the mode of the beams has to be very smooth in order to get good molasses. Every few days it becomes necessary to clean optics to prevent interference fringes from forming in the beams.

We have found that it is possible to transfer 20 to 30% of the atoms from the MOT to the magnetic trap (only 33% are transfered in the ideal case, since the atoms in the MOT are evenly distributed among three m_F spin states, and only one of them is weak-field seeking). Within the first few milliseconds after the magnetic trap is turned on, however, a large fraction of these atoms are lost. This is probably due to the large size of the atom cloud and the proximity of the radial

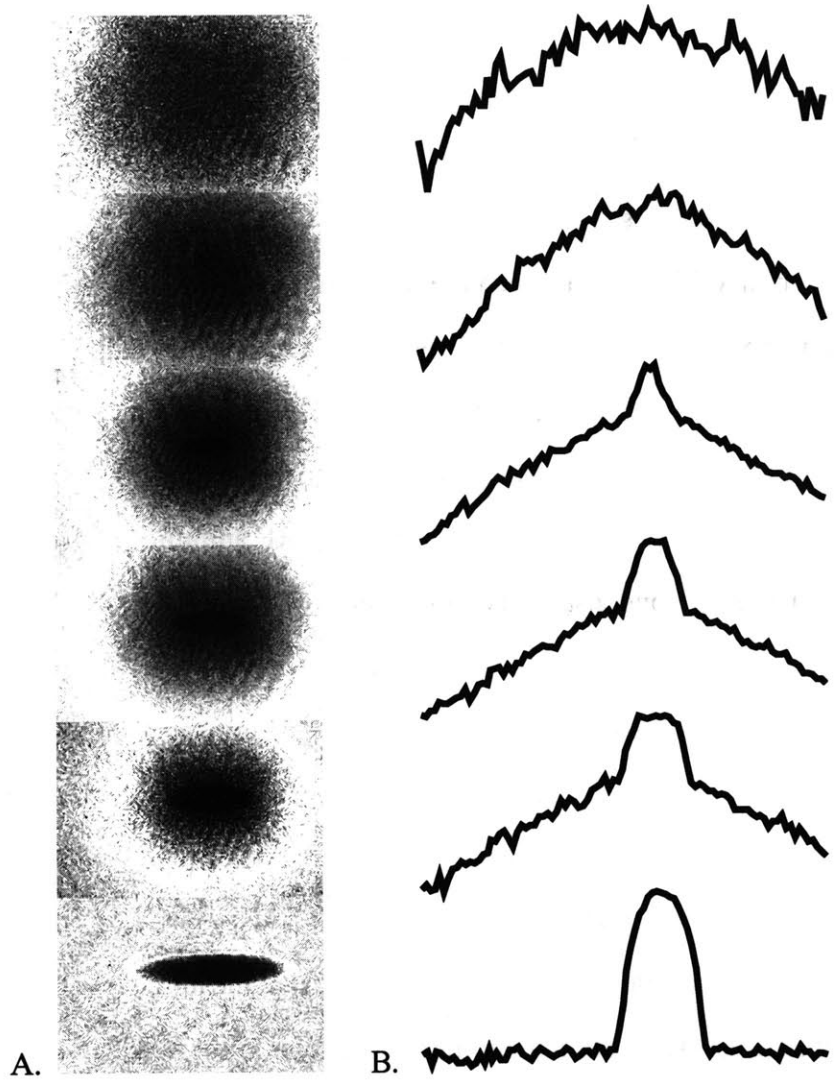

Figure 2.2: BEC transition viewed in time-of-flight. Figure (A) shows a series of time-of-flight absorption images of clouds at various temperatures near the BEC transition. Each picture was taken by evaporating a magnetically trapped cloud of sodium to a different RF frequency (a lower frequency corresponds to a lower temperature—see Section 5.6), and then allowing the cloud to expand ballistically for several milliseconds. The final evaporation frequency for each frame is 2.5, 2.2, 2.15, 2.1, 2.0, and 1.7 MHz (from top to bottom). The appearance of a cold, asymmetrically expanding cloud of atoms as the temperature is lowered is an indication that the BEC transition temperature has been reached. The bottom image shows a nearly pure Bose condensate consisting of 5 million atoms. The time-of-flight for each image was 30 ms, with the exception of the bottom most image for which it was 50 ms, and the width of each frame is 3.2 mm. Figure (B) shows the vertical atom density profile for each frame in figure (A). The presence of a Bose condensate in the four coldest clouds is easily seen in the bimodality of their profile.

instability in the magnetic trap (see Section 5.2). This loss was minimized, but not eliminated, by adjusting the initial currents used to generate the magnetic trap. After this initial loss, the trap becomes very stable, decaying with a reasonably long time constant of about 1 minute due to collisions with background gases.

Careful optimization of the rate at which the RF frequency is swept during evaporation has not seemed to make an impact in the size of condensates produced. Some improvement was obtained by increasing the RF power at the beginning of the evaporation from the values initially used. Beyond a certain point, however, increasing the power further no longer had an effect. The total RF power requirements are not very high, and when our RF amplifier malfunctioned, we found that we could produce Bose condensates using just the power which was output directly from our digital frequency synthesizer ($\sim$ 10 mW).

Due to the high field gradients produced by our magnetic trap, and to the limited frequency of our RF synthesizer (30 MHz), RF evaporation starts off by spin-flipping atoms which are closer to the middle of the cloud rather than at the edge. But when we tried to extend evaporation to higher initial frequencies with the use of a higher frequency RF oscillator, we saw no improvement in the condensate number. After several attempts, we concluded that the initial evaporation was probably limited by the spontaneous evaporation of atoms over the saddle point in the magnetic trap (Section 5.3).

One more important optimization to the experiment is the decompression of the atoms near the end of evaporation. Although decompression was not necessary in our first BEC apparatus, the higher field gradients in our new trap create high enough atom densities in the condensate to make inelastic three-body collisions a significant loss mechanism. As a result, we found that it is beneficial to weaken the trap during the last 250 to 500 ms of evaporation.

Now that the apparatus has been debugged and peaked up, it functions quite reliably. On a reasonably good day the lasers can be warmed up and aligned within three or four hours, after

which the apparatus requires only a minimal amount of re-adjustment throughout the rest of the run, allowing us to concentrate fully on the experiment at hand.

Chapter 3

The Vacuum System

The vacuum system is a very important part of any alkali gas BEC experiment. The vacuum system must be leak tight, sturdy, and bakeable to high temperatures in order to create and maintain the ultra-high vacuum environment necessary for BEC. Care must be taken to assure that the vacuum conductance between the trapping chamber and the vacuum pumps is high, and that the conductance between the trapping chamber and regions of poorer vacuum is low. Also, since the vacuum chamber is the largest single piece of a typical BEC apparatus, the physical size and shape of the vacuum chamber must be designed to accommodate the magnetic coils and optics which surround it. In this chapter, the vacuum requirements for BEC will be discussed, followed by a description of the vacuum system implemented in the new BEC apparatus.

3.1 Vacuum Requirements

During the creation and study of a Bose-Einstein condensate, the atoms must be thermally insulated from the outside world. This is accomplished by holding the atoms in a trap, isolated from any room-temperature walls by a very good vacuum. No type of insulation is ever perfect, however,

and all sources of heating must be understood and controlled.

There are essentially three ways in which the atoms are coupled to the environment: through blackbody radiation, through fluctuations in the magnetic trap, and through collisions with residual gases in the vacuum chamber. Blackbody radiation can be neglected because the scattering of thermal photons occurs at a *very* low rate (about one photon per $\sim 10^{20}$ years!). The coupling of thermal energy through the magnetic trap is also negligible. A typical magnetic trap has a resistance of a few ohms and an RC time constant much longer than a microsecond. This results in an RMS thermal (Johnson) noise current of much less than a nanoamp, a trivial current compared to the typical DC currents in the trap of tens or hundreds of amps (current ripple or other "technical" instabilities in the trap current can, however, lead to significant heating — see Appendix C.3.1). Heating due to collisions with background gases, however, can be a severely limiting factor.

Several groups have reported that a background pressure in the low 10^{-11} mbar range is necessary to obtain magnetic trap lifetimes of one minute. This low required pressure is a result of the extremely low temperature necessary for Bose condensation; in a hotter gas, heating due to grazing incidence collisions are less significant and can be neglected. The rate of background collisions per trapped atom can be written as $R_{bg} = n_{bg}\sigma_{bg}v_{bg}$, where n_{bg} and v_{bg} denote the density and thermal velocity of the background gas molecules. Assuming a collision rate of 1 per minute, a background pressure of 10^{-11} mbar, and a background gas temperature near room temperature, we find that the collisional cross-section, σ_{bg}, is on the order of 10^3 Å^2, *much* larger than the length scale of the inter-atomic potentials. This reflects the major contribution of the long-range van der Waals potential and small-angle scattering to the total cross-section. Background collisions are discussed in detail in [64, 65].

3.2 Layout of the Apparatus

In order to load the atom trap quickly, a high flux oven must be used. The temperature of the oven, and the regular introduction of contamination into the vacuum each time it is reloaded with sodium, creates a background gas pressure near the oven which is unacceptable for magnetic trapping and evaporative cooling (the pressure near the oven in our new apparatus typically ranges from a few times 10^{-8} Torr to as high as 10^{-7} Torr).

We have gone to great lengths to reconcile the need for a fast loading atom trap with the need for an extreme UHV environment. The result is the vacuum system shown in Fig. 3.1. The vacuum system is segmented into four parts: the oven chamber, the slower, the quartz cell, and the pumping chamber. The narrow, low conductance tubes connecting each part of the vacuum system, and pumps placed in three different places along the chamber assure that the vacuum pressure in the trapping region is about 1500 times lower than the pressure in the oven chamber. A detailed analysis of vacuum conduction in our new apparatus is given in Appendix A.

The oven chamber, shown in Fig. 3.2, contains the sodium oven (the source of the sodium beam which loads the trap) as well as the cold plate/differential pumping assembly which collimates the sodium beam and limits the flow of gases between the oven chamber and the slower. The oven chamber is pumped by a 60 liter/second ion pump (protected from alkali poisoning by a water cooled baffle) supplemented by a titanium sublimation pump. In order to isolate the relatively dirty vacuum of the oven chamber from the UHV vacuum of the slower and trapping regions, a nitrogen cooled cold plate and differential pumping tube are employed. The cold plate is a piece of copper which can be cooled with liquid nitrogen to collect hot sodium atoms and other vacuum contaminants. The differential pumping tube, detailed in Fig. 3.3, is a 7" long copper tube connected to the back of the cold plate. The 3/16" diameter opening in the end of the differential pumping tube nearest the oven serves to collimate the sodium beam. The other side of the tube

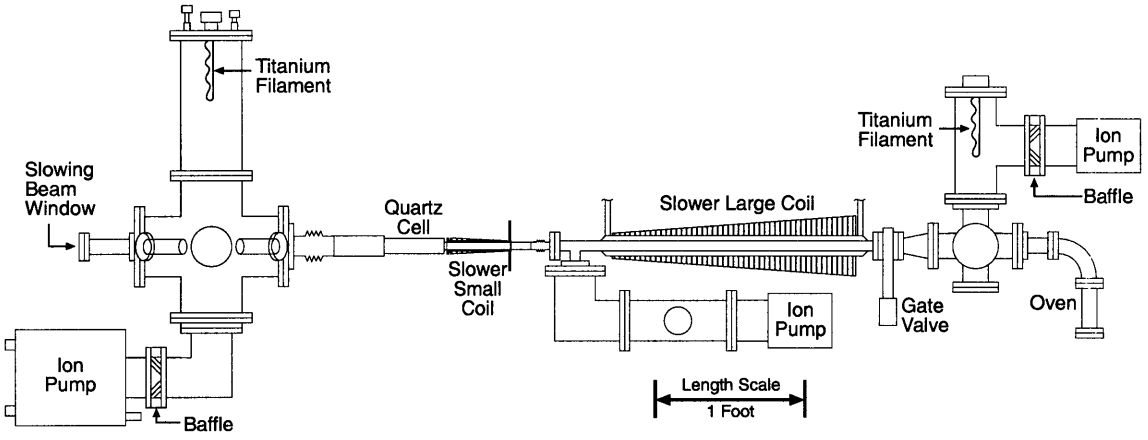

Figure 3.1: Side view of the vacuum system. The major gas load in the experiment is produced by the oven (on the far right). The oven, pumped by a 60 liter/second ion pump in combination with a Ti sublimation pump, is separated from the slower by a narrow differential pumping tube (see Figs. 3.2 and 3.3) with a vacuum conductance of just 0.14 liters/second. As a result, background molecules are several thousand times more likely to find themselves stuck in the ion or Ti sublimation pumps than to enter the slower. The relatively narrow slower tube is pumped on by a 60 liter/second ion pump, adding an additional stage of differential pumping. The trapping region is evacuated through the UHV pumping chamber, pumped on by a 120 liter/second ion pump and a Ti sublimation pump. The result is a vacuum of order 10^{-11} Torr in the quartz cell where the atoms are trapped. A detailed vacuum conduction analysis is contained in Appendix A.

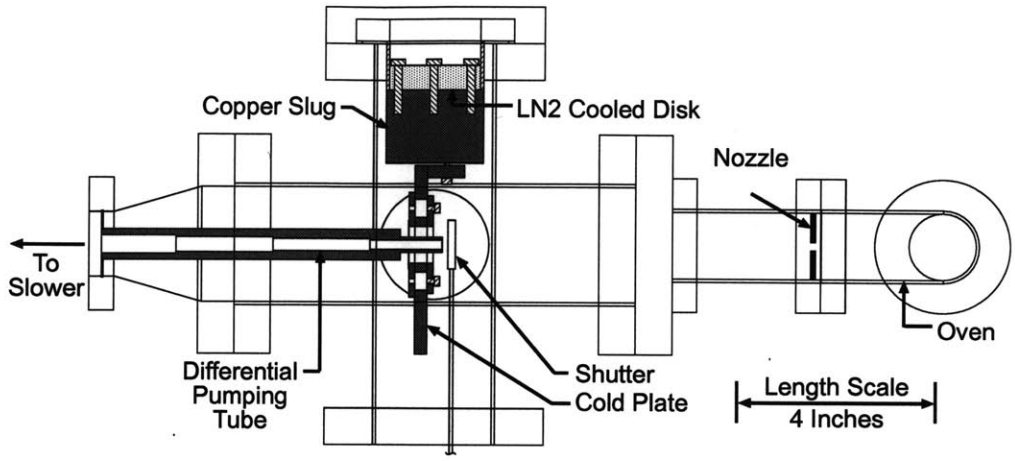

Figure 3.2: Oven/Cold plate assembly (top view). Sodium escaping through the oven nozzle is collimated by the end of the differential pumping tube attached to the cold plate. The cold plate is heat sunk to a copper slug which conducts heat out of the vacuum system to a liquid nitrogen cooled disk, allowing the cold plate and the differential pumping tube to be cryogenically cooled. The differential pumping tube (detailed in Fig. 3.3) has a stepped inner bore to minimize its conductance without clipping the radially expanding sodium beam.

is sealed to the vacuum chamber with a thin stainless steel membrane, such that the only channel between the oven chamber and the slower is through the narrow bore of this tube. This limits the rate at which gas from the oven chamber streams into the slower. The steel membrane was made as thin as possible to lower the heat conduction from the vacuum chamber walls to the nitrogen cooled differential pumping tube.

The inner bore of the differential pumping tube changes in several steps, approximating a cone. This design provides the lowest possible conductance without occluding any part of the conically expanding atomic beam. This is an improvement over the cylindrical tube used in our first experiment, because not only are the atoms which strike the side of the tube lost from the beam, but they also have the tendency to bounce around inside the tube, colliding with other atoms and choking the atomic beam.

Cooling of the cold plate/differential pumping tube assembly is accomplished by conduction

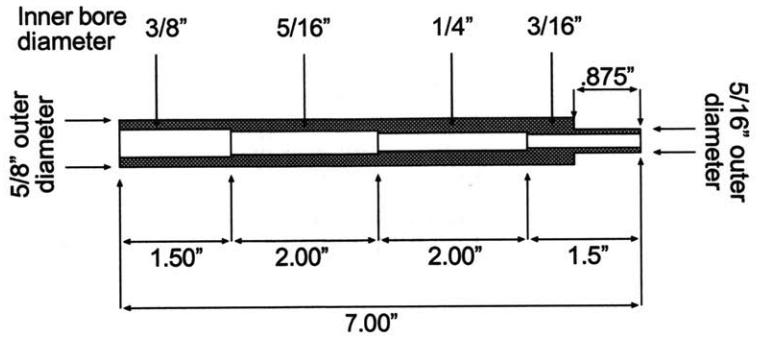

Figure 3.3: Differential pumping tube.

through a two inch diameter copper slug which protrudes out of the vacuum chamber. This method was chosen over direct liquid nitrogen feedthroughs into the vacuum to allow the possibility of cooling the cold plate with a commercial closed loop chiller. We have found that, unlike our older BEC apparatus, once the sodium oven has been degassed, liquid nitrogen cooling is not necessary. This is probably a result of the conically tapered differential pumping tube; since very little sodium strikes the inner walls of this tube, the enhanced sticking coefficient for sodium which cryogenic cooling provides is not necessary to prevent sodium from bouncing around in the tube.

The oven chamber can be isolated from the slower with the gate valve shown in Fig. 3.1. This allows us to vent the oven chamber (to reload the sodium, for example) without losing the ultra-high vacuum in the slower and trapping chambers. The o-ring in the gate valve is made of Kalrez, allowing the valve to be baked to 250° C. The slower (essentially consisting of a one meter long tube surrounded by magnetic coils) is divided into two parts, as described in Section 4.2.4. A tee placed between the two halves is attached to a 60 liter per second ion pump, providing an additional stage of differential pumping and limiting the flow of gas from the slower into the trapping region.

Immediately after the tee is a glass to metal seal to which the quartz cell is attached. The first segment of the cell is a narrow (3/4 inch) quartz tube. The second slower coil is wound around a brass spool surrounding this tube. At the end of the tube, the cell opens up into a square

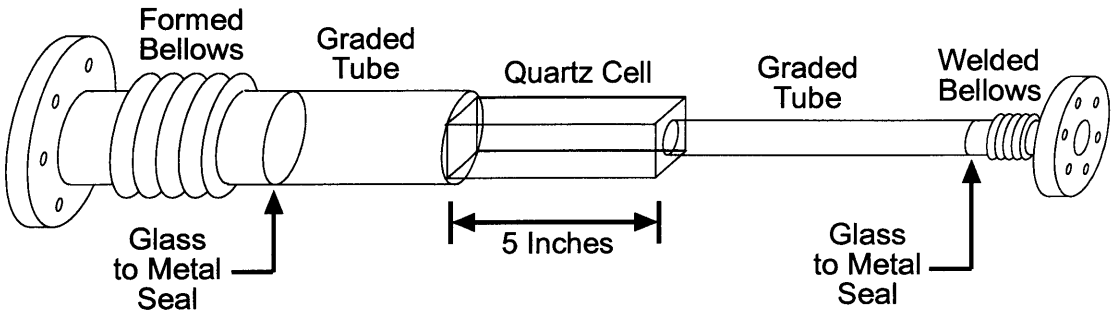

Figure 3.4: The quartz cell. The central part of the quartz cell has a square cross-section, 32 mm on a side. This section is 5 inches long, and is made of four optically flat quartz plates diffusion bonded together. The square section of the cell is connected to round tubes on either side. Both the larger 2" diameter tube and the smaller 3/4" diameter tube are fused to quartz-to-metal seals utilizing graded glass which gradually changes from pure quartz on one side, to a special glass with a thermal expansion coefficient matched to the metal seal on the other. Both glass-to-metal seals are welded to a conflat flange through a bellows to help isolate the fragile cell from mechanical stress.

cross-section. The rectangular part of the cell was made to be 32 mm on a side and 5" long to accommodate the magnetic trap and allow for the entrance of large MOT beams. At the far end of the trapping region, the quartz cell opens up further into a 2" diameter round tube to optimize vacuum conduction into the pumping chamber. A scale diagram of the quartz cell is given in Fig. 3.4.

It should be noted that while many experiments employ fragile glass cells, very few of them are connected between two large steel chambers. This made the assembly of the vacuum system an extremely delicate process. To lower the risk of breaking the quartz cell, both ends of the cell are connected to the vacuum chamber through vacuum bellows. A stiffer formed-bellows on the large diameter tube provides a small amount of flexibility while the chamber is being assembled. Under vacuum, this bellows is completely compressed, providing a rigid mounting point for the glass cell. The small end of the cell is connected to a very flexible welded-bellows which prevents small amounts of relative motion between the slower and the UHV pumping chamber from cracking the

cell.

When first installing the quartz cell into the vacuum chamber, the cell cracked twice. The cell was then modified by reinforcing its three weakest points. Two of the weak points were the unions between the square section of the cell and the round tubes on either side of it. Since diffusion bonding creates weaker joints than traditional heat fused joints, and since these junctions are a fair distance from the center of the trap, it was possible to have a glass blower strengthen these joints by fusing them with a torch without destroying the optical quality of the quartz near the center of the magnetic trap. The third weak point was a molybdenum quartz-to-metal seal. Flexible molybdenum seals allow quartz to be sealed directly to stainless steel, minimizing the length of the 2" diameter tube between the square part of the cell and the UHV pumping chamber. We found, however, that molybdenum seals can be very fragile. This seal was replaced with a quartz-to-metal seal which utilizes a section of graded glass tube, varying across its length from pure quartz on one side to a glass formulated to match the thermal expansion of the metal tube on the other. This length of graded glass tubing added a few inches to the length of the cell, reducing vacuum conduction by a small amount, but resulted in a more durable seal. Since these modifications were made, we have installed the cell twice without problems, and the cell has been in use for over a year now without mishap.

The pumping chamber is a large, many-port chamber to which an ion and titanium sublimation pump are connected. Other ports on this chamber are filled with ion gauges, a residual gas analyzer, and a bakeable all-metal valve (to which a turbo pump can be attached when pumping down and baking out the chamber). The port opposite the quartz cell contains a window to allow the slowing laser beam to enter the apparatus. The window is set off from the chamber with a nipple so that it can be heated (to prevent fogging of the window from the sodium beam) with minimal heat conduction to the rest of the chamber. UHV vacuum is obtained in the new apparatus by pumping on the chamber with a turbo pump while the chamber is baked at 250°C for several days.

Chapter 4

Laser Cooled Atoms

The first stage of cooling in our apparatus is done with near resonant laser light. In the past decade, laser cooling has become one of the most important tools for atomic physics, and three pioneers in this field were recently honored with the Nobel prize [66, 67, 68]. The relative ease and versatility of laser cooling allowed experimenters to bypass many of the difficulties experienced by groups working with cryogenic pre-cooling methods. As a result, BEC was achieved with alkali gases three years before a much longer standing effort with hydrogen was completed.

It should be noted that despite the advantages of laser cooling, BEC has never been achieved with laser cooling alone. The simplest methods of laser cooling are limited in temperature to the Doppler limit ($T_{Doppler} = \frac{\hbar\Gamma}{2k_B}$ for a two-level atom with a linewidth Γ)[69]. More refined methods can cool atoms down to the recoil limit, the point at which the average momentum of an atom is equal to the recoil of a single photon: $T_{recoil} = \frac{\hbar^2 k^2}{2Mk_B}$ (where k is the wavenumber of resonant light, and M is the mass of the atom) [70]. At the recoil temperature, Bose condensation requires that the atoms be spaced by less than an optical wavelength. At these densities, light scattering can no longer be considered as an interaction between a photon and a single atom, causing traditional laser cooling to fail. Methods have been demonstrated which illustrate the possibility to laser cool atoms

to temperatures much colder than the recoil temperature [71, 72], but these methods have so far been limited to small dilute samples. Laser cooling processes in general tend to break down when the sample becomes optically dense. The purpose of laser cooling in current BEC experiments is, then, to bridge the gap between thermal sources of atoms (at hundreds of Kevin) and magnetic traps (with trap depths of only $\sim$ 1 mK).

As previously mentioned, the production of laser cooled atoms in the new apparatus is done in several phases. First, a thermal beam of atoms is produced in an oven. This beam is then slowed by a Zeeman slower to a velocity of about 60 m/s, small enough to be captured in a magneto-optical trap (MOT). The MOT confines, compresses, and cools the atoms to a density on the order of 10^{11} atoms/cm^3 and a temperature just below 1 mK. A final stage of laser cooling is done by suddenly switching off the MOT and performing a few milliseconds of polarization gradient cooling. This lowers the temperature of the atoms to $\sim$ 100 μK, at which point the atoms are cold enough to be held in the magnetic trap.

4.1 The Oven

As with all current alkali gas BEC experiments, the source of the atomic beam used in the new apparatus is nothing more complicated than an oven. Due to the low melting point of sodium, it is possible to create reasonably large vapor pressures at relatively low temperatures. This makes very simple oven designs possible. The oven in the new apparatus is simply a conflat elbow, heated externally by band heaters and heater tape, which contains a lump of metallic sodium. One side of the elbow is connected to a short conflat nipple with a copper disk brazed in one end (see Fig. 3.2). The copper disk has a 4 mm diameter hole in its center (the "nozzle") which allows atoms to escape from the oven. The escaping atoms are collimated by a 3/16 inch hole in the copper cold plate/differential pumping assembly positioned 7.5 inches from the nozzle. The position of the

nozzle and the differential pumping tube are radially adjustable so that the atomic beam can be aligned with the bore of the slower. Between the oven and the cold plate is a beam flag which can be moved up and down to shutter the sodium beam.

The oven is typically heated to 270 C, producing a $6 \cdot 10^{-3}$ Torr sodium vapor in the oven. The nozzle is typically heated to 350 C to prevent sodium from condensing on and clogging it. The nozzle is brazed into a water cooled nipple to both prevent the nozzle heater from heating the rest of the vacuum chamber and to help the sodium which strikes the wall of the nipple to stick to it, thereby limiting the sodium build-up in the rest of the oven chamber. Atoms in the thermal beam have a most probable velocity of 762 m/s, and the beam produces a flux of $2.9 \cdot 10^{12}$ atoms per cm^2 per second at the center of the MOT (1.6 meters from the oven nozzle) [73]. A considerable amount of work has been done in our lab to develop and test a more complicated "recirculating" oven, similar to the one discussed in [74]. The recirculating oven, documented in Appendix B, has not yet been implemented on the experiment.

4.2 The Zeeman Slower

The hot beam of atoms from the oven must be slowed before it can be used to effectively load a MOT. Atoms which pass through the beams of the magneto-optical trap with too high a velocity will be out of resonance with the MOT beams and will not be stopped. Increasing the size of the MOT beams allows faster atoms to be captured, since the larger capture volume gives the atoms passing through the MOT more time to be slowed through off-resonance scattering. Increasing the MOT volume to increase the maximum capture velocity is only effective to a certain point, however. In the limit of high velocity atoms (with Doppler shifts much greater than the linewidth), the scattering rate scales with the atom's velocity as v^{-2}, the number of photons which must be scattered to catch an atom scales with v, and the time taken by an atom to pass through a given

capture volume scales as v^{-1}. So in this limit, doubling the capture range of a MOT involves expanding one dimension of the MOT by a factor of 16. As a result, the velocity capture range of a MOT is practically limited to Doppler shifts of several times the linewidth of the slowing transition. The maximum capture velocity of our sodium MOT is $\sim$ 60 m/s, compared to the $\sim$ 800 m/s RMS velocity of atoms emitted by the oven. Therefore a beam slower is necessary to efficiently load the MOT.

4.2.1 Atomic Beam Slower Theory of Operation

Atoms are slowed in a beam slower by scattering photons out of a counter-propagating laser beam. This laser beam is detuned from the natural resonance of the atoms to compensate for the Doppler shift of the fast moving atoms. But as fast moving atoms scatter photons and slow down, their Doppler shift decreases, causing them to get out of resonance. If nothing is done to compensate for the changing Doppler shift, the atoms will be lost from the slowing process. So rather than slowing a large number of atoms to a low velocity, atoms with a Doppler shift within several linewidths of the laser detuning are simply "burned out" of the Maxwell-Boltzmann distribution.

The changing Doppler shift can be compensated for with one of two commonly used methods. In a chirped slower, atoms are kept in resonance as they slow down by raising the frequency of the laser beam. With each chirp of the laser, a large number of atoms can be pushed to low velocities. Using this type of slower to load a MOT, however, has two major drawbacks. First, in chirped slowers the slow atoms are produced along the entire length of the slower. Slow atoms produced far from the MOT are more likely to end up striking the walls of the vacuum chamber than to actually enter the MOT's capture region. Second, the fastest atoms which can be slowed are only in resonance periodically when the laser is at its lowest frequency. Fast moving atoms which enter the slower just after a chirp has begun will not be slowed until the next chirping cycle. As a result,

an additional length, equal to the velocity of the fastest atoms to be slowed times the duration of a chirp, must be added to the optimum length of the slower given in Eq. (4.1). Otherwise only a fraction of the faster moving atoms will be slowed.

Zeeman slowers overcome these two drawbacks of chirped slowers. In Zeeman slowers the laser beam is kept at a constant frequency, and an inhomogeneous magnetic field is used to keep the atoms in resonance with the slowing light. As the atoms pass through the slower, the inhomogeneous magnetic field causes them to have a changing Zeeman shift which compensates for their changing velocity, keeping the atoms on resonance as they slow down. Since the resonant condition depends on the position of the atoms in the slower rather than time within a frequency chirp, Zeeman slowers do not suffer from the limited duty cycle of a chirped slower. Furthermore, the velocity of the slowed atoms is determined by the local magnetic field, and is therefore a function of position along the slower rather than a function of time within a chirp. As a result, the atoms reach their slowest velocity at the end of the slower, near to the capture region of the MOT.

4.2.2 Spin-flip Zeeman Slowers

The counter-propagating laser beam in a Zeeman slower is usually circularly polarized to drive either a $\Delta m_F = 1$ or $\Delta m_F = -1$ transition (typically a cycling transition). By switching the handedness of the polarization, either of these two types of transitions can be driven, allowing one to select a transition which either increases or decreases in energy with increasing magnetic field. This makes it possible to implement Zeeman slowers in which the magnetic field either decreases or increases in magnitude along its length. In the literature, these two types of slowers are usually referred to as $\sigma-$ and $\sigma+$ slowers, or increasing field and decreasing field slowers. Both types of slowers have advantages and disadvantages.

In an increasing field slower, the atoms entering the slower initially encounter a very low mag-

netic field. This means that the laser beam must be detuned by the entire Doppler shift of the fastest atoms to be slowed. In our case this is technically undesirable since our laser system is locked near to resonance (for the purpose of magneto-optic trapping and probing of the atoms), and the fastest atoms we wish to slow have Doppler shifts of ~ 1.5 GHz. This requires either the use of a high frequency AOM (which typically have very low conversion efficiencies), or the use of a second laser. A more serious disadvantage of the increasing field slower is that the magnetic field has a large maximum at the far end of the slower closest to the MOT. As a result, the MOT has to be placed a considerable distance downstream from the slower to avoid interference from the slower's fringing fields (in our first BEC apparatus, this distance was about half a meter, resulting in the loss of about 90% of the slowed beam). On the plus side, the longitudinal derivative of the magnetic field, dB/dx, changes sign at the end of the slower, causing the atoms leaving an increasing field slower to move out of resonance quickly. As a result, the scatter in the final velocity of the slow atoms is only a few times the linewidth of the optical transition (see Fig. 4.1).

A decreasing field slower has a small field at its exit, allowing it to be placed in closer proximity to the MOT. It can also be operated with a laser beam much closer to the zero velocity resonance of the atoms. The main drawback of this type of slower is that the fringing fields at its end decay in a way which tends to extend the slowing process (see Fig. 4.1) resulting in a broadening of the velocity distribution of the slowed atoms.

A spin-flip slower capitalizes on the strengths of both of the standard types of Zeeman slowers by combining the two in series. The atoms are first slowed with a decreasing field slower. Since this first segment only does a fraction of the slowing, the field at the beginning of the slower can be reduced. The first coil still provides significant compensation for the Doppler shift, however, and only relatively small laser detunings are required. At the end of this first stage of slowing, a second coil produces a field which increases in the opposite direction. This small increasing field slower produces a fringing field at the end which is small and decays quickly, allowing the slower to be

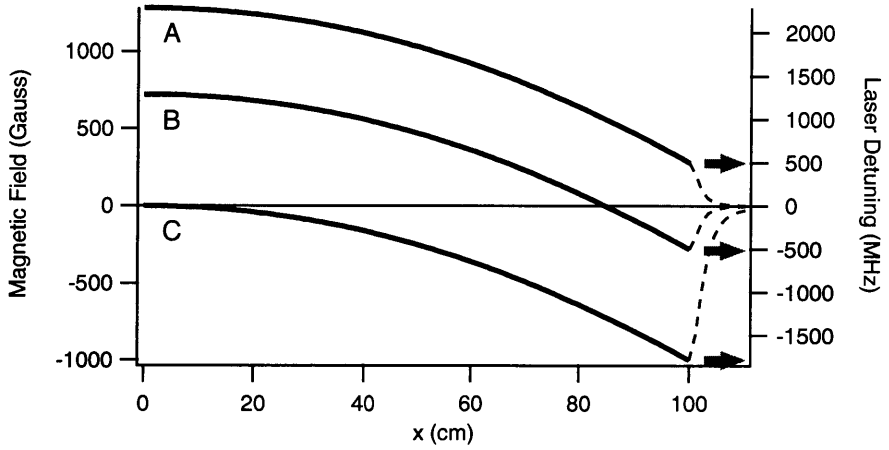

Figure 4.1: Comparison of Zeeman slowers. The magnetic field along the length of three hypothetical sodium beam slowers is shown. The required laser detuning for each slower is indicated by the short horizontal arrows on the right. Regardless of the type of Zeeman slower used, the optimum rate at which the field changes along the length of the slower is the same, with each type of slower simply differing by a constant magnetic field offset. In a decreasing field slower (A), high fields at the beginning of the slower compensate for most of the Doppler shift, and only small laser detunings are needed. The fall off of the magnetic field at the end of the slower (indicated by a dotted line), however, causes some atoms to continue to be slowed after exiting the slower, resulting in a large spread in the velocities of the slow atoms. In an increasing field slower (C), the field starts at zero, and all of the Doppler shift of the fast atoms has to be compensated for with a large detuning of the laser. In addition, the large field at the far end of the slower requires the MOT to be placed a considerable distance away. As a result, only a small fraction of the transversely expanding beam of slow atoms actually reaches the capture region of the MOT. The spin-flip slower (B) overcomes these limitations, requiring relatively small laser detunings, having a small field at the end of the slower, and having fringing fields which fall off in such a way that the slow atoms quickly get out of resonance with the slowing light.

placed in close proximity to the MOT. Furthermore, the fringing fields at the end of the slower decay in a way such that the atoms exiting the slower immediately get out of resonance with the slowing beam.

The term "spin-flip slower" can be somewhat confusing. It should be noted that the orientation of an atom's spin does not change, as measured in the lab frame, in a spin-flip slower. It is the magnetic field, rather, which changes direction. So the "spin-flip" only occurs in the frame of the magnetic field, changing the handedness of the transition driven by the slowing laser. In fact, keeping the spins aligned in the same direction as the atoms pass between the two slower coils is the one experimental difficulty of spin-flip slowers, since the change in direction of the magnetic field between the coils has a tendency to destroy the polarization of the atoms. This problem is solved by adding several inches of constant low magnetic field between the two coils to give the atoms time to be re-polarized by the circularly-polarized slowing light.

4.2.3 Designing an Optimized Slower

The most important parameter to consider when designing a slower is its length. Since atoms can only be slowed optically at a finite rate, a longer slower is necessary to capture faster moving atoms. The maximum acceleration rate a_s of an atom is determined primarily by three things. First, the momentum transfered to the atom per scattered photon simply depends on the wavelength of the photon: $p_{recoil} = h/\lambda$. Second, given the recoil momentum of a single photon, the change in velocity per scattered photon scales inversely with the mass of the atom: $v_{recoil} = p_{recoil}/M$. Third, with a given laser intensity, the atoms scatter photons at a finite rate R (in the limit of infinite laser intensity, R is simply the saturated scattering rate, $\Gamma/2$). Given the maximum acceleration rate and the maximum velocity v_{max} of atoms to be slowed, the length of the slower must be at least:

$$L_s = \frac{1}{2}\frac{v_{max}^2}{a_s} \tag{4.1}$$

This equation holds for both Zeeman and chirped slowers. In practice, the slower has to be somewhat longer than L_s to allow for imperfections in the slower; if the field/chirp changes too rapidly at any point/time, atoms will not stay in resonance and will leave the slowing process. On a side note, L_s scales inversely with a_s which in turn scales inversely with the mass of the atoms. However, at a given temperature, the mean velocity squared also scales as M^{-1}. As a result, for a given oven temperature L_s is mass independent.

An optimized slower should balance the benefits and drawbacks of increasing or decreasing the slower's length. Equation (4.1) tells us that a longer slower can capture a larger portion of the thermal velocity distribution. However, by making the slower longer, we also decrease the solid angle of the atomic beam, reducing the fraction of the atoms leaving the oven which reach the end of the slower. (In addition, since scattered photons not only slow the beam longitudinally, but also heat the beam transversely, atoms which are slowed from higher velocities will be transversely hotter at the end of the slower.) If we are only interested in the highest flux of slow atoms from an ideal thermal source, we simply compare the gain obtained by catching more of the thermal velocity distribution versus the loss in solid angle from a longer slower. For simplicity I will assume that the oven is a point source of atoms positioned at one end of the slower, and that the flux of hot atoms per unit solid angle is a constant across the beam. This is a good approximation if the distance between the oven and the slower is very small compared to the length of the slower, and if the end of the slower subtends an angle much less than 1 radian (viewed from the oven nozzle). In this case, the flux of hot atoms through the slower as a function of velocity is simply

$$Q_{hot}(v)dv = Q_0(v)\Omega dv = Q_0(v)\frac{A_s}{L_s^2}dv \tag{4.2}$$

where A_s and L_s are the cross-sectional area and length of the bore of the slower, and $Q_0(v)$ is the flux per unit solid angle from the oven. If the beam is collisionally dilute, $Q_0(v)$ goes as

$$Q_0(v) \propto v^3 e^{-Mv^2/2k_BT} \tag{4.3}$$

where M is the mass of the atom, k_B is the Boltzmann constant, and T is the temperature of the oven. The flux of slow atoms is just the integral of $Q_{hot}(v)$ from zero velocity up to the maximum capture velocity:

$$Q_{slow} = \int_0^{v_{max}} Q_{hot}(v)dv \propto \frac{A_s}{L_s^2}\left[1 - \left(1 + \frac{M}{2k_BT}v_{max}^2\right)e^{-\frac{Mv_{max}^2}{2k_BT}}\right] \tag{4.4}$$

Combining this with Eq. (4.1), we find that the flux of slow atoms as a function of the slower's length goes as

$$Q_{slow} \propto \frac{A_s}{L_s^2}\left[1 - \left(1 + \frac{Ma_s}{k_BT}L_s\right)e^{-\frac{Ma_s}{k_BT}L_s}\right] \tag{4.5}$$

This is a monotonically decreasing function of L_s, which implies that the best slower is one of zero length. It must be remembered, however, that the production of a dense atomic vapor involves the use of a hot, usually somewhat dirty, oven. Some method must be used to isolate the UHV trapping region from the contaminant gases coming from the source. In addition, if the dense thermal beam is allowed to collide with atoms in the magnetic trap, the magnetic trap's lifetime will be severely limited. So an effective method of shuttering the atomic beam is necessary. Both of these details require the atomic beam source to be separated from the trapping region by some finite distance, and putting a small slower into the gap between the oven and the MOT will maximize the flux of slow atoms.

There are two reasons why the slower should be made longer than just the minimum size of the gap between the oven and the trap. First of all, the above derivation assumed that the output of the

oven abuts the end of the slower. But in reality there is usually some distance between the oven and the slower. For example, in our new apparatus there is a beam shutter, a cold plate to collimate the sodium beam, a differential pumping tube, and a gate valve which all have to be placed between the oven and the slower (see Chapter 3). If we include the space between the slower and the oven, denoted as L_{os}, in the above derivation, Eq. (4.5) turns into

$$Q_{slow} \propto \frac{A_s}{(L_s + L_{os})^2} \left[1 - \left(1 + \frac{M a_s}{k_B T} L_s \right) e^{-\frac{M a_s}{k_B T} L_s} \right] \tag{4.6}$$

It is readily apparent that the maximum of Q_{slow} is no longer at $L_s = 0$. For a sodium beam from a 550 Kelvin oven placed about 45 cm from the slower (as is the case in our apparatus), the flux is optimized if the slower is about 40 cm long. Of course this assumes that the atoms are slowed at the maximum possible rate, leaving no room for any flaws in the slower. A more cautious design in which the atoms are slowed at half the maximum acceleration is optimized at a length of about 60 cm.

The second reason that a longer slower might be favorable is due to collisions in the atomic beam. The flux of atoms from an oven can be increased simply by turning up the oven temperature. Beyond a certain point, however, the beam will become collisionally dense, resulting in the narrowing of the velocity distribution of the beam as faster atoms collide into the back of slower atoms, transferring some of their momentum. This process of beam "jetting" removes atoms from the lower and higher velocity tails of the distribution and pushes them towards the mean atom velocity. While the flux from a short slower is diminished by this process, a slower which is long enough to capture atoms faster than the mean velocity of the distribution ($\bar{v} = \frac{3}{2}\sqrt{\frac{\pi k_B T}{2M}}$) will actually *gain* from jetting, since the narrowing of the velocity distribution about the mean results in a higher fraction of atoms which are below the maximum velocity which can be slowed. For sodium at 550 K, the mean velocity is about 800 m/s and the required slower length is about 40 cm long.

Once again, to allow for flaws in the construction of the slower, we should only attempt to slow atoms at a fraction of the maximum acceleration, resulting in a slower which is somewhat longer.

4.2.4 Design and Construction Details of the New Slower

As shown in the preceding sections, an optimal slower for our design constraints would be a spin-flip Zeeman slower about 60 to 80 cm long, and should be capable of slowing atoms traveling faster than the 800 m/s mean velocity of the thermal beam. The slower which was designed and built consists of two coils with a total length of 80 cm plus a 23 cm gap between them (to allow space for a vacuum pumping port, and to allow time for the atoms to be re-polarized between coils). The slowing transition used is the $F = 2$, $m_F = 2$ to $F = 3$, $m_F = 3$ cycling transition. The slowing beam is circularly polarized to prevent off-resonant scattering out of the $F = 2$, $m_F = 2$ ground state. Sidebands at 1.71 GHz are added to the slowing laser beam to generate $F = 1$ to $F = 2$ light to pump atoms into the $F = 2$ ground state before entering the slower.

The laser detuning was selected to be 500 MHz (50 linewidths). This detuning was chosen to be as small as possible without greatly perturbing the MOT. Given this detuning, the initial field of the slower was chosen to be 720 gauss in order to slow atoms with velocities up to 888 m/s. The final field of the slower was designed to be -300 gauss, giving the atoms a final velocity of 50 meters per second. In practice, we found that the loading of the MOT was optimized when the final field was somewhat smaller (about 280 gauss), and when the laser detuning was slightly larger (508 MHz) indicating that the capture range of our MOT is larger than we originally expected.

One of the primary concerns when building the slower was heat management. In order to effectively cool the slower, the large coil was wound around a nested pair of (non-magnetic) stainless steel tubes. The inner, 1" outer-diameter tube constitutes the UHV vacuum bore of the slower, and is welded to conflat flanges at either end. A concentric 2" outer-diameter tube, around which the

magnet coil is wound, surrounds the inner tube. Water is flushed between the inner and outer tubes to carry away heat generated by the coils. In addition, a loosely packed winding of water cooled copper tubing was wound around the outside of the large coil to assist in the cooling of the coil.

The small coil was also wound around a pair of tubes, with the gap between them being flushed with cooling water. The center bore of this tube, however, is not part of the vacuum system, but fits around the small quartz tube of the quartz cell, allowing the small slower coil to be in close proximity to the trapping region (see Fig. 3.1). In order to minimize the radius of the small slower coil (and thereby decrease the length scale over which the fringing fields decay), the brass tubes used are just large enough to fit over the quartz tube. As a result, this coil had to be installed on the quartz cell before the cell was fused to the vacuum flanges.

Both coils were made with materials capable of withstanding a 250° C vacuum bake-out. In addition to the two main slower coils, an extra "compensation coil" is wound around a spool which covers the end of the small slower coil. This coil was designed to cancel the residual bias and gradient produced by the slower at the center of the magneto-optical trap.

4.2.5 Performance of the New Slower

The performance of the new slower was evaluated by measuring the absorption of a laser beam as it passed through the slowed atomic beam. The probe beam was sent through at 45° from the axis of the atomic beam, making the probe sensitive to the Doppler shift of the atoms, and the frequency of the laser was swept to map out the velocity distribution of the atoms in the slowed beam. In order to correct for noise due to laser power fluctuations, a second reference beam was split off and measured on a second photodiode. This simple setup allowed us to measure absorption to better than one percent. The results of the measurements are shown in Fig. 4.2. From these measurements we found that the slower was able to produce a flux of a few times 10^{11} slow atoms per second at the center of the MOT.

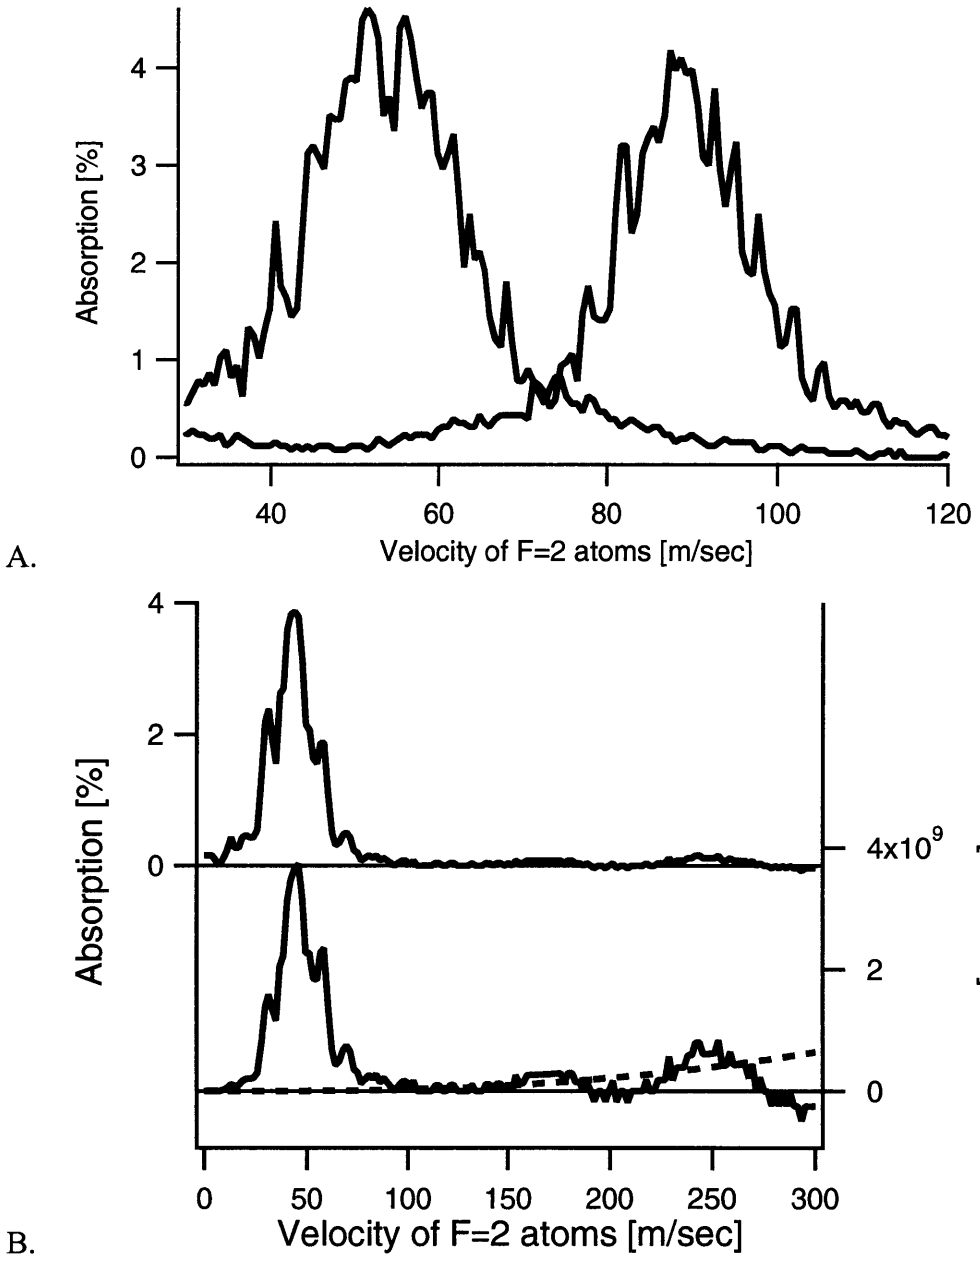

Figure 4.2: Absorption spectra of slowed atoms. Figure (A) shows two absorption spectra of atomic beams slowed to two different final velocities by modifying the current in the small slower coil. Figure (B) shows a larger portion of an absorption spectra, showing that atoms are swept from higher velocities in the thermal distribution down into a narrow distribution of slower velocities. The upper trace in figure (B) is the raw absorption data. In the lower trace, absorption has been converted to atom flux. The dotted line represents the theoretical velocity distribution of an unslowed thermal beam.

4.3 Magneto-Optical Trapping and Cooling

Atoms from the slower are collected, cooled, and compressed by a dark SPOT magneto-optical trap. The MOT parameters in the new apparatus are almost identical to those of our first BEC apparatus, described in [75, 76]. The implementation has been altered somewhat in order to improve stability, increase the capture and trapping volumes, and to simplify operation. Detailed overviews of magneto-optical trapping can be found in [69, 75].

The maximum number of atoms which can be held in a MOT is limited by "radiation trapping," a process in which light scattered by one atom is rescattered by another atom, generating an effective repulsive force between the atoms [77]. In order to maximize the number of atoms which can be captured in the MOT, a dark SPOT configuration is used. In a dark SPOT MOT, atoms are kept in a dark state which doesn't interact with the trapping beams. A repumping beam is then used to pump the atoms into the bright state and cause them to interact with the trapping light. A dark shadow at the center of the repumping light is created by passing the repumping light through a glass slide with an opaque spot on it, and imaging the spot into the center of the MOT. As a result, atoms in the center of the trap are only repumped weakly, and spend only a small fraction of their time in the bright state. This effectively lowers the optical density of the cloud, allowing more atoms to be trapped in a given volume. Radiation trapping is still the limiting factor in a dark SPOT MOT. The limit, however, is greatly reduced; the dark SPOT has been shown to increase the density of trapped atoms by two orders of magnitude over a conventional MOT [78].

4.3.1 Implementation

One of the most important parameters in a magneto-optical trap is the diameter of the laser beams. The quartz cell in our experiment measures 32 mm on a side (with an optically flat clear "window" of $\sim$ 25 mm), allowing the MOT beams to be quite large. The large beams present a longer capture

region which is capable of catching slightly faster atoms from the atomic beam. More importantly, large MOT beams provide a larger trapping volume. This is important because radiation trapping sets an effective limit on the optical density of the atoms in the trap. By increasing the linear dimension of the trap, the number of atoms which can be confined at a given optical density scales up; since the density n scales as N/r^3, and the optical density scales as nr, (where N is the number of atoms in the MOT, and r is the radius of the MOT), at constant optical density the number of atoms scales as r^2.

The MOT is made from six independent laser beams (four in the horizontal plane, at ± 45 degrees to the axis of the slower, plus two vertical beams). Most MOT setups recycle some of the beams; in our first BEC apparatus, two beams are retro-reflected to produce counter propagating beams. Although this approach makes more efficient use of available laser light, we chose to use six independent beams to avoid shadows in the beams and to limit deterioration of the light's polarization. A MOT beam traveling through a dense cloud of atoms can suffer from a large amount of absorption. If this beam is retro-reflected, the retro-reflected beam will be weaker, and the MOT will not be balanced. This can be compensated for by focusing down the diameter of the retro-reflected beam. The shadow of the atoms, however, creates unavoidable spatial variations in the retro-reflected beam. In addition, since the horizontal beams pass through the glass cell at 45 degrees from normal incidence, the circular polarization of these beams is degraded. A retro-reflected beam would have to pass through two additional 45 degree surfaces before returning to the atoms, further degrading the beam's polarization.

Before entering the apparatus, each of the six beams passes through a telescope assembly which expands them to a Gaussian diameter of about 1.5 inches, a bit larger than the glass cell. By over-expanding the beams we are able to throw away light in the noisy wings of the laser mode, using just the center of each beam where the mode is the best. The telescope assemblies each contain a quarter-wave plate which produces the appropriate circular polarization for each beam. The

lenses and polarizer in each telescope are mounted into dust-tight lens tube assemblies to lower the required maintenance of the MOT.

The six MOT beams are split from a single beam arriving from the laser table with the use of polarizing beam splitters. A half-wave plate in front of each beam splitter allows us to control the splitting ratio of each beam splitter. The beam is first split into a horizontal and vertical beam. The horizontal beam is then split into two, and then each of these three beams is split again to produce three pairs of beams. Each beam in a pair is routed to opposite sides of the experiment to produce counter-propagating MOT beams. This setup is somewhat more cumbersome than splitting the beam into a left-hand side and right-hand side beam, and then further splitting the beams after they have been routed to the correct side of the experiment. But this complication is necessary to give us independent control of the balance between each pair of counter-propagating MOT beams. Fine balance between the intensities of counter propagating beams is necessary to produce a large, stable MOT and for the implementation of polarization gradient cooling (Section 4.3.2).

As each pair of counter-propagating beams is routed around the experiment, care is taken so that the number of reflections in one beam of a counter-propagating pair differs from the number of reflections in the other beam of the pair by an odd number (including the reflection in the beam splitter). Furthermore, the length that each beam travels is made to be roughly the same. This assures that any jitter in the beam pointing or in the mode of the beam will be of a similar nature and in a common direction for both counter-propagating beams. This helps to stabilize the MOT against beam pointing and mode fluctuations.

Due to the proximity of the reflective surfaces of the quartz cell to the atoms (and due to the lack of an anti-reflective coating on the cell), light from the dark SPOT repumping beam can be scattered into the center of the trap, reducing the "darkness" of the dark SPOT. This problem was solved by using linear polarized repumping light sent in at a 45° angle to the surface of the quartz cell. By sending the beam in near Brewster's angle, reflections were minimized (we couldn't send

it *exactly* at Brewster's angle because the MOT beam telescopes were in the way). In addition, at this angle the small amount of light that is reflected misses most of the MOT trapping volume. To get the beam in at this angle, it was necessary to overlap the repumping light with one of the MOT beams using an edge mirror (placed near the focus between the MOT beam's two main telescoping lenses). We also found that to get the best MOT, we had to mask the repumping light with a pair of razor blades to prevent the edges of the beam from scattering off of the pinch coils of the magnetic trap and into the MOT.

The slowed atomic beam enters the trapping region with a velocity of about 60 m/s and a density on the order of 10^9 atoms per cm^3. When the MOT is filled, the trapped atoms have been compressed to a density of $\sim 10^{11}$ cm^{-3} at a temperature of a few hundred microkelvin.

4.3.2 Polarization Gradient Cooling

The final phase of laser cooling consists of a short burst (3-5 ms) of polarization gradient cooling. Polarization gradient cooling is a sub-Doppler method of cooling, and is discussed in [69, 70]. It is implemented in our experiment using the same beams used for the magneto-optic trap. After the MOT is filled to capacity, the magnetic quadrupole field is suddenly switched off, the MOT repumping light intensity is lowered, and the MOT beams are shifted a few MHz further from resonance. In this configuration, the atoms experience a "dark molasses" which quickly cools them to ~ 100 μK, a bit more than a factor of two below the Doppler limit for sodium.

Three sets of Helmholtz coils tuned to cancel the earth's magnetic field, as well as any stray fields, are on during the entire cooling process. These coils are required because the presence of magnetic fields breaks the symmetry for absorption of $\sigma_+$and σ_- light, causing the atoms to be pushed into a moving frame by the polarization gradient cooling laser beams. In addition to stray magnetic fields, polarization gradient cooling is also very sensitive to the intensity bal-

ance and alignment of counter-propagating laser beams. In order to increase the stability of our MOT/polarization gradient beams, our first attempt at magneto-optical trapping on the new apparatus was done using light from six single-mode fibers. By rigidly fixing the output of each fiber less than half a meter from the trap, beam jitter and mode distortions were eliminated. Unfortunately, reflections at the ends of the fibers caused etaloning to occur, so the transmission of the fibers was extremely sensitive to mechanical noise and drift. As a result, it was impossible to get polarization gradient cooling to work consistently, and the fibers were soon replaced with mirrors. It is possible that the problems with fibers could be fixed, however, by using angle cleaved fibers which prevent reflections from being coupled back into the fiber.

4.4 The Laser System

All of the sodium resonant light used for laser cooling, as well as for optical pumping and probing, is provided by a rhodamine 590 ring dye laser pumped by a large frame argon ion laser[1]. Light from the dye laser is split off into many beams. Some of the beams are used for diagnostic purposes and to lock the laser to the sodium resonance. The other beams are shifted to an appropriate frequency, often with the addition of sidebands, for use in some stage of the experiment. The dye laser is a Coherent 899 model, modified by the addition of a high pressure dye nozzle and dye recirculation pump. These modifications allow the laser's dye jet to run at higher pressures, resulting in higher laser output power. In addition, the modified nozzle provides greater stability and simplifies alignment of the laser. The dye is cooled to 5° C to increase the viscosity of the dye, preventing the onset of turbulence at high pressures. In typical operation, the dye laser produces 1.5 watts of power with 10 watts of multi-line visible pump power from the argon ion laser.

[1] Coherent model Inova 100

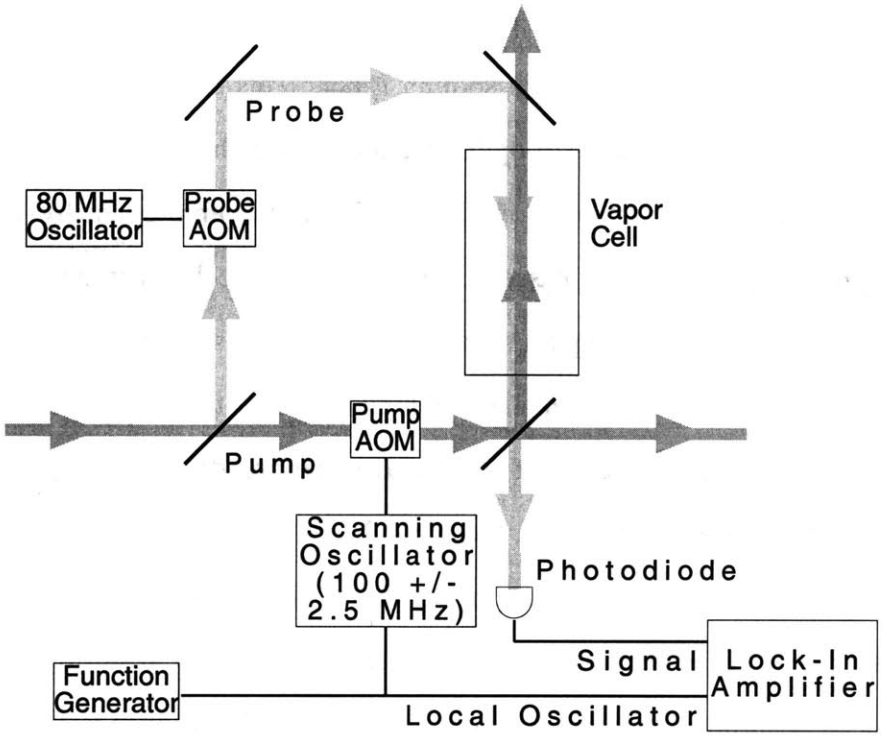

Figure 4.3: Saturated absorption laser locking setup. Using saturated absorption spectroscopy, the laser is locked to the sodium $F = 2$ to $F = 3$ ground state transition. The frequency of the pump beam is dithered, allowing lock-in detection of the probe beam absorption. This both enhances the signal-to-noise and generates a dispersive zero crossing at the lock point (making the lock point insensitive to laser intensity fluctuations).

4.4.1 Locking the Laser

The dye laser, in its off-the-shelf configuration, is locked to a temperature stabilized Fabry-Perot cavity, giving the laser an instantaneous linewidth of about 0.5 to 2 MHz. Although this is less than the natural linewidth of the sodium atoms (10 MHz), over longer time scales the laser tends to drift in frequency by a much larger amount. Frequency drift is prevented by locking the laser to a sodium vapor cell using saturated absorption spectroscopy. This locking method is illustrated in Fig. 4.3.

Saturated absorption spectroscopy is a common technique which produces spectral lines with

the Doppler free linewidth from Doppler broadened gases. In our setup, part of the light from the dye laser is sent to a cell containing a sodium vapor. This beam is split into two parts: a pump beam and probe beam. An acousto-optic modulator shifts the frequency of the probe beam up by 80 MHz. A second acousto-optic modulator shifts the pump beam by 100 $\pm$2.5 MHz. The pump and probe beams are then passed through the vapor cell, overlapping but passing in opposite directions. The more intense pump beam pumps the atoms which are Doppler shifted into resonance with it into a different ground state, making the vapor transparent to this frequency of light. If the weaker probe beam is resonant with the Doppler broadened vapor, but is not resonant with the same velocity group of atoms as the pump beam, it will be heavily attenuated by the vapor. Otherwise, the probe will "pass through" the "hole" generated by the pump beam with little attenuation. As a result, if the laser is scanned through resonance, the resulting spectra is a large Doppler broadened absorption dip with transmission peaks at the resonance frequencies of each ground state transition, each with a linewidth comparable to the natural linewidth of the transition. In addition, there are several other "cross-over" peaks which correspond to atoms pumped from one ground state and probed in another.

The laser is locked to the saturated absorption cell using a lock-in technique. This is accomplished by dithering the frequency of the pump beam by $\pm$2.5 MHz at a rate of 27 kHz. The dither signal, as well as the photodiode signal, are sent to a lock-in amplifier. Not only does the lock-in method improve signal to noise, but it also turns the transmission peaks into dispersive zero crossings, making the lock point insensitive to laser power fluctuations. The output of the lock-in amplifier is integrated with a time constant of 3 ms and sent to the laser's external frequency scan input. The result is a laser which is locked 90 MHz below the sodium $F = 2$ to 3 transition.

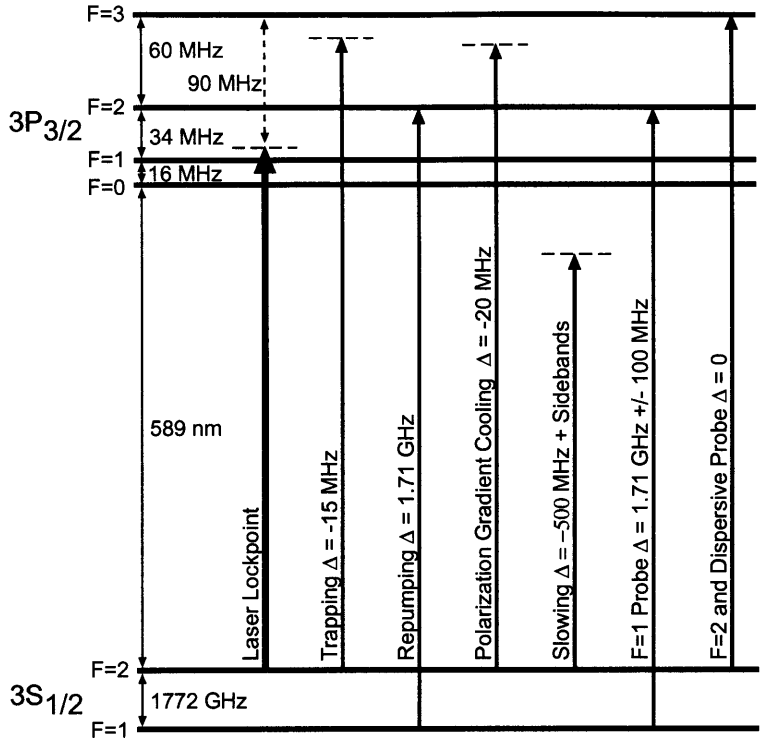

Figure 4.4: Energy level diagram of the $3S_{1/2}$ to $3P_{3/2}$ transition in sodium. The lock point of our laser and the atomic transitions used to trap, cool, and probe the atoms are depicted. To generate each of the required frequencies, light from the laser must be shifted by the amount Δ listed on the diagram (the detuning from the $F = 2$ to 3 transition) plus the 90 MHz detuning of the laser from the $F = 2$ to 3 transition. In addition to frequency shifts, 1.71 GHz sidebands are placed on the slowing light to pump atoms into the $F = 2$ state before entering the slower.

4.4.2 Generating Frequencies

The new apparatus requires laser beams at many different frequencies in order to trap, cool, prod, and probe the atoms. Figure 4.4 shows all of the frequencies of light which are necessary to run the experiment. To generate all of these different beams from a single laser, several acousto- and electro-optical modulators are used. A block diagram of our laser setup is shown in Fig. 4.5.

In addition to the near resonance beams, other lasers are often used to generate far detuned light, depending on the experiment being performed. In particular, light from an infrared Nd-

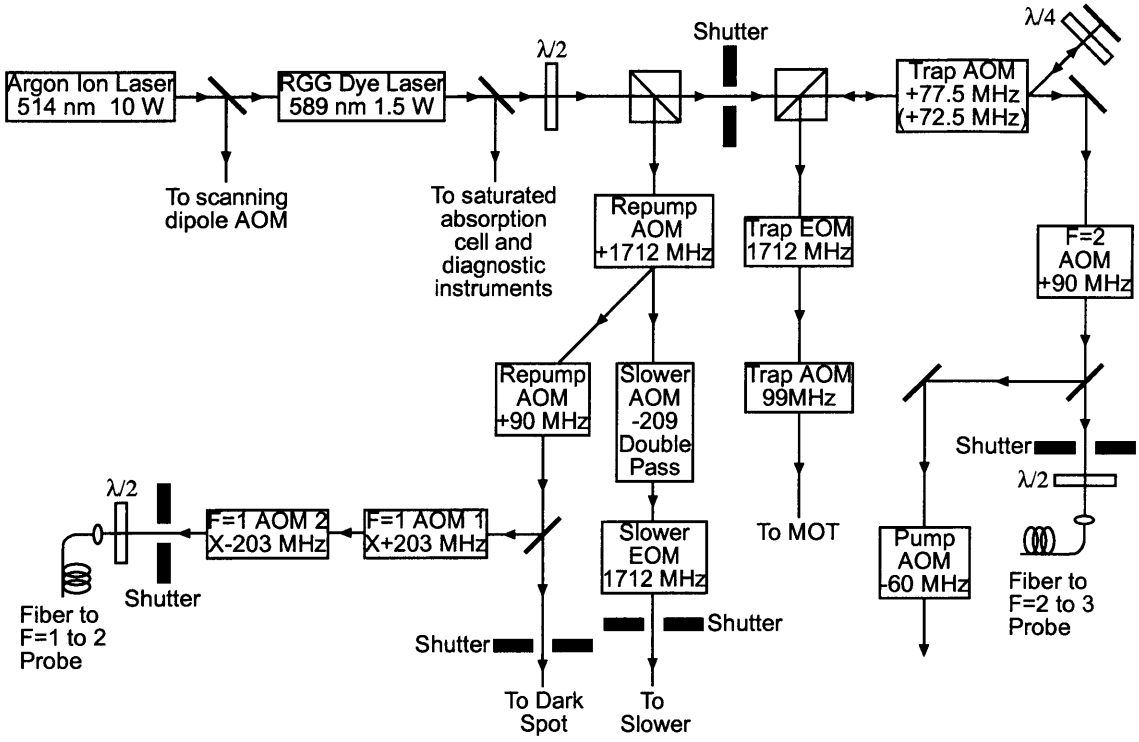

Figure 4.5: Laser setup. Various AOMs are used to generate the frequencies of light needed for the experiment. Before entering the dye laser, some of the light from the argon laser can be split off and used to generate dipole potentials (see Chapter 7). The light produced by the dye laser (locked 90 MHz below the $F = 2$ to 3 transition) is split up and shifted in frequency using several acousto-optic modulators (AOMs). In most cases, the first order of each AOM is used, and the zeroth order discarded. The two exceptions are the "Trap AOM" (in which case the zeroth order is sent on to the "F=2 AOM", and the first order is retro-reflected for a second pass through the AOM and then separated from the incoming beam with a polarizing beam splitter) and the 1712 MHz "Repump AOM" (in which case the zeroth order is sent on to the "Slower AOM" to generate the beam used by the slower, and the first order is sent to the 90 MHz "Repump AOM" for use as repumping light in the dark SPOT trap and for probing of atoms in the $F = 1$ ground state). The two AOMs labeled "F=1 AOM" are used to generate a scanning probe. The AOMs are aligned in such a way that as the AOMs are scanned, the change in deflection angle of one AOM is compensated for by the deflection in the other AOM. By shifting the frequency of one of them up, and the other one down, the probe light can be detuned while keeping the beam aligned on the fiber coupler.

YAG laser and light from our argon ion laser have been used to generate far-off-resonance dipole potentials. The use of far-off-resonance lasers is discussed in Chapter 7.

4.5 Controlling the Experiment

Many coordinated events must occur during each phase of an experiment. Magnetic coil currents and laser beams have to be switched on and off at the right junctures, laser beams need to be shifted in frequency at several times, and RF magnetic fields have to be swept in frequency during evaporation (see Section 5.6). All of the switching and timing operations are controlled by a single computer. The computer has a pair of I/O boards which allow it generate 32 digital (TTL level) signals, as well as ten 12-bit analog signals. The state of each channel can be changed with microsecond precision. An arbitrary function generator board provides an additional four analog channels, used to control the scanning dipole trap discussed in Section 7.2. External frequency synthesizers, such as the one used for evaporative cooling, are programmed by the control computer through a GPIB bus. The software for the control computer was written specifically for this experiment, and has been implemented in a way which allows us to very simply and quickly alter the fashion in which an experiment is performed.

Chapter 5

Magnetic Trapping

During the final stage of cooling the atoms are confined by an inhomogeneous magnetic field. Magnetic traps are ideally suited for the last cooling stage for two reasons. First, magnetic traps can be made with very low heating rates such that they don't compete with the cooling process. Second, magnetic traps facilitate a very effective method of cooling — RF evaporation. In the following chapter, several details of magnetic trapping and the magnetic trap designed for our new apparatus will be discussed. More general information on magnetic trapping is contained in Appendix C.

5.1 Overview of Magnetic Trapping

The trapping potential felt by an atom in a magnetic trap is generated by the interaction of the atom's magnetic dipole moment with the magnetic fields of the trap. A classical magnetic dipole interacting with a magnetic field has a potential energy $U_m = -\boldsymbol{\mu}_m \cdot \mathbf{B} = -\mu_m B \cos\theta$ (quantum-mechanically, the energy levels in a magnetic field are $E(m_F) = g\mu_B m_F B$, where g is the g-factor and m_F is the quantum number of the z-component of the angular momentum, and the classical

term $\cos\theta$ can be related to the quantum mechanical quantity m_F/F). A dipole can lower its energy by moving to a lower field, or by changing its orientation such that it is more closely aligned with the magnetic field.

The angle between the dipole moment and the magnetic field is constant when an atom is exposed to a homogeneous field ($\boldsymbol{\mu}_m$ simply precesses around the magnetic field at constant θ). As a dipole moves through an *inhomogeneous* field, it experiences a local magnetic field which is changing in both magnitude and direction. If the field changes slowly enough in the frame of the atom, however, adiabaticity requires that the angle θ be unchanged (or, in a quantum picture, the atom must remain in one m_F state). Therefore an isolated atom moving slowly enough through an inhomogeneous magnetic field feels a conservative potential simply proportional to the magnitude of the field. To trap atoms in such a potential, a local minimum of the magnetic potential energy $E(m_F)$ must be created. For $gm_F > 0$ (weak-field seeking states) this requires a local magnetic field minimum. Strong-field seeking states ($gm_F < 0$) cannot be trapped by static magnetic fields, since a magnetic field maximum in free space is not allowed by Maxwell's equations [79, 80].

Collisions Since magnetic traps only confine atoms in the weak-field seeking state, atoms making a transition to a strong-field seeking state will be lost from the trap. In addition, transitions between different weak-field seeking states can also be problematic, since these transitions can release internal state or magnetic field interaction energy as kinetic energy, thereby heating up the atoms. As discussed above, transitions to other m_F states will not occur for slowly moving, isolated atoms. In the presence of other atoms, however , it is possible for collisions to occur which change the atom's internal state.

By trapping atoms which are spin-polarized along the field axis ($m_F = \pm F$), angular momentum conservation prevents the simplest m_F changing collisions in which two atoms exchange internal angular momentum. Other spin-changing processes are, however, still possible, the most

notable being dipolar relaxation (in which orbital angular momentum is converted into internal angular momentum) and inelastic three body collisions (which become limiting at high densities). Fortunately, for spin-polarized sodium ($F = 1, m_F = \pm 1$) in the density range (10^{11} to 10^{14} atoms per cm^3) and temperature range (50 nK to 100 μK) typical for magnetically trapped atoms in our experiment, these processes are slow compared to the rate of elastic, spin-conserving collisions (see Fig. 2.1). This allows evaporative cooling to work efficiently (see Section 5.6), and makes it possible to cool the atoms through the BEC transition before the sample has decayed from the trap.

Majorana flops If an atom moves too quickly through an inhomogeneous field, non-adiabatic transitions to other m_F states can occur. An atom's magnetic moment will adiabatically follow the direction of the magnetic field only if the rate of change of the field's direction θ (in the atom's moving reference frame) is slower than the precession of the magnetic moment: $\frac{d\theta}{dt} < \frac{\mu_m |\mathbf{B}|}{\hbar F} = \omega_{Larmor}$. If the field's direction changes faster, atoms will make so called "Majorana" spin-flips [81] and be lost from the trap. Since the precession rate of an atom scales with the magnitude of the local magnetic field, adiabaticity is most easily violated in the center of the trap where $|\mathbf{B}|$ is smallest.

In the frame of a moving atom, the rate of change in the magnetic field direction is given by $\frac{d\theta}{dt} = v\frac{d\theta}{dl}$ where v is the velocity of the atom, and $\frac{d\theta}{dl}$ is the change in magnetic field direction with distance in the direction in which the atom is moving. Although the atoms in the trap have a thermal distribution of velocities, and the derivative of the field direction $\frac{d\theta}{dl}$ is not the same at all points and along all trajectories in the trap, an upper bound on the susceptibility of a harmonic trap to Majorana flops can be obtained by assuming that the direction of the magnetic field rotates by a maximum of π radians as an atom travels from one side of the trap to the other. Knowing that an atom cannot traverse this distance faster than the trapping frequency, this gives a maximum field rotation of $\frac{d\theta}{dt} = \frac{\omega_{trap}}{2}$ and a condition for trap stability against Majorana spin-flips of $\omega_{Larmor} >$

$\frac{\omega_{trap}}{2}$ (where ω_{trap} is the harmonic frequency of the trap). For the field configuration of our new trap (an Ioffe-Pritchard configuration — see Section 5.2), the field only changes direction when the atoms oscillate radially, so the radial trap frequency is the relevant ω_{trap} to use in these equations.

5.2 Types of Magnetic Traps

Many types of magnetic traps have been used in the field of Bose-Einstein condensation of dilute gases. A description and comparison of many trap geometries are given in Appendix C. There are several figures of merit to consider when designing a magnetic trap. One figure of merit is the strength of confinement of the trap. This is important because more tightly confining traps can compress atoms to higher densities, thereby increasing collision rates within the trapped gas. With higher collision rates, evaporative cooling can be done more quickly and efficiently (see Section 5.6), and heating and trap loss due to collisions with residual gases in the vacuum can be avoided (see Section 3.1).

Linear traps As discussed in Appendix C.2.1, lower order potentials are more effective at confining atoms than are higher order potentials, with the optimum confinement being provided by a linear potential. Three-dimensional linear potentials can only be generated around a zero-crossing in the magnetic field (as shown in Fig. 5.1), and as a result, Majorana spin-flips are a serious concern in linear traps. For hotter atoms, the region where Majorana spin-flips occur is small relative to the size of the atom cloud, and the rate at which atoms are lost from the trap is minimal. However, Majorana flops in linear traps become more problematic as the atoms are cooled to lower temperatures. This can be explained roughly as follows; although the average velocity of the trapped atoms drops with temperature (decreasing the probability that a Majorana spin-flip will occur with each pass through the center of the trap), the colder atoms also settle into a smaller region near the

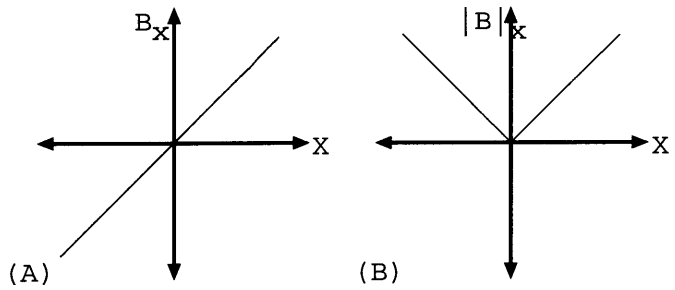

Figure 5.1: Linear magnetic traps. Although Maxwell's equations require that magnetic fields in free space have continuous derivatives, linear magnetic traps with a cusp at the origin are possible, due to the fact that the trapping potential is proportional only to the magnitude, not the direction, of the field. Figure (A) shows a magnetic field which varies linearly with displacement from the origin. Figure (B) shows the magnitude of the field (which is proportional to the trap potential), which has a cusp at the zero field crossing.

center of the trap (where the oscillation frequency is higher), causing them to pass through the trap center more frequently. A more rigorous model, discussed in [52] and in Appendix C.2.1, shows that the loss rate in a spherical quadrupole potential scales as T^{-2}.

To avoid Majorana losses, two approaches have been taken; higher order traps which do not have a zero crossing have been used, and the zero-field crossing in linear traps has been "patched" with AC magnetic fields [1] or optical potentials [2]. In optically plugged linear traps, the type of trap used in our first demonstration of BEC in sodium, a focused laser repels atoms from the low field region of the trap. As a result, only the very bottom of the of the potential is weakened — hot atoms are still tightly confined in a linear potential. Furthermore, the harmonic trapping frequencies characterizing the "rounded out" bottom of the trap are typically somewhat higher than than those produced in comparable higher order magnetic traps (the geometric mean of the frequencies for the optically plugged trap described in [2] is 415 Hz, compared to 200 Hz in our new Ioffe-Pritchard trap). There are a few added technical complications of an optically plugged trap; the beam has to be carefully aligned from day to day, and if care is not taken, fluctuations in beam pointing and intensity can induce heating. The tight confinement of an optically plugged

trap is especially useful for lighter atoms such as lithium, or atoms with weak magnetic moments.

In the other common "patched" linear trap, the TOP trap, a rotating magnetic bias is added to the linear potential, creating a parabolic time-averaged potential. The TOP trap is discussed in [1] and in Appendix C.2.1.

Ioffe-Pritchard traps The lowest order (and therefore tightest) trap which has a finite bias field is a harmonic trap. A magnetic trap with a finite bias field along the z-direction has an axial field of $B_z = B_0 + B''z^2/2$. Given this axial field, Maxwell's equations allow the leading term of the transverse fields to be linear. Applying Maxwell's equations (and assuming axial symmetry) leads to the following field configuration [82]:

$$\mathbf{B} = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2+y^2) \end{pmatrix} \tag{5.1}$$

The parabolic trap was first suggested and demonstrated for atom trapping by Pritchard [83, 84], and is similar to the Ioffe configuration developed earlier for plasma confinement [85]. Although the axial confinement of this trap is parabolic, for large hot clouds ($k_B T > \mu_m B_0$) the radial confinement is linear, characterized by the radial gradient B'. For smaller, cooler clouds ($k_B T < \mu_m B_0$), since the magnetic potential is created by the quadrature sum of the linearly increasing radial field and the axial bias field, the radial confinement is harmonic with a trapping frequency of $\omega_r = \sqrt{\frac{\mu_m B_r''}{M}}$, where B_r'' is the "effective" radial curvature, given by:

$$B_r'' = \frac{B'^2}{B_0} - \frac{B''}{2} \tag{5.2}$$

At a certain axial displacement the trap has a saddle point, beyond which the radial confinement

vanishes. This occurs because the radial components of the curvature term ($-\frac{B''}{2}xz$ and $-\frac{B''}{2}yz$) cancel the confinement provided by the the radial gradients ($B'x$ and $-B'y$). Since the sign of the gradient term is different in x and y, the instability in each axis occurs at opposite displacements along the z axis. From Eq. (5.1) the points of instability z_{inst} occur at:

$$z_{inst} = \pm \left(\frac{B'}{B''} - \frac{1}{2}\frac{B_0}{B'} \right) \tag{5.3}$$

When $B_r'' \leq 0$, the instability is at the origin, and there is no radial confinement at all. The saddle point at z_{inst} requires special attention when loading large clouds into an Ioffe-Pritchard trap, as discussed in Appendix C.2.3.

Comparison of Ioffe-Pritchard traps to TOP traps Appendix C.2.5 compares the confinement of Ioffe-Pritchard-type traps to the confinement of a TOP trap. It is found that the maximum compression for both traps is similar. Since the TOP trap is more complicated to build and use, it would seem that the Ioffe-Pritchard configuration has a certain advantage. As a result, the trap which we designed and built for the new apparatus is an Ioffe-Pritchard type trap.

5.3 A Modified Ioffe-Pritchard Trap

For reasons discussed above and in Appendix C, we chose to use an Ioffe-Pritchard type trap in the new apparatus. There are, however, many ways to wind coils which produce an Ioffe-Pritchard type of geometry. Several simple and effective winding patterns, such as the baseball [86, 82], yin-yang [82], three-coil [87, 88], and four-dee [89] trap, all lack the flexibility to tune the axial curvature and the radial gradient independently (although the radial confinement can still be varied by augmenting the bias field B_0 with an additional set of Helmholtz coils). The more complicated

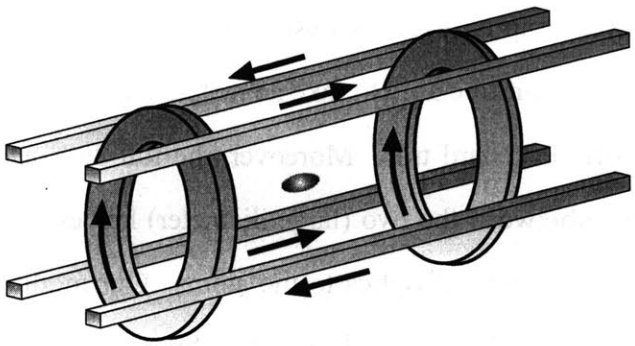

Figure 5.2: Traditional Ioffe-Pritchard trap. Four "Ioffe bars" produce a radial gradient field, confining the atoms radially. Two "pinch coils" produce a bottle field which provides the axial confinement.

cloverleaf trap, which we first demonstrated in 1995, has become popular due to its flexibility, and due to the fact that it provides 360° optical access to the atoms [53]. However, when compared to the traditional Ioffe-Pritchard winding pattern, this design is somewhat less efficient at producing the high radial gradients necessary for tight radial confinement. After considering the possibilities, we settled upon a design very similar to the traditional Ioffe-Pritchard winding pattern.

The traditional Ioffe-Pritchard trap consists of two pinch coils and four Ioffe bars, as depicted in Fig. 5.2. The Ioffe bars generate the radial gradient field B', while the pinch coils produce a bias field and the curvature term B''. Since the pinch coils and the Ioffe bars are very close to the atoms, this configuration is very efficient at producing a tight trapping potential. This traditional winding pattern has been used in several BEC experiments [90, 43].

The trap which we designed and built for the new apparatus varies slightly from the traditional Ioffe-Pritchard winding pattern. One modification we made was to elongate the pinch coils to allow clearance for the four horizontal MOT beams, each crossing at 45° to the axis of the trap (see Figs. 5.3 and 5.4). Although the elongated curvature coils do not provide the 360° optical access of a cloverleaf trap, they do offer more than 90° access on either side, in addition to allowing access between the coils for the atomic beam and the slowing laser beam. In practice, since much of the

optical access of a cloverleaf trap is typically obscured by the vacuum chamber, an experiment with a cloverleaf trap usually won't be able to take advantage of any more optical access than we have with our modified Ioffe-Pritchard trap. Moreover, the full 360° access to a cloverleaf trap is through the narrow space between the two (large diameter) halves of the trap (see Fig. C.2 in Appendix C), limiting the proximity of optics to the atoms. Our modified IP trap, on the other hand, allows for f/1 imaging in both radial and axial directions. Furthermore, since Ioffe bars are more efficient than cloverleafs, our modified Ioffe-Pritchard trap is able to produce radial field gradients about two times larger than a comparable cloverleaf trap. The trap which we designed is depicted in Figs. 5.3 and 5.4.

The elongated pinch coils break the axial symmetry of the trap, modifying the first order expansion of the magnetic fields given in Eq. (5.1). The more general expression is given by

$$\mathbf{B} = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -\alpha xz \\ -\beta yz \\ z^2 - \frac{1}{2}(\alpha x^2 + \beta y^2) \end{pmatrix} \tag{5.4}$$

where α and β depend on the asymmetry of the trap and are related by the expression $\alpha + \beta = 2$. This results in a trapping potential which is slightly elliptical in the xy plane, characterized not by a single effective radial curvature (see Eq. (5.2)), but by two different effective curvatures. The resulting "small displacement" potential is given by:

$$U \simeq \frac{\mu_m}{2}\left[B''_x x^2 + B''_y y^2 + B'' z^2\right] \tag{5.5}$$

$$B''_x = \frac{B'^2}{B_0} - \frac{\alpha B''}{2} \tag{5.6}$$

$$B''_y = \frac{B'^2}{B0} - \frac{\beta B''}{2} \tag{5.7}$$

Figure 5.3: Modified Ioffe-Pritchard Trap. In this rendering of our modified IP trap, each of the different coils can be seen. The shorter "racetrack" coils running left and right are the elongated pinch coils. The two long coils above and below the pinch coils make up the gradient producing Ioffe bars. The large round coils are divided into two sections, differentiated by different colored wire. The inner, lighter colored windings make up the "anti-bias" coils, and the darker outer windings make up an extra set of coils which can be used to generate large bias fields if they are required by an experiment. Bias fields produced by these extra coils are also used to weaken radial confinement of the trap (see Section 5.4).

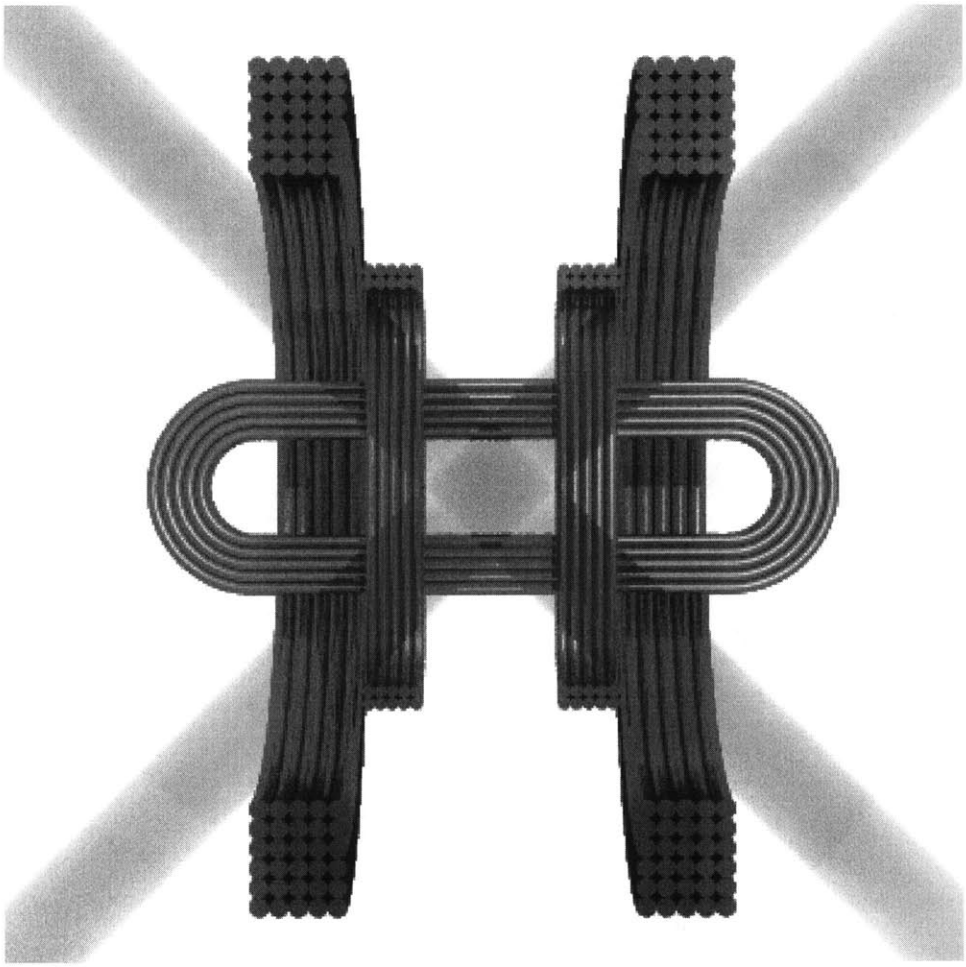

Figure 5.4: Cross-section of the modified IP trap. This is a cross-sectional view of the coils depicted in Fig. 5.3, showing how clearance is provided for inch-wide MOT beams crossing at 90°. With this coil design, optics can be placed above or below the gradient (Ioffe bar) coils, approaching within 0.8 inches of the center of the trap, with an aperture set by the coils which is slightly greater than one inch — allowing the possibility of f/1 imaging along the vertical axis. Horizontally, lenses can be placed just outside of the pinch coils, just over an inch away from the center of the trap, with an inch of vertical clearance — also allowing an f/# of about 1 (although our currently implemented imaging system is diffraction limited with an f/# of 4, limited by the size of the front lens rather than the clearance of the trap). For a sense of scale, the inner diameter of the large round coils is 6 inches.

In our trap, the B'^2/B_0 term dominates the $B''/2$ terms, and the difference in the effective x and y curvatures is small (about 10 percent when the trap is expanded to catch atoms from the MOT, and less than 0.1 percent when the trap is run at full current).

At large displacements from the origin, where the potential derived from Eq. (5.4) cannot be simply characterized by three harmonic frequencies, the asymmetry of the pinch coil has a larger effect on the trap. One result is a change in the position of the two saddle points (see Eq. (5.3)). Since the asymmetric pinch coils create a different radial curvature term in x and y, the instabilities in x and y are at different displacements along the z axis. If the pinch coil is elongated in the x dimension (resulting in $\alpha < \beta$), the point at which the confinement in x vanishes moves further from the center of the trap, while the instability in y moves closer to the origin (see Fig. 5.5). Unfortunately, when calculating the position of the unstable points from Eq. (5.4), the asymmetry creates a pair of quadratic terms (which cancel in the symmetric case), resulting in more complicated expressions. Simplicity is restored if we assume that we are only interested in points along the z axis for which $|z| \ll \sqrt{2\left(\frac{B'^2}{B''^2} - \frac{B_0}{B''}\right)}$. For the full range of operating parameters for our new trap, this restriction is valid if we only consider points which are displaced along the z axis by much less than 2 cm. This is *almost* a good assumption, since we typically catch clouds out of the MOT with radii of about 0.5 cm. In this approximation, the location of the saddle points are given by:

$$z_{instx} = \quad + \left(\frac{B'}{\alpha B''} - \frac{B_0}{2B'}\right) \tag{5.8}$$

$$z_{insty} = \quad - \left(\frac{B'}{\beta B''} - \frac{B_0}{2B'}\right) \tag{5.9}$$

In the extreme limit (a pinch coil which is stretched to be infinitely long while keeping B'' constant) we can find the instabilities without the above approximation. In this case, $\alpha = 0$ and

$\beta = 2$ (for our coils, α and β differ from these values by less than one percent), and the unstable points are found at:

$$z_{instx} \rightarrow \infty \tag{5.10}$$

$$z_{insty} = -\frac{2B'}{B''}(1-\gamma) \tag{5.11}$$

where

$$\gamma = \sqrt{\frac{1}{2} + \frac{B_0 B''}{2B'^2}} \tag{5.12}$$

Values for γ are generally very near to $\sqrt{1/2}$ (for our trap γ ranges from 0.78 when trap is weakened to catch atoms from the MOT, to 0.7 when it is running at full current). It is interesting to compare the instabilities in our asymmetric trap to a symmetric IP trap when the traps are weakened to catch atoms from the MOT (this is when the instabilities are most problematic). While the instability in x for our asymmetric trap is extremely far from the atoms, the unstable point in y is just one centimeter from the center of the trap, about a factor of two closer than it would be for a symmetric trap characterized by the same magnetic field bias, curvature, and gradient (this effect is illustrated in Fig. 5.5). As a result, extra care has to be taken to assure that the large cloud of atoms captured from the MOT sits well within this closer unstable point.

To tighten the radial confinement of the trap, the large (760 gauss) bias field generated by the pinch coils is canceled down to about 1.5 gauss by a pair of additional "anti-bias" coils. By lowering B_0, radial confinement is increased (see Eqs. (5.6) and (5.7)). A small bias is left to suppress Majorana spin-flips. The bias cancellation must be done carefully. If the bias field is under-compensated, the radial confinement of the trap will be weakened. If it is over-compensated, the field will cross zero and Majorana flops will occur (Fig. C.1). And if the bias field is not stable,

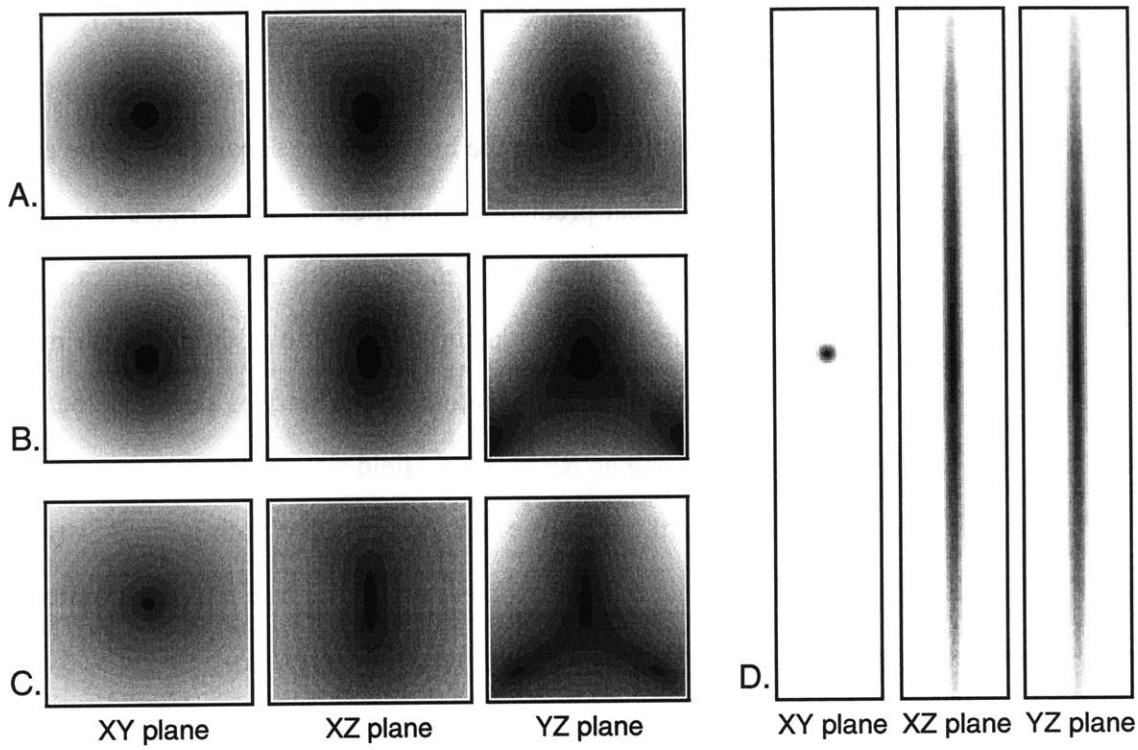

Figure 5.5: Magnetic field profiles of symmetric and asymmetric Ioffe-Pritchard traps. Rows (A) and (B) compare the magnetic potentials of symmetric and asymmetric Ioffe-Pritchard traps. Both plots were calculated using the typical parameters (B_0, B', and B'') of our trap when configured for catching atoms out of the MOT. In (A), the expansion for a symmetric IP trap was used (see Eq. (5.1)), while (B) includes the radial asymmetry of our modified IP trap (see Eq. (5.4)). The unstable point for confinement in y (the point at which the equipotential lines in the yz plane are perpendicular to the z axis) are just out of the field of view in (A), but are readily apparent much closer to the origin in (B). The unstable point for confinement in x is, once again, just out of the field of view in (A). In (B), this unstable point has moved far away, and the xz plane shows no visible asymmetry. Row (C) shows the field profile after the trap has been stiffened by raising currents to their maximum values. After RF cooling in the stiffened potential, the size of the cloud near the BEC transition temperature is much smaller, and the atoms only sample the very center of the potential, depicted in row (D). Although (D) is simply a close-up view of (C), we see that for small clouds, the asymmetry of the pinch coils has very little effect on the shape of the confining potential. The field of view in rows (A), (B), and (C) is 4 x 4 cm, and in row (D) is 100 x 700 μm. In rows (A) and (B) the grey-scale runs from 90 gauss (black) to 300 gauss (white). In image (C) the scale runs from 0 to 1500 gauss, and in image (D) runs from 0 to 0.2 gauss.

the trap parameters will be modulated, resulting in heating. Since the two sets of coils cancel each other to a few parts in 10^3, differential current ripple between the coils is magnified by this same factor. To prevent differential current noise, both the pinch coils and the bias-canceling coils are powered in series by the same power supply. To tune the bias, first the separation of the anti-bias coils was carefully set according to numerical predictions and measurements made during bench testing of the magnetic trap. The residual bias field was then measured, and fine adjustments were made by adding an additional large diameter three-loop coil to one side of the trap. We made the bias measurement in-situ using the atoms as our magnetometer. This was done by loading the magnetic trap with atoms, and then sweeping an RF magnetic field slowly down in frequency until all of the atoms had been expelled from the trap (see Section 5.6). By noting the RF frequency at which the trap was emptied, the magnetic field at the bottom of the trap was determined. Like the larger bias-reducing coil, the extra trimming coil was added in series with the curvature coils, allowing us to ramp current in the curvature coils to change the axial confinement without having to worry about disproportionate changes in the bias field.

5.4 Subtleties of Magnetic Trapping

Transfer from the MOT and adiabatic compression After the atoms are loaded into the magnetic trap, the cloud is compressed to increase the collision rate of the atoms by ramping the currents in the magnetic trap. Higher collision rates result in faster and more efficient RF evaporation (see Section 5.6). Ideally, when transferring atoms from the MOT to the magnetic trap, the magnetic trap is switched on with currents configured to generate the same potential felt by the atoms in the MOT. Otherwise, phase space density will be lost in the transfer, resulting in a loss in collision rate when the atoms are compressed. An equation describing the loss in collision rate resulting from a non-"mode matched" transfer is given in Appendix C.2.3.

In order to match the trapping parameters of the MOT, the magnetic trap must be much weaker and more symmetric than it is when running with full currents. To produce a weaker trap, the currents in the Ioffe bars and the pinch/anti-bias coils are reduced. The gradients can only be lowered to a certain point, however, since as they are lowered the unstable points move inwards (see Eqs. (5.3) and (5.11)). So to further weaken the radial confinement, an additional bias field is added using the trap's extra bias coils (see Fig. 5.3) .

Adiabatic decompression By the time the atoms have been cooled below the BEC transition temperature, the density of atoms in the magnetic trap has increased significantly — to much more than 10^{14} atoms/cm^3. At these high densities, trap loss and heating due to three-body collisions begin to limit the lifetime of the trapped atoms. Furthermore, the condensates formed in the compressed trap are very small, often making it difficult to image or manipulate the condensate accurately. These problems can be overcome by ramping down the magnetic currents to decompress the atoms.

Care must be taken to prevent excitations of the condensate during decompression. In our trap we found that decreasing the currents in the Ioffe bars and in the pinch coils independently caused the center of the magnetic trap to move. This is due to slight misalignments of the coils, giving each pair of coils a slightly different axis of symmetry, and to the small displacement of the trap minimum from the center of the coils due to the the fact that the Ioffe-bar coils aren't infinitely long (and therefore produce a field gradient in the z direction at the center of the coils). Similar misalignments in the cloverleaf trap of our old experiment were used to excite dipolar oscillations in our earliest studies of collective excitations [54]. In general, however, it is desirable to decompress the cloud *without* creating excitations. Trap displacement during decompression places an added restriction on the way in which the currents can be ramped down. In a perfectly symmetric trap, to achieve adiabatic decompression it is only necessary to limit the *rate* at which

the current is ramped down. But if trap displacement is a concern, the transition from a time independent trap to a trap which is slowly being weakened can also result in a transition from a trap which is not moving to, suddenly, one which is moving. So in addition to limits on the ramp rate, the *derivative* of the ramp rate must also be small.

We found that best results were achieved if we ramped currents in both sets of coils at a similar rate. We also found that below a certain current, oscillations became much harder to avoid. At this point, rather than decreasing currents in the Ioffe bars and the pinch coils, we had better success decompressing the condensate through the addition of a bias field.

Effects of gravity In a stiff magnetic trap gravity plays only a minor part in determining the trap potential. But due to the narrow linewidth of the RF transition, gravity could have a major effect on RF evaporation (see Section 5.6). As a result of the small displacement of the trap center by gravity, trap equipotentials are no longer surfaces of constant magnetic field strength, and as a result, the surface over which the RF is resonant is not concentric with the atom cloud. This can result in atoms being evaporated only near the point where an equipotential surface is tangent to the surface of constant $|\mathbf{B}|$ where the atoms are in resonance with the RF field, resulting in slow, inefficient evaporation. In our first attempts to achieve BEC in a cloverleaf trap, we suspected that this effect could be a limiting factor, and to prevent this problem, the trap was rotated such that the tight radial confinement (as opposed to the much weaker axial confinement) was opposing gravity. This same preventive measure has been taken with our new trap.

Effects of current noise Noise in the magnetic coil currents can translate into heat in the trapped atoms. The acceptable level of current ripple is discussed in Appendix C.3.1. Due to the large inductance of the coils, the susceptibility of the magnetic trap to high frequency noise and voltage ripple in the power supply is decreased. When selecting power supplies we found that there were

many commercially available power supplies with sufficiently low noise, and other than carefully selecting a good power supply, no additional measures had to be taken to stabilize the currents in the coils.

5.5 Design and Construction of the New Magnetic Trap

Our modified Ioffe-Pritchard trap was optimized using a numerical simulation of the fields generated by the coils. Several constraints had to be included in the optimization. Physical space constraints included the need for clearance for 1" MOT beams and clearance for high resolution imaging lenses. In addition, it had to be possible to adequately cool the coils with water pumped through them at a reasonable pressure. Furthermore, the coils had to generate the tightest confinement possible given the finite power provided by the power supplies. An additional constraint was that the z component of the gradient produced at the center of the Ioffe bars (as a result of the finite length of these coils) had to be small enough that differential noise in the curvature and gradient power supplies would modulate the center of the cloud by much less than the cloud diameter.

In the numerical simulation, the thickness of each wire was taken into account. But to simplify the calculations it was assumed that all of the current was concentrated in an infinitely thin channel running through the center of the wire. The calculations were further simplified by approximating the pinch coils and Ioffe-bar coils as having rectangular (rather than "race-track") windings. The optimized simulation resulted in the coils depicted in Fig. 5.3. The dimensions of the optimized coils are listed in Table 5.1, and the magnetic field parameters predicted by the calculation are listed in Table 5.2. Figure 5.6 shows the predicted magnetic potential for the current settings typically used when catching atoms from the MOT.

In addition to creating the magnetic trapping potential, the coils are also used to generate the spherical quadrupole field for the MOT. A spherical quadrupole field is created by adding the field

Table 5.1: Numerically optimized dimensions of the magnetic coils. The wire thickness denotes the length of each side of the square tubing (including the thickness of the insulation). The radial and axial windings indicates the number of layers of wire in each coil in the radial and axial directions.

	Ioffe-bars	pinch coils	anti-bias	extra bias
wire thickness (inches)	0.145	0.145	0.2	0.2
small inner diameter (inches)	1.25	1	6	7.2
large inner diameter (inches)	8	4.25	6	7.2
minimum separation (inches)	1.58	1.4	2.69	2.69
radial windings	4	2	3	2
axial windings	3	5	5	5

from the Ioffe-bars to the field produced by the pinch coils (wired up in anti-Helmholtz configuration). Numerical calculations showed that a current of 6.1 amps in the pinch coils and a current of 5.7 amps in the Ioffe-bars produced a very good spherical quadrupole field, with gradients of 5.1, 4.3, and -9.4 gauss/cm. The calculated MOT fields are shown in Fig. 5.7.

5.5.1 Selection of Materials and Construction

The "magnet wire" used to wind the coils was fiberglass insulated copper tubing. The use of hollow tubing allowed the coils to be cooled simply by running water through the bore of the magnet wire. The tubing used has a square cross-section. The square cross-section allowed for more efficient packing of the coil loops (without the air gaps produced by the round tubing used in our cloverleaf coils). In addition, the flat surfaces of the wire made it easier to wind the coil neatly. When selecting the gauge of wire to use, we considered multiple factors. In principle, it is possible to build the same trap with many different wire thicknesses. A trap with thiner wires will require more windings and will have more resistance, but will also require less current to produce the same magnetic fields as a trap with thicker wires. Both traps would in principle consume the same amount of power — one drawing less current but with a higher voltage drop.

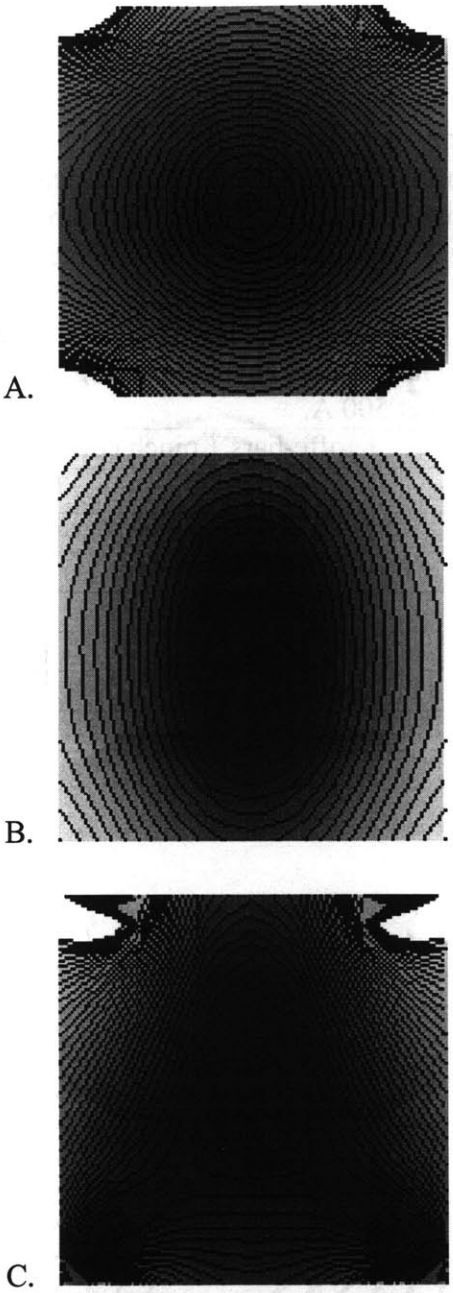

Figure 5.6: Calculated magnetic fields in the modified IP trap. The magnitude of the magnetic field is plotted in the xy (A), xz (B) and yz (C) planes. These plots were calculated using the typical currents employed when catching atoms from the MOT ($I_{Gradient}$ = 150 A, $I_{Curvature/Anti-Bias}$ = 100 A, and I_{Bias} = 60 A). These images, calculated with the full numerical simulation of the coils, are very similar to Fig. 5.5B which was calculated using the lowest order expansion of the magnetic potential given in Eq. 5.5. The field of view in each image is 4 cm by 4 cm, and the color scale runs from 0 gauss (black) to 300 gauss (white).

Table 5.2: Calculated figures of merit for the magnetic trap. All parameters listed in this table were computed assuming a current of 500 A in each coil (all field parameters for each coil scale linearly with current, with the exception of α and β which are geometric factors and are independent of the current). Although the fields for the Ioffe-bars and the pinch coils were calculated assuming rectangular coils, the calculated voltage drop was done more accurately by taking into account the rounded ends of the coils. The voltage drop calculations also included an additional 4 meters of wire to account for the voltage drop across the leads. B'_z is the gradient in the z direction at the center of the Ioffe-bars, resulting in a displacement of the trap minimum by z_0 when all of the coils except the extra bias coils are run at 500 A.

	Ioffe-bars	pinch coils	anti-bias	extra bias
voltage drop (volts)	24.8	13.0	12.7	10.4
B_0(gauss)		760.7	-760.2	481.2
B' (gauss/cm)		413		
B'' (gauss/cm^2)		279	-4	1
α		-0.01	1	1
β		2.01	1	1
B'_z (gauss/cm) / z_0 (μm)	5.2 / 189			

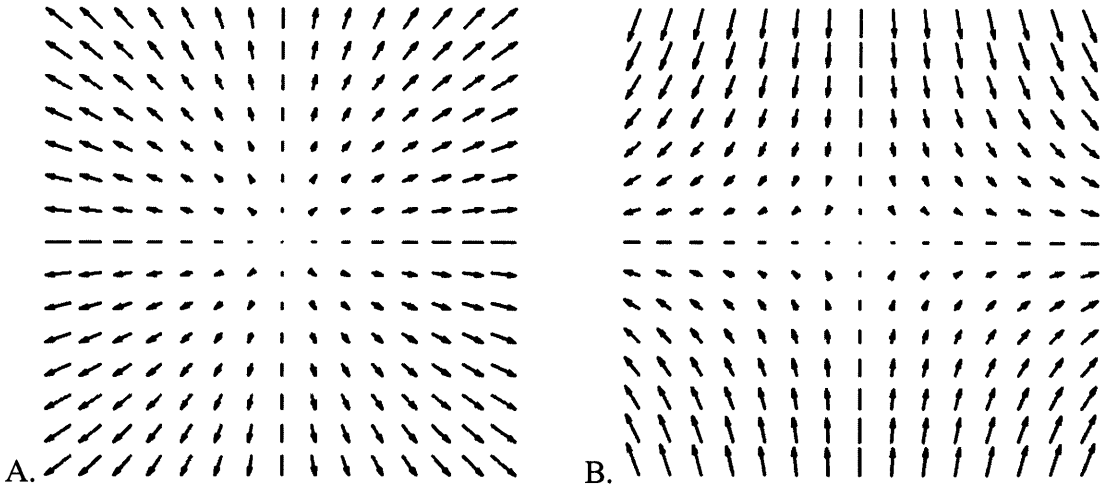

Figure 5.7: Calculated MOT fields. These plots of the calculated MOT field in the xz (A) and xy (B) planes illustrate that the field generated is, to a very good approximation, a spherical quadrupole field. The field of view in these plots is 2 x 2 cm.

Thinner wires are more flexible and easier to wind, and in addition, the leads of a trap with many windings of thin wire carry a smaller fraction of the traps total "amp turns" than do the leads of a trap with fewer windings, each of which carrying more current. Thus the leads of a trap wound with thinner wire will produce lower stray fields. Using wire of too fine a gauge, however, can be problematic. A trap wound with very fine wire would require a voltage drop of hundreds of volts or more. Working with higher voltages increases the danger involved and requires more sophisticated means of dealing with voltage spikes generated during trap switch-off (see Appendix D). Also, winding many layers of very fine wire can be tedious; it can be difficult to give each winding of a coil with hundreds of turns the same care as one would give the windings of a ten turn coil. Nevertheless, magnetic traps for BEC experiments have been successfully made from a wide variety of wire gauges. We set a limitation on the wire thickness for our trap when we decided to use hollow "tubing" rather than solid magnet wire to facilitate water cooling of the coils. Other methods have been used to cool magnetic traps, but many of them rely on the conduction of heat through several layers of wire and insulation, allowing parts of the coil to be many times hotter than the cooling water.

We selected two types of magnet wire for the trap. Thinner wire, 1/8" inch thick, was used to wind the curvature and gradient coils. This was the thinnest wire which could be made with a center bore large enough to allow a reasonable water flow through the coils (the center bore of this wire is square, 0.06" on a side). Since the anti-bias and bias coils are further from the atoms and take up less important space on the experiment, we decided to use thicker 0.2" wire (with a circular center bore 0.1" in diameter) for these coils. This resulted in coils which are larger, but which consume less power. Both types of wire used were made of copper with a soft temper, making the wire easier to bend (the temper of the copper has almost no effect on its conductivity). When selecting wire, we tended towards wire with a larger center bore to increase the water conductivity through the coil at the expense of higher electrical resistance. This is advantageous because water

conductivity scales with a higher power of the cross-section (proportional to A^2 for laminar flow, and approaching $A^{5/4}$ as the flow becomes turbulent [91]) than does electrical conductivity (which scales linearly with A). So by stealing some of the cross-sectional area from the flow of electrons, and dedicating it to the flow of water, the cooling power increases more than the joule heating does. The result is a coil which burns off more power at a given current, but which is capable of safely burning off *even more* power, provided that your power supply is capable of delivering it.

The magnet wire is insulated with a thin but durable double-layer fiberglass insulation. Durability is very important, since the process of winding the coils submits the wire to a fair amount of abrasion (during construction of our cloverleaf trap, despite the extreme care taken, several coils had to be discarded due to shorts which developed while they were wound). Each coil was carefully wound around a form and then potted in epoxy. This produced coils which were stiff enough that they required no additional support.

5.5.2 Power Supplies and Switching

Of equal importance to the coils themselves are the power supplies and switching circuitry which drive the coils. In order to limit heating in the trap, the currents in the coils have to be stable to about one part in 10^3 (see Section 5.4). In addition, the current must be able to switch on quickly to catch atoms from the MOT before they have a chance to expand and lose phase space density. Since the $\sim$ 1 cm diameter cloud of atoms out of the MOT has a thermal velocity of about 30 cm/s after polarization gradient cooling, a switch-on time for the magnetic trap of 2 ms implies that the cloud diameter expands by less than 10%. We must also be able to switch the magnetic trap off quickly (compared to the $\sim$ 500 Hz radial trap frequency) to allow time-of-flight studies of the atoms. The current must also be able to be smoothly ramped to compress or decompress the atoms, or to be modulated to drive excitations. In addition, it must also be possible to tap the current in

several places and drive it backwards in one coil in order to produce a spherical quadrupole field for the MOT. A wiring diagram of the setup which we created to meet all of these requirements is contained in Appendix D.

The magnetic trap is powered by three high current power supplies, allowing us to run the trap with up to 500 amps of current. The specific power supplies chosen have a specified noise ripple of 60 mV (out of 30 V), sufficient to prevent significant heating of the atoms. They also have relatively fast recovery and a fast programming rate, allowing the trap current to be turned on or modulated quickly. The current produced by each of these power supplies is controlled by an analog voltage signal from the experiment's control computer. To facilitate fast switch-on and fast switch-off, each power supply is connected through an insulated gate bipolar transistor (IGBT). The gate of each IGBT is powered on or off by a control circuit based upon a TTL signal received from the control computer. Although the IGBT is capable of switching hundreds of amps in microseconds, they can easily be destroyed by the inductive voltage spikes produced when the magnetic trap is suddenly turned off. These voltage spikes are suppressed by varistors placed across the coil leads. We found, however, that even the short leads between the varistors and the IGBT were inductive enough to cause spikes lethal to the IGBT. So to further protect the IGBTs, a lower power varistor was also wired directly across each IGBT.

Four additional low-current supplies are used to generate the spherical quadrupole field for the MOT. In the cloverleaf trap used in our older experiment (Fig. C.2), the MOT field can be simply generated by powering up just one curvature coil on one side of the trap, and one anti-bias coil on the other side. In our new experiment, however, the symmetry axis of the MOT is rotated by 90° from the symmetry axis of the magnetic trap. To generate the required field, the pinch coils are powered up in anti-Helmholtz configuration (with the current in one coil running opposite to the sense of current used in the magnetic trap) along with the Ioffe-bar coils. The pinch coils, when wired this way, approximate a spherical quadrupole field with a symmetry axis aligned with the

magnetic trap's axis. The addition of a two-dimensional gradient field produced by the Ioffe bars serves to rotate the symmetry axis of the quadrupole by 90° (see Fig. 5.7). Using two separate power supplies for each set of coils allows us to move the center of the quadrupole field toward or away from a coil just by changing the current in each power supply individually, giving us the ability to align the MOT with the magnetic trap in two dimensions (adjustments in the third dimension are done by tweaking the current in the slower compensation coil).

The fields produced by the coils were mapped out during bench testing. At each point in space, the fields were measured to be within a few percent of the predicted values. This assured us that no shorts had developed while winding the coils. The switching characteristics of the trap were also tested, and we found that the coils could be switched on in about 3 milliseconds, and switched off in about 100 microseconds.

5.6 Evaporative Cooling

Once caught in the magnetic trap, the atoms can be very conveniently cooled by RF forced evaporation. RF evaporation works by ejecting the most energetic atoms by spin-flipping them into an untrapped internal state. The spin-flip is induced by an oscillating magnetic field generated by a pair of small "antenna" coils near the quartz cell. The spin-flip can be applied selectively to the energetic atoms due to the sensitivity of the spin-flip resonance to magnetic fields. RF evaporation is discussed more fully in Appendix C.3 and in reference [92].

Rf cooling in the new apparatus In our new apparatus, the RF antenna coils consist of two square single loop coils tucked inside the Ioffe bars above and below the quartz cell. Single loop coils were used to lower the inductance of the antenna, allowing them to produce bigger magnetic fields at higher frequencies. Since the current in the antenna typically doesn't exceed 200 mil-

liamps, it is possible to make the RF antenna out of very thin wire (thereby minimizing the bulk of the coil and the resulting loss in optical access to the atoms). However, since the self-inductance of a loop of wire can depend greatly on the diameter of the wire used (diverging in the limit of infinitely thin wire), we did careful calculations to assure that this wasn't a problem. Calculation of the self inductance of a loop is usually done by separating the inductance into two parts, one due to the flux through the loop which is external to the wire, and one due to the internal flux passing directly through the thickness of the wire. We found that for 0.25 mm diameter wire wound in a 1" diameter loop, the internal flux term was completely negligible. In typical use, the antenna produces an oscillating field of from 10 to 100 milligauss. The coils are driven with higher voltages at the beginning of evaporation due to the increased impedance of the coils at higher frequencies.

RF cooling is started by driving the antenna at 30 MHz with the output of a computer controlled function generator. As the atoms cool and settle lower in the magnetic potential, the frequency is swept down. During most of the sweep, the RF field is kept in resonance with atoms sitting at potential energies several times $k_B T$, resulting in efficient cooling. Bose condensation typically occurs in our experiment when the frequency has been swept down to about 1.2 MHz. By sweeping to lower frequencies, and observing at which frequency the last atoms are expelled from the trap, the magnetic bias field at the center of the trap can be inferred (see Section 5.3).

Chapter 6

Imaging

Optical probes have so far been the only means used to probe dilute gas Bose condensates. Due to the relatively small number of atoms in alkali gas Bose condensates (on the order of 10^7 atoms in our experiments), the use of macroscopic ($\sim 10^{23}$ atoms) contact probes is not possible — rather than the probe coming into equilibrium with the condensate, the condensate would come into equilibrium with the probe! As a result, all of the data which has been collected on dilute gas condensates has been taken by optical means. A detailed account of the imaging methods used in our experiments is given in Appendix E. In this chapter a brief overview of standard imaging methods is given, additional experimental details are described, and the details of the application of these methods on the new apparatus are discussed. The extraction of physical parameters from an image of a condensate is discussed in depth in reference [47].

6.1 Absorptive Imaging Methods

Imaging methods can be divided into two classes depending on whether the probe light passing through the atoms is primarily scattered by large angle Rayleigh scattering, or whether the probe

light is primarily refracted dispersively. In the former case, imaging can be done by either catching and focusing the scattered photons (fluorescence imaging), or by imaging the shadow left in the probe light (absorption imaging). These two methods are depicted in Fig. 6.1. Both methods measure the same scattering events. But since absorption imaging can detect almost all of the "missing" photons, and fluorescence imaging only detects those which are scattered into the small solid angle subtended by the imaging lens, absorption imaging is typically a few hundred times more sensitive. A quantitative analysis of absorption imaging is contained in Appendix E.1.2.

6.1.1 Limits of Absorption Imaging

One problem with absorption and fluorescence imaging is their limited dynamic range — absorption can only be as high as 100 percent! This is a major limitation for optically dense clouds. Since the value we wish to measure is the integrated number density of atoms, which is proportional to the optical density of the clouds, we have to convert the measured absorption A at each point in the image to an optical density $\tilde{D}$ with the formula: $\tilde{D} = -ln|1 - A|$. If A is near unity, small errors are magnified in the data reduction. For typical Bose condensates in our apparatus, the peak optical density is $\sim$300 with a peak absorption of $A = 1 - 5 \cdot 10^{-131}$, resulting in images which are completely blacked out to within the noise floor of the image over the entire condensate. This effect is illustrated in Fig. 6.2.

6.1.2 Time-of-flight Imaging

The imaging of optically dense alkali condensates can be accomplished by allowing the atoms to freely expand, in the absence of a trapping potential, until the optical density of the atoms has dropped below unity. This "time-of-flight" method of imaging was used in the first detection of BEC in dilute gases [1] as well as in our first detection of a sodium Bose condensate [2]. In

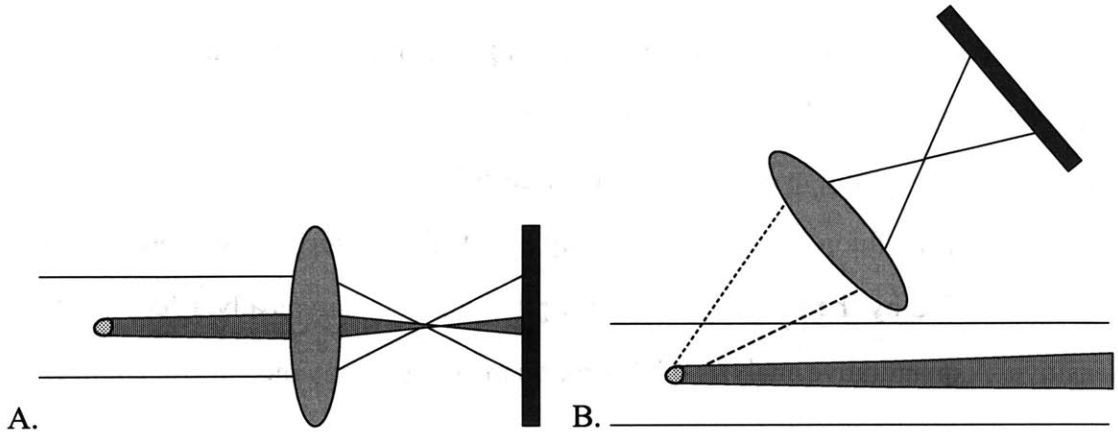

Figure 6.1: Absorption (A) and fluorescence (B) imaging setup. Probe light from the left illuminates the atoms. Some of this light is scattered by the atoms over large angles. In absorptive imaging, the shadow created by the "missing photons" is imaged. In fluorescence imaging, some of the scattered photons are collected to form an image of the atoms.

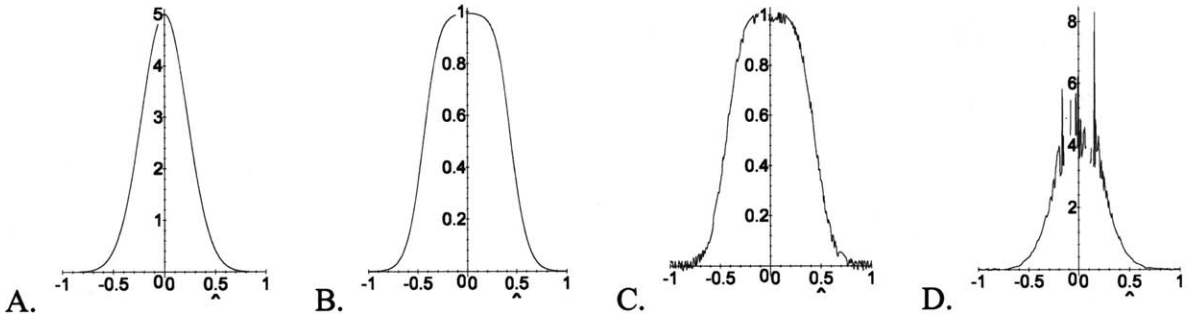

Figure 6.2: Susceptibility to noise in blacked out absorption images. Frame (A) represents the optical density (OD) profile of a hypothetical cloud of atoms, plotted as a function of position. When taking an image, this, rather than the absorption, is the value which we wish to measure, since it is directly proportional to the integrated atomic density. Frame (B) is the absorption profile which this optically dense cloud would produce in an ideal imaging system. Notice that the absorption near the center of the cloud is nearly 100 percent. In frame (C), four percent noise has been added to the absorption profile of frame (B). In frame (D), the optical density of the cloud has been extracted from the noisy absorption profile. Notice that in the wings, where the cloud's optical density is low, the profile is extracted very clearly despite the noise. In the center, where the OD is high, the extracted profile has been completely distorted.

addition to reducing the optical density, time-of-flight methods also increase the size of the atom cloud, making it easier to resolve the cloud with the imaging system. However, while it is possible to reconstruct a picture of the trapped condensate from a time-of-flight image, perturbations caused by a trap which doesn't turn off instantaneously, and complications due to collisions between atoms during the early part of the expansion can limit the accuracy of a time-of-flight measurement. Another disadvantage of time-of-flight imaging is that it involves the destruction of the condensate.

6.2 Dispersive Imaging Methods

When faced with optically dense samples, one natural thing to do is to detune the light from resonance, thereby reducing the effective optical density of the cloud. When detuning from resonance, however, care has to be taken to assure that dispersive refraction of the detuned light by the atoms doesn't bend light out of the imaging system. As discussed in Appendix E.1.2, wide angle dispersive scattering caused by optically dense clouds require that the detuning be either very small (resulting in blacked out absorption images) or very large (resulting in an off-resonant optical density which scales as the *inverse* of the resonant optical density). In the case of large detuning, the imaged cloud is transparent to the probe light and cannot be discerned on the camera.

Transparent objects do, however, distort the phase of the probe light. An image of a dense cloud can be made, then, by converting this phase modulation into an intensity modulation which can be detected by a CCD camera. This conversion is done by Fourier filtering of the image. Since the phase shift induced by the atoms causes the probe light to refract, the scattered and unscattered light is separated in the Fourier plane of the first lens of the imaging system (see Fig. 6.3). This makes it possible to operate on the scattered and unscattered photons separately in a way which will produce an intensity modulated image on the camera.

The simplest form of Fourier filtering is done by placing a small opaque spot in the center of the

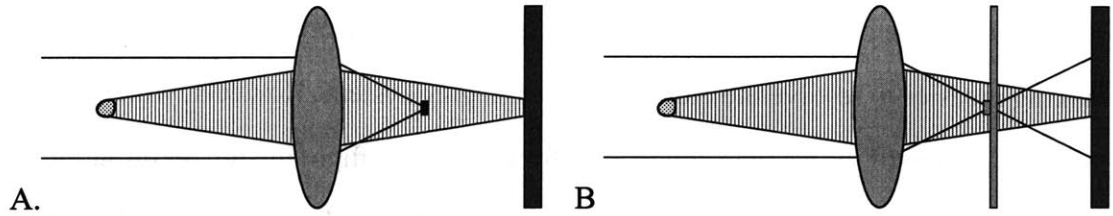

Figure 6.3: Dark-ground (A) and phase-contrast (B) imaging setup. Off-resonance probe light from the left illuminates the atoms. Some of this light is scattered dispersively by the atoms. In the Fourier plane of the first imaging lens, the scattered and unscattered components of the probe light are spatially separated. In dark-ground imaging, unscattered light is blocked by an opaque object in the center of the Fourier plane. In phase-contrast imaging, the unscattered light is shifted by 90° by a phase plate placed in the Fourier plane.

Fourier plane of the first lens of the imaging system (see Fig. 6.3). This spot (actually a thin wire in our case) blocks the unscattered light, resulting in a dark-ground image at the camera which, to lowest order, has an intensity proportional to the square of the phase shift induced by the atoms (see Appendix E.1.2).

Our first attempts at dark-ground imaging produced excellent images of dense Bose condensates. We found, however, that the less dense thermal clouds were easily distorted by low levels of stray probe light reaching the camera. We quickly began work to implement a more advanced imaging method: phase-contrast imaging. In phase-contrast imaging, the zeroth order light in the Fourier plane (the unscattered light) is not blocked. Rather, it is shifted in phase by 90° by a phase plate (an optically flat piece of glass with a small, $\lambda/4$ thick bump in the center). This causes the unscattered light to act as a local oscillator, interfering with the scattered light and producing a "homodyned" signal. The result is an image in which the intensity at a point is, to lowest order, linearly proportional to the phase shift of the light. This linear signal, compared to the quadratic signal of dark-ground methods, allows clearer imaging of the dilute thermal wings of a trapped cloud, as shown in Fig. 6.4. A detailed analysis of dark-ground and phase-contrast imaging can be found in Appendix E.

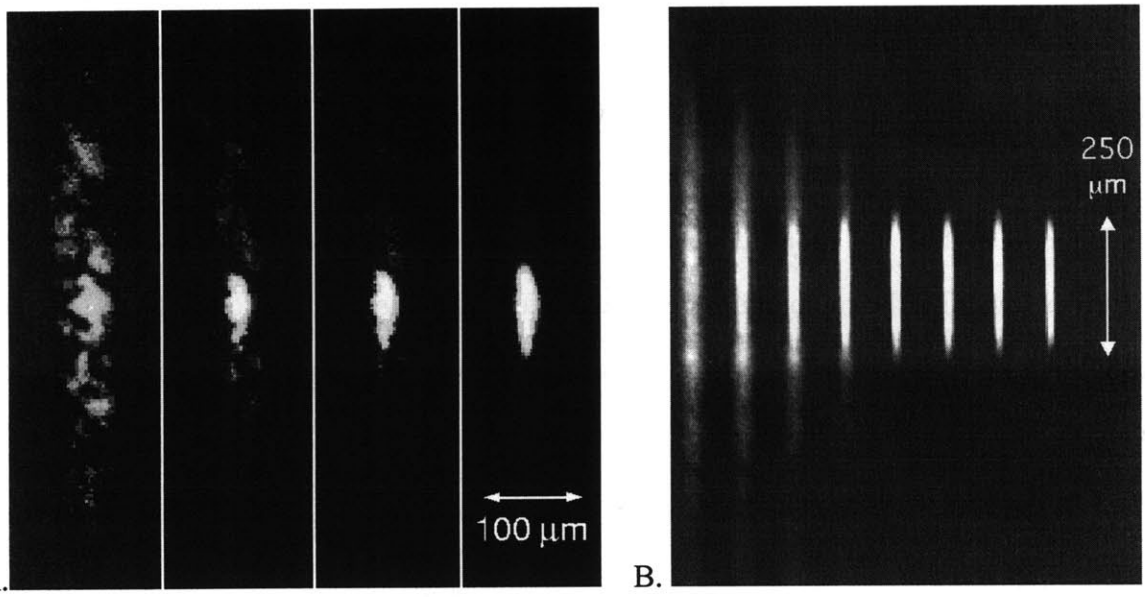

Figure 6.4: Comparison of dark-ground and phase-contrast images of Bose condensates and thermal clouds. These two sets of images were taken of clouds of trapped atoms at various temperatures near the critical temperature for BEC (temperature decreases as you move to the right in both sets). These are typical dispersive images showing dense condensates as well as less dense "thermal" atoms. Set (A) was taken during our early work with dispersive imaging using dark-ground techniques, while image (B) was taken after implementing the more sensitive phase-contrast imaging technique. In both sets of images, the condensate is clearly visible. In set (A), however, the dilute thermal cloud is riddled with noise, while in image (B), the thermal cloud is cleanly imaged.

6.2.1 Non-destructive Imaging

As discussed in the above paragraphs, dispersive imaging methods overcome the limited dynamic range of absorptive methods when imaging optically dense condensates. But in addition to getting around problems with optical density, dispersive imaging has one added benefit: dispersive imaging makes it possible to image the condensate non-destructively. This is due partially to the fact that dispersively scattered photons, scattered at smaller angles than Rayleigh scattered photons, impart less momentum to the cloud. Moreover, dispersive scattering is a coherent process, and therefore induces coherent, rather than thermal, excitations in the atoms. If the probe light is on for much longer than the oscillation frequency of an atom in the trap, these coherent excitations cancel out, further lowering the momentum deposited into the atoms by the probe light. The reduced heating makes it possible for us to take tens of images of a Bose condensate without heating the condensate appreciably. The non-destructive nature of dispersive imaging is discussed further in Appendix E.1.4.

6.2.2 Rapid Sequences of Images

Using non-destructive dispersive imaging, we have been able to take "movies" of Bose condensates, observing the evolution of a condensate in real time. Before developing the ability to take rapid sequences of images, dynamics of dilute gas Bose condensates were studied exclusively by destructive methods in which a new condensate was created for each image taken. Our condensate "movies" have several advantages over these methods. First of all, data rates are improved considerably. Whereas before the determination of the frequency of a mode of oscillation in the condensate required on the order of an hour and the production of tens of individual condensates, we can now use a rapid sequences of images taken non-destructively and map out the oscillation frequency in just the time required to make and probe a single condensate (less than one minute

in our experiment). In addition, since a single condensate is used, shot to shot fluctuations are eliminated.

Rapid image sequences are taken using a fast CCD camera operating in kinetics mode. In kinetics mode, all but a small strip of pixels on the CCD chip are masked off, and the rest of the chip is used as an effective "high speed film" where images can be transfered and stored for later readout. Images are taken by illuminating the exposed pixels, and then rapidly transferring the charge acquired in those pixels, row by row, until the image has moved under the masked region of the CCD. Since no readout is involved, the images can be shifted along the chip very rapidly. After each picture this shifting occurs until the chip is full of images, at which point the entire chip is read out slowly to minimize read noise.

Since the CCD chip[1] in our cameras is sealed in an evacuated vacuum encasing (to both limit heat conductivity and allow the chip to be cooled without the risk of water condensation) it is impossible to put a mask directly in front of the chip. The masking of the probe light is done instead with a razor blade in an intermediate image plane of the imaging system (see Fig. 6.6).

Using large chips, we have been able to take more than twenty pictures of a single condensate. The chips used in our cameras have a maximum shift rate of 3 μs per row. We found, however, that the images blurred if the chips are configured to shift faster than about 30 μs per row. Combined with the typical width of the images taken (50 pixels — giving a field of view of 140 μm at the maximum resolution of our imaging system) this results in a maximum image rate of just over 600 Hz. An example of a non-destructive rapid sequence of images is shown in Fig. 6.5.

[1] EEV model CCD02-06

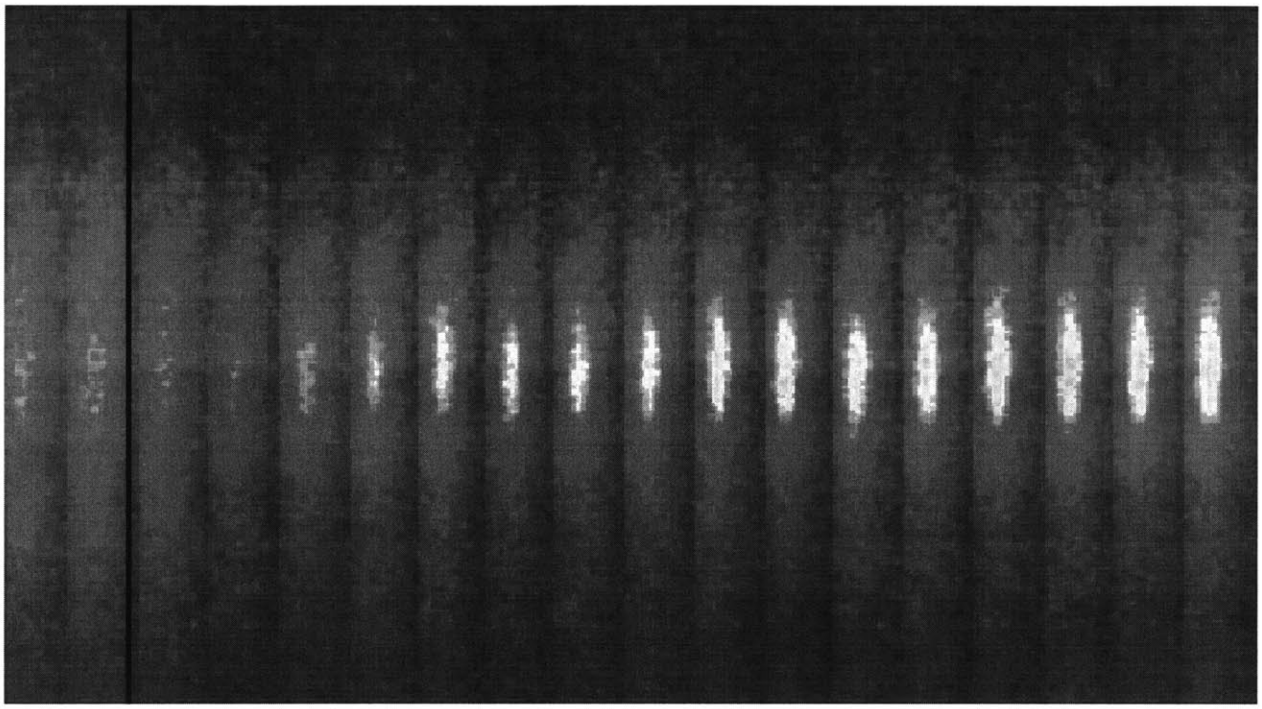

Figure 6.5: "Movie" of condensate formation. One application of non-destructive rapid-sequence imaging was the study of condensate formation. These studies were done by cooling a cloud of magnetically trapped atoms with RF evaporation until the cloud was just above the transition temperature for BEC. Then over 10 ms, the RF was quickly swept below the transition. After this shock cooling, we observed the relaxation of the cloud to equilibrium. In the image above, a sequence of pictures taken at 13 ms intervals shows the formation of a dense condensate out of the shock cooled thermal cloud (the black line indicates the point in time at which the shock cooling sweep was begun). In these experiments [48], we were able to detect Bosonic stimulation during condensate formation, illustrating one more "laser-like" property of Bose condensates. The length of the images is 630 μm.

6.3 Specifics of Our Imaging Layout

We have implemented two nearly identical imaging systems on the new apparatus, allowing us to image the atoms both from the side (along the symmetry axis of the magnetic trap) and from the top. The imaging systems on the new apparatus are made from off the shelf standard achromats. This has allowed us a great deal of flexibility (in fact, the details of the imaging systems have been altered many times in the past months). The layout of the horizontal imaging system is depicted in Fig. 6.6. The vertical imaging system is functionally identical, using identical lenses. Both imaging systems have an f/# of about 4. The absolute resolution of both imaging systems was measured using standard mil-spec test patterns, and determined to be about 6 μm. The front lens of both imaging systems is mounted on a micrometer adjustable translation stage, allowing the imaging system to be carefully focused on the atoms. In addition, the first two lenses of the horizontal imaging system can be translated vertically with micrometer screws, allowing us to translate the axis of the imaging system in order to image atoms which have been allowed to fall under gravity.

Both imaging systems allow three magnifications settings: 8x, 2x, and 0.4x. In addition, pixels on the camera can be binned in groups of 1, 4, or 16 pixels. This gives us nine effective magnifications ranging from 0.1x to 8x. Low magnification allows us to characterize large clouds of atoms (such as atoms in the MOT or atoms which have been expanded in time-of-flight). The high magnification settings allow us to resolve fine features such as surface waves on a Bose condensate (see Section 9.3) or the interference fringes formed by overlapping condensates (Chapter 8).

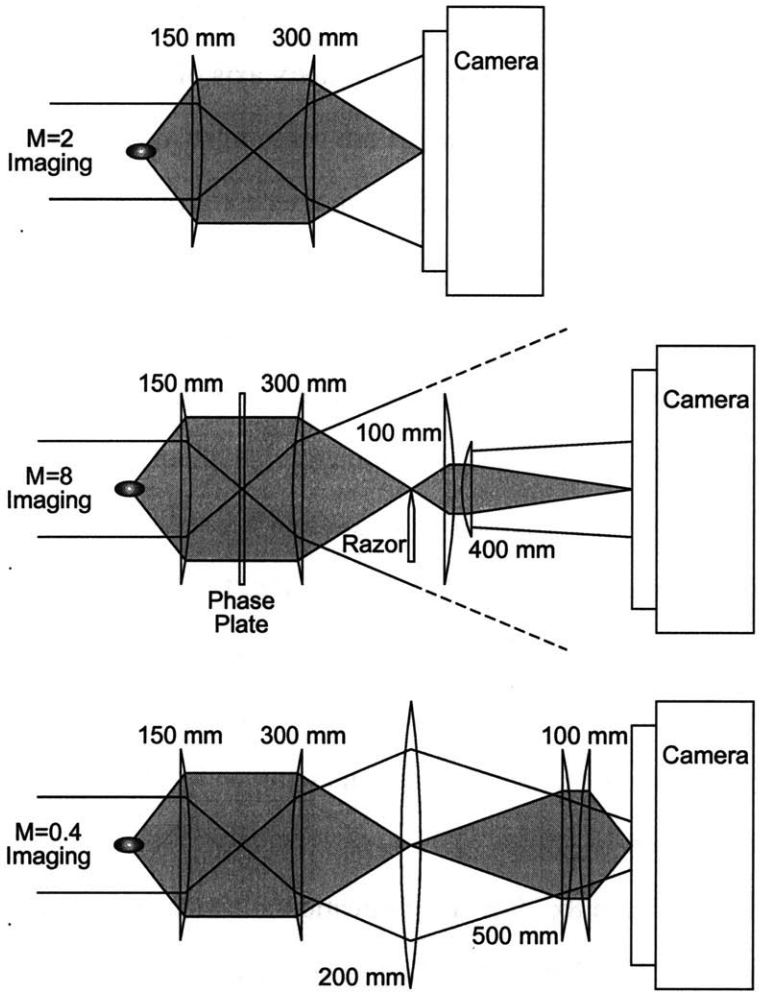

Figure 6.6: Horizontal imaging system. The horizontal imaging system can be configured for three different magnifications. M=2 imaging is done with a pair of achromats. M=8 and M=0.4 imaging is done by re-imaging the focused image of the M=2 lenses. An additional field lens is added at the intermediate image plane in the M=0.4 configuration. In this position, the lens has little effect on the imaging system other than reducing vignetting. Space between the two first lenses gives us access to the Fourier plane of the first lens, allowing the implementation of phase-contrast imaging. In order to mask the probe light for kinetics mode imaging, a razor blade can be placed in the intermediate focus of the M=8 set-up. The focal length is indicated next to each lens.

Chapter 7

Optical Manipulation of Bose-Einstein Condensates

Although the magnetic trap is an essential part of our apparatus, magnetic traps do have many limitations. One problem is the limit that Maxwell's equations place on the possible shapes of magnetic potentials. In particular, given the distance of the magnetic coils from the atoms, it is difficult to make small features in the trap. The large inductance of the magnetic trap's coils also make it difficult to make changes to the potential on a time scale faster than a few kHz. To overcome these problems we turned to the use of optical dipole forces. By adding the potential generated by a focused off-resonance laser beam to the magnetic trapping potential, we are able to "sculpt" the trap and add small features to it.

Early on we began using green light from an argon laser (with a wavelength 50 million linewidths "blue" of the sodium resonance) to repel atoms. Our first successful attempt at Bose-Einstein condensation used such a beam to keep atoms away from the low-field "hole" in a quadrupole trap (see Section 5.2 and Appendix C.2.1). We later used a focused sheet of argon laser light to split a magnetic trap in two — resulting in the first matter-wave interference pattern and the first evidence

of long range coherence in a Bose condensate [58]. The sudden introduction of a similar light sheet into a condensate was used to generate localized "Greens-function" types of density perturbations for the study of sound pulse propagation in a condensate [60, 61].

Another limitation of magnetic traps is that they can only confine atoms in weak-field seeking states. This limits studies of any atomic properties which depend crucially on the hyperfine state of the atom, and makes studies of multi-component spinor condensates impossible. Also, since the magnitude and spatial variation of the magnetic field is determined by the trapping requirements, it is not possible to do experiments in which the magnetic field must be very low or very high, or in which the field must be homogeneous. To get around these and other limitations, members of our group pursued the goal of confining a condensate in a *purely* optical potential. Unlike earlier attempts to *create* a Bose condensate in an optical potential, in this work the Bose-Einstein condensate was first formed by evaporative cooling in a magnetic trap. And then, only after the condensate was formed was it transfered into an optical dipole trap [63]. This approach bypasses many barriers to optical cooling. And since the Bose condensed atoms are a thousand times cooler than atoms right out of a MOT, only a few milliwatts of laser power were needed to make the trapping potential (dipole traps for hotter atoms typically require many Watts of laser power [93, 94]). The success of this approach resulted in several striking experiments [95, 51, 96], and the ability to confine atoms in optical potentials continues to be one of the most valuable tools used on our original BEC apparatus.

In the new apparatus we have extended the application of optical dipole forces with the use of a rapidly scanning dipole beam, allowing us to create "arbitrary" time-averaged perturbations to magnetic traps (making possible the studies described in Chapter 9). Furthermore, scanning a laser at slower speeds made it possible to study the transition from dissipation free to dissipative flow (see Chapter 10).

In this chapter, a general description of optical dipole forces will first be given. Then the

implementation of rapidly scanned optical dipole forces in our experiment will be described, and certain limitations which we encountered will be addressed. The way in which we bypassed some of these limitations with the use of scanned optical "perturbations" to a magnetic trap will also be mentioned, followed with several examples of possible future applications of scanned potentials.

7.1 Theory of Optical Dipole Trap Operation

A very general description of optical dipole potentials is given in this section. Further details on the theory and operation of optical dipole traps can be found in [97, 47].

The dipole potential Optical dipole potentials are generated due to the Stark shift of the atoms in the presence of a laser beam. The AC Stark shift for a two level atom is given by the expression:

$$U(\mathbf{r}) = -\frac{\hbar\omega_R^2(\mathbf{r})}{4}\left(\frac{1}{\omega_0 - \omega_L} + \frac{1}{\omega_0 + \omega_L}\right) \tag{7.1}$$

where ω_0 is the resonant frequency of the atom, ω_L is the laser frequency, and $\Delta = \omega_L - \omega_0$ is the laser's detuning. The Rabi frequency $\omega_R(\mathbf{r})$ depends on the local intensity of the laser, and is defined by the relation

$$\frac{2\omega_R^2(\mathbf{r})}{\Gamma^2} = \frac{I(\mathbf{r})}{I_{\mathrm{SAT}}} \tag{7.2}$$

where Γ is the natural linewidth of the optical transition, and I_{SAT} is the saturation intensity:

$$I_{\mathrm{SAT}} = \frac{2\pi^2\hbar\Gamma c}{3\lambda_0^3} \tag{7.3}$$

with λ_0 being the wavelength of resonant light. For sodium, I_{SAT} is 6 mW/cm^2 (assuming an oscillator strength of unity). In the rotating-wave approximation, Eq. (7.1) becomes:

$$U(\mathbf{r}) \simeq \frac{\hbar\omega_R^2(\mathbf{r})}{4\Delta} = \frac{\hbar\Gamma^2}{8\Delta}\frac{I(\mathbf{r})}{I_{\mathrm{SAT}}} \tag{7.4}$$

Conveniently, this potential energy is simply proportional to the intensity of the laser light. The sign of the potential energy shift is the same as the detuning; blue detuned light creates a positive energy shift (thereby repelling the atoms), and red detuned light makes a negative shift in the energy (attracting the atoms). Therefore, the radially decreasing intensity of a red detuned Gaussian beam creates a potential which is lowest in the center of the beam (where the light is most intense), providing radial confinement of the atoms. By focusing the laser, the light intensity is made to vary axially as well as radially, and three dimensional confinement of the atoms is possible. The distribution of light in the focus of a Gaussian beam with total power P is given by :

$$I(r,z) = \frac{2P}{\pi w_0^2(1+(z/z_R)^2)} \exp\left(-\frac{2r^2}{w_0^2(1+(z/z_R)^2)}\right) \tag{7.5}$$

where w_0 is the $1/e^2$ beam waist radius of the focus, z_R is the Rayleigh range, and r and z are the displacement from the center of the focus along the radial and axial directions [98]. The potential generated by a red detuned laser with this intensity distribution is harmonic to lowest order, and is described by radial and axial trapping frequencies of :

$$\nu_r = \sqrt{\frac{3\Gamma P\lambda_0^3}{8\pi^5|\Delta|Mcw_0^4}} \tag{7.6}$$

$$\nu_z = \nu_\rho \cdot \frac{w_o}{z_R}\sqrt{\frac{1}{2}} \tag{7.7}$$

Rayleigh heating Besides generating a conservative potential $U(\mathbf{r})$, the laser light also induces heating via Rayleigh scattering. The rate of Rayleigh scattering is given by :

$$\gamma_R = \left(\frac{\omega_L}{\omega_0}\right)^3 \frac{\Gamma^3}{8} \left(\frac{1}{\omega_0 - \omega_L} + \frac{1}{\omega_0 + \omega_L}\right)^2 \frac{I}{I_{\mathrm{SAT}}}$$

Neglecting counter-rotating terms ($|\Delta| \ll \omega_0$), the potential depth scales as I/Δ while the scattering rate scales as I/Δ^2. As a result, heating through Rayleigh scattering can be decreased while keeping trap parameters constant by detuning the laser further from resonance and increasing the laser's intensity.

7.2 Time Dependent "Arbitrary" Potentials

Using optical dipole forces, it is possible to generate a wide variety of trap geometries simply by shaping the mode of the laser beam. Work has been done to generate novel potentials for laser cooled atoms by such methods as the use of axiconical optics [99] and by steering the laser with moving mirrors. These efforts, however, have been aimed at generating very specific types of potentials with a specific laser field pattern. In our new apparatus we have extended these techniques with a scanning dipole beam which can generate "arbitrary" two-dimensional optical potentials which can be varied in time in an arbitrary way. This was accomplished by the use of two scanning acousto-optic modulators (AOMs) to deflect a laser beam in two dimensions, allowing the position and intensity of the scanned beam to be controlled by a set of computer programmable arbitrary function generators. By scanning the laser much faster than any of the natural oscillation frequencies of the atoms in the trap, the atoms can be made to see a time-averaged potential.

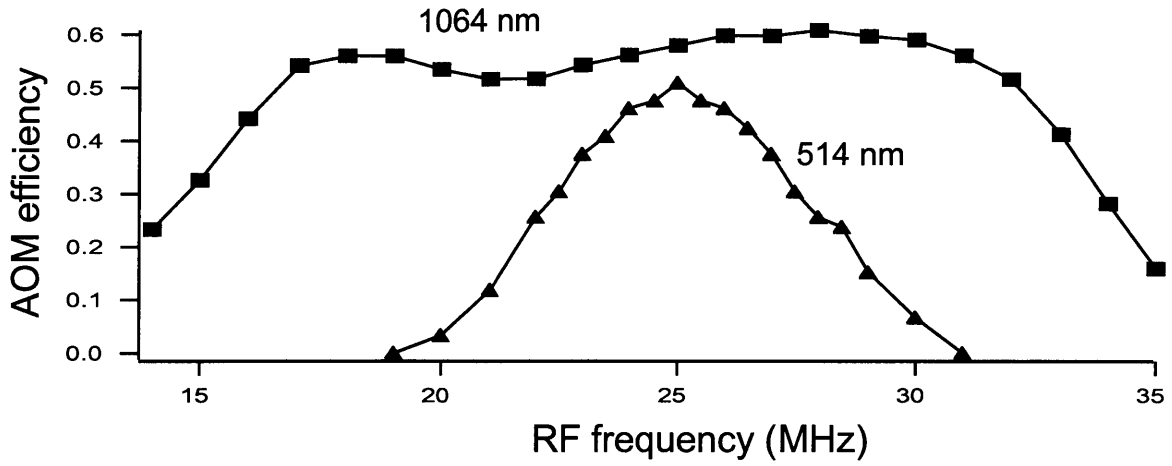

Figure 7.1: AOM Bragg diffraction efficiency. The efficiency for Bragg scattering into the first order of the AOM is plotted as a function of the RF frequency used to drive the AOM for both Nd:YAG and argon ion laser beams.

7.2.1 Implementation of Scanned Dipole Potentials

AOM beam deflector In order to make time-averaged potentials, the dipole beam must be scanned much faster than any of the resonances of the trapped atom cloud. Rapid xy scanning is achieved in our lab with the use of a pair of scanning acousto-optic modulator beam deflectors. The AOMs employed[1] were engineered specifically for scanning beam applications using Nd:YAG lasers, and have a nearly constant efficiency for scattering into the first Bragg order over a deflection range of 28 mrad (Fig. 7.1).

The maximum scan speed of an AOM is limited by the speed of sound in the AOM crystal — when the drive frequency changes, it takes a finite amount of time for the altered acoustic wave to travel across the width of the laser beam. The effective speed of an AOM is, then, also limited by the diameter of the laser beam. And, given the finite range of deflections of the device $\Delta\theta$, the diameter of the laser beam propagating through the AOM also determines the maximum number of resolvable spots across the AOM's scan range:

[1] Intraaction model DTD-274HA6

$$N_{spots} \simeq \frac{\Delta\theta\, D}{\lambda} \tag{7.8}$$

where D is the diameter and λ is the wavelength of the laser beam. As such, the scan speed of the AOM is specified as an "access time" per resolvable spot. Our AOMs are specified at 6.4 μs per spot, allowing a 4 x 4 point pattern to be scanned at about 40 kHz. The active size of the crystal is 4 x 4 mm, allowing patterns to be generated with a maximum resolution of 114 x 114 points at a specified scan rate of 1.4 kHz. In practice we found that we were able to scan the beam slightly faster than the specified rate.

The RF signals which power the AOMs are generated by a set of voltage controlled oscillators. The voltage input to these oscillators is produced by a pair of arbitrary function generators which are programmed by the experiment's control computer, allowing us to scan the AOMs in arbitrary patterns. In addition to steering the beam, the AOMs are also used to modulate the power of the laser beam. This is done by controlling the amount of RF power fed into the AOM, thereby changing the efficiency of Bragg refraction into the deflected beam. The RF power in each AOM is controlled by a DC mixer placed before the amplifier for each RF signal. In typical operation, the control input to one of the mixers is manually controlled with a potentiometer, allowing us to set an overall maximum power in the deflected beam. The other mixer input is either set to a constant voltage with a potentiometer, or is controlled by an additional arbitrary function generator, allowing the laser beam intensity to be controlled by the computer.

Scanning Nd:YAG beams Attractive optical potentials are produced in our experiment using light from a Nd:YAG laser[2] emitting at 1064 nm, with an output power of 470 mW. Due to the low transmission efficiency of the optics used in our setup, only about 30 mW can be delivered to the atoms. The laser is first passed through a single-mode, polarization preserving fiber to filter the

[2]CrystaLaser model IRCL-400-1064S

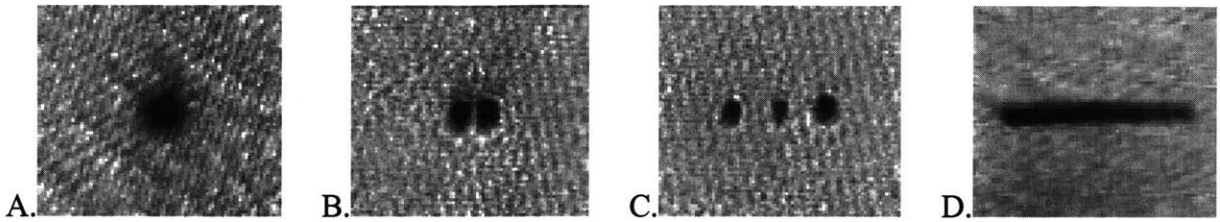

Figure 7.2: Scanning dipole potentials. These in-situ absorption images show Bose condensates trapped in scanning dipole potentials created with light from a Nd:YAG laser. In (A), the dipole beam is held at a fixed position, creating a single axially-symmetric trap. In (B) and (C), the beam is made to hop quickly from point to point, generating a time averaged potential with two (B) and three (C) local minima. These types of trap geometries could be used for the interferometric experiments described in Section 8.2. Image (D) shows atoms trapped in a sheet potential (the visible width of the sheet is 120 μm, with a similar dimension into the page). Problems with heating and trap loss have so far prevented us from loading atoms into more complicated potentials (Section 7.2.2).

mode of the laser, and is then collimated, telescoped, and passed through the two AOMs. Since the diameter of the beam in the AOMs affects the maximum speed at which the beam can be scanned, in the experiments described in Chapter 9 it was adjusted to be about 1 mm, allowing a scan rate of 10 kHz. The deflected beam from the AOMs is telescoped up by a factor of 10, and then focused onto the atoms with a 145 mm focal length achromatic lens. The beam waist of the optical trap is set by changing the effective f/# of the final lens, either by adjusting the final telescope or aperturing the beam. Typical beam waists at the trap have a diameter of 10 to 20 μm. Figure 7.2 demonstrates some of the potentials which have been used to trap condensates with our scanning IR trap.

Special care has to be taken when aligning the optics for the optical dipole trap. The presence of a small amount of astigmatism in the focus, for example, can create a potential with two (weakened) minima, and can limit the minimum achievable beam waist. Furthermore, aligning the optical trap to overlap with the magnetic trap requires very fine control of the position of the focus. Fine alignment was done by imaging the focus of the IR beam onto our CCD camera with

the same imaging setup that is used to probe atoms in the condensate. To account for chromatic aberration, we first imaged a test pattern, illuminating it first with yellow and then with IR light, and noted the amount we had to adjust the focus of the imaging system, and the distance that the pattern was translated in the plane of the image when we switched between the two wavelengths of light. We then took an absorption image of a small condensate to find the location of the center of the magnetic trap. Finally, after adjusting the imaging system for IR light, we would image the focused IR beam, carefully tilting the final focusing lens and mirror to remove astigmatism and to align the focus over the center of the magnetic trap.

Scanning argon ion beams Repulsive optical potentials are produced using green light from an argon ion laser. This light is produced by splitting a $\sim$ 500 mW beam off of the argon laser beam used to pump our dye laser (see Section 4.4). Since the argon laser is configured for multi-line operation, before sending the light through the AOMs it is first sent through a prism, allowing us to single out the 514 nm line. This beam is then passed through the AOMs, telescoped up, and focused into the condensate. No fiber was necessary to spatially filter the green light, due to the quality of the mode from the argon laser. Since the experiment done with green light required that the beam propagate perpendicular to the condensate axis, it was necessary to send this beam along the same path as one of the MOT beams. This was done by mixing the two beams with a glass plate beam splitter just before the final 150 mm focal length lens.

Our AOM beam deflectors were designed to be used with IR light, and as a result, when working with green argon light, the power in the deflected beam is a strong function of the deflection angle (see Fig. 7.1). To overcome this problem, the Bragg diffraction efficiency for green light was measured as a function of deflection angle. Then, with these measurements, it was possible to compensate for the changing efficiency by adjusting the RF power sent to the AOMs as the beam was scanned, resulting in a roughly constant light level.

Alignment of the green optical dipole beam was done in a similar manner as the IR light alignment. First the imaging system was analyzed with a test pattern, illuminated first with yellow light and then with green light, to allow us to correct for chromatic aberration. Then a condensate was imaged absorptively to find the position of the trap center. Finally, the argon beam was viewed on the camera and steered into position.

7.2.2 Limitations of Scanned Potentials

Scan speed / resolution limits When trying to load atoms into complicated optical potentials, we encountered very large heating and trap loss rates. In one example, when scanning in one dimension to create a sheet trap (see Fig. 7.2D), the heating seemed to begin very suddenly when the scan width of the laser was increased to more than 3 or 4 beam diameters, possibly indicating that the "micro-motion" induced by the scanning beam was driving some resonance of the trapped cloud. It is very likely that these effects could be overcome by simply increasing the frequency at which the beam is scanned — in our experiments the beam was scanned at 10 to 50 kHz, compared to the 100 to 500 Hz oscillation frequency of the atoms in a typical single spot IR trap. Unfortunately we were unable to scan the beam faster to test this hypothesis.

It should be noted that, as stated above, the scan speed of an AOM can be increased by reducing the size of the laser beam passing through it, but only at the expense of reduced scan resolution (Eq. (7.8)). As a result, the required beam diameter in the AOMs is limited by the complexity of the pattern to be produced, and may not be scaled down to be arbitrarily small to increase the AOM scan speed.

AOMs as phase gratings It might be possible in future experiments to overcome AOM scan rate limitations by driving an AOM with more than one frequency at a time. For example, N dipole traps could be made in a line by feeding a signal into an AOM which is the sum of N different RF

frequencies. Since the resulting potential would be a static rather than a time-averaged potential, problems with AOM scan rates would be eliminated. Extending this method to two dimensions should also be possible. It should be simple to create potentials of the form $U(x,y) = U_x(x) \cdot U_y(y)$ resulting from one beam being split into many beams horizontally by the first AOM, and then each of these beams being split into many beams vertically by the second AOM. These types of potentials include the multiple-well point and sheet potentials useful for condensate interferometry and studies of two-dimensional Bose condensates. For more complicated 2-D patterns, such as ring-shaped potentials, the two perpendicular 1-D phase gratings must act coherently as one thick 2-D phase grating. This would require that the drive to both AOMs be phase locked and that the relative position and orientation of the two AOMs be stable to less than an optical wavelength.

7.2.3 Scanning Perturbations

Most of the complications discussed above can often be avoided simply by generating the desired potential with a combination of magnetic and optical fields. In the experiments discussed in Section 9.3, the potential of the scanning IR beam was used only as a small perturbation to the magnetic trapping potential. As such, the heating due to micro-motion was correspondingly small. Furthermore, since the magnetic trap provided all of the confinement, the optical field could be switched off after driving an excitation, allowing the atoms to freely oscillate with no optical field present to induce heating.

In the experiments described in Section 10.3, a scanning AOM was used to move a focused argon laser beam through the condensate at a relatively slow rate, presenting the atoms with a single beam potential which moved in time rather than a time-averaged potential. As a result, no heating due to micro-motion of the atoms in the scanned beam had to be taken into account in the analysis of the experiment. Once again, confinement of the atoms was provided by a conservative

magnetic potential.

7.3 Future Experiments

We have only begun to explore the possibilities of scanning optical perturbations to magnetic traps. The experiments described in Chapters 9 and 10 illustrate the ability of scanning dipole beams to excite several unexplored types of dynamics in a Bose condensate. In addition, if problems with purely optical "painted-potential" traps can be solved, these traps will be useful for a wide variety of new experiments. As discussed in Section 10.3, studies of persistent currents and vortices would benefit from a potential consisting of a ring trap, plus a second round trap in the center of the ring. The ring shaped trap would create a multiply-connected condensate in which persistent currents should be stable, and the additional trap in the center of the ring would facilitate the detection and measurement of the circulation in the ring (see Section 8.3). Using scanning beam optical potentials, it is also possible to create multiple-welled potentials (Fig. 7.2) which could be used to overcome mechanical stability issues in interferometric studies of Bose condensates (see Section 8.2). In addition, it may be possible to use scanning dipole traps to make a condensate interferometer in which the two halves of the condensate are widely separated. Unlike interference experiments which split a magnetically trapped condensate in two with a fixed light-sheet dipole potential [58], or which split the condensate into individual wells of a standing wave optical trap [100], a scanning beam dipole trap would be able to spatially separate the two halves of the condensate wavefunction by *many* cloud diameters.

Sheet potentials like the one shown in Fig. 7.2D would allow the study of condensates in potentials in which one dimension approximates a box potential. This would result in a condensate with number dependent collective excitation frequencies (due to the fact that the length of the condensate would be fixed, but the density, and therefore the speed of sound, would depend on

the number of atoms in the condensate). Furthermore, it has been suggested that if the sheet can be made thinner than the healing length of the atoms, an effectively two-dimensional condensate could be made (although BEC does not occur in two-dimensions in the thermodynamic limit, it *can* occur in inhomogeneous systems). As shown in Appendix F, however, this may be difficult to achieve.

The use of scanning dipole potentials gives experimenters a very precise way to sculpt and manipulate condensates. As mentioned, the use of scanned perturbations has already been demonstrated in two ground-breaking experiments in our lab. And in the future, scanned optical perturbations and optical painted potentials are sure to be used to make many more experiments possible. If the heating problems with painted potentials turn out to be difficult to overcome, several of the experiments described above could be performed by a combination of static magnetic and optical potentials with the addition of a small perturbative scanning optical dipole beam. Due to the fine level of control which scanning dipole forces provide, it is possible that new experiments and applications which take advantage of this technique will continue to be performed for many years to come.

Chapter 8

Interference of Bose Condensates

A Bose-Einstein condensate is described by a single coherent wavefunction with a single unique quantum phase across its entire volume. This simple fact gives rise to all of the non-classical behavior of Bose condensates, including superfluidity and superconductivity, Josephson oscillations, and quantized vortices. The detection of any of these phenomena is, then, indirect proof of a macroscopic quantum phase. In dilute gas Bose condensates, however, the existence of a macroscopic phase can be measured *directly* by interference. The first observation of interfering Bose condensates was achieved in our lab in the winter of 1996 [58]. We performed this experiment by creating two separate condensates in a double well potential (made by slicing a magnetic trap with a blue detuned laser light sheet). After the condensates were formed, the trap was suddenly turned off, allowing the condensates to expand into one another and interfere. The observed interference pattern demonstrated that dilute gas Bose condensates, like optical lasers, have long-range first order coherence.

In addition to serving as a probe of coherence properties, interference patterns can be used as a tool to study other features of the condensate. In this chapter, the basic concepts for two possible types of interference studies are discussed. The first idea presented is a way in which mechanical

stability limitations in interference experiments can be overcome by the use of multiple (more than two) condensates. This technique is discussed in the context of phase diffusion experiments, but applies as well to more general interference experiments. The second study considered is the use of interference to detect superfluidity and quantized vortices. This prospect was proposed independently in references [101] and [102]. I will discuss the application of this approach to a specific trap geometry which has certain advantages over the original geometry proposed in these references.

8.1 Condensate Interferometry

In our first interference experiments, the barrier between the two Bose condensates was higher than their chemical potentials. As a result, the two condensates were formed independently, and no consistent phase relationship between them was to be expected. The creation of multiple condensates with a definite phase relation could be accomplished, however, by splitting a single condensate into two or more pieces, ensuring that the phase of each piece is initially the same. After performing any number of operations on each of the pieces, the final phase relation between them could be probed by allowing them to expand and overlap, and then measuring the resulting interference pattern. Such an experiment would constitute a separated beam atom interferometer.

8.1.1 Phase Diffusion Measurements

One of the simplest interference experiment which could be performed with a condensate is to split the condensate in two, wait for a time, and then allow the two clouds to interfere. For short hold times, a fringe pattern with a reproducible phase would be expected. At longer times, however, the phase relation between the condensates would disappear. By probing condensates held for

different amounts of time, the loss of phase correlation between two condensates could be studied.

The relative phase between two condensates can change for several reasons. If the trapping conditions are different for the two condensates, a difference in their chemical potential could cause an oscillation in their differential phase. As long as the experimental conditions are reproducible, this phase oscillation should also be reproducible and should not interfere with measurements of incoherent mechanisms for phase diffusion.

Even if the experimental parameters are well controlled, however, there is still an unavoidable uncertainty in the splitting of the atoms between the two wells, since the splitting ratio of atoms constitutes an observable which is conjugate to the relative phase of the condensates, satisfying the uncertainty relation $\Delta N\,\Delta\phi \geq 1$. By establishing a well defined phase between the two condensates, by definition we are introducing uncertainty into the number of atoms in each condensate, resulting in an uncertainty in their relative chemical potentials. In the Thomas-Fermi approximation, a small difference in the number of atoms in each condensate ΔN causes a difference in the chemical potentials of the two clouds given by $\Delta\mu = \frac{2}{5}\mu\frac{\Delta N}{N}$, where μ is the average chemical potential of the two condensates. This results in a rate of phase diffusion which is given by [103, 104, 105, 106]:

$$\gamma_\phi = \frac{2\mu}{5\hbar}\frac{\Delta N}{N} \tag{8.1}$$

For a condensate split symmetrically in two, $\Delta N \simeq \sqrt{N}$ [107]. For typical parameters for our condensates ($N \sim 10^7$, $\mu \sim 200$ nK), this results in a loss of phase correlation at a rate of a few Hz.

Interactions with thermal atoms create an additional source of phase diffusion. In a process similar to the phase diffusion of an optical laser, the phase of the condensate changes slightly each time a thermal atom scatters incoherently into the condensate. Since coherent "Bose-stimulated"

scattering into the condensate occurs N times more frequently than does incoherent scattering into the ground state, it is expected that the time scale for this process should scale inversely with the number of atoms in the condensate. Furthermore, since each small phase shift produced by an incoherently scattered atom adds up in a random way, the time scale should also scale with the square root of the collision rate between condensate and thermal atoms. This process is most pronounced for small condensates (where N is small, and the thermal cloud completely penetrates the condensate). If a small condensate of 1000 atoms is created by evaporating just below T_c under experimental conditions typical for our apparatus (T_c being about 1 μK, and the elastic collision rate at the transition being about one per condensate atom per millisecond), the phase diffuses on a time scale of one second. At lower temperatures, where N is larger and the collision rate is lower (due to the decrease in the number of thermal atoms and the fact that thermal atoms are repelled from the condensate), this phase diffusion process occurs over even longer time scales. In the case of a nearly pure Bose condensate, this process should not compete with the process described in Eq. (8.1).

In an experiment conducted in Boulder, measurements of the repeatability of the initial phase of two split condensates, as well as the oscillation of the relative phase of two condensates in time, were made by putting the condensate into a superposition of two different internal states with a $\pi/2$ pulse, allowing the two states to evolve for some time, and then causing them to interfere with a second $\pi/2$ pulse [108]. In a related experiment at Yale, a standing wave dipole trap was used to trap a multitude of condensates in a 1-D optical lattice, resulting in Josephson oscillations between the wells of the lattice. The Yale experiment detected the dephasing of the condensates by observing the time evolution of the Josephson oscillations [109]. Although this is a measurement of the reproducible dephasing due to the differing chemical potentials of each trapped condensate, the large number of condensates involved resulted in only partial revivals of the Josephson oscillations.

Our first interference experiments could not measure a reproducible phase because the relative

positions of the camera, the condensate, and the "cutting" laser beam were not mechanically stable to within the 10 μm scale of the interference fringes [110]. But, as discussed below, even if it is not possible to measure the phase of an interference pattern, it is still possible to measure phenomena such as phase diffusion by observing the qualitative features of an interference pattern of three, rather than two, condensates.

8.2 Interference Patterns from Multiple Point Sources

The interference of two identical, fully coherent condensates results in a pattern with only one free parameter which can be measured — the relative phase of the two condensates. The relative phase can be determined by measuring the distance from the center of the interference pattern to the first interference maximum. As mentioned above, due to shot to shot drift in the camera alignment and the trapping potential, early experiments were unable to accurately make this measurement. When more than two condensates are made to interfere, however, extra phases are introduced to the problem, resulting in additional observable parameters which can be measured without a precise knowledge of the camera-trap alignment.

The interference pattern generated by the sudden release of atoms from several ideal point sources, each with a given phase at the time of their release, can be simply calculated. After the atoms have been allowed to freely expand for a time t, the phase that a given atom has acquired is correlated with the distance over which it traveled during the expansion. In other words, in order for the atom to have reached a certain place in a given time t, it must have had a given kinetic energy, resulting in an accumulated phase of $Et/\hbar$. So the wavefunction of a collection of coherent atoms freely expanding from a single point source at the origin is given by

$$\Psi = A(\mathbf{r})e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t+\phi_0)} = A(\mathbf{r})e^{i\left(\frac{Et}{\hbar}+\phi_0\right)} \tag{8.2}$$

$$= A(\mathbf{r})e^{i\left(\frac{P^2t}{2M\hbar}+\phi_0\right)} \tag{8.3}$$

$$= A(\mathbf{r})e^{i\left(\frac{M}{2\hbar t}\left(x^2+y^2+z^2\right)+i\phi_0\right)} \tag{8.4}$$

where $A(\mathbf{r})$ is a function of position and time which describes the spatial distribution of the particles (resulting from their thermal kinetic energy distribution).

Interference of collinear condensates The interference of N collinear point sources can be described by the sum of N wavefunctions of the form shown in Eq. (8.4). If the point sources are spaced along the x axis at regular intervals, each term in the sum is of the form

$$\Psi_n = A_n(\mathbf{r})e^{i\left(\frac{M}{2\hbar t}\left((x-nD)^2+y^2+z^2\right)+i\phi_n\right)} \tag{8.5}$$

where n is an integer (for an odd number of condensates) or a half integer (for an even number of condensates), and nD is the position of the nth point source on the x axis. For two condensates, the interference pattern is just $\left|\Psi_{-1/2}+\Psi_{+1/2}\right|^2$. If we ignore the A_n terms, this yields the cosine pattern observed in our first interference experiment.

Three condensate interference For three condensates positioned at $x = -D$, $x = 0$, and $x = +D$, the probability amplitude of the interference pattern is given by:

$$\Psi = \Psi_{-1} + \Psi_0 + \Psi_{+1} \tag{8.6}$$

To calculate the density of atoms at any point in the interference pattern, we simply evaluate $|\Psi|^2$. To simplify matters, I will assume that the atoms are probed after a long enough expansion time such that the three clouds are completely overlapping. This allows us to set $A_{-1}(\mathbf{r}) = A_0(\mathbf{r}) = A_{+1}(\mathbf{r}) \equiv A$. The equation is further simplified if we define the phase of the center condensate to be zero. We can now express the the shape of the interference pattern in the following form:

$$|\Psi|^2 = A\left[3 + C\cos\left(\kappa\left(x - x_0\right)\right) + 2cos\left(2\kappa\left(x - x_0\right)\right)\right] \tag{8.7}$$

where

$$C = 4\cos\left(\frac{MD^2}{2\hbar t} + \phi_S\right) \tag{8.8}$$

$$\kappa = \frac{MD}{\hbar t} \tag{8.9}$$

$$x_0 = \phi_A \frac{\hbar t}{MD} \tag{8.10}$$

and where

$$\phi_S = \frac{1}{2}\left(\phi_{+1} + \phi_{-1}\right) \tag{8.11}$$

$$\phi_A = \frac{1}{2}\left(\phi_{+1} - \phi_{-1}\right) \tag{8.12}$$

are the symmetric and antisymmetric variations of the phase of the $+1$ and -1 condensates with respect to the middle condensate.

The resulting interference pattern is simply the sum of two cosines, one with twice the spatial frequency of the other, plus a constant. The antisymmetric phase shift simply translates the whole pattern along the x axis. This translation is difficult to measure, for the same reasons that the phase of a two condensate interference pattern is hard to measure. The symmetric phase shift leaves the

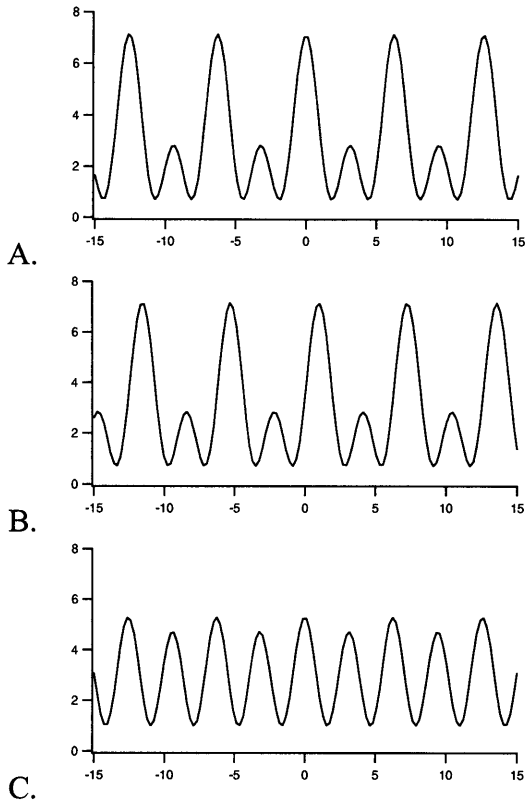

Figure 8.1: Calculated cross-section of a three condensate interference pattern. In these three plots, the density of atoms is shown as a function of position. (A) shows the interference of three condensates which all have the same phase. If the phase of the left and right condensates slip by equal amounts but in opposite directions, the interference pattern simply shifts in space (B). If the two outer condensates shift in phase in the same direction, a qualitative change occurs in the interference pattern (C).

center of symmetry of the pattern unchanged, but changes the amplitude of the lower frequency cosine. As a result, the symmetric phase shift makes a qualitative change in the appearance of the interference pattern, and can therefore be measured without any knowledge of the the origin of the trap (see Fig. 8.1).

8.3 Interference of a Superfluid Ring or Vortex State

Since a Bose condensate is a pure quantum state, the flow velocity at any point in a condensate is simply characterized by the gradient of the quantum phase. This puts certain limits on the types of flow patterns which can be sustained by a condensate. In particular, a "phase flow" is characterized by quantized vorticity; because the phase cannot be multiple valued, the line integral of the phase gradient around any closed loop must be an integer times 2π. In general, since the local fluid flow and the local phase gradient of the condensate are simply proportional to each other, interferometric measurements of the phase of the condensate can be used to measure the dynamics of a condensate.

The use of interference as a probe for vortices in Bose condensates was suggested independently in [101] and [102]. Both of these references suggest that a rotating condensate could be probed by interference with a stationary "reference" condensate sitting next to it. If a singly charged vortex is present in the rotating condensate, it would result in a bifurcation in the interference pattern. Prompted by these ideas, I have calculated what the interference pattern would look like in a different, possibly more useful geometry.

Since it may be possible in the near future to create arbitrary trap shapes (see Section 7.2), we have considered doing studies of vortices and persistent currents in a combination of a ring trap plus an additional round trap in the center of the ring. In this geometry, currents could be generated in one of the condensates, and the other condensate could be used as a "phase standard" for interferometric detection of the currents. This geometry has certain advantages over two side-by-side condensates. First of all, it allows us to explore the stability of vortices in simply-connected versus multiply-connected geometries. Theories have predicted that a vortex in a simply-connected inhomogeneous condensate will not be stable at zero temperature, since it can always lower its energy by moving towards the edge of the condensate. Numerical studies, however, have shown that the lifetime of such a vortex should be very long [111]. Furthermore, it has been shown that

finite temperature effects can stabilize vortices [112, 113]. Currents flowing around a ring shaped trap, however, can be thought of as having a vortex which is "pinned" by the potential barrier in the center of the trap, resulting in a stable supercurrent [112]. The stability of vortices could be studied in this trap geometry by comparing the interference patterns produced when currents are driven in the outer ring condensate and the central core condensate is used as the phase reference, to the reverse case in which the core atoms are stirred and the ring is used as the reference.

Another advantage of this proposed geometry is the proximity of the two condensates. The fringe spacing in an interference pattern is inversely proportional to the initial spacing between the interfering condensates (Eq. (8.9)). So by getting the rotating condensate closer to the reference condensate, larger fringe patterns, which can be more accurately imaged and measured, will be produced. In this proposed ring pattern, the ring and the core can be separated by as little as the beam waist of the scanning laser used to generate the trap. Two side-by-side round condensates, on the other hand, could only be this close at their nearest point, but with an average separation which is much larger.

The experiment proposed above would be performed as follows. First, a Bose condensate would be formed in a magnetic trap. The condensate would then be transfered into an optical trap, which would then be deformed and broken into a round trap surrounded by a ring shaped trap, pulling approximately half of the condensate into each. Then either the ring or core trap would be deformed slightly to break axial symmetry, and the deformation would be rotated to impart angular momentum to the trapped atoms. Finally, the deformation would be removed, the condensates allowed to evolve freely for a given time, and then the trap would be switched off suddenly, causing the two condensates to ballistically expand into each other. If no supercurrent is present, the resulting interference pattern will be an axially symmetric pattern of concentric rings. If the outer ring is rotating with one quanta of circulation, there will be a 2π phase shift around the ring. As a result, if we travel around one interference maxima, when we get to the other side we

expect it to join up with the next interference maxima, resulting in a spiral pattern moving out one de Broglie wavelength per turn. If the ring trap is given N quanta of circulation, we would expect an interference pattern with N arms spiraling out (see Fig. 8.2).

Calculation of the exact interference pattern is made simple if we assume that atoms in the ring trap only expand perpendicular to the surface of the ring. This is a good assumption if the thickness of the ring is several times smaller than the circumference of the ring, and if the mean field energy of the ring-shaped condensate is much larger than its kinetic energy. In this case, the interference pattern can be calculated independently for each polar angle θ using the equation for a three condensate interference pattern calculated above (see Eq. 8.7)). The entire pattern can then be written, in polar coordinates, as:

$$|\Psi(r,\theta)| = A\left[3 + C\cos\left(\kappa\left(r - r_0\right)\right) + 2cos\left(2\kappa\left(r - r_0\right)\right)\right] \tag{8.13}$$

where

$$C = 4\cos\left(\frac{MD^2}{2\hbar t} + \phi_0\right) \tag{8.14}$$

$$\kappa = \frac{MR}{\hbar t} \tag{8.15}$$

$$r_0 = -q\theta\frac{\hbar t}{MD} \tag{8.16}$$

and where ϕ_0 is the relative phase of the atoms in the ring at $\theta = 0$ with respect to the atoms in the spot at the center of the ring, and q is the number of quanta of circulation in the rotating trap. The predicted interference pattern, given reasonable experimental parameters, is plotted in Fig. 8.2 for several values of q.

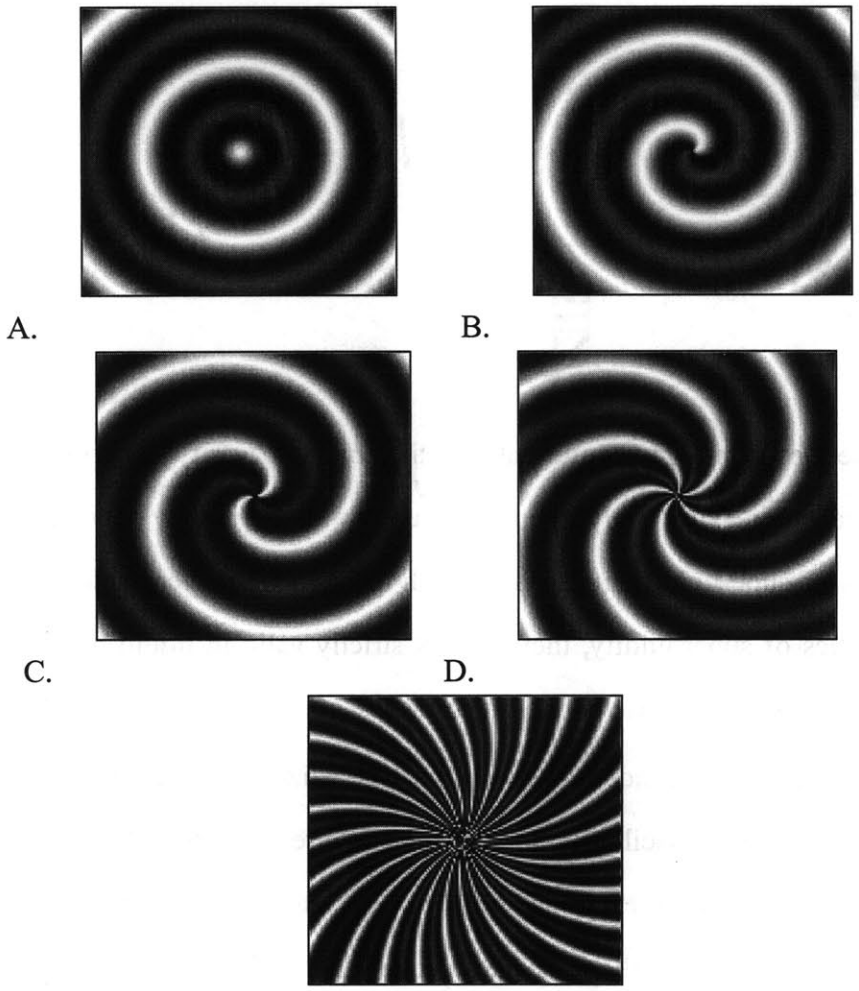

Figure 8.2: Interference of a point with a superfluid ring. Shown is the interference pattern which would result from the ballistic expansion of a point condensate surrounded by a ring condensate with (A) zero, (B) one, (C) two, (D) five, and (E) 25 quanta of angular momentum.

Chapter 9

Surface Waves

One exciting aspect of dilute gas condensates is the ability to verify experimentally many long standing theoretical studies of weakly interacting superfluids (originally developed to gain insight into the dynamics of superfluid helium). Although these theories were able to reproduce many of the characteristics of superfluidity, they are not strictly valid in liquid helium. They should, however, be valid in dilute gas Bose condensates.

Soon after alkali gas Bose condensates were finally realized, studies which compared the fundamental collective modes of oscillation to zero temperature mean-field theories were performed [54, 114]. These early collective excitation experiments, carried out by modulating the magnetic trap parameters at the theoretically predicted frequencies [115], showed good agreement with mean-field theory. These experiments, however, were limited by the geometry of the magnetic trap to exciting modes with cylindrical symmetry. The first experiment done in our new apparatus extended the techniques of these early studies with the use of the optical dipole potential of a scanning laser beam (Section 7.2), allowing modes with finite angular momentum to be excited. The use of scanning optical potentials allows for a broad range of excitations to be created with a great deal of mode selectivity (the modes which are excited are selected by both the modulation

frequency *and* the geometry of the optical dipole potential).

Using this new tool, we have been able to excite and study standing and rotating "surface wave" oscillation modes. Oscillations such as these could be useful for the study of the formation of vortices and vortex stability [111], and are a close counterpart to surface excitations in mesoscopic liquid Helium droplets [116, 117]. Also, as will be discussed, these types of oscillations might prove useful for very low temperature thermometry and studies of the transition from hydrodynamic to single particle excitations. In this chapter a brief theoretical description of the collective modes studied will be given, followed by a detailed description of the experiments which were performed and the results which we obtained.

9.1 Hydrodynamic Excitations at Zero Temperature

The behavior of dilute gas condensates at absolute zero temperature can be described exactly by the nonlinear Schrödinger equation (NLSE):

$$i\hbar\frac{\partial\Psi}{\partial t} = -\frac{\hbar^2}{2M}\nabla^2\Psi + V_0|\Psi|^2\Psi + V_{ext}\Psi \tag{9.1}$$

In this equation, the interaction between atoms is described by a mean-field term proportional to the particle density and to the constant $V_0 = \frac{4\pi\hbar^2 a}{M}$ (where a is the s-wave scattering length). Although the NLSE had been studied in great detail for many years (since long before the advent of dilute gas condensates), most of these studies were aimed at understanding the properties of superfluid helium, and as such they were mainly carried out in the thermodynamic limit. Since the realization of dilute gas condensates, these theories have been extended to describe gases in inhomogeneous potentials.

Static solutions to the NLSE In the thermodynamic limit the static solution of the NLSE is trivial; the mean-field term simply assures that the condensate density is uniform throughout. In an inhomogeneous trapping potential, however, the mean-field repulsion is countered by the confining potential, resulting in an inhomogeneous particle density. In the trivial limit in which the mean-field energy is much smaller than the zero-point energy of the trap, the condensate wavefunction is approximated by the single particle harmonic oscillator ground state. In the opposite limit, in which the mean-field energy dominates the zero-point energy of the condensate (the Thomas-Fermi limit), the density of particles is simply distributed such that the local mean-field energy exactly cancels the local trapping potential, resulting in a density profile which mimics the inverted shape of the trap potential:

$$\rho(\mathbf{r}) = \rho_0 - \left(\frac{M}{4\pi\hbar^2 a}\right) V_{ext}(\mathbf{r}) \tag{9.2}$$

For the condensates described in this chapter, the Thomas-Fermi approximation is valid everywhere except the very edges of the condensate, allowing a very simple and accurate prediction of the condensate density profile.

Dynamic solutions to the NLSE By performing the Madelung transformation (in which the wavefunction is written as $\Psi(\mathbf{r}) \equiv \sqrt{\rho(\mathbf{r})}e^{iS(\mathbf{r})}$, where $\rho(\mathbf{r})$ is the local density of the condensate, and $S(\mathbf{r})$ is related to the local flow velocity by the relation $\mathbf{v}(\mathbf{r}) = \frac{\hbar}{M}\nabla S(\mathbf{r})$), the NLSE can be mapped into a hydrodynamic equation in which the local density ρ and the local flow field velocity $\mathbf{v}$ appear explicitly in place of the wavefunction Ψ. With the assumption that the kinetic energy pressure is negligible in comparison to the mean-field and trapping potential terms (in a manner similar to the Thomas-Fermi approximation in the static case), the hydrodynamic equation takes the form [118]:

$$M\frac{\partial}{\partial t}\mathbf{v} + \nabla\left(V_{ext} + \frac{4\pi\hbar^2 a}{M}\rho + \frac{Mv^2}{2}\right) = 0 \tag{9.3}$$

where a is the s-wave scattering length of the atoms. For the most part, the assumption used to obtain this equation is valid for typical condensates and excitations produced in our apparatus. The approximation breaks down, however, for $r \gtrsim R$ where R is the Thomas-Fermi radius of the cloud (the displacement from the center of the cloud at which Eq. (9.2) predicts zero local atom density), because near this point the mean-field energy becomes small. It also breaks down for modes in which the local density varies on a length scale smaller than the local healing length $\xi = (8\pi\rho a)^{-1/2}$.

In reference [115], Eq. (9.3) is solved for an isotropic harmonic oscillator potential ($V_{ext} = \frac{1}{2}M\omega_0^2 r^2$) . In this spherically symmetric case, the density oscillations have the form of a polynomial in r times a spherical harmonic in θ and ϕ. As such, the oscillations can be described by the quantum numbers n, l, and m, where n is the number of radial nodes in the oscillation, and l and m describe the angular momentum and the z-projection of the angular momentum of the oscillation mode. The frequencies of the modes are given by

$$\omega(n,l) = \omega_0\sqrt{2n^2 + 2nl + 3n + l} \tag{9.4}$$

Hydrodynamics in cylindrically symmetric harmonic traps The experiments described in this chapter were not done in a completely symmetric trap, but in a trap with axial symmetry ($V_{ext} = \frac{1}{2}M\omega_r^2 r^2 + \frac{1}{2}M\omega_z^2 z^2$), complicating somewhat the hydrodynamic equation. Each mode in an asymmetric harmonic potential can, however, be linked to a corresponding mode of the symmetric case, allowing the same terminology to be used. In the special case of $n = 0$ and $m = \pm l$, the oscillation frequencies in an axially symmetric trap are [115]:

$$\omega_{(m=\pm l)} = \omega_r \sqrt{l} \tag{9.5}$$

where $\omega_r = 2\pi\nu_r$, ν_r being the radial trapping frequency. In this case, the oscillating density perturbation has the form:

$$\delta\rho \propto r^l Y_{lm}(\theta, \phi) e^{-i\omega t} \tag{9.6}$$

9.2 High Angular Momentum Modes

The r^l dependence of the density perturbation in Eq. (9.6), combined with the knowledge that this solution breaks down near the Thomas-Fermi radius ($R = \sqrt{\frac{2\mu}{M\omega_r^2}}$, with μ being the the chemical potential of the condensate), implies that as l is made larger, $\delta\rho$ becomes increasingly localized near the Thomas-Fermi radius. (In the full solution, without the above approximation, $\delta\rho$ should quickly approach zero beyond the Thomas-Fermi radius, since the static density rapidly approaches zero beyond this point.) These $n = 0$, $m = \pm l$ excitations, then, constitute surface waves, the study of which could reveal many interesting properties of the condensate. Surface waves could, for example, facilitate precise measurements of interactions between condensate and non-condensate atoms (even at very low temperatures, at which very few non-condensate atoms remain). This is due to the fact that the non-condensate atoms feel a stronger mean-field repulsion from the condensate than do atoms which are part of the condensate, resulting in a combined magnetic trap plus mean-field potential with a minimum near the Thomas-Fermi radius. For temperatures below the condensate's chemical potential, this combined potential localizes the non-condensate atoms to a thin shell around the surface of the condensate, causing them to interact strongly with surface waves.

Another interesting aspect of surface waves is that for modes with sufficient angular momen-

tum, the surface oscillation will be localized to a shell thinner than the local healing length. As such, surface waves could facilitate the study of the transition from hydrodynamic to single-particle excitations. For single-particle excitations, the dispersion relation should be similar to that of an ideal, interaction free Bose gas ($\omega_{IG} = \omega_r \cdot [2n + l]$) rather than the relations given in Eqs. (9.4) and (9.5). By studying collective modes for which $n = 0$, the transition from collective modes (for which the frequency scales as $\sqrt{l}$) to free particle excitations (for with the frequency scales as l) can be studied. This transition is expected to occur for $l \geq 2^{1/3}(R/a_{HO})^{4/3}$ where $a_{HO} = \sqrt{\hbar/M\omega_r}$ is the characteristic radial size of the harmonic trapping potential [118]. For the condensates used in the experiments described in this chapter, this transition should occur at $l \gtrsim 26$. In these experiments we were only able to excite modes with l as high as four, due to the size of the beam waist of our optical dipole beam as compared to the size of the condensate. Exploration of the crossover to single particle excitations might be possible in our apparatus, but will require expanding the condensate to be several times larger, and focusing the dipole beam more tightly.

9.3 Experimental Realization of Surface Waves

The use of optical dipole forces has allowed us to excite $l = 2$ and $l = 4$ collective excitations modes. Each of these experiments was done by first creating a condensate of about 25 million atoms in a magnetic trap characterized by radial and axial trapping frequencies of $\nu_r = 90.1 \pm 0.5$ Hz and $\nu_z = 18$ Hz, and then superimposing a time-averaged optical dipole potential on top of the magnetic trapping potential. By modulating the intensity of the optical potential, or by rotating the optical potential about the trap axis, we have been able to create and study standing and rotating surface waves. In each of these experiments, the optical potential was generated with the focused beam of a Nd:YAG laser (emitting at 1064 nm).

9.3.1 $l = 2$ "Quadrupole" Modes

Standing waves To generate a standing $l = 2$ surface wave, an optical dipole potential was generated with a focused laser beam propagating parallel to the symmetry axis of the condensate. The laser was focused to a beam waist diameter of 12 μm (compared to the $\sim$ 50 μm Thomas-Fermi diameter of the condensate in the radial direction). The Rayleigh range of the focused beam was 1 mm (several times larger than the 220 μm axial length of the condensate); as a result, the laser mainly generated radial perturbations to the magnetic trap. The laser focus was rapidly scanned in space between two positions, jumping from one side of the condensate to the other at a frequency of 10 kHz (about 100 times higher than the radial trapping frequency). This was done using the AOM beam deflectors described in Chapter 7. The scanning laser focus created a time-averaged potential consisting of two dimples on either side of the trap, each at a distance of 20 μm from the trap center (see Fig. 9.1).

To excite a standing $l = 2$ surface wave, the laser power was modulated sinusoidally between 0 and 80 μW (creating dimples in the potential which had a maximum depth of $\sim$ 20% of the condensate's 200 nK chemical potential). This modulation was done at a frequency of 130 Hz (near the frequency predicted by Eq. (9.5) for the $l = 2$ modes), exciting a superposition of the $l = 2$, $m = 2$ and the $l = 2$, $m = -2$ modes. After five cycles the laser beam was turned off, and the condensate was left to oscillate in the magnetic trap for a variable amount of time. The atoms were then probed in time-of-flight (see Fig. 9.2). By measuring the aspect ratio of the expanding clouds as a function of the free oscillation time in the magnetic trap, the frequency of this excitation was determined to be 130.5 $\pm$ 2.5 Hz = $1.45\nu_r$, very close to the value of $\sqrt{2}\nu_r$ predicted by hydrodynamic theory. The damping time of this excitation was measured to be about 600 ms.

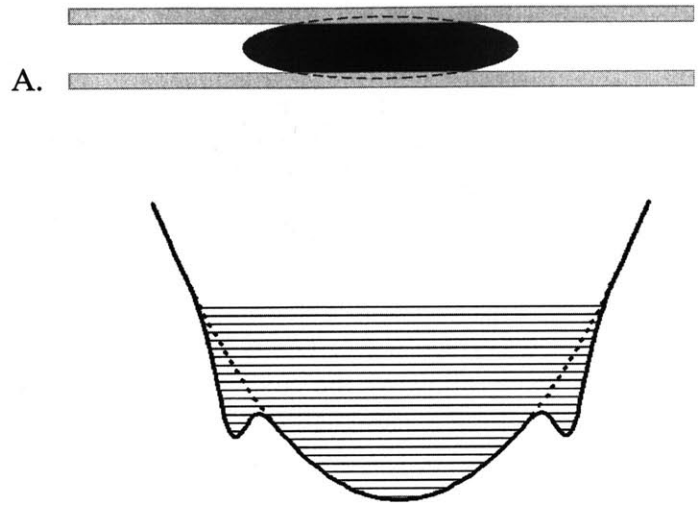

Figure 9.1: Optical perturbations to the magnetic trap. Modes with angular momentum of $l = 2$ and $l = 4$ were driven by adding a time-averaged optical perturbation, generated with a scanning focused laser beam, to a magnetic trap. With this technique, perturbations with "arbitrary" geometries could be made. To excite the $l = 2$ mode, the dipole beam was switched rapidly from one side of the condensate to the other (A), generating a time averaged optical potential consisting of two dips on either side of the condensate, as shown in figure (B). Excitations were driven by modulating the intensity of these dips or by rotating the dips around the trap axis.

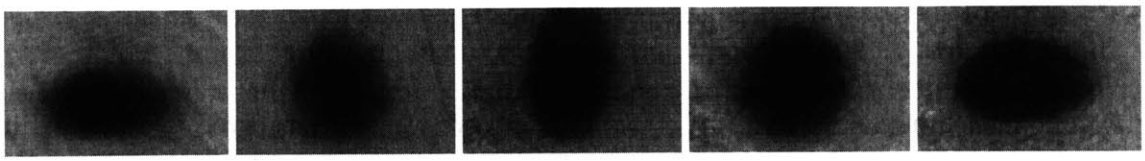

Figure 9.2: Time-of-flight images of a freely oscillating $l = 2$, $|m| = 2$ mode. The images shown were taken of condensates which were driven and then allowed to freely oscillate for 3.5, 5.25, 7, 8.75, and 10.5 ms (left to right). After this free oscillation time the magnetic trap was turned off, and the clouds were allowed to expand ballistically for 20 ms before being imaged. Each frame is 720 μm wide.

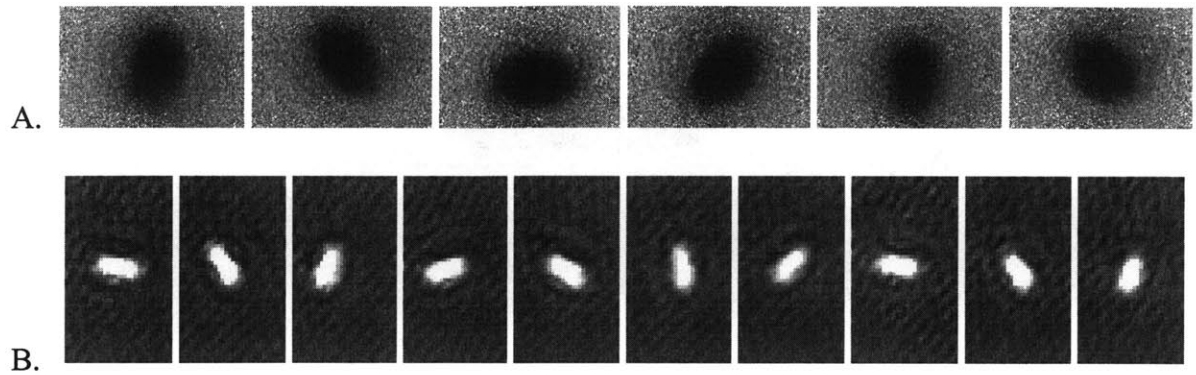

Figure 9.3: Rotating $l = 2$, $m = 2$ surface wave. (A) is a series of time-of-flight images taken after exciting a rotating surface wave, allowing the condensate to oscillate freely for a variable amount of time, and then letting the cloud expand ballistically for 20 ms. The free oscillation time of each condensate was (from left to right) 0, 1.9, 3.8, 5.7, 7.6, and 9.5 ms. In (B) a rotating surface wave was observed in-situ by taking multiple images of a single oscillating condensate using non-destructive phase contrast imaging (see Chapter 6). The images were taken after the optical driving potential was turned off at a rate of one frame every 2 ms. Each image is 720 μm wide in (A) and 140 μm wide in (B).

Rotating waves A pure $l = 2$, $m = 2$ rotating excitation was driven by rotating the two-point optical dipole perturbation described above (at constant laser intensity) around the perimeter of the condensate. The rotation was done at half the measured quadrupole frequency, resulting in the rotating surface waves shown in Fig. 9.3.

9.3.2 $l = 4$ "Octopole" Standing Waves

A higher angular momentum $l = 4$, $|m| = 4$ standing wave was also observed in our experiment. This mode was driven with an intensity-modulated optical dipole perturbation similar to the one used to generate the $l = 2$ standing wave discussed above. In this case the laser was scanned in two dimensions, creating an optical potential consisting of four dimples at the corners of a square (40 μm on a side). After creating the condensate in the magnetic trap, the optical potential was ramped on as before. Then the depth of the optical potential was modulated at the frequency predicted by

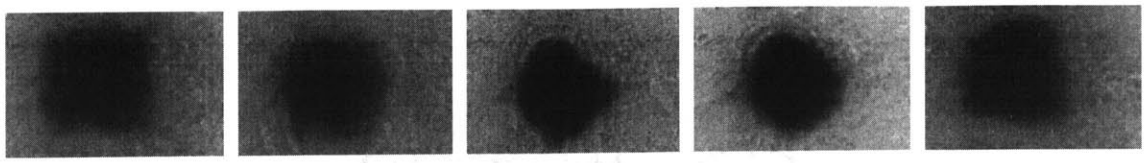

Figure 9.4: Time-of-flight images of $l = 4$ mode oscillations. To produce these images a condensate was driven with a four-point optical perturbation at the frequency predicted for the $l = 4$, $|m| = 4$ mode and then allowed to freely oscillate for 1, 2, 3.5, 4.5, and 6.5 ms (from left to right). They were then released from the trap and observed after 20 ms of free expansion. Each of the images is 720 μm wide.

Eq. (9.5) for the $l = 4$, $|m| = 4$ modes. After a few cycles, the optical potential was switched off, and the condensate was left to ring freely in the magnetic trap for a variable amount of time. Time-of-flight images were then taken (see Fig. 9.4), revealing an oscillating velocity distribution with the same symmetry as the driving potential.

The time-of-flight images were analyzed by extracting the Fourier components of the mean radius $r_0(\theta)$ of each image. To make our analysis more sensitive to surface effects and less sensitive to noise, each image was analyzed using only the data which fell between the contours of 30 and 60 percent transmission of the probe light (see Fig. 9.5). The dominant Fourier components in this analysis were found to be the $l = 4$ sine and cosine components. The evolution of the $l = 4$ cosine component in time is plotted in Fig. 9.6. From a fit to an exponentially decaying sine wave, the frequency of this mode was determined to be 177 ± 5 Hz $= 1.96\nu_r$, very close to the $2\nu_r$ prediction from hydrodynamic theory, and the damping time was determined to be 9.5 ± 2.2 ms. It is interesting to note that this mode damped much more quickly than the $l = 2$ mode, although both modes should be stable in the limit of zero temperature. This is probably an indication of the extreme sensitivity of surface waves to interactions with the thermal cloud mentioned earlier in this chapter.

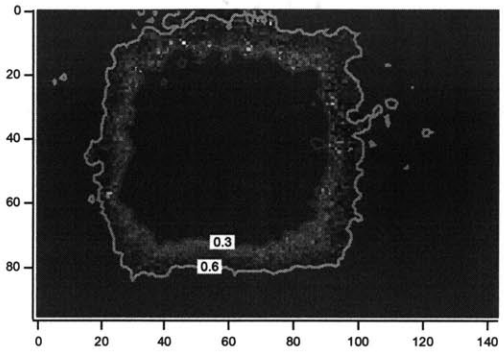

Figure 9.5: Analysis of data for the $l = 4$, $|m| = 4$ excitation. The time-of-flight images of the standing $l = 4$ mode were analyzed by extracting the Fourier components of the expanded cloud's mean radius. In order to emphasize the contribution from the surface of the condensate, and to reduce sensitivity to noise present in the center of the cloud (where the optical density was high — see Section 6.1.1), only data falling between the contours of 30% and 60% transmission of the probe light was considered in the analysis.

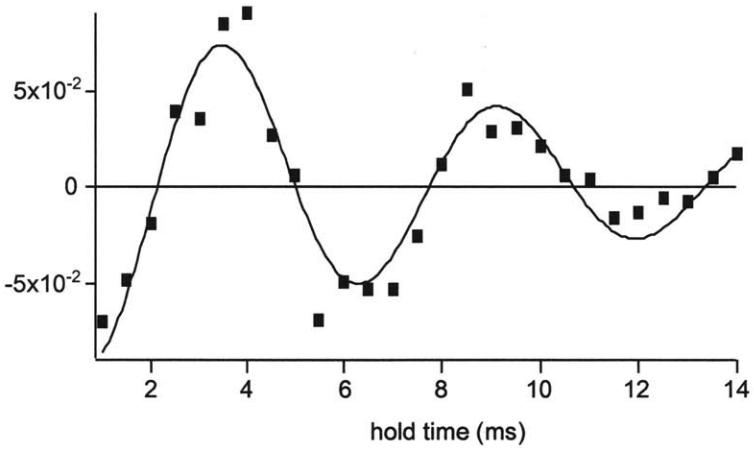

Figure 9.6: Time evolution of the $l = 4$ cosine term of the angular Fourier spectra.

9.4 Conclusion

The types of oscillations studied in these experiments could be used in further experiments to explore several interesting aspects of condensate dynamics. As mentioned in Section 9.2 above, the localization of a high angular momentum surface wave to a thin shell would make it possible to study the transition from hydrodynamic to single-particle excitations. Using optical dipole forces, it is difficult to drive excitations with a wavelength smaller than the beam waist of the laser — limiting us to $l = 4$ excitations in the above experiments. In order to excite higher l oscillations the laser must be focused to a smaller spot, and the cloud must be decompressed further to increase the wavelength of high l surface waves.

A factor of two in the laser's beam waist could be achieved without changing the f/# of our optics simply by using a shorter wavelength blue-detuned laser, like the one used in the studies described in Chapter 10. An additional factor of 12 could be gained if the dipole beam were focused with unity f/# optics. Further expansion of the cloud has so far been difficult due to the tendency of decompression to create dipole excitations in the condensate. This is a technical limitation, however, which we hope to soon overcome. In addition to increasing the wavelength of high l excitations, expanding the atom cloud also results in condensates which have a lower maximum density. As a consequence, R/a_{HO} is smaller, allowing the transition to single particle excitations to be explored at lower l.

Rotating surface waves also have an interesting connection to studies of vortex formation. While it is tempting to equate the observed $l = 2$, $m = 2$ oscillation in Fig. 9.3 to experiments with liquid helium in a rotating vessel, this mode is best pictured as two out-of-phase radial oscillations rather than by a static state in a rotating frame. In a manner similar to the way in which linear momentum is carried by a wave on a wire that is vibrating transversely to the wave's propagation, the angular momentum contained in these oscillations is carried by the wave rather than by

individual particles, with the condensate atoms oscillating radially with no circulation.

The strongest evidence that what we observed was a collective excitation, rather than a rigid body rotation enabled by the formation of vortices, is that the rotating wave could only be excited with drive frequencies near the frequency predicted for the $l = 2$, $m = 2$ surface wave. The creation of vortices, on the other hand, should begin at rotation frequencies which are much lower. In the case of an ideal, non-interacting gas, the condition analogous to a single charged vortex is the state in which each particle in the condensate has one quanta of orbital angular momentum. This state is created by rotating the system at exactly the radial trap frequency ν_r. The presence of interactions reduces the critical rotation frequency for vortex formation such that it is far from the surface wave resonance. In [113], the Thomas-Fermi approximation is used to derive the minimum rotation frequency for vortex formation in a condensate held in an axially symmetric harmonic trap:

$$\nu_c = \nu_r \frac{5}{2} \left(\frac{\nu_z}{\nu_r}\right)^{1/3} \left(\frac{\bar{a}_{HO}}{R_r}\right)^2 \ln\left[\frac{0.67 R_r}{\xi}\right] \tag{9.7}$$

In this equation, $\bar{a}_{HO} = \sqrt{\frac{\hbar}{2\pi M(\nu_z \nu_r^2)^{1/3}}}$ is the harmonic oscillator length associated with the mean trapping frequency, ξ is the healing length, and R_r is the Thomas-Fermi radius of the condensate in the radial direction. For the the condensates used in this study, $a_{ho} = 2.9$ μm, $R_r = 25$ μm, and $\xi = 0.28$ μm, resulting in a critical velocity for vortex formation of $0.08\nu_r$, about 9 times lower than the drive frequency used to create rotating surface waves. If the oscillations we observed were to be attributed to vortex formation, a small increase in the drive frequency should merely result in the creation of a few more vortices, with very little change in the observed dynamics. However, when we varied the rotation frequency up or down by 10%, no observable excitation occurred. This sharp resonance does not *rule out* the formation of vortices. It does, however, indicate that majority of the dynamics observed were due to a collective excitation mode, and that vortices, if present, accounted for an immeasurable contribution to the density oscillation.

There are, however, certain links between vorticity and $m = \pm l$ excitations. For example, the frequency of these types of excitations are affected by the presence of vortices, and it may be possible to detect vortices by observing a difference between the frequencies of the $m = +l$ and $m = -l$ modes [119, 118] in a manner similar to the way in which quantized vortices were first measured in liquid helium (see Section 10.2).

Chapter 10

Dissipation, Critical Velocities, and Vortices

Perhaps the most striking aspect of superfluid liquid helium is the existence of a critical velocity, below which fluid flow is completely dissipationless. While dissipationless flow was one of the first things which caught the attention of early liquid helium experimenters (the word "superfluid" was, in fact, coined specifically to describe the "abnormally low viscosity" of liquid helium below the λ-point [13]), superfluids are more than just "ideal friction free" fluids. Phenomena such as second sound, Josephson oscillations, and quantized excitations cannot be explained simply by eliminating the viscosity term of the Navier-Stokes equation. In fact, the *existence* of a critical velocity, below which flow is dissipationless, is at least as interesting as the dissipationless flow itself.

In this chapter, an experiment which resulted in evidence for the existence of a critical velocity in a sodium vapor Bose condensate is discussed. This experiment was done by stirring the condensate using the repulsive force of a focused blue-detuned laser beam, and observing the temperature rise in the condensate as a function of the velocity of the stirring beam. Before giving the experimental details, a short discussion of Landau's prediction of a critical velocity in liquid helium, and the relationship between critical velocities and vortices will be given, and two unsuccessful meth-

ods to directly measure vorticity and supercurrents will be mentioned. The experiment described in Section 10.4 was begun while I was in the process of writing my thesis, and as such, several of the more recent developments will not be discussed here. The results of this experiment will be published soon [120], and I refer the reader to this reference for more details.

10.1 Critical Velocities in Landau's Model of Damping

The origin of a critical velocity in superfluid helium has been the subject of many experimental and theoretical studies. The concept of a critical velocity was introduced by Landau in 1941 [121, 122]. In Landau's theory of damping, an object moving through a superfluid generates quasiparticle excitations in the fluid (phonons and rotons), resulting in dissipation. Phonons and rotons with a given momentum $\hbar\mathbf{k}$ have a specific energy $\epsilon(\mathbf{k})$, given by the phonon and roton dispersion relations. As a result, conservation of energy and momentum puts a limit on what types of excitations can be generated by an object moving through a superfluid:

$$\frac{1}{2}Mv^2 = \frac{1}{2}Mv'^2 + \epsilon(\mathbf{k}) \tag{10.1}$$

$$M\mathbf{v} = M\mathbf{v}' + \hbar\mathbf{k} \tag{10.2}$$

where $\epsilon(\mathbf{k})$ is the energy of the emitted quasiparticle, $\mathbf{k}$ is the quasiparticles wave-vector, and M is the mass of the moving object. These two equations can be combined into a single relation which must be satisfied in order for quasiparticles to be created:

$$\hbar\mathbf{v}\cdot\mathbf{k} = \epsilon_{\mathbf{k}} + \frac{\hbar^2k^2}{2M} \tag{10.3}$$

This equation is satisfied for the lowest possible velocity when $\mathbf{v}$ and $\mathbf{k}$ are in the same direction. If we assume that the object's mass is large (such that a single quasiparticle removes only a small fraction of the object's momentum), this results in a critical velocity for dissipation $v_c = \min\left(\frac{\epsilon(\mathbf{k})}{\hbar k}\right)$. If an object moves through the superfluid at a velocity below v_c, it cannot excite quasiparticles and the fluid will flow around it without dissipation (see Chapter 7 of [6] for an introductory level discussion of the critical velocity for Landau damping). Above the critical velocity, it becomes possible to remove momentum and energy from the fluid in the correct proportions to allow the creation of quasiparticles, thus allowing dissipation to occur.

Landau's prediction was proven to be correct when a critical velocity was experimentally measured in helium. Early measurements of the critical velocity, however, were much lower than the one predicted by Landau [123]. It seemed from these measurements that there must be another important mechanism for damping which can occur at lower flow velocities.

10.2 Relationship Between Critical Velocities and Vorticity

As pointed out by Feynman in 1955, vortices in a flowing classical fluid with low viscosity are unstable, resulting in the onset of turbulence and the conversion of kinetic energy into heat. Similar behavior, he points out, may occur in superfluid liquid helium [124]. The theory of superfluid vortices and their relation to dissipation in fluid flow has a long history. In 1949, Lars Onsager suggested that, since the mass flow of a pure quantum state is simply proportional to the gradient of the local quantum phase, and since the quantum phase is single valued, the hydrodynamic circulation ($\oint \mathbf{v} \cdot d\mathbf{l}$) must be quantized (see Chapter 2, Section 3 of reference [125]). In the case of a cylindrically symmetric flow pattern, quantization of circulation forces the velocity field of the flowing superfluid component to be $\mathbf{v}_s = \frac{q\hbar}{Mr}\hat{\theta}$, where M is the particle mass, r is the displacement from the axis of the flow, and q is an integer. Note that for any value of q other than zero (the

static case), this flow pattern has a finite vorticity ($\nabla \times \mathbf{v}_s$), violating the assumptions of Tisza's two-fluid model of liquid helium (which explicitly assumed that the vorticity of the superfluid component is *everywhere* zero). As a result, it was predicted that the superfluid component of a liquid in a cylindrically symmetric geometry should remain stationary, regardless of the rotation of the container.

Experiments by Osborne in 1950 challenged the assumption of zero vorticity [126]. In these experiments, a cylinder of liquid helium was rotated about its axis, and the shape of the meniscus of the helium surface was measured. In a classical rotating fluid, the shape of the meniscus can be easily calculated by balancing the potential energy due to gravity with the centrifugal potential. According to the two fluid model, however, only the normal component should participate in the flow. So the height difference between the center and the edge of the container should be reduced by the fraction of normal fluid density to the total fluid density. The experimental results, however, were in close agreement with *classical* predictions — the superfluid fraction seemed to be undergoing classical rigid-body rotation! This experiment gave an early indication of the existence of vorticity in the superfluid component, and as a result, theoretical models had to be altered to include a critical velocity above which vorticity was allowed [17].

In the early fifties, after the report of Osborne's experimental results, there were several theoretical conjectures as to the nature of superfluid vorticity. In 1955 Feynman suggested a model in which the vorticity was concentrated in vortex filaments of atomic dimensions [124] (reviewed in [127] and [125]). It was already known that quantized circulation could be sustained, while satisfying the curl free conditions of the two fluid model, in multiply-connected geometries. Feynman proposed that the superfluid could be somehow expelled from the vortex core, punching a hole in the superfluid component and thereby generating a multiply-connected fluid. In this paper he went on to state the possibility that the critical velocity for superfluid flow was determined by the onset of these vortices, and showed that a good estimate for the critical velocity of superfluid

helium could be made based on this assumption. He also proposed several mechanisms which could convert the energy absorbed by the creation of a vortex into heat, including collisions between vortices and quasiparticles, interactions of the vortices with the walls of the helium dewar, and the repeated subdivision of vortices into smaller vortices (resulting, finally, in a multitude of rotons—the microscopic analog to vortices).

Another mechanism by which vortices can lead to heating in *dilute gas* condensates can be inferred from the work of Pitaevskii. In a paper written in 1961, Pitaevskii considers the properties of vortices in a weakly interacting gas described by the nonlinear Schrödinger equation [128]. In this paper it is shown that once a vortex is created, vibrations of the vortex line constitute a new branch of elementary excitations. Furthermore, he shows that it is energetically favorable for high wavenumber excitations of this type to decay into multiple lower wavenumber excitations. This behavior could lead to an additional damping mechanism which is mediated by the creation of vortices.

The first detection of quantized vortices was done by measuring the vibration frequencies of a wire strung along the axis of a rotating cylinder of liquid helium [129, 130]. By measuring the splitting of the degeneracy between right- and left-handed circularly polarized vibrations on the string, the presence of quantized vortices were clearly demonstrated. Soon after the detection of quantized vortices, studies of ion drift through liquid helium explored many of the properties of superfluid vortices [131, 132]. These studies measured critical velocities for dissipation which were orders of magnitude lower than the velocity predicted by Landau's theory, demonstrating that under normal circumstances, the critical velocity for dissipation in He II is set by the conditions for vortex generation, and not by the phonon/roton spectrum. Landau's critical velocity was eventually measured in ion drift experiments [133, 134], but not without going to great lengths to avoid the generation of vortices. For further discussion of the link between vortex formation and the existence of a critical velocity, I refer the reader to Chapter 2 of reference [17] and the references

contained therein.

10.3 Direct Search for Persistent Currents and Vortices

Bose condensates in dilute vapors have long been expected to exhibit superfluidity in a manner similar to liquid helium. Although an ideal gas exhibits no hydrodynamic behavior (quantum or otherwise), numerous theoretical and numerical studies have been performed which suggest that the mean-field term in the nonlinear Schrödinger equation is sufficient for the existence of a critical velocity [135], and the existence of stable [112, 113] or metastable [111] vortices in trapped dilute Bose gases. Experimental studies which can detect superfluidity or quantized vortices, however, have taken longer to realize. Several experiments have been suggested or tried, but the experimental observation of vortices and persistent currents has turned out to be somewhat difficult — partially due to the difficulty in *detecting* supercurrents and vortices once they are produced.

In an early study by our group, a beam from an argon ion laser, focused through the center of a harmonic magnetic trap, was used to generate a toroidal trapping geometry. The condensate was then spun up by displacing the center of the magnetic trap with respect to the focused laser beam in a circular pattern $\boldsymbol{r}_0 = \hat{x} r_0 \cos(\omega t) + \hat{y} r_0 \sin(\omega t)$. After the cloud was driven, and then allowed to oscillate freely for some time, the condensate was then imaged in time-of-flight. It was hoped that, if a supercurrent was present, the center of the time-of-flight image would be depleted (see [118] and the references therein). The results of this experiment were inconclusive. No depleted core was detected, and it was unclear whether the core, if present, would have a high enough contrast to be detected above the background noise in the image.

In our new apparatus, we had hoped to create and detect persistent currents with a more complicated toroidal trapping geometry, generated with a scanning optical dipole beam (see Section 7.2). The flexibility of the scanning potential would allow us to "cut" the torus. By rotating the

location of this cut, the condensate could be pushed along the ring to generate a current. Probing of the condensate would also be facilitated by the scanning dipole potential. By generating a second trap minimum in the center of the ring, a second "test" condensate could be held. After driving the ring condensate, both traps would be turned off, allowing the two condensates to expand into one another and interfere. The resulting interference pattern would give a definite indication of supercurrents in the ring (see Section 8.3). Although this approach has the potential to overcome the uncertainties of our earlier attempt, the work has been thwarted for the time being by heating problems in the scanning dipole trap as discussed in Section 7.2.2.

10.4 Evidence of a Critical Velocity in a Dilute Bose Gas

Evidence for the existence of a critical velocity in dilute gas condensates was recently observed in our lab. This was done by observing the heat generated when a focused laser beam was "dragged" through the center of the condensate. To perform this experiment, a Bose condensate was first created in a magnetic trap and then decompressed adiabatically until the condensate had expanded to a length of 150 μm and a radial width of 45 μm. The expansion process was not completely adiabatic, and as a result the cloud was heated, reducing the condensate fraction to about 60% of the total number of trapped atoms. The expanded cloud had radial and axial trapping frequencies of 65 and 18 Hz, respectively, and a peak density of $1.5 \cdot 10^{14}$ atoms per cm^3.

After the condensate was formed and expanded an argon laser beam (514 nm), traveling perpendicular to the symmetry axis of the condensate and focused on the center of the atoms, was adiabatically turned on. The position of this beam could be controlled using the scanning beam deflector described in Section 7.2. The beam had a $2w_0$ Gaussian beam waist diameter of 13 μm and a total power of 0.4 mW, resulting in an optical potential height of 700 nK (compared to the 110 nK chemical potential of the condensate). The focused beam repelled atoms, creating a cylin-

drical hole in the cloud of trapped atoms (see Fig. 10.1). The argon beam was then rastered back and forth along the axis of the condensate at a constant velocity by applying a triangle waveform to the beam deflector. After stirring the condensate for 900 ms the laser was turned off, and the condensate was allowed to equilibrate for 100 ms. The magnetic trap was then turned off, and the condensate was absorptively probed after a 35 ms time-of-flight. The temperature of the expanding atoms was evaluated by comparing the number of atoms in the rapidly expanding cloud of non-condensate atoms to the number of atoms remaining in the slowly expanding condensate remnant.

The experiment was repeated at many different raster velocities by scanning the laser over a different distance (from 0 to 60 μm) at a constant frequency. To assure that the measured heating was not due to resonant collective excitations of the condensate, the experiment was done at three different scan frequencies: 56 Hz, 83 Hz, and 167 Hz (resulting in linear scans which lasted 9, 6, or 3 ms, respectively, before the beam changed direction). The results of this experiment, shown in Fig. 10.2, are striking. It was found that if the condensate was stirred below a given velocity, the amount of heating was lower than the noise floor of our measurements. Above this velocity, the heating increased steadily from zero with increasing stirring velocity. The fact that the data taken at different scan rates all have the same dependence on velocity indicates that the onset of dissipation is *not* determined by discrete resonances of the trapping potential, but by the bulk properties of the condensate. This measurement is the first evidence of the existence of a critical velocity for dissipative flow in dilute gas Bose condensates.

Several "control" experiments were run by applying the same "stirring" process to clouds of atoms above the critical temperature for Bose-Einstein condensation. These experiments were performed at two different temperatures, 530 and 800 nK (the lower of the two being $\sim$ 20 nK above the transition temperature). The experiments were performed with scan velocities up to 14 mm/s with no measurable heating. The low heating rate in thermal clouds above T_c is due to

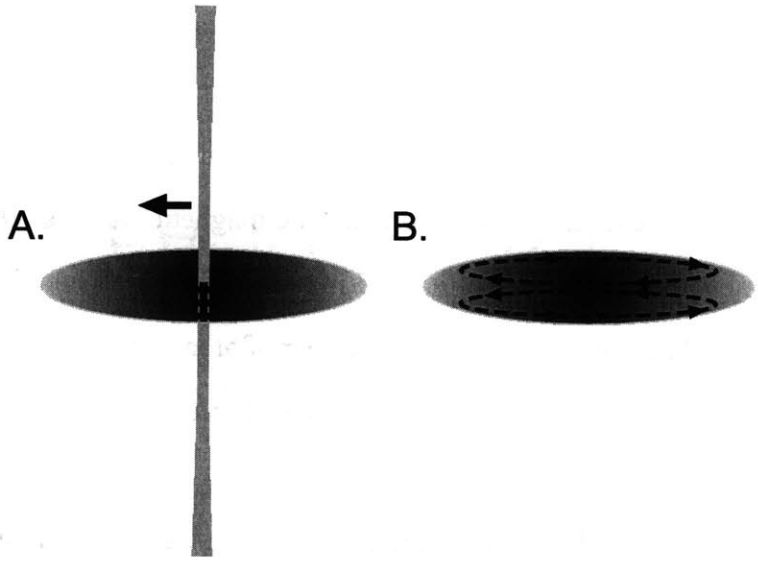

Figure 10.1: Stirring the condensate with a focused laser beam. To measure the critical velocity for dissipation, light from an argon laser was focused through the condensate, punching a hole in the density distribution (A). The beam was then scanned back and forth, forcing the condensate to flow away from its path, generating flow patterns such as the ones shown in (B).

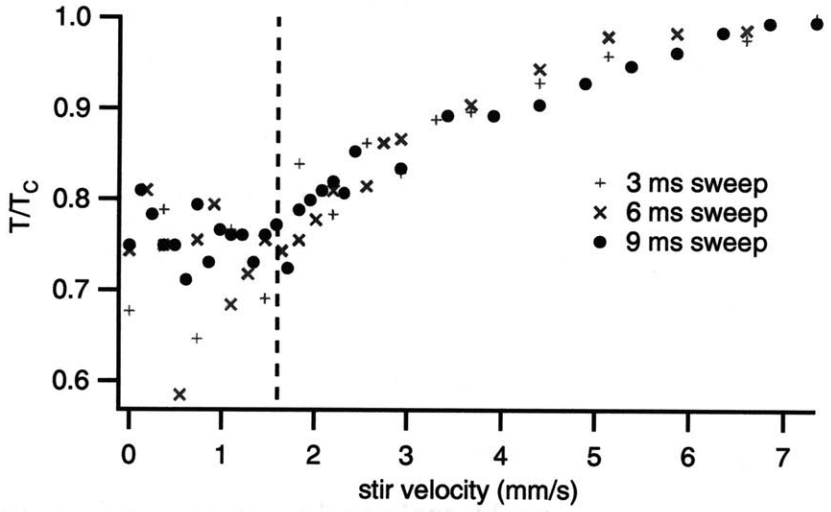

Figure 10.2: Evidence of a critical velocity for dissipation. The temperature of magnetically trapped clouds, after being stirred for 900 ms with a focused argon ion laser beam, is plotted vs. the velocity of the stirring beam. At a velocity of about 1.6 mm/s there is a sudden change in the character of the heating, going from a state in which the heating is undetectable within the noise, to a state where the heating increases with velocity. The data taken with a 9 ms raster has the least amount of scatter, but the critical point is evident in the other data sets as well.

the small overlap of the large thermal cloud with the small area over which the stirring laser is scanned. An ideal gas model was then used to calculate the heating of the dilute atoms in the thermal cloud due to collisions with a moving object, allowing the upper bound on the heating rate for clouds measured above T_c to be scaled to predict the heating due to the non-condensate atoms below T_c. The result was an upper bound on the heating due to the non-condensate atoms which could account for, at most, 15% of the observed temperature rise in the stirred condensates. This is evidence that the observed heating was due to the dynamic flow of the condensate, and not simply due to heat imparted to non-condensate atoms.

The measured critical velocity is consistent with zero-temperature theoretical studies. Numerical studies in [135] predicted the critical velocity for dissipation in the case of a homogeneous, weakly interacting fluid. In this study, the flow around a disk moving through an otherwise stationary fluid was considered. As pointed out in this reference, dissipation begins when any part of the fluid exceeds the velocity of sound. At this point, since superfluids cannot support a shock wave, a wake of vortices are emitted. The circulation of these vortices forms a "bridge" between parts of the fluid with relative speeds in excess of the sound velocity. It is also pointed out in [135] that the flow of an incompressible fluid around a disk has a maximum velocity (with respect to the rest frame) at the points where the disk's cross-section is widest. At these points, the fluid flows at twice the velocity at which the disk is dragged through the fluid. As a result, one would expect that the critical velocity at which the disk would cause dissipation would be at 1/2 the speed of sound. A value for the critical velocity of a disk dragged through a weakly interacting *compressible* fluid was derived by finding the velocity above which a steady state solution to the NLSE does not exist, resulting in a critical velocity of:

$$v_c = \sqrt{\frac{2}{11}}\, c_s \tag{10.4}$$

where c_s is the speed of Bogoliubov sound in the unperturbed fluid. This relation was then verified with numerical simulations.

Although derived for a homogeneous gas in the thermodynamic limit, Eq. (10.4) is in reasonably good agreement with the results of our experiment. The speed of sound in the center of the condensate, measured from a time-of-flight expansion, was found to be $c_0 = 6.2$ mm/s. With this speed of sound, Eq. (10.4) predicts a critical velocity of 2.6 mm/s, just 63% more than our measured value of 1.6 mm/s. There are several reasons why we might expect that the critical velocity would be slightly lower in our experiment, including the finite size of the condensate and the inhomogeneous condensate density distribution.

It is interesting to note that, in the numerical study referred to above, dissipation above the critical velocity was seen to take the form of vortex formation, the energy of the vortices eventually being "radiated" into phonons. The fact that we see a critical velocity which agrees with these simulations, in combination with the association of vortices with dissipation in liquid helium, and the known mechanism for vortex mediated dissipation in dilute gases, as discussed in Section 10.2, raises the question as to whether the critical velocity measured in this experiment is an indication of vortex nucleation. To further pursue this possibility, the effects of the finite size and inhomogeneous density of the condensate on the critical velocity must be carefully examined. There are some indications that, even in inhomogeneous fluids, vortex formation is governed completely by the *local* sound velocity, in which case Eq. (10.4) tells the whole story (this point of view is supported by the numerical results in [111], in which it was found that the size of a vortex core depends primarily on the local density with only a weak dependence on the geometry of the condensate). On the other hand, one can relate the critical velocity to the critical angular rotation for vortex formation discussed in [113, 136], resulting in a critical velocity which depends on the size of the condensate as well as the peak speed of sound. It is possible that the correct picture will only be revealed by a generalization to the inhomogeneous case of studies such as the one described in [135].

Appendix A

Vacuum Conduction Analysis

The use of a quartz cell in our experiment makes it impossible to install a vacuum gauge in the vicinity of the trapped atoms. We have to *infer* the pressure in the trapping region from the ion gauge near the pump between slower coils and the ion gauges in the pumping chamber at the far end of the experiment. A full analysis of vacuum pumping and conduction was performed to understand how the vacuum pressures in each part of the experiment relate to one another. In this analysis, the pressure at five points of interest were considered (see Fig. A.1): the oven chamber, the slower ion pump, the UHV pumping chamber, the tee between the slower coils, and the center of the quartz cell. Of these five points, the first three represent the places at which ion gauges are present, allowing us to directly measure the vacuum pressure. The fourth point is important because it is a junction between three parts of the vacuum system. The last point, the center of the quartz cell, is important since this is where the atoms are magnetically trapped.

In order to do this analysis, I made the assumption that the oven chamber was the source of all gases in the system. This assumption simplifies matters a great deal; it would be impossible to consider the out-gassing of each piece of the vacuum assembly correctly and independently. This assumption is known to be approximately correct due to the fact that the pressure measured at the

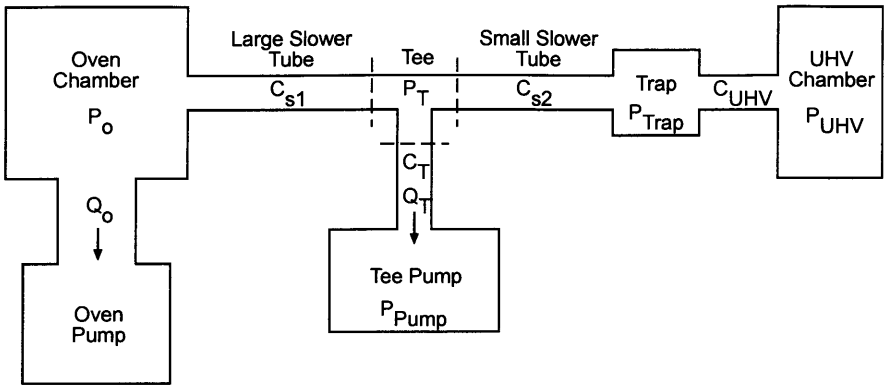

Figure A.1: Block diagram of the vacuum system. Pressures are denoted by the letter P , conductance by the letter C , and pump rates by the letter Q. For a drawing of the vacuum system which is to scale, see Fig. 3.1.

pump between slower coils is a factor of three to four higher when the gate valve between the oven and the slower is open as compared to the reading when the gate valve is closed.

With this assumption it is easy to put together a system of coupled equations which can be solved for the vacuum pressure at each important point in the experiment. These coupled equations can be simply written down by considering the mass flow in and out of each part of the chamber; in equilibrium the total rate of gas flow into any part of the chamber minus the flow out must be equal to zero. Using the variables defined in Fig. A.1, we arrive at the following set of coupled equations:

$$Q_o P_o = (P_T - P_o)C_{s1} + R_{source}P_o \tag{A.1}$$

$$0 = (P_o - P_T)C_{s1} + (P_{Pump} - P_T)C_T + (P_{Trap} - P_T)C_{s2} \tag{A.2}$$

$$Q_T P_{Pump} = (P_T - P_{Pump})C_T \tag{A.3}$$

$$0 = (P_T - P_{Trap})C_{s2} + (P_{UHV} - P_{Trap})C_{UHV} \tag{A.4}$$

The known quantities in this analysis are the pumping rates Q_o and Q_T, the pressures P_{Pump},

Table A.1: The known parameters in the vacuum evaluation. Pumping rates are taken from the specifications given by the pump's manufacturer, conductances are calculated from the known geometry of the vacuum chamber, and pressures represent typical pressures measured by ion gauges at three points in the vacuum. Only an upper limit on P_{UHV} can be given, since the pressure in the UHV chamber is lower than the x-ray limit of the ion gauge. The conductance C_{s1} is mainly determined by the conductance of the differential pumping tube (see Fig. 3.3), and is only weakly determined by the bore of the large slower coil. The effective value of C_{s1} can be reduced by a small amount by cooling the differential pumping tube between the oven and the slower with liquid nitrogen, causing some gases to "stick" to its inner surface rather than bounce through.

Parameter	Value
Q_o	60 l/s
Q_T	60 l/s
P_{Pump}	$2 \cdot 10^{-10}$ Torr
P_o	$6 \cdot 10^{-8}$ Torr
P_{UHV}	$< 5 \cdot 10^{-12}$ Torr
C_{s1}	0.14 l/s
C_{s2}	2.2 l/s
C_T	46 l/s
C_{UHV}	14 l/s

P_o, and P_{UHV}, and the conductance of each part of the chamber, C_{s1}, C_{s2}, C_T, and C_{UHV}. The conductance and pump rates, as well as typical values for the known pressures are summarized in Table A.1. The unknowns are the pressures at the tee and in the center of the glass cell, P_T and P_{Trap} , and the rate at which residual gases enter the oven chamber R_{source}.

From the last two equations of the set, we can find the pressure in the center of the glass cell as a function of known parameters:

$$P_{Trap} = \frac{1}{C_{s2} + C_{UHV}} \left(\frac{Q_T + C_T}{C_T} C_{s2} P_{Pump} + P_{UHV} C_{UHV} \right) \tag{A.5}$$

Since we don't know P_{UHV} precisely, due to the limit of our ion gauge, a worse-case estimate is obtained by assuming that P_{UHV} is at the upper possible limit of $5 \cdot 10^{-12}$ Torr. In this case, even with the lowest pressure measurements which we have made for P_{Pump}, the P_{Pump} term in A.5 is

still more than an order of magnitude higher than the P_{UHV} term. As a result, for the rest of this analysis I will simplify matters by assuming that $P_{UHV} = 0$. This gives a simplified relation for the pressure in the trap:

$$P_{Trap} = \frac{Q_T + C_T}{C_{s2} + C_{UHV}} \frac{C_{s2}}{C_T} P_{Pump} \tag{A.6}$$

Plugging in the values in Table A.1, we find that the pressure in the trap is about three times lower than the pressure measured in the slower ion pump. It is still useful, however, to find the trap vacuum in terms of just the pressure in the oven. This gives us an idea of how low the oven pressure must be in order to run the experiment, and allows us to test the consistency of this analysis. Using the last three equations of the set, we find that:

$$P_{Trap} = \frac{C_{s1}C_{s2}(Q_T + C_T)}{(Q_T(C_{s1} + C_{s2} + C_T) + C_T(C_{s1} + C_{s2}))(C_{s2} + C_{UHV}) - C_{s2}^2(Q_T + C_T)} P_o \tag{A.7}$$

By putting in the known conductances and the known pump rate, we find that the pressure in the trap is about 1500 times lower than the pressure in the oven. Given the values in Table A.1, it is obvious that many of the terms in A.7 can safely be neglected, giving us a simpler equation to remember:

$$P_{Trap} \simeq \frac{C_{s1}C_{s2}}{C_{s2} + C_{UHV}} \frac{Q_T + C_T}{Q_T C_T} P_o \tag{A.8}$$

Using Eqs. (A.6) and (A.7) and the measured pressure in the oven and in the slower ion pump, we get two estimates of the pressure in the trap. Both estimates indicate that the pressure in the cell is a few times 10^{-11} Torr.

Appendix B

Recirculating Ovens

Only a small fraction of the sodium leaving the oven travels into the small solid angle defined by the bore of the slower. The rest of the sodium ends up trapped in the oven vacuum chamber or in the slower. In the past the oven has been typically loaded with 10 grams of sodium, allowing us to operate the experiment for only about 120 hours, after which the sodium must be renewed. Reloading sodium is a process which takes several hours. Furthermore, after a sodium change the experiment must often be shut down for one or two days while the oven is baked to restore acceptable vacuum levels. This is compounded by the fact that the sodium in our oven has been known to run out at inopportune times, often abruptly ending an experimental run. And since very little of the sodium actually ends up in the thermal beam used to load the trap, the bulk of the sodium sticks to the walls of the oven vacuum chamber, where it can clog the collimation hole in the cold plate and act as a vacuum contaminant. For these reasons we explored the possibility of using a recirculating oven which would recapture most of the uncollimated sodium and return it to the oven for reuse.

Several designs were tried, including two "candlestick" designs based on the oven described in [74], and one much simpler design similar to the oven described in [137]. All of these designs use

the surface tension of sodium in a porous material to "wick" the uncollimated sodium back into the oven. In the simplest design the oven is connected to the experiment through a narrow tube lined with several layers of stainless steel mesh. Heat conducted from the oven generates a temperature gradient along the length of the tube. Atoms which leave the oven at large angles will strike the porous material lining the tube and will be wicked back into the oven. This design suffers from the fact that the wick is directly connected to the hot oven. As a result, liquid sodium present in the tube near the oven is heated to a temperature comparable to the oven's temperature, causing it to source a vapor pressure comparable to the vapor pressure in the oven. This effect limits the flux of atoms which can be generated with this type of oven. As the temperature of the oven is increased, the vapor pressure in the collimation tube increases as well, presenting a collisional barrier for atoms leaving the oven. After a point, increasing the oven temperature does not increase the flux from the oven. It simply causes the effective source of the atoms to move down the tube to a cooler point where the mean free path is longer. After a few short experimental tests and calculations, it was determined that this type of oven could not produce the flux which we desired, and this design was abandoned in favor of the candlestick design.

The design of our final candlestick oven is shown in Fig. B.1. In this design, sodium is wicked up a "candlestick" from a pool of sodium at the bottom of the oven by a piece of porous material running through the center of the candlestick. At the top of the candlestick is a small chamber which is heated to $\sim$ 300° C to produce a dense sodium vapor. Sodium atoms stream out from a small hole in this chamber. Atoms going in the correct direction continue through a somewhat larger hole in the side of the oven and form a collimated atomic beam. Atoms which miss the hole in the side of the oven strike a lining of porous material and are wicked back to the pool of sodium at the bottom of the oven. In the candlestick design, the lining of porous material can be kept at a temperature just above the melting point of sodium, such that the vapor pressure in the oven is relatively low.

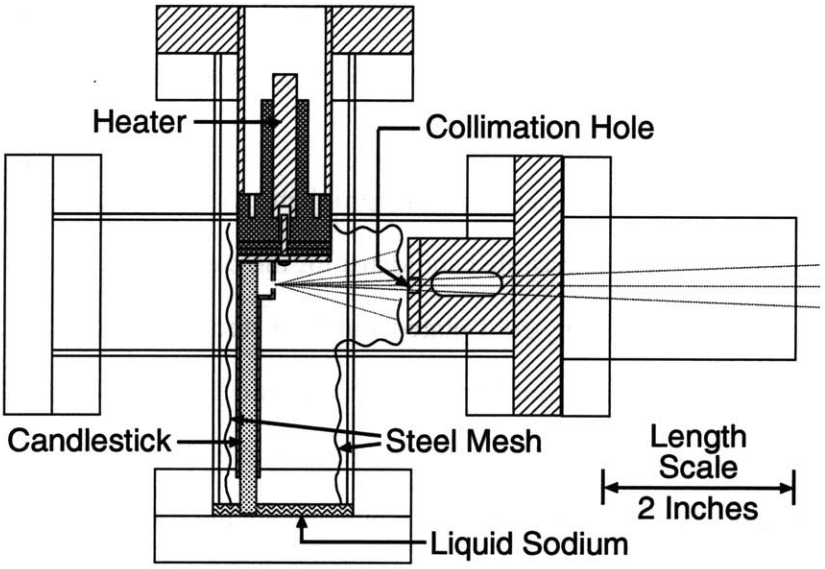

Figure B.1: Candlestick recirculating oven. The entire conflat cross which contains the oven is heated just above the melting point of sodium, producing a pool of liquid sodium at the bottom of the cross. Sodium from this pool is wicked up the candlestick by the capillary action of a tightly rolled piece of steel mesh. In the small heated chamber at the top of the candlestick a thick vapor is formed. Atoms escape from this chamber through a small hole. Atoms leaving the candlestick at very shallow angles pass through the collimation hole and form an atomic beam. Atoms exiting at larger angles strike the steel mesh surrounding the collimation hole and are wicked back to the sodium pool at the bottom of the oven.

One additional advantage of the candlestick design is that only a small part of the oven is heated to high temperatures, thereby reducing the amount of background gas emitted from the oven (although the large surface area of the mesh requires more careful cleaning and vacuum bake-out procedures to achieve this improved vacuum). In addition, since the sodium vapor pressure is low outside of the candlestick, the plate with the collimation hole can be set off from the vacuum chamber, forming a baffle which obscures the direct line of sight out of the oven from the hole in the candlestick, but which allows background gas to be pumped out of the oven with high conductance.

In the end, the recirculating oven was limited by problems with the wicking material used. Several porous materials were tried as possible candidates for the candlestick's "wick." This included tight weave stainless steel wire cloth, gold plated stainless steel wire cloth, and porous sintered tungsten. We observed that sodium would only wet to the gold plated stainless steel cloth. Unlike other alkali metals which have been reported to wick directly to stainless steel, it seems that sodium cannot easily penetrate or wet to the oxide layer which covers the surface of the steel. In the gold plating process, the oxide is removed, and a layer of gold prevents the oxide from returning. When exposed to liquid sodium, the much milder gold is dissolved, and the sodium is free to wick to the clean steel surface. If too much gold is present, however, gold-sodium alloys with high melting points will form. We found that the gold plating done by a typical plating company either produced samples which had regions of uncoated steel, or regions where the gold was too thick. The best samples that we tried *did* successfully wick sodium and produce a beam. However, the beam flux was about a factor of 3 lower than the flux expected from a conventional oven running at the same temperature.

For the time being, work on the recirculating oven has stopped. In its place, a cold plate assembly which is capable of sustaining greater sodium build up was implemented. In addition, a conflat nipple was added at the bottom of the oven elbow to simply allow larger amounts of sodium to be loaded into our conventional oven at once.

Appendix C

Magnetic Trapping and Evaporative Cooling - Varenna Summer School 1998

The following appendix is a selection reprinted from our contribution to the proceedings of the 1998 Varenna summer school on Bose-Einstein condensation [47]. It has been modified from the original only in the sense that cross references, citations, and section headings have been renumbered and clarified to be consistent with the rest of this thesis.

C.1 Conservative Atom Traps

Conservative atom traps fulfill two essential roles in BEC experiments: they keep the atoms tightly compressed during cooling, and hold the condensate for study. In principle, any trap that has a sufficiently small heating rate could be used. Conservative trapping potentials have been realized with dc magnetic fields (sect. C.2), ac magnetic fields [138], microwave fields [139] and far-off-resonant laser beams (sect. 8 of [47]).

The requirements for the trap during cooling are more stringent than they are for holding condensates. First, the time for cooling (typically thirty seconds for evaporative cooling) is usually much longer than the time for performing experiments on BEC, requiring low heating and trap loss rates. Furthermore, for cooling, the trap needs a sufficiently high trap depth and trapping volume to hold the initial (pre-cooled) cloud, and must accommodate a cooling scheme able to reach BEC temperatures. So far, only the combination of magnetic trapping (sect. C.2) and evaporative cooling (sect. C.3) has accomplished this. Evaporative cooling has also been observed in an optical dipole trap [93], but with only a factor of 30 gain in phase-space density. Rf-induced evaporation

is particularly effective and simple to implement in a magnetic trap, and perhaps this combination will become the workhorse of the nanokelvin temperature range, just as the MOT has in the microkelvin range. However, there is always room for improvement, in particular for atoms which don't have the collisional properties necessary for evaporative cooling in a magnetic trap (sect. 2.6 of [47]). Considerable progress is being made on far-off resonant trapping using blue-detuned [140, 141, 142], near-infrared [143], and CO_2 lasers [144, 145].

After cooling, trap requirements are different, and other options are available for holding the condensate. Our recent experiments on optically confined BEC (sect. 8 of [47]) demonstrate that optical dipole traps confine Bose-Einstein condensates more easily than much hotter atoms. Traps for condensates can be much weaker than for laser cooled atoms, making them easier to implement and greatly reducing heating due to beam jitter, intensity fluctuations, and spontaneous emission [146].

There are several traps which have so far not been pursued beyond their first demonstration: A microwave trap originally suggested for hydrogen atoms [147] has been realized with laser-cooled cesium atoms [139]. An ac magnetic trap for strong-field seeking atoms was suggested in 1985 [148], and realized in 1991 [138]. Finally, ac electric fields offer another possibility to trap strong-field seekers [149, 150]. All these traps are rather weak and don't seem to offer obvious advantages over the combination of optical and magnetic forces. A recent resource letter contains many references on atom traps [151].

C.2 Magnetic Trapping

The major role of the magnetic trap in a BEC experiment is to accommodate the pre-cooled atoms and compress them in order to achieve high collision rates and efficient evaporative cooling. The steps that must be taken to obtain high collision rates, including "mode-matched" transfer and anisotropic compression, are discussed in this section. The collision rate after compression is suggested as the most important figure of merit for a magnetic trap, and will be related to the magnetic-field parameters.

Magnetic trapping of neutral atoms was first observed in 1985 [152]. Shortly afterwards, orders of magnitude improvements in density and number of trapped atoms were achieved at MIT and in Amsterdam using superconducting traps and different loading schemes [84, 153, 154]. Important aspects of magnetic trapping are discussed in [155, 82, 44, 156, 92].

Magnetic forces are strong for atoms with an unpaired electron, such as the alkalis, resulting in magnetic moments μ_m of the order of a Bohr magneton μ_B. However, it is worth pointing out that magnetic confinement was first observed for neutrons, despite their thousand-times smaller magnetic moment [157].

The interaction of a magnetic dipole with an external magnetic field is given by $-\boldsymbol{\mu} \cdot \mathbf{B} = -\mu_m B \cos\theta$. Classically, the angle θ between the magnetic moment and the magnetic field is constant due to the rapid precession of $\boldsymbol{\mu}_m$ around the magnetic-field axis. Quantum-mechanically, the energy levels in a magnetic field are $E(m_F) = g\mu_B m_F B$, where g is the g-factor and m_F the

quantum number of the z-component of the angular momentum F. The classical term $\cos\theta$ is now replaced by m_F/F; the classical picture of constant θ is equivalent to the system remaining in one m_F quantum state.

An atom trap requires a local minimum of the magnetic potential energy $E(m_F)$. For $gm_F > 0$ (weak field seeking states) this requires a local magnetic-field minimum. Strong field seeking states ($gm_F < 0$) cannot be trapped by static magnetic fields, because Maxwell's equations don't allow a magnetic-field maximum in free space [79, 80].

Because magnetic traps only confine weak field seeking states, atoms will be lost from the trap if they make a transition into a strong field seeking state. Such transitions can be induced by the motion in the trap because an atom sees a field in its moving frame which is changing in magnitude and direction. The trap is only stable if the atom's magnetic moment adiabatically follows the direction of the magnetic field. This requires that the rate of change of the field's direction θ must be slower than the precession of the magnetic moment:

$$\frac{d\theta}{dt} < \frac{\mu_m |\mathbf{B}|}{\hbar F} = \omega_{Larmor} \tag{C.1}$$

The upper bound for $d\theta/dt$ in a magnetic trap is the trapping frequency. This adiabatic condition is violated in regions of very small magnetic fields, creating a region of trap loss due to spin-flips to untrapped states. These spin-flips are referred to as "Majorana flops" [81].

C.2.1 Quadrupole-type Traps

There are two basic types of static magnetic traps: those in which the minimum is a zero crossing of the magnetic field, and those which have a minimum around a finite field [82]. Traps with a zero-field crossing usually generate a linear potential characterized by the gradient of the magnetic field: $B_x = B'_x x$, $B_y = B'_y y$, $B_z = B'_z z$. Maxwell's equations require that $B'_x + B'_y + B'_z = 0$. The case of axial symmetry is a spherical quadrupole field in which $B' \equiv B'_x = B'_y = -B'_z/2$. This configuration is realized with two coils in "anti-Helmholtz" configuration. This was the configuration which was first used to trap neutral atoms magnetically [152].

A linear trapping potential offers superior confinement compared to traps with a parabolic potential minimum (sect. C.2.2). This follows from a simple argument. Coils which are a distance R_{coil} away from the trapped cloud and generate a field B_{coil} at the coil typically produce a field gradient of $B' \approx B_{coil}/R_{coil}$ and a curvature of $B'' \approx B_{coil}/R^2_{coil}$. A cloud of size r in a linear potential with gradient B' would be confined to the same size in a parabolic potential with a curvature equal to B'/r. This allows us to define an "effective curvature" of linear confinement: $B''_{eff} = B'_{coil}/r$. This exceeds the curvature of a parabolic trap by R_{coil}/r, which is usually an order of magnitude or more.

When linear traps were employed for the first demonstrations of evaporative cooling with alkali atoms [158, 159], trap loss due to Majorana spin-flips [81, 152, 160, 155, 161] near the zero of the magnetic field was encountered. For atoms moving at a velocity v, the effective size of this "hole"

in the trap is $\sqrt{2\hbar v/\pi\mu_m B'}$ which is about 1 μm for $\mu_m = \mu_B$, $v = 1$ m/s and $B' = 1000$ G/cm. As long as the hole is small compared to the cloud diameter, the trapping time can be long (even longer than a minute), and evaporative cooling in such a trap was used to increase phase-space density by more than two orders of magnitude [162]. However, as the temperature drops, the trap loss due to the hole becomes prohibitive for further cooling. Although the size of the hole depends on the thermal velocity of the atoms, and therefore shrinks as the atoms are cooled, the diameter of the atom cloud shrinks even faster with temperature, resulting in a T^{-2} dependence of the loss rate [162, 163].

Two methods have been demonstrated to plug the hole. One solution is to add a rotating magnetic bias field B_0 to the spherical quadrupole field. The frequency of the rotating field is much higher than the orbiting frequency of the atoms, but much lower than the Larmor frequency. The resulting time-averaged, orbiting potential (TOP) trap is harmonic, but much tighter than what could be obtained by dc magnets of the same size [163]. The time-averaged potential can be written as

$$U_{TOP} = \frac{\mu_m}{2}\left(B''_\rho \rho^2 + B''_z z^2\right) \tag{C.2}$$

$$B''_\rho = \frac{B'^2}{2B_0}, B''_z = \frac{4B'^2}{B_0} \tag{C.3}$$

where $\rho^2 = x^2 + y^2$ is the radial coordinate.

The rotating field moves the zero of the magnetic field around in a circle (the "circle of death") of radius $r_D = B_0/B'$. Due to Majorana flops, this limits the depth of the potential to $U_{TOP}(r_D) = \mu_m B_0/4$. A large circle of death requires either large B_0 or small B', both of which lead to weaker confinement. The TOP trap was used in the first demonstration of BEC [1]. An interesting variant of the TOP trap has been proposed, where a rotating quadrupole field avoids the circle of death [164].

Another solution is to plug the hole using the optical dipole forces of a tightly focused blue-detuned laser beam to repel atoms from the center of the trap (see also sect. 2.5.2 of [47]) [2]. The "optically plugged trap" achieves very tight confinement corresponding to a curvature of about B'/x_0, where x_0, the separation of the potential minimum from the zero of the magnetic field, is typically $\sim 50\mu$m. BEC in sodium was first achieved with this trap [2].

C.2.2 Ioffe-Pritchard Traps

The lowest order (and therefore tightest) trap which can have a bias field is a harmonic trap. A magnetic trap with finite bias field along the z-direction has an axial field of $B_z = B_0 + B''z^2/2$. The leading term of the transverse field component B_x is linear, $B_x = B'x$. Applying Maxwell's

equations (and assuming axial symmetry) leads to the following field configuration [82]:

$$\mathbf{B} = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2+y^2) \end{pmatrix} \tag{C.4}$$

The parabolic trap was first suggested and demonstrated for atom trapping by Pritchard [83, 84], and is similar to the Ioffe configuration discussed earlier for plasma confinement [85]. We refer to any trap which has this field configuration as a Ioffe-Pritchard (IP) trap.

The Ioffe-Pritchard trap has two different regimes. For temperatures $k_B T < \mu_m B_0$, the cloud experiences the potential of a 3D anisotropic harmonic oscillator. In the case of $k_B T > \mu_m B_0$, the potential is predominantly linear along the radial directions ($U_\rho = \mu_m B' \rho$) and harmonic along the axial direction ($U_z = \mu_m B'' z^2/2$). As a consequence, the loss in confinement compared to a linear trap is not as severe as implied in section C.2.1. For small clouds (and all condensates) the trapping potential is very well approximated by an anisotropic harmonic oscillator potential

$$U \simeq \tfrac{\mu_m}{2} \left[B''_\rho \rho^2 + B'' z^2 \right] \tag{C.5}$$

$$B''_\rho = \tfrac{B'^2}{B_0} - \tfrac{B''}{2} \tag{C.6}$$

Beyond a certain axial displacement, the radial confinement vanishes. This occurs because the radial component of the curvature term $-B''/2\, xz$ interferes destructively with the radial gradient B'. From eq. (C.4) we find that this point of instability z_{inst} occurs at:

$$z_{inst} = \pm \left(\frac{B'}{B''} - \frac{1}{2} \frac{B_0}{B'} \right) \tag{C.7}$$

When $B''_\rho \leq 0$, the instability is at the origin, and there is no radial confinement at all. The saddle point at z_{inst} requires special attention when loading large clouds into an Ioffe-Pritchard trap, as discussed in section C.2.3.

The Ioffe-Pritchard trap has been used in many BEC experiments. The most straightforward implementation consists of two pinch coils and four Ioffe bars [82]. The Ioffe bars generate the radial gradient field B', while the pinch coils produce a bias field and the curvature term B''. Most of the bias field is usually canceled by a pair of additional "anti-bias" coils. By lowering B_0, radial confinement is increased (eq. (C.6)). The bias field should be just high enough to suppress Majorana flops. A typical value is 1 G, although we have used values as low as 0.4 G, and smaller values are probably still stable. If the bias field is over-compensated, the field will cross zero and Majorana flops will occur (fig. C.1). Since both the pinch coils and the anti-bias coils can produce fields on the order of 500 to 1000 G, the cancellation has to be carefully controlled. We achieved best results when both sets of coils were powered in series by the same power supply, and fine bias adjustments were made by carefully moving a few loops of the anti-bias coils.

Since the pinch coils and the Ioffe bars are very close to the atoms, this configuration is very

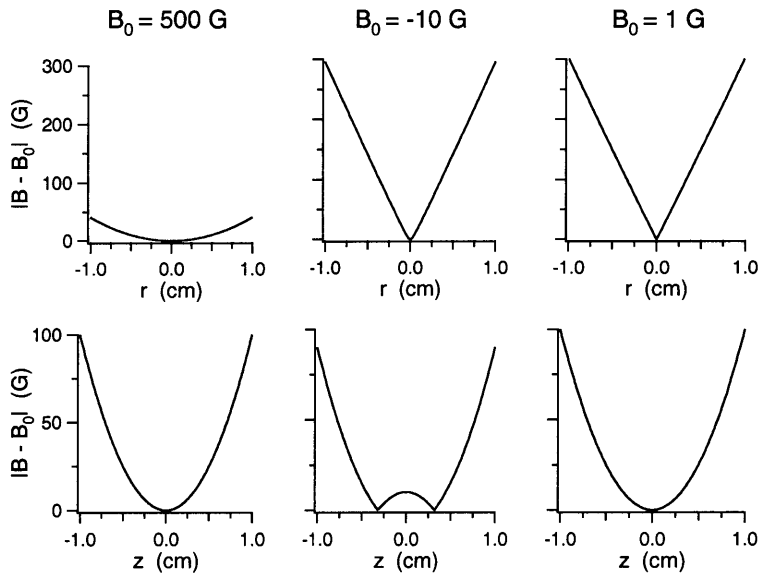

Figure C.1: Bias field compensation in an Ioffe-Pritchard trap. Bias field compensation in an Ioffe-Pritchard trap is important for tight radial confinement. The magnetic field in an IP trap characterized by a radial gradient of 300 G/cm and an axial curvature of 200 G/cm^2 is shown for three bias fields B_0. The upper row displays radial cuts, and the lower row displays axial cuts of the magnetic-field profile. In the first column, radial confinement is softened as a result of the large bias field. In the second column, the bias field is over-compensated, resulting in a pair of zero-field crossings along the axis of the trap. In the third column, the bias field is tuned correctly, resulting in tight radial confinement and no zero field crossings.

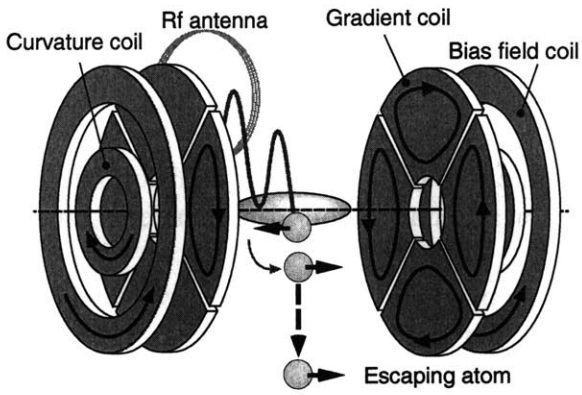

Figure C.2: Cloverleaf trap. In a cloverleaf trap, Ioffe bars are replaced by eight "cloverleaf" coils surrounding the pinch coils, providing 360 degree optical access. Evaporation is done by selectively spin-flipping atoms into untrapped states with rf radiation.

efficient in producing a tight trapping potential. It has been used in several BEC experiments [90, 43]. The optical access can be improved by elongating the pinch coils, or using a cloverleaf configuration [53] (fig. C.2) where the Ioffe bars are removed, and the radial field gradient is created by a pair of four coils surrounding each pinch coil. This configuration is used in several experiments because it has open, 360° access in the x-y plane, and the coils may be placed outside the vacuum chamber. Some simpler winding patterns (baseball [86, 82], yin-yang [82], three-coil [87, 88], four-dee [89]) don't allow for independent control of B' and B'', although the radial confinement can still be varied through B_0. Other implementations of the Ioffe-Pritchard trap include permanent magnet traps [165, 3], and traps with ferromagnetic pole pieces [166].

C.2.3 Mode Matching

Evaporative cooling can increase phase-space density by many orders of magnitude as long as the atoms rethermalize quickly compared to the trap lifetime. This makes the collision rate during cooling even more important for BEC than the initial phase-space density of the pre-cooled atoms. In order to maximize the collision rate, atoms are transferred into an optimized ("mode matched") magnetic trap, and then adiabatically compressed.

"Mode matching" between an atom cloud and a trap is achieved when the transfer of atoms maximizes phase-space density. This also optimizes the elastic collision rate which will be achieved after compressing the trap (sect. C.2.4), because all properties of the compressed cloud, including the collision rate, are completely determined by the number of atoms N and the phase-space density $\mathcal{D}$. Adiabatic compression conserves N and $\mathcal{D}$, so a loss in phase-space density when loading the magnetic trap corresponds to an equal loss in the compressed trap. For a power law potential $U(r) \propto r^{d/\delta}$ in d dimensions, the volume of a cloud at temperature T scales as $V \propto T^{\delta}$. The

collision rate Γ_{el} for atoms with a collisional cross section σ and a thermal velocity v is equal to:

$$\Gamma_{el} = n\sigma v \propto \mathcal{D}^{\frac{\delta-1/2}{\delta+3/2}} N^{\frac{2}{\delta+3/2}} \tag{C.8}$$

In particular, for a 3D harmonic trap ($\delta = 3/2$), the collision rate is proportional to $\mathcal{D}^{1/3}$. This implies that the loss in collision rate due to non-ideal transfer is much less severe than the loss in phase-space density.

Atoms are transferred into a magnetic trap by suddenly switching off the MOT, applying polarization gradient cooling for a few milliseconds, and then suddenly turning on the magnetic trap and building a new potential around the atoms. When transferring a spherical, Gaussian shaped cloud of atoms with an rms radius r_0 and temperature T into a harmonic trap, phase-space density is conserved if the potential $\frac{1}{2}\kappa r^2$ has a stiffness $\kappa = \kappa_0 = k_B T/r_0^2$. This will ensure that the atoms maintain their volume and temperature. If the trap is too tight, the atoms will be heated by the transfer. If it is too loose, the atoms will expand non-adiabatically. In either case phase-space density will be lost. For ideal loading we have used the term "mode matching," in analogy with the coupling of a laser beam into a single mode fiber, where the efficiency suffers from focusing which is either too weak or too tight.

If the magnetic trap is not mode-matched, i.e. $\kappa \neq \kappa_0$, then the phase-space density $\mathcal{D}$ is reduced from its initial value $\mathcal{D}_0$ [52]:

$$\frac{\mathcal{D}}{\mathcal{D}_0} = \frac{8\left(\frac{\kappa}{\kappa_0}\right)^{3/2}}{\left(1+\frac{\kappa}{\kappa_0}\right)^3} \tag{C.9}$$

This loss is only a weak function of mode mismatch. For example, the phase-space density decreases only by half when κ is four times larger or smaller than κ_0. More importantly, the collision rate is still 80% of its maximum value.

Mode matching into a Ioffe-Pritchard trap is done with a high bias field ($\mu_m B_0 > k_B T$). This requires (eq. (C.6)):

$$B'' = \frac{\kappa_0}{\mu_m} \tag{C.10}$$

$$\frac{B'^2}{B_0} = \frac{3\kappa_0}{2\mu_m} \tag{C.11}$$

It is important that the initial cloud radius r_0 be much smaller than the distance to the trap instabilities (eq. (C.7)), otherwise, atoms will spill out of the trap. Inserting eqs. (C.10) and (C.11) into eq. (C.7) gives an instability distance of $z_{inst} = \sqrt{B_0(2\mu_m/3\kappa_0)} = B'(2\mu_m/3\kappa_0)$. This shows that both a large bias field and a large gradient are necessary, which is unfortunate because this requires a fast rise time of the high-current power supplies.

For $\mu_m B_0 < k_B T$, the radial confinement is linear, and mode-matching would require $\mu_m B' \approx$

$k_B T/r_0$. With this condition we find that $z_{inst} \approx r_0$, and therefore mode matching cannot be done at low bias field.

C.2.4 Adiabatic Compression

Adiabatic compression plays a crucial role in BEC experiments because it increases the collision rate before evaporative cooling. In our first BEC demonstration [2] adiabatic compression increased the collision rate by a factor of 20, resulting in BEC after only 7 seconds of evaporative cooling!

If, after loading, the potential (again characterized by a volume $\propto T^\delta$) is adiabatically increased by a factor α, the temperature rises by a factor $\alpha^{2\delta/(2\delta+3)}$ and the density by $\alpha^{3\delta/(2\delta+3)}$. Phase-space density is constant, but the elastic collision rate increases by a factor $\alpha^{4\delta/(2\delta+3)}$ as long as the cross-section is constant.

Adiabaticity requires

$$\frac{d\omega_{trap}}{dt} \ll \omega_{trap}^2 \tag{C.12}$$

Violation of adiabaticity, however, doesn't have severe consequences. If the compression is done suddenly in a harmonic trap ($\delta = 3/2$), the collision rate increases by a factor $2\alpha^{3/2}/(1+\alpha)$ rather than α. For a sudden compression by a factor of five, one still reaches 75% of the maximum possible collision rate. Note that this is the same loss we would see if we transferred atoms into a trap which was mode-mismatched by a factor α (eq. (C.9)). Non-mode matched transfer is equivalent to a mode-matched transfer plus a sudden (de-)compression to the same potential strength.

Compression in an Ioffe-Pritchard trap involves one more complication. Initially, the aspect ratio is adjusted to ≈ 1 to ensure mode-matched loading of a spherical cloud. This requires a high bias field B_0. Subsequent compression includes lowering the bias field to achieve the maximum radial confinement. This process elongates the cloud into a cigar shape. In addition to the usual adiabaticity criterion (C.12), the compression has to be slow compared to the elastic collision rate. Otherwise, the anisotropic compression would lead to an anisotropic temperature and subsequent loss of phase-space density during equilibration.

Anisotropic heating is hard to avoid when the initial collision rate is low. Fortunately, nature is forgiving. Radial compression by a factor of α causes the temperature to rise by a factor of $\alpha^{1/3}$ if done adiabatically, and by a factor of $(2\sqrt{\alpha}+1)/3$ (after rethermalization) if done suddenly. Since the elastic collision rate is inversely proportional to temperature in a given 3D harmonic oscillator potential, one still obtains 94% of the maximum possible collision rate when a 2D compression by a factor of $\alpha = 5$ is done quickly with respect to the thermalization rate.

Furthermore, the radial compression in a Ioffe-Pritchard trap usually takes us to the small bias field regime ($B_0 < k_B T/\mu_m$), changing the radial confinement from harmonic to linear. As pointed out by Pinkse *et.al.* [167] and discussed in sect. 8.2 of [47], this increases the phase-space density by a factor of e. This increase is due to collisions which change the population of energy levels, but conserve entropy.

When the bias field B_0 is lowered adiabatically, the collision rate Γ increases by

$$\frac{\Gamma_{final}}{\Gamma_{initial}} = \frac{e^{1/2}}{2}\left[\frac{\mu_m B_{0initial}}{k_B T_{initial}}\right]^{1/2} \tag{C.13}$$

We can apply estimates for adiabatic compression in a harmonic trap to linear potentials by noting that B' confines a cloud at the same temperature and density as a harmonic trap with magnetic-field curvature $B''_{equiv} = B'/r = B'^2/k_B T$, where r characterizes the size of the cloud. If we use this equivalent curvature to characterize the compressed trap, we get the same result as eq. (C.13) — apart from the factor $e^{1/2}/2$. This emphasizes that lowering the bias field beyond a value $k_B T/\mu_m$ does not compress the cloud any further.

C.2.5 Figure of Merit for Magnetic Traps

The major goal when designing magnetic traps for BEC experiments is to get the highest collision rate after compression. Therefore, we regard this as the figure of merit of a trap. In a 3D harmonic oscillator, the collision rate increases due to adiabatic compression in proportion to the geometric mean of the three curvatures: $(B''^2_\rho B''_z)^{1/3}$.

In a Ioffe-Pritchard trap, after full compression to a final temperature T, the equivalent radial curvature is $\mu_m B'^2/k_B T$. Knowing that the final temperature scales with the geometric mean of the trapping frequencies, we find that the collision rate will depend on $B'B''^{1/4}$, implying that it is much more important to have strong field gradients than curvatures.

Let's compare this result to a TOP trap, assuming that the TOP trap is compressed until the trap depth due to the circle of death is $5k_B T$, a typical barrier for evaporative cooling. In this case, the effective radial curvature is

$$B''_{TOP,r} = \frac{\mu_m B'^2}{40 k_B T} \tag{C.14}$$

This implies that the TOP trap corresponds to an IP trap with a $\sqrt{40}$ smaller value of B'. Along the axial direction, the effective curvature B'' of a TOP trap is typically 100 times larger than that of an IP trap. The figure of merit, $B'B''^{1/4}$, is slightly higher for an IP trap, but this does not take into account the fact that one can usually obtain higher gradients in spherical quadrupole traps than in IP traps.

If we wish to compare the confinement for condensates after cooling, the effective radial curvatures of TOP traps and IP traps are $B'^2/2B_0$ and B'^2/B_0 respectively. The maximum confinement depends on how small B_0 can be, as determined by the adiabatic condition which ensures stability of the trap (eq. (C.1)). Usually B_0 is smaller in IP traps than in TOP traps, because the use of time-dependent fields in the TOP trap requires larger Larmor frequencies to prevent non-adiabatic spin-flips. However, experiments have not yet explored the lower limits for B_0 in either type of trap (see [168] for a theoretical treatment).

The bottom line is that both traps work very well. The TOP trap might have advantages in

studying vortices due to the built-in rotation, and does not require careful bias-field cancellation. The advantages of the IP trap are the variable aspect ratio and the use of only dc fields. Several groups have now built magnetic traps using room temperature electromagnets which provide sufficient confinement for evaporation to BEC. Even tighter confinement, and therefore faster evaporation, could be achieved with permanent magnets or superconducting magnets, but at the price of less flexibility.

C.3 Evaporative Cooling

Gaseous Bose-Einstein condensates have so far only been obtained by evaporative cooling. Evaporative cooling is done by continuously removing the high-energy tail of the thermal distribution from the trap. The evaporated atoms carry away more than the average energy, which means that the temperature of the remaining atoms decreases. The high energy tail must be constantly repopulated by collisions, thus maintaining thermal equilibrium and sustaining the cooling process. Evaporative cooling is a common phenomenon in daily life — it's how hot water cools down in a bathtub or in a cup of coffee. Evaporative cooling of trapped neutral atoms was developed at MIT as a method for cooling atomic hydrogen which had been pre-cooled by cryogenic methods [169, 170, 171]. The first suggestion by Hess [169] was soon followed by an experimental demonstration [170]. Evaporative cooling has been reviewed in [156, 92], and we only briefly summarize the basic aspects here.

The essential condition for evaporative cooling is a long lifetime of the atomic sample compared to the collisional thermalization time. Trapped atom clouds are extremely dilute (about ten orders of magnitude less dense than a solid or a liquid) and collisional thermalization can take seconds. A major step was taken in May 1994 when the MIT and JILA groups combined laser cooling with evaporative cooling, extending the applicability of evaporative cooling to alkali atoms [158, 159].

In these experiments, the evaporation of atoms was controlled by radio frequency radiation (rf induced evaporation). This technique was proposed by Pritchard [172] and Walraven [173] and first demonstrated by our group when spatial truncation of magnetically trapped atoms was observed [174]. Increases in phase-space density were reported by the MIT and Boulder groups at IQEC in May 1994 and at ICAP-XIV [158, 159] and published after further progress had been achieved [162, 163]. Other early work on evaporation of alkali atoms was done at Rice [175] and Stanford [93].

In rf-induced evaporation, the rf radiation flips the atomic spin. As a result, the attractive trapping force turns into a repulsive force and expels the atoms from the trap (fig. C.2). This scheme is energy-selective because the resonance frequency is proportional to the magnetic field, and therefore to the potential energy of the atoms. In the case of transitions between magnetic sublevels m_F, the resonance condition for the magnetic field strength B is $|g|\mu_B B = \hbar\omega_{\rm rf}$. Since the trapping potential is given by $m_F g\mu_B[B(r) - B(0)]$, only atoms which have a total energy $E > \hbar|m_F|(\omega_{\rm rf} - \omega_0)$ will evaporate (ω_0 is the rf frequency which induces spin-flips at the bottom

of the trap). Rf induced evaporation has several advantages over other evaporation methods [92]. First, the evaporation process can be completely separated from the design of the magnetic trapping potential. In particular, there is no need to weaken the trapping potential in order to lower its depth. This makes it easier to reach runaway evaporation where the decrease in temperature leads to a net increase in density and collision rate despite the loss of atoms. Furthermore, atoms evaporate from the whole surface where the rf resonance condition is fulfilled. This make the evaporation three-dimensional in velocity space, and therefore very efficient.

Let us briefly summarize the major developments in the three years since our earlier review [92]. First, evaporative cooling has become a standard technique of atomic physics. The complete list of experiments in 1995 comprised only six groups — since then, thirteen more groups have implemented evaporative cooling and achieved BEC, and there are more to come! Second, the evaporation process has been simulated in more detail [176, 177, 178, 179, 180]. In particular, Monte Carlo simulations can account for realistic experimental situations [180] providing insight on how to optimize the cooling cycle. A careful study at Amsterdam demonstrated that saddle-point evaporation is one-dimensional and was the bottleneck for achieving BEC in hydrogen [181]. This problem was finally overcome when rf evaporation was implemented [43]. Several groups reported evaporation in TOP traps by the circle of death [182, 183, 184]. This scheme should be one- or two-dimensional depending on the orientation of the circle of death, but no striking detrimental effects due to that have been reported.

Sympathetic cooling between rubidium atoms in different hyperfine states has been observed [86]. Combining sympathetic cooling with cryogenic buffer gas loading [62] should vastly extend the number of species which can be cooled. Finally, cooling by adiabatic deformation of the trapping potential has been achieved [167, 185]. The relation of this reversible, adiabatic scheme to irreversible evaporative cooling is discussed in section 8.2 of [47].

C.3.1 Trap Loss and Heating

Evaporative cooling requires a favorable ratio of elastic to inelastic collisions. In the case of alkali atoms, as shown in fig. 6 of [47], the dominant loss mechanism over a large range of temperatures and densities is background gas collisions. Therefore, experimentalists have made major efforts to reconcile the loading and trapping techniques with an extreme UHV environment. The standard technique is to use a combination of an ion pump and a titanium sublimation pump, and to implement differential pumping to isolate the trapped atoms from the high gas load of the atom source. Our atom beam source (the atomic beam oven) is separated from the magnetic trap by about one meter, and an additional 10 cm-long narrow tube ensures good differential pumping. After loading the MOT for about 2 seconds, we close a shutter near the atomic beam oven, isolating our sample in a perfect UHV environment. More compact experiments have to achieve the separation between a region of high vapor pressure and UHV with less space, making differential pumping more difficult to implement.

Several groups have reported that a background pressure in the low 10^{-11} mbar range is nec-

essary to obtain magnetic trap lifetimes of one minute. The background collision rate per trapped atom is $n_{bg}\sigma_{bg}v_{bg}$, where n_{bg} and v_{bg} denote the density and thermal velocity of the background gas molecules. The collision cross section, σ_{bg}, is typically 10^3 angstrom, much larger than the size of the molecules. This reflects the major contribution of the long-range van-der-Waals potential and small-angle scattering to the total cross-section. Background collisions are discussed in detail in refs. [64, 65].

Another possible source of heating is magnetic field fluctuations which modulate the trapping potential. The simple model of ref. [146] predicts that the characteristic time t_{heat} for the exponential growth of the energy of the trapped atoms is given by

$$\frac{1}{t_{heat}} = \pi^2 \nu_{trap}^2 S(2\nu_{trap}) \tag{C.15}$$

where ν_{trap} is the trapping frequency and $S(\nu)$ is the power spectrum of the fractional magnetic field noise. An estimate for $S(\nu)$ is given by $\epsilon^2/\Delta\nu$ where ϵ denotes the rms fractional fluctuations of the magnetic field in a bandwidth of $\Delta\nu$. Assuming an inductance limited bandwidth of 1 kHz and a trapping frequency of 100 Hz, we obtain

$$t_{heat} = \left[\frac{10^{-2}\text{sec}}{\epsilon^2}\right] \tag{C.16}$$

Current fluctuations of 1% result in $t_{heat} = 100$ seconds. Therefore, commercial power supplies with a current stability of 10^{-3} to 10^{-4} are fully adequate for achieving BEC. Stability requirements are more severe at higher trapping frequencies as in optical traps.

Another source of heating is mechanical noise, but it seems not to be very critical. As a precaution we usually don't use any mechanical pumps close to the experiment, but there were occasions where these pumps were running without affecting Bose condensation. However, a fan-cooled, high-power rf amplifier which was on the experimental table prevented us from getting BEC until we isolated it with a foam pad. So although BEC is usually forgiving, it is still wise to consider sources of mechanical vibration as potential problems.

Appendix D

Magnetic Coils Wiring Diagram

The magnetic coils are wired up according to the diagram in Fig. D.1. The magnetic coils are powered by three high current power supplies. The power supply which drives current through the curvature/anti-bias coils, and the power supply which drives the gradient coils can each generate up to 500 amps of current at a voltage of up to 30 volts[1]. These power supplies have a maximum rated voltage ripple of 60 millivolts when run at full power. The power supply connected to the extra bias coils is a similar but lower power model, capable of producing 250 amps at 20 volts[2]. Four additional low current (10 amp) power supplies[3] are used to generate the spherical quadrupole field for the MOT. These power supplies are are protected from the high current supplies by diodes.

The currents in the trap have to be switched on and off rapidly in order to transfer atoms from the MOT into the magnetic trap, and in order to do time-of-flight studies of the condensate. The high speed switching of the high current power supplies is done with insulated gate bipolar transistors (IGBTs), and the lower current supplies are switched with MOSFET solid state switches. The sudden switch-off of the magnetic trap causes large inductive voltage spikes. The spikes are

[1]EMI model ESS 30-500
[2]EMI model EMS 20-250
[3]Kepco model ATE25-10M

minimized by running the positive and negative leads close together. Nevertheless, the spikes generated are still easily capable of destroying the IGBTs and MOSFETs. The solid state switches are protected by several large varistors which absorb the power in the spike. We found, however, that even with carefully placed wires, there was still enough inductance in the few feet of wire *between* the varistors and the transistor switches to produce voltage spikes capable of killing them. This problem was solved by putting an additional varistor, smaller in size but with a Zener voltage equal to the larger varistors used, directly across each IGBT and MOSFET switch.

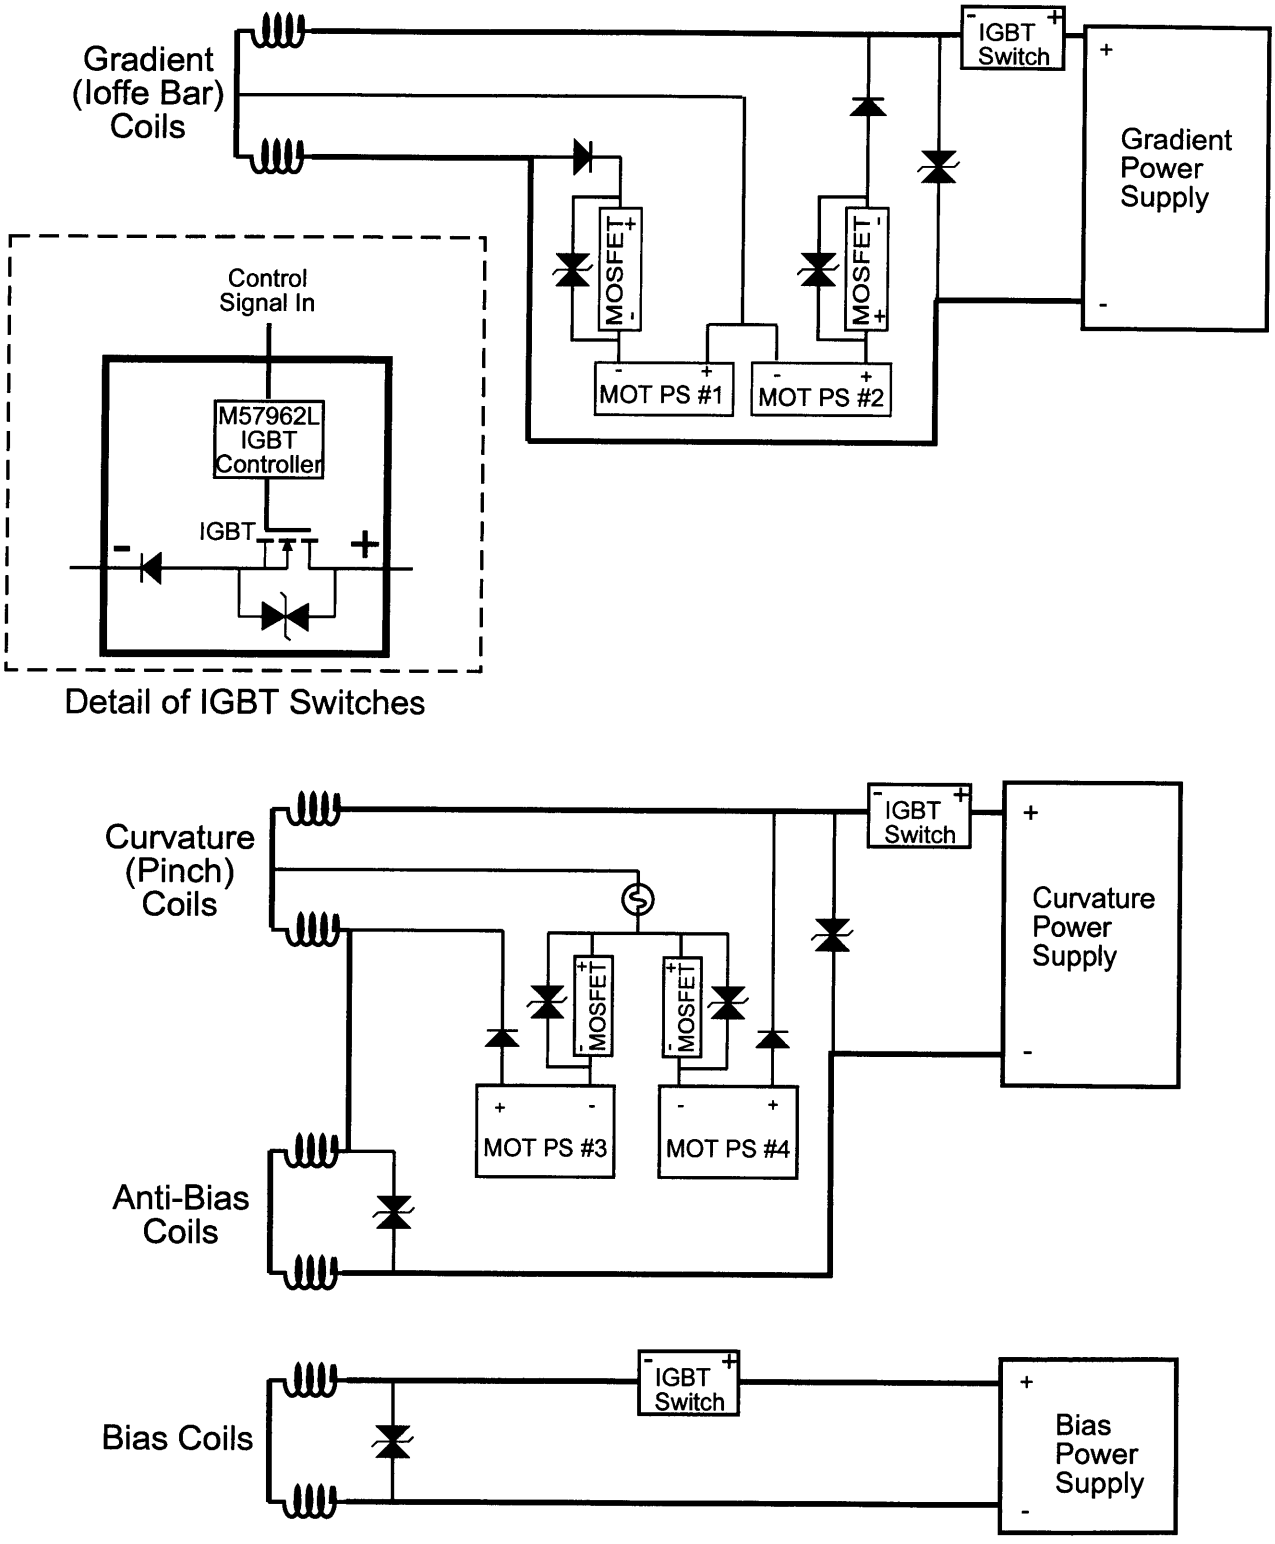

Figure D.1: Magnetic trap wiring diagram.

Appendix E

Imaging - Varenna Summer School 1998

The following appendix is a selection reprinted from our contribution to the proceedings of the 1998 Varenna summer school on Bose-Einstein condensation [47]. It has been modified from the original only in the sense that cross references, citations, and section headings have been renumbered and clarified to be consistent with the rest of this thesis.

E.1 Techniques to Probe Bose-Einstein Condensates

Everything we know about gaseous Bose condensates has been obtained by optical diagnostics. "Contact probes" cannot be used because the samples are much smaller ($\sim 10^7$ atoms) than even a 10 μm sized probe ($\sim 10^{13}$ atoms), which would cause the atoms to equilibrate with the probe rather than the opposite. With the "cooling power" of rf evaporation we can cool at most 10^8 atoms per minute, and it would take several months just to cool the sensor tip. Fortunately, optical diagnostics are very versatile, and the ease with which light scattering methods are implemented for dilute atomic samples is a major advantage over condensed-matter systems.

The two most important techniques for observing Bose-Einstein condensates are in-situ and time-of flight imaging. In both cases, one obtains an image which reflects the density distribution of the atoms either in a trapped state or in ballistic expansion. In this section we discuss the basic interactions between atoms and light and the physical principles of absorptive and dispersive imaging. Finally, we refer the interested reader to appendix A of [47] for a discussion of techniques for processing absorption and phase-contrast images.

E.1.1 Atom - Light Interactions

The interaction of atoms with a beam of light involves three processes: absorption of photons, re-emission of photons, and shifting the phase of the transmitted light. These properties are used in absorptive, fluorescence, and dispersive imaging methods, respectively.

The interaction can be described by the complex index of refraction of the atoms $n_{ref} = \sqrt{1 + 4\pi n\alpha}$, where α is the atomic polarizability, and n is the density of atoms. Assuming $n_{ref} - 1 \ll 1$, the index of refraction for a two-level system can be written (in the rotating wave approximation):

$$n_{ref} = 1 + \frac{\sigma_0 n\lambda}{4\pi}\left[\frac{i}{1+\delta^2} - \frac{\delta}{1+\delta^2}\right] \tag{E.1}$$

where σ_0 is the resonant cross-section (for a two level atom $\sigma_0 = 6\pi\lambda^2$), and $\delta \equiv \frac{\omega-\omega_0}{\Gamma/2}$ is the detuning in half linewidths. If more than two levels are involved (e.g. Zeeman sublevels, see sect. E.1.5), several resonances contribute to the polarizability.

Equation (E.1) assumes the limit of a weak probe laser beam. Saturation may be non-negligible for near-resonant absorption imaging. In this case, one has to replace δ^2 by $\delta^2(I/I_{\mathrm{SAT}})$ in the expression for the Lorentzian lineshape. I_{SAT} is the saturation intensity for the transition, and the intensity I depends on x, y (beam profile) and z (due to absorption).

Assuming that light enters and exits the cloud at the same (x, y) coordinate (the "thin lens" approximation applied to the atom cloud), the atoms simply attenuate and phase shift the probe light:

$$E = tE_0 e^{i\phi} \tag{E.2}$$

The transmission coefficient t and phase shift ϕ depend on the product of the column density $\tilde{n} = \int n \cdot dz$, and σ_0 (eq. (E.1)).

$$t = e^{-\tilde{D}/2} = \exp\left(-\frac{\tilde{n}\sigma_0}{2}\frac{1}{1+\delta^2}\right) \tag{E.3}$$

$$\phi = -\delta\frac{\tilde{D}}{2} = -\frac{\tilde{n}\sigma_0}{2}\frac{\delta}{1+\delta^2} \tag{E.4}$$

where $\tilde{D} = \tilde{n}\sigma_0/(1+\delta^2)$ is the off-resonance optical density. Section E.1.3 discusses a quantum mechanical treatment of light scattering.

Most BEC experiments have relied on spatially imaging the real or imaginary part of the index of refraction. We discuss various imaging methods in this section. An alternative and in many aspects complementary method is high resolution spectroscopy, where information is obtained from line shifts and line broadenings. Such methods have so far only been applied to hydrogen condensates (see Kleppner and Greytak's article in [186]), but can be extended to alkalis in the form of Doppler sensitive Raman or Bragg scattering.

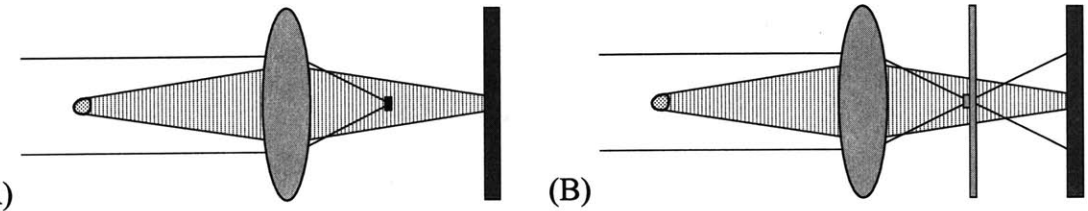

Figure E.1: Dark-ground (A) and phase-contrast (B) imaging set-up. Probe light from the left is dispersively scattered by the atoms. In the Fourier plane of the lens, the unscattered light is filtered. In dark-ground imaging (A), the unscattered light is blocked, forming a dark-ground image on the camera. In phase-contrast imaging (B), the unscattered light is shifted by a phase plate (consisting of an optical flat with a $\lambda/4$ bump or dimple at the center), causing it to interfere with the scattered light in the image plane.

E.1.2 Absorptive and Dispersive Methods

Absorption imaging is done by illuminating the atoms with a laser beam and imaging the shadow cast by the atoms onto a CCD camera. Because photosensors aren't sensitive to phase, the absorption image shows the spatial variation of t^2 (eq. (E.3)).

To image a transparent object, information encoded in the phase shift of the light must be converted into intensity information which can be detected by a photosensor. Several methods are commonly used in microscopy. The Schlieren method was introduced by A. Toepler in 1864 and the phase-contrast method by Fritz Zernike in 1934 (earning him the 1953 Nobel prize in physics). All dispersive methods rely on the ability to separate scattered and unscattered components of the probe light and manipulate them independently. This is usually done by spatially filtering in a Fourier plane of the imaging system where the unscattered probe light comes to a focus.

Dark-Ground Imaging: The simplest form of spatial filtering is to block the unscattered light by placing a small opaque object into the Fourier plane (fig. E.1).

The probe light field after passing through the atoms (fig. E.2) can be separated into the scattered and unscattered radiation.

$$E = tE_0 e^{i\phi} = E_0 + \Delta E \tag{E.5}$$

Blocking the unscattered light gives the dark-ground signal:

$$\langle I_{dg} \rangle = \frac{1}{2} |E - E_0|^2 = I_0 \left[1 + t^2 - 2t\cos\phi\right] \tag{E.6}$$

where $I_0 = E_0^2/2$. For small ϕ the dark-ground signal is quadratic in ϕ.

Phase-contrast imaging: Phase-contrast imaging can be regarded as a homodyne detection scheme in which the unscattered light acts as the local oscillator and interferes with the scattered radiation. This is accomplished by shifting the phase of the unscattered light by $\pm\pi/2$ in the Fourier plane of the imaging lens (fig. E.1). This is done with a "phase plate" which is an optical

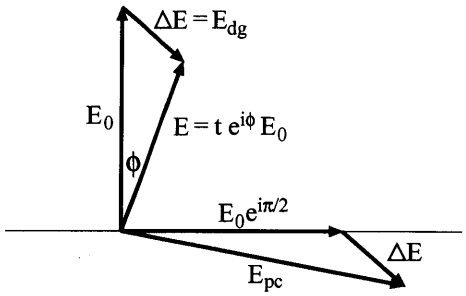

Figure E.2: Phasor diagram of dark-ground and phase-contrast imaging. A ray of incident light with an electric field E_0 is scattered by the atoms, causing the light to be attenuated and shifted in phase, resulting in the electric field E. The dark-ground method images $\Delta E = E_{dg}$, the difference between incident and scattered electric fields. Phase-contrast methods cause ΔE and E_0 to interfere by rotation the phase of E_0 by 90°, resulting in the field E_{pc}.

flat with a small bump or dimple in the center. The result of this phase shift is shown in the phasor diagram in fig. E.2.

The intensity of a point in the image plane is then

$$\langle I_{pc} \rangle = \frac{1}{2} \left| E + E_0 \left(e^{\pm i \frac{\pi}{2}} - 1 \right) \right|^2 = I_0 \left[t^2 + 2 - 2\sqrt{2} t \cos \left(\phi \pm \frac{\pi}{4} \right) \right] \tag{E.7}$$

where the $\pm$ sign corresponds to a phase plate which changes the phase of the unscattered light by $\pm\pi/2$. For small ϕ one obtains

$$\langle I_{pc} \rangle \simeq I_0 \left[t^2 + 2 - 2t \pm 2t\phi \right] \tag{E.8}$$

which is linear in ϕ. This makes phase-contrast imaging superior to dark-ground imaging for small signals.

Polarization-contrast imaging: Another dispersive imaging method takes advantage of the anisotropic polarizability of spin-polarized atoms [187, 3]. It relies on different phase-shifts for orthogonal polarization components of the probe light. A polarizer before the camera can block the unscattered light, resulting in a dark-ground image. Rotating the polarizer leads to interference of the two polarizations, resulting in a phase-contrast image.

In the case of high magnetic fields, the degeneracy between σ_+, σ_-, and π transitions is lifted, and it is possible to choose a probe detuning such that only one polarization component is close to an atomic resonance and experiences a phase shift. The other polarization component then serves as a uniform local oscillator.

This method has the advantage that it does not require spatial filtering in the Fourier plane, making it easier to correct for stray light (see appendix A in [47]). However, this method is not universally applicable. The standard imaging geometry for a Ioffe-Pritchard trap involves a probe beam propagating perpendicular to the axis of a low magnetic bias field. For our typical detunings

(much larger than the excited state hyperfine structure), the phase shift is the same for both polarizations leading to zero signal. However, the independence of the phase shift from polarization is advantageous for phase-contrast imaging, making it insensitive to the polarization of the probe beam.

Imaging dilute clouds: Dispersive and absorptive methods can be discussed together by generalizing the previous treatment to a generic phase plate with a central spot which retards the unscattered light by an arbitrary phase γ and transmits a fraction τ^2. Absorption ($\tau = 1$, $\gamma = 0$), dark-ground ($\tau = 0$), and phase-contrast ($\tau = 1$, $\gamma = \pm\frac{\pi}{2}$) imaging are obtained as special cases.

The intensity at a point in the image plane is:

$$\langle I_{gen} \rangle = I_0 \left[1 + t^2 + \tau^2 + 2t\tau \cos(\phi - \gamma) - 2t \cos\phi - 2\tau \cos\gamma\right] \tag{E.9}$$

If we expand this around $\tilde{n} = 0$ (i.e. $t = 1$, $\phi = 0$) and keep only the first order term we get

$$\langle I_{gen} \rangle = I_0\tau^2 - I_0\sigma_0\tau \left[\frac{\delta}{1+\delta^2} \sin\gamma + \frac{1}{1+\delta^2} \cos\gamma\right] \tilde{n} \tag{E.10}$$

This shows that for a given detuning δ, the signal is maximized when $\tau = 1$ and $\tan\gamma = -\delta$. The absolute maximum signal is obtained by choosing $\delta = \gamma = 0$, the case of resonant absorption imaging. However, the maximum phase-contrast signal ($\gamma = \pm\pi/2$), obtained when $\delta = 1$, is only a factor of two lower. Dark-ground imaging is far from optimum at any detuning

High sensitivity is usually only an issue when imaging expanded clouds in time-of-flight, or for imaging the wings of the spatial distribution [89], because trapped condensates are generally optically dense.

Imaging optically dense clouds: The typical resonant optical densities of our condensates (in the radial direction) is on the order of 300. When such a dense cloud is probed with resonant light, the image is "blacked out," making it extremely difficult to extract atomic column densities from the image.

Usually, an optimum detuning for absorption imaging gives about 50% absorption. However, for non-zero detuning, the atomic susceptibility has a finite real component which causes the atoms to refract light like a lens. This does not affect the images as long as all of the refracted light is collected by the imaging system. Otherwise, the refracted light appears as a false absorptive signal. Since the magnitude of this effect varies across the cloud, it affects both relative and absolute density measurements.

The refraction angle for a cloud of atoms with diameter d and a maximum phase shift ϕ can be estimated as $2\lambda\phi/\pi d$. As long as this phase shift is less than $\pi/2$, the refraction angle is smaller than the diffraction angle λ/d due to the finite size of the object. For larger phase shifts, refraction leads to a noticeable additional divergence of the light after passing through the atoms. For a diffraction limited imaging system, diffraction of the smallest resolvable object just fills the solid angle extended by the lens. Therefore, a dispersively dense object (i.e. phase shifts larger than $\pi/2$) will degrade the spatial resolution for quantitative absorption imaging.

For a cloud with a resonant optical density $\tilde{D}_0 = \tilde{n}\sigma_0 \gg 1$, the highest contrast for absorption

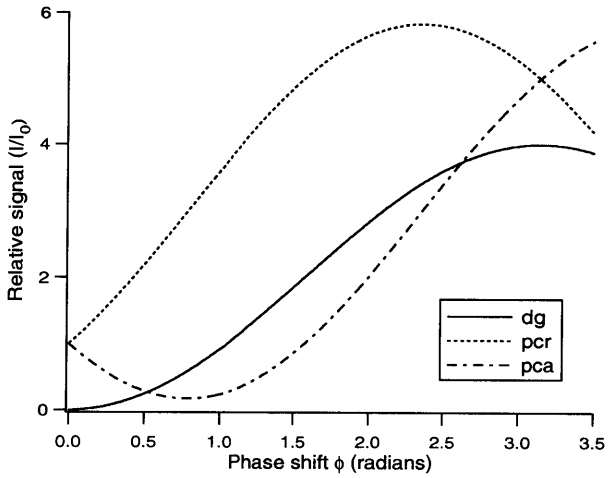

Figure E.3: Comparison of signal for dark-ground and phase-contrast imaging. Relative signal as a function of phase shift for dark-ground imaging (dg) and phase-contrast imaging with a phase plate which either retards (pcr) or advances (pca) the phase of the unscattered light. For negative phase shifts (as would be generated with a blue detuned probe), (pcr) and (pca) are swapped.

imaging is obtained with a detuning $\delta = (\tilde{D}_0)^{1/2}$ at which the optical density drops to unity. However, at this detuning, the cloud is dispersively dense with a phase shift of $\phi \approx (\tilde{D}_0)^{1/2}/2$. For large clouds, quantitative absorption imaging is still possible and has been applied [162, 89], but for small clouds, refraction is a limiting factor.

These limits are avoided by larger detunings. However, when the phase shift is reduced to $\pi/2$ by choosing $\delta = \tilde{D}_0/\pi$, the optical density has dropped to $\tilde{D} = \pi^2/\tilde{D}_0$. For typical condensates of optical density $\tilde{D}_0 = 300$, the absorption signal is far too small to be detected. In contrast, dispersive imaging still features high signal levels.

Equations (E.6) and (E.7) show that dark ground and phase-contrast signals are periodic in ϕ. So unlike absorption methods, dispersive imaging does not saturate at large ϕ, and arbitrarily high phase shifts can be measured by counting fringes. However, large phase shifting should be avoided since the interpretation of images is easier if the dispersive signal never "rolls over." For image analysis, an approximately linear relationship between signal and phase is desirable. Figure E.3 shows that a wide range of approximate linearity is achieved for phase-contrast imaging with a phase plate which retards the phase of the unscattered light ($\omega t \rightarrow \omega t + \pi/2$) for red-detuning, or advances the phase ($\omega t \rightarrow \omega t - \pi/2$) for blue-detuning of the probe light. For this choice of phase plate, the maximum signal before rolling over is 5.8 times the incident light level. This gives phase-contrast imaging a larger dynamic range than absorption imaging.

E.1.3 Quantum Treatment of Light Scattering

Further insight into the interaction between the probe light and the atoms is obtained from a quantum-mechanical treatment [188]. For a single atom in state $|i>$, the differential cross section for light scattering involves a sum over all possible final states $|f>$:

$$\frac{d\sigma_R}{d\Omega} = C \sum_f \left| \langle i | e^{i\Delta \mathbf{k}\cdot\mathbf{r}} | f \rangle \right|^2 \tag{E.11}$$

This is simply the differential cross section for Rayleigh scattering. We consider only the low intensity limit, i.e. the central peak of the Mollow triplet [189] and assume that Doppler shifts are small, i.e. $|\mathbf{k}_f| = |\mathbf{k}_i|$, where $|\mathbf{k}_f - \mathbf{k}_i| = |\Delta \mathbf{k}| = 2k \sin(\theta/2)$ is the change of the photon's wavevector due to scattering by an angle θ. The atomic matrix elements and fundamental constants are contained in C.

We now extend the treatment to N atoms in a trap, where each trap level $|j>$ is initially populated with N_j atoms, still assuming that the cloud is optically dilute. The relevant matrix elements are those of the operator $e^{i\Delta \mathbf{k}\cdot\mathbf{r}}$ with the symmetrized many-body eigenstates, characterized by the populations N_j. For bosonic atoms, one obtains [188]

$$\frac{d\sigma}{d\Omega} = C \left| \sum_i N_i \langle i | e^{i\Delta \mathbf{k}\cdot\mathbf{r}} | i \rangle \right|^2 + C \sum_{i \neq f} N_i (N_f + 1) \left| \langle i | e^{i\Delta \mathbf{k}\cdot\mathbf{r}} | f \rangle \right|^2 \tag{E.12}$$

The scattering is naturally divided into coherent and incoherent scattering. The first term (the coherent part) represents elastic scattering in which the state of the atom is unchanged. It is related to the Fourier transform of the density distribution:

$$\sum_i N_i \langle i | e^{i\Delta \mathbf{k}\cdot\mathbf{r}} | i \rangle = \int n(\mathbf{r}) e^{i\Delta \mathbf{k}\cdot\mathbf{r}} d\mathbf{r} \tag{E.13}$$

For scattering angles smaller than the diffraction angle, $\Delta k < 1/d$, the intensity of the scattered light is N^2 times the Rayleigh scattering of a single atom. This term creates the real part of the index of refraction and generates the signal in dispersive imaging.

The second term in (E.12) represents inelastic scattering. We can separate the "spontaneous part" from the part due to bosonic stimulation:

$$\left(\frac{d\sigma}{d\Omega}\right)_{incoh} = C \sum_{i \neq f} N_i \left| \langle i | e^{i\Delta kr} | f \rangle \right|^2 + \sum_{i \neq f} N_i N_f \left| \langle i | e^{i\Delta kr} | f \rangle \right|^2 \tag{E.14}$$

The first term is N times the single atom Rayleigh scattering (E.11), if we assume that the sample is much larger than an optical wavelength so that we can drop the condition $i \neq f$. This term represents incoherent scattering into a solid angle of 4π. Since photons scattered by a large angle are "missing" from the probe beam, it is this term which gives rise to the absorption signal.

The second term, describing Bose-enhanced scattering, is only non-vanishing near the phase transition. Politzer distinguishes two contributions, scattering in and out of the condensate, and scattering between excited states [188]. This scattering is limited to angles smaller than the ratio of the optical wavelength to the thermal de Broglie wavelength. Because this contribution is much weaker than the diffractive signal, it might be difficult to observe it. It is neglected in the quantitative analysis of images. The effects of quantum statistics can also be included in the index of refraction [190]. This treatment confirms that non-classical light scattering effects are small, are noticeable only near the phase transition, and vanish as $T \to 0$.

E.1.4 Non-destructive Imaging

It is obvious from the previous discussion that the different imaging techniques do not excite the atoms in a different way — they only "collect" different parts of the scattered electric field. Whether the atom cloud is heated during the probing or not only depends on the intensity and duration of the probe light.

A discernible image with 30×30 pixels and 100 detected photons per pixel involves 10^5 photons. Even if each photon would "knock" an atom out of the condensate, this would be "non-perturbative" if the sample had many millions of atoms. Therefore, the large hydrogen Bose condensates could be observed non-destructively with absorption spectroscopy [43].

An important figure of merit of the different imaging techniques is the ratio of signal to heating. In this regard, dispersive imaging often has a big advantage. In absorption imaging, each absorbed photon (i.e. photons scattered at large angles) heats up the cloud by about one recoil energy. The dispersive signal is based on the coherent term (E.13) which leaves the atom in its initial quantum state, and therefore does not contribute to the heating. The strength of the signal in dispersive scattering is best estimated from the number of photons collected in dark-ground imaging. The total number of forward scattered photons (E.13) is larger than the total number of absorbed photons (E.14) by a factor which is approximately N times the solid angle of the diffraction cone, λ^2/d^2. Since

$$N\frac{\lambda^2}{d^2} \simeq \tilde{n}\lambda^2 \tag{E.15}$$

this factor is about the resonant optical density $\tilde{D}_0$, identifying some common roots between forward scattering and superradiance [191]. More quantitatively, the ratio of the signal in dark-ground imaging (E.6) to the number of Rayleigh scattered photons $(1 - t^2)$ in the limit of large detunings is $\tilde{D}_0/4$. For typical values of $\tilde{D}_0 \simeq 300$ this means that we can obtain two orders of magnitude more signal from dispersive imaging than from absorption imaging for the same amount of heating. With absorption methods, we could take only a single "non-destructive" image, but with dispersive imaging we could take real-time movies consisting of up to 100 frames, more than the storage capacity of our camera.

The elastic scattering of photons at small angles is analogous to the Mössbauer effect — it is recoil-free scattering of photons where the momentum transfer to the photons is absorbed by

the trap and not by the individual atoms. This picture is valid when the pulse duration of the probe light is longer than the trapping period. Rays of light passing through the atoms are bent in opposite directions on opposite sides of the atom cloud. If the pulse duration is longer than the trapping period, atoms have time to travel across the width of the cloud during the probe and the recoil is averaged to zero. For short pulses, atoms on opposite sides of the cloud experience opposite momentum transfer. However, since the scattering angles are small, typically about 10 mrad, the recoil heating is 10^4 times lower than for a Rayleigh-scattered photon. Therefore, even for short probe pulses, the dominant source of heating is large-angle Rayleigh scattering (as long as $\tilde{D}_0 \leq 10^4$).

Since the relative figure of merit of dispersive vs. absorptive imaging is $\tilde{D}_0/4$, dispersive imaging is usually only applied to very dense clouds, either trapped or in the early phase of ballistic expansion. Absorption imaging is used for time-of-flight imaging with sufficiently long expansion times so that the resonant optical density has dropped to values around unity. There is no obvious advantage of dispersive imaging over absorptive imaging for dilute clouds.

E.1.5 Other Aspects of Imaging

Fluorescence imaging So far we have not discussed fluorescence imaging. In fluorescence and absorption, one detects the same photons either as missing or as counted photons. In comparing signal strengths, however, fluorescence imaging is inferior by a factor of typically 100, since only a small fraction of the scattered photons are collected by the imaging system. Fluorescence methods usually have the advantage of being background-free, in contrast to absorption methods, where a small signal might be invisible due to fluctuations of the laser intensity. However, in BEC experiments, one can usually choose a detuning where the absorptive or dispersive signal is comparable to the incident light intensity.

Maximum light intensity in single-shot imaging Above we discussed heating by the probe pulse as the signal-limiting factor in non-destructive imaging. The situation is different if the goal is to take only one picture of a condensate. All groups reported their first observations of BEC using this "single-shot" technique. In this case, the figure of merit is not how much the condensate is heated during the probe pulse, but rather by how much is the image blurred due to recoil-induced motion.

If an atom in the sample scatters N_p photons in a time Δt, it acquires a velocity of $v_{rec}N_p$ (where v_{rec} is the recoil velocity) along the direction of propagation of the probe beam due to absorption, and a random (approximately isotropic) velocity $\sqrt{N_p}v_{rec}$. This random velocity v_{rms} leads to a random position of $r_{rms} = v_{rms}\Delta t/\sqrt{3} = \sqrt{N_p/3}\, v_{rec}\Delta t$ [192]. A large number of photons N_p per atom can be scattered if the pulse duration is short. This strategy finds its limit when the atomic transition is saturated. The acceptable amount of blurring depends on the required spatial resolution which is higher for a trapped cloud than for a cloud in ballistic expansion. Typical numbers for sodium are that 100 photons scattered in 50 μs give a final $v_{rms} = 30$ cm/s and

$x_{rms} = 9\,\mu$m.

Optical pumping The scattering of many photons per atom gives high signal levels and results in a very good signal-to-noise ratio, making the optical system less susceptible to stray light. However, it can only be applied on a cycling transition. If the atoms are trapped in a different state, they can be transferred to the cycling transition by optical pumping.

Optical pumping was used by our group for two other reasons. In the imaging of spinor condensates (sect. 8.4 of [47]) atoms trapped in different hyperfine states were pumped into the same state for probing, allowing easy comparison of relative atom numbers without correcting for different transition matrix elements. Spatially selective optical pumping can also be used to select only a part of the cloud for subsequent absorption imaging. We used this method to get around the line-of-sight integration in normal absorption methods: a thin light sheet selected atoms by optical pumping in a tomographic way, and only these atoms absorbed the probe light which propagated perpendicular to the light sheet (fig. 27 of [47]). This method was necessary to image the fringes created by two interfering Bose condensates [58].

Imaging in inhomogeneous magnetic fields When the temperature of the cloud is larger than the natural linewidth, the spectroscopic nature of light scattering becomes important. In this regime, the inhomogeneous trapping potential can "tune" the atoms in and out of resonance. This was the situation for the experiment on Lyman-α spectroscopy of hydrogen in Amsterdam [193], and also for our first evaporative cooling experiments [162].

Figure E.4 shows the variety of spectroscopic absorptive shadows cast by identical clouds in a spherical quadrupole trap at different probe detunings. Some shadows are much larger than the rms size of the cloud. These figures nicely show resonant shells where the atoms are Zeeman-shifted into resonance with the probe light.

Temperature and density can be extracted by comparing experimental images to simulations. These simulations solve for the propagation of the complex electric field through a medium characterized by a susceptibility tensor which depends on the local atom density and magnetic-field direction. Simulations done by Michael Andrews [101] were based on the method used by the Amsterdam group [193] and explained all observed features (fig. E.4).

It was surprising that some spectroscopic images had a lower symmetry than the ellipsoidal cloud and exhibited swirl-like structures. These structures are due to the optical activity of the trapped cloud which rotates the polarization of the probe light. The symmetry of the cloud in the quadrupole trap would suggest that the intensity distribution of the transmitted light has an angular dependence of $a_0 + a_1 \cos 2\phi$ (see eq. (21) in ref. [193]). However, there is a complex phase of the transmitted electric field which changes the angular dependence to $a_0 + a_1 \cos(2\phi - \delta)$ causing the swirls seen in fig. E.4. This angle δ was not considered in [193] but did not affect their angle-averaged results.

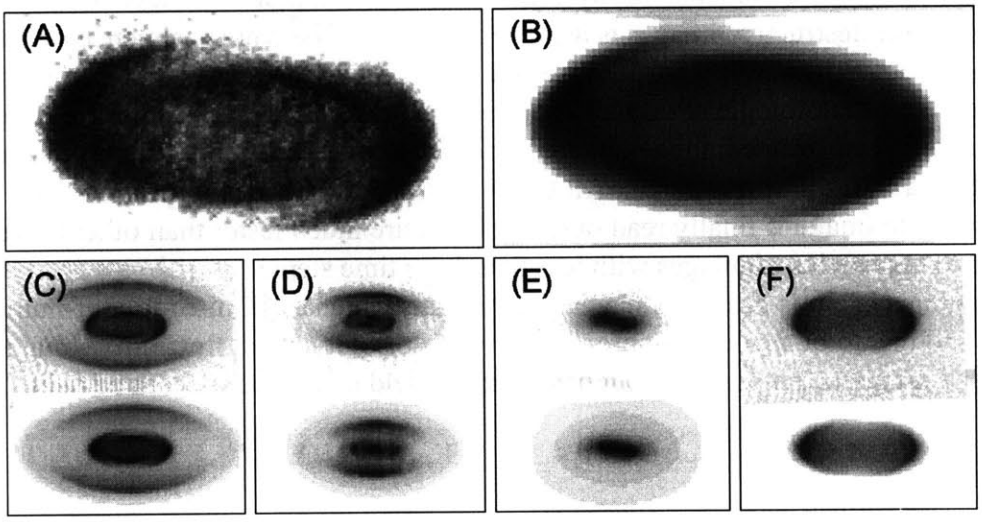

Figure E.4: Absorption images of clouds of sodium in a spherical quadrupole trap showing "swirls" and "resonant shells" (see text). Image (A) shows a cloud at $\sim 80\mu$K in a quadrupole field with $B' = 500$ G/cm probed with linearly polarized light detuned -97 MHz from the F=1 to F'=2 transition. Image (B) is a simulation of image (A). Images (C), (D), (E), and (F) show clouds at $\sim 500\mu$K in a quadrupole field with $B' = 250$ G/cm at different detunings. The experimental data is on top, and the simulation below. Image (C-F) were taken with -80, -40, 0, and +80 MHz detuning respectively. The true rms size of the cloud is similar to the black area in (E). The width of images (A) and (B) is 1.5 mm, and the width of images (C-F) is 5 mm. Figures taken from ref. [101].

Experimental aspects The current generation of low-noise CCD cameras meets all the needs of current BEC experiments. With about 5 electrons rms noise per pixel, the camera is already shot-noise limited for more than about 50 detected photons per pixel. Usually the detection sensitivity is limited by bad probe beam profiles and imperfect image normalization (see appendix A of [47]) rather than sensor performance. Image-intensified cameras have nominally higher sensitivity for very small signals, but for more than ~ 50 incident photons per pixel, the higher quantum efficiency of CCD sensors (40%) compared to image intensifiers (15%) guarantees better signal-to-noise ratios.

The low noise of the camera usually requires a slow readout over several seconds. Rapid sequencing of non-destructive images is achieved by storing the whole sequence on the chip, and then reading it out slowly. For this purpose, the sensor is divided into an area for exposure and a much larger area for storage. The division is done by a mask, either on the sensor, or for more flexibility, in an intermediate image plane of the optical system. After exposure, the accumulated charge is shifted into the storage area. All CCD chips should be able to perform this shift (since this is the way the image is finally read out), but some are much faster than others. We have used this method to take series of images with less than 1 ms time steps.

Our imaging system is based on standard achromats which are limited to an f-number of four. We have found this to be a good compromise between resolution and flexibility. Imaging systems with smaller f-numbers are possible, but have smaller field of view and depth of field.

Appendix F

Optically Trapped Two-Dimensional Condensates

It has been suggested that the use of sheet potentials, created with scanned optical dipole beams, could possibly be used to study two-dimensional Bose condensates. In order to understand the feasibility of such experiments, we can perform the following evaluation. If we assume that a sheet condensate has a given width, length, and thickness, then two-dimensional behavior of the condensate is ensured if the density of the condensate is small enough such that the healing length is longer than the condensate's thickness. This density limit results in a limit on the maximum number of atoms in the 2-D condensate.

In order to increase the maximum number of atoms allowed in a two-dimensional condensate, it is useful to decrease the laser power used to generate the trap — thereby expanding the volume of the condensate and allowing it to hold more atoms. Increasing the size of the condensate, however, means that the cloud's thickness increases. So to maintain the two-dimensional character of the cloud, the density must be reduced in order to increase the healing length of the condensate by the same factor. If we assume that the laser generates a harmonic potential with an infinite well depth,

we find that as the laser power is reduced, the maximum healing length of a cloud with a given number of atoms increases at a faster rate than the thickness of the cloud. As a result, a weaker trap made with lower laser power can satisfy the criterion for two-dimensionality with a larger number of atoms.

The number of atoms allowed in a two-dimensional trap cannot, however, be made infinitely large simply by reducing the laser power to zero. This is due to the fact that the laser potential has a finite well depth, and if the power is reduced too far, the trap will become unstable. The maximum number of atoms in a two-dimensional potential can be calculated, then, by assuming that the healing length is just equal to the thickness of the condensate, and that the optical potential has a depth which is f times the chemical potential of the trapped cloud of atoms. Although the first assumption implies that the condensate is no longer in the the Thomas-Fermi limit, the Thomas-Fermi approximation can still be applied to obtain an order-of-magnitude estimate of the maximum number of atoms which can be held in a two-dimensional condensate. In the absence of gravity, we find that the maximum number of atoms is:

$$N = s \cdot \frac{z_R}{4\pi a}\sqrt{\frac{1}{f}} \tag{F.1}$$

In this equation, s is the ratio of the sheet condensate's width to the condensate's thickness (determined by the range over which the laser beam is scanned), a is the s-wave scattering length of the atoms, and $z_R = \pi w_0^2/\lambda$ is the Rayleigh range of the optical dipole beam (w_0 being the beam waist radius of the dipole beam focus). This optimum trap is generated with a total laser power of:

$$P = sf^2 \cdot \frac{2\pi^3\hbar^2|\Delta|c}{3M\Gamma\lambda_0^3} \tag{F.2}$$

Note that the optimum power is independent of the beam waist of the laser. If we choose $f = 10$ (a reasonable number, just large enough to prevent rapid "evaporation" of the atoms over the edge

of the trap) and $s = 10$, we find that an optimized 2-D trap for sodium can be created with a beam from our Nd:YAG laser with a total power of $\simeq$ 1 mW. Furthermore, if we desire to have at least 10^5 atoms in the trap, the criterion for a two-dimensional trap is achieved by focusing the laser to a beam waist radius of 17.8 μm, requiring a lens with an f/# of 28. These parameters would result in a cloud with a density of $2.5 \cdot 10^{11}$ atoms / cm^3 having dimensions of 8 x 80 x 630 μm, and having a healing length equal to 8 μm. The potential would be described by harmonic confinement in two axes (along the longest and shortest dimensions of the trap), characterized by the trapping frequencies $\nu_x = 0.17$ Hz and $\nu_y = 14$ Hz, and a box potential in the third axis, characterized by the 80 μm width of the condensate, resulting in a frequency spectrum of $\nu_{1\to n} = (n^2 - 1) \cdot 2.1$ Hz. The optical potential would have a trap depth of just 1.6 nK.

These parameters are within the range of our current setup. The extremely low trap depth could, however, make the trap extremely sensitive to beam-jitter induced heating [146] and to background gas collisions (see Section 3.1). Furthermore the potential gradient due to gravity, which was neglected in this analysis, is 150 times stronger than the maximum gradient of the dipole potential in the example given above. As a result, such a configuration would not be able to hold atoms against gravity.

If we add the additional criterion that the maximum potential gradient generated by the optical field be j times larger than Mg, we find that the power given in Eq. (F.2) must be focused to a beam waist of

$$w_0 = \left(\frac{sf^2\hbar^2}{2jM^2g} \right)^{1/3} \tag{F.3}$$

and we find that the maximum number of atoms is then:

$$N = \frac{s}{4a\lambda} \sqrt{\frac{1}{f}} \left(\frac{sf^2\hbar^2}{2jM^2g} \right)^{2/3} \tag{F.4}$$

For sodium trapped with 1064 nm light from a Nd:YAG laser, using the parameters $s = f = j = 10$, this requires a beam waist radius of 3.4 μm, requiring an f/5.3 focusing lens. With some care, this setup could be realized without much difficulty. Unfortunately, this setup gives a maximum number of just 3500 atoms, making observation of the 2-D condensates very difficult.

Alternatively, if the trapped atoms are spin polarized, gravity can be "nullified" by a magnetic field gradient such that $\mu B' = Mg$. For sodium in the $F = 1$ ground state, cancellation of gravity requires a magnetic field gradient of 8 gauss/cm. Maxwell's equations, however, do not allow the creation of a field gradient in just the y dimension (defined to be the dimension in the direction of gravity). As a result, a field gradient in y must necessarily affect confinement in x and z. The variation in the magnetic potential in x and z can be diminished without effecting the field gradient in y, however, by the addition of a large bias field in the y direction. In the case of an axially symmetric field configuration, this would result in a field of the form

$$\mathbf{B} = B_0\hat{y} + B'\left(y\hat{y} - \frac{1}{2}x\hat{x} - \frac{1}{2}z\hat{z}\right)$$

resulting in the magnetic potential (to lowest order in x, y, and z) of

$$U_B = \mu\left(B_0 + B'y\right) + \frac{1}{2}m\omega^2\left(x^2 + z^2\right)$$

where $\omega \equiv \sqrt{\frac{\mu B'^2}{4B_0 M}}$ is the harmonic oscillator frequency associated with the magnetic potential. The effect of the magnetic potential on the confinement of the atoms can be characterized by the magnetic potential's harmonic oscillator length $a_{HO} = \sqrt{\frac{\hbar}{M\omega}}$. In the example above, in order to produce a sodium condensate with 10^5 atoms, a condensate was discussed which had a length of 630 μm. In order to assure that the harmonic oscillator length of the magnetic potential is greater than this length, an incredibly large bias field of about 4000 tesla is required!

Bibliography

[1] M.H. Anderson, J.R. Ensher, M.R. Matthews, C.E. Wieman, and E.A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269:198, 1995.

[2] 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(22):3969–3973, November 1995.

[3] C.C. Bradley, C.A. Sackett, and R.G. Hulet. Bose-Einstein condensation of lithium: Observation of limited condensate number. *Phys. Rev. Lett.*, 78:985, 1997.

[4] A. Einstein. Zur quantentheorie des idealen gases. *Sitzungsber. Preuss. Akad. Wiss.*, Bericht 3:18, 1925.

[5] K. Huang. *Statistical Mechanics*. Wiley, New York, 1987.

[6] Charles Kittel and Herbert Kroemer. *Thermal Physics*. W. H. Freeman and Company, New York, 2nd edition, 1980.

[7] H. Kamerlingh Onnes. Further experiments with liquid helium. *Konink. Akad. Wetensch. Amsterdam, Proc.*, 13:1093–1113, April 1911.

[8] L.I. Dana and H.K. Onnes. Further experiments with liquid helium. Preliminary determinations of its latent heat of vaporisation and specific heat. *Konink. Akad. Wetensch. Amsterdam, Proc.*, 29(8):1051–1068, 1926.

[9] W.H. Keesom and M. Wolfke. The two different liquid states of helium. *Comptes Rendus*, 185:1465–1467, December 1927.

[10] W.H. Keesom and M. Wolfke. Two liquid states of helium. *Konink. Akad. Wetensch. Amsterdam, Proc.*, 31(190b):90–94, 1928.

[11] E.F. Burton. Viscosity of helium I and helium II. *Nature*, 135:265, February 1935.

[12] J.F. Allen and A.D. Misener. Flow of liquid helium II. *Nature*, 141(3558):75, January 1938.

[13] P. Kapitza. Viscosity of liquid helium below the λ-point. *Nature*, 141(3558):74, January 1938.

[14] J.F. Allen, R. Peierls, and M. Zaki Uddin. Heat conduction in liquid helium. *Nature*, 140:62–63, 1937.

[15] J.F. Allen and H. Jones. New phenomena connected with heat flow in helium II. *Nature*, 141(3562):243–244, February 1938.

[16] F. London. The λ-phenomenon of liquid helium and the Bose-Einstein degeneracy. *Nature*, 141(3571):643–644, April 1938.

[17] R.J. Donnelly. *Experimental Superfluidity*. University of Chicago Press, Chicago and London, 1967.

[18] David R. Tilley and John Tilley. *Superfluidity and Superconductivity*. John Wiley and Sons, New York, 1974.

[19] William E. Keller. *Helium-3 and Helium-4*. Plenum Press, New York, 1969.

[20] J. Bardeen, L.N. Cooper, and J.R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108(5):1175–1204, December 1957.

[21] D.D. Osheroff, R.C. Richardson, and D.M. Lee. Evidence for a new phase of solid He 3. *Phys. Rev. Lett.*, 28(14):885–888, April 1972.

[22] D.D. Osheroff, R.C. Richardson, and D.M. Lee. Evidence for a new phase in solid He 3. In *Bulletin of the American Physical Society*, January 1972.

[23] L.L. Chase and N. Peyghambarian. Evidence for Bose-Einstein condensation of biexcitons in CuCl. *Phys. Rev. Lett.*, 42(18):1231–1234, April 1979.

[24] N. Peyghambarian and L.L. Chase. Bose-Einstein statistical properties and condensation of excitonic molecules in CuCl. *Phys. Rev. B*, 27(4):2325–2345, February 1983.

[25] J. Kogut and J.T. Manassah. π^- condensation and neutron star cooling. *Phys. Lett.*, 41A(2):129–131, September 1972.

[26] D.B. Kaplan and A.E. Nelson. Strange goings on in dense nucleonic matter. *Phys. Lett. B*, 175(1):57–63, July 1986.

[27] D.B. Kaplan and A.E. Nelson. Errata: Strange goings on in dense nucleonic matter. *Phys. Lett. B*, 179(4):409, October 1986.

[28] J. Ambjørn, M. Laursen, and M.E. Shaposhnikov. Baryon asymmetry of the universe. A Monte Carlo study on the lattice. *Phys. Lett. B*, 197(1,2):49–54, October 1987.

[29] Laszlo Tisza. The theory of liquid helium. *Phys. Rev.*, 72(9):838–854, November 1947.

[30] N. Bogoliubov. On the theory of superfluidity. *J. Phys.*, XI(1):23–32, 1947.

[31] Oliver Penrose and Lars Onsager. Bose-Einstein condensation and liquid helium. *Phys. Rev.*, 104(3):576–584, November 1956.

[32] S.T. Beliaev. Application of the methods of quantum field theory to a system of bosons. *Soviet Phys. JETP*, 34(2):289–299, August 1958.

[33] S.T. Beliaev. Energy-spectrum of a non-ideal Bose gas. *Soviet Phys. JETP*, 34(2):299–307, August 1958.

[34] N.M. Hugenholtz and D. Pines. Ground-state energy and excitation spectrum of a system of interacting bosons. *Phys. Rev*, 116(3):489–506, November 1959.

[35] Kerson Huang and C.N. Yang. Quantum-mechanical many-body problem with hard-sphere interaction. *Phys. Rev.*, 105(3):767–775, February 1957.

[36] T.D. Lee, Kerson Huang, and C.N. Yang. Eigenvalues and eigenfunctions of a Bose system of hard spheres and its low-temperature properties. *Phys. Rev.*, 106(6):1135–1145, June 1957.

[37] T.D. Lee and C.N. Yang. Many body problem in quantum mechanics and quantum statistical mechanics. *Phys. Rev.*, 105:1119–1120, 1957.

[38] Kerson Huang, C.N. Yang, and J.M. Luttinger. Imperfect Bose gas with hard-sphere interaction. *Phys. Rev.*, 105(3):776–784, February 1957.

[39] T.D. Lee and C.N. Yang. Low-temperature behavior of a dilute Bose system of hard spheres. I. Equilibrium properties. *Phys. Rev.*, 112(5):1419–1429, December 1958.

[40] T.D. Lee and C.N. Yang. Low-temperature behavior of a dilute Bose system of hard spheres. II. Nonequilibrium properties. *Phys. Rev.*, 113(6):1406–1412, March 1959.

[41] Charles E. Hecht. The possible superfluid behaviour of hydrogen atom gases and liquids. *Physica*, 25:1159–1161, 1959.

[42] W.C. Stwalley and L.H. Nosanow. Possible "new" quantum systems. *Phys. Rev. Lett.*, 36(15):910–913, April 1976.

[43] Dale G. Fried, Thomas C. Killian, Lorenz Willmann, David Landhuis, Stephen C. Moss, Daniel Kleppner, and Thomas J. Greytak. Bose-Einstein condensation of atomic hydrogen. *Phys. Rev. Lett.*, 81(18):3811–3814, November 1998.

[44] T.J. Greytak. Prospects for Bose-Einstein condensation in magnetically trapped atomic hydrogen. In A. Griffin, D.W. Snoke, and S. Stringari, editors, *Bose-Einstein Condensation*. Cambridge University Press, 1995.

[45] D. Kleppner, T.J. Greytak, T.C. Killian, D.G. Fried, L. Willmann, D. Landhuis, and S.C. Moss. Bose-Einstein condensation of atomic hydrogen. In *Bose-Einstein Condensation* [186].

[46] E. Arimondo, W.D. Phillips, and F. Strumia, editors. *Proceedings of the International School of Physics "Enrico Fermi"*, number CXVIII, Laser Manipulation of Atoms and Ions. Italian Physical Society, North-Holland, 1992.

[47] W. Ketterle, D.S. Durfee, and D.M. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In *Bose-Einstein Condensation* [186].

[48] H.-J. Miesner, D.M. Stamper-Kurn, M.R. Andrews, D.S. Durfee, S. Inouye, and W. Ketterle. Bosonic stimulation in the formation of a Bose-Einstein condensate. *Science*, 279:1005, 1998.

[49] W. Ketterle and N.J. van Druten. Bose-Einstein condensation of a finite number of particles trapped in one or three dimensions. *Phys. Rev. A*, 54:656, 1996.

[50] D.S. Hall, M.R. Matthews, J.R. Ensher, C.E. Wieman, and E.A. Cornell. The dynamics of component separation in a binary mixture of Bose-Einstein condensates. *Phys. Rev. Lett.*, 81:4531, 1998.

[51] J. Stenger, S. Inouye, D.M. Stamper-Kurn, H.-J. Miesner, A.P. Chikkatur, and W. Ketterle. Spin domains in ground state spinor Bose-Einstein condensates. *Nature*, 396:345, 1998.

[52] M.-O. Mewes. *Bose-Einstein Condensation of Sodium Atoms*. PhD thesis, Massachusetts Institute of Technology, February 1997.

[53] M.-O. Mewes, M.R. Andrews, N.J. van Druten, D.M. Kurn, D.S. Durfee, and W. Ketterle. Bose-Einstein condensation in a tightly confining DC magnetic trap. *Phys. Rev. Lett.*, 77(3):416–419, July 1996.

[54] M.-O. Mewes, M.R. Andrews, N.J. van Druten, D.M. Kurn, D.S. Durfee, C.G. Townsend, and W. Ketterle. Collective excitations of a Bose-Einstein condensate in a magnetic trap. *Phys. Rev. Lett.*, 77(6):988–991, August 1996.

[55] M.R. Andrews, M.-O. Mewes, N.J. van Druten, D.S. Durfee, D.M. Kurn, and W. Ketterle. Direct, non-destructive observation of a Bose condensate. *Science*, 273:84, 1996.

[56] D.S. Durfee and W. Ketterle. Experimental studies of Bose-Einstein condensation. *Optics Express*, 2:299, 1998.

[57] M.-O. Mewes, M.R. Andrews, D.M. Kurn, D.S. Durfee, C.G. Townsend, and W. Ketterle. Output coupler for Bose-Einstein condensed atoms. *Phys. Rev. Lett.*, 78(4):582–585, January 1997.

[58] 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, January 1997.

[59] C.G. Townsend, N.J. van Druten, M.R. Andrews, D.S. Durfee, D.M. Kurn, M.-O. Mewes, and W. Ketterle. *Ultracold Atoms and Bose-Einstein-Condensation, 1996*, page 2. OSA Trends in Optics and Photonics Series, Vol. 7. Optical Society of America, Washington D.C., 1996.

[60] M.R. Andrews, D.M. Kurn, H.-J. Miesner, D.S. Durfee, C.G. Townsend, S. Inouye, and W. Ketterle. Propagation of sound in a Bose-Einstein condensate. *Phys. Rev. Lett.*, 79(4):553–556, July 1997.

[61] M.R. Andrews, D.M. Stamper-Kurn, H.-J. Miesner, D.S. Durfee, C.G. Townsend, S. Inouye, and W. Ketterle. Erratum: Propagation of sound in a Bose-Einstein condensate. *Phys. Rev. Lett.*, 80(13):2967, March 1998.

[62] Jinha Kim, Bretislav Friedrich, Daniel P. Katz, David Patterson, Jonathan D. Weinstein, Rober DeCarvalho, and John M. Doyle. Buffer-gas loading and magnetic trapping of atomic europium. *Phys. Rev. Lett.*, 78(19):3665–3668, May 1997.

[63] D.M. Stamper-Kurn, M.R. Andrews, A.P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle. Optical confinement of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 80:2072, 1998.

[64] C.J. Myatt. *Bose-Einstein Condensation Experiments in a Dilute Vapor of Rubidium*. PhD thesis, University of Colorado, Boulder, Colorado, 1997.

[65] E. Cornell, J.R. Ensher, and C.E. Wieman. Experiments in dilute atomic Bose-Einstein condensation. In *Bose-Einstein Condensation* [186].

[66] Steve Chu. The manipulation of neutral particles. *Rev Mod Phys*, 70(3), July 1998. (Nobel Lecture).

[67] Claude N. Cohen-Tannoudji. Manipulating atoms with photons. *Rev Mod Phys*, 70(3), July 1998. (Nobel Lecture).

[68] William D. Phillips. Laser cooling and trapping of neutral atoms. *Rev Mod Phys*, 70(3), July 1998. (Nobel Lecture).

[69] W.D. Phillips. Laser cooling and trapping of neutral atoms. In Arimondo et al. [46], pages 289–343.

[70] S. Chu. Laser cooling and manipulation of atoms, and selected applications. In Arimondo et al. [46], pages 239–287.

[71] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji. Laser cooling below the one-photon-recoil energy by velocity-selective coherent population trapping. *Phys. Rev. Lett.*, 61:826, 1988.

[72] J. Lawall, S. Kulin, B. Saubamea, N. Bigelow, M. Leduc, and C. Cohen-Tannoudji. Three-dimensional laser cooling of helium beyond the single-photon recoil limit. *Phys. Rev. Lett.*, 75:4194, 1995.

[73] N.F. Ramsey. Beams of atoms and molecules. In Arimondo et al. [46], pages 1–28.

[74] Lene Vestergaard Hau and J.A. Golovchenko. A new atomic beam source: The "candlestick". *Rev. Sci. Instrum.*, 65(12):3746–3750, December 1994.

[75] Michael Arnold Joffe. *Trapping and Cooling of Atoms at High Densities*. PhD thesis, Massachusetts Institute of Technology, November 1993.

[76] Kendall Bruce Davis. *Evaporative Cooling of Sodium Atoms*. PhD thesis, Massachusetts Institute of Technology, May 1995.

[77] T. Walker, D. Sesko, and C. Wieman. Collective behavior of optically trapped neutral atoms. *Phys. Rev. Lett.*, 64:408, 1990.

[78] Wolfgang Ketterle, Kendall B. Davis, Michael A. Joffe, Alex Martin, and David E. Pritchard. High densities of cold atoms in a dark spontaneous-force optical trap. *Phys. Rev. Lett.*, 70(15):2253–2256, April 1993.

[79] W.H. Wing. On neutral particle trapping in quasistatic electromagnetic fields. *Prog. Quant. Electr.*, 8:181, 1984.

[80] W. Ketterle and D.E. Pritchard. Trapping and focusing ground state atoms with static fields. *Appl. Phys. B*, 54:403, 1992.

[81] E. Majorana. Atomi orientati in campo magnetico variabile (atomic orientation in variable magnetic fields). *Nuovo Cimento*, 9:43, 1932.

[82] T. Bergeman, G. Erez, and H. Metcalf. Magnetostatic trapping fields for neutral atoms. *Phys. Rev. A*, 35:1535, 1987.

[83] D.E. Pritchard. Cooling neutral atoms in a magnetic trap for precision spectroscopy. *Phys. Rev. Lett.*, 51:1336, 1983.

[84] V.S. Bagnato, G.P. Lafyatis, A.G. Martin, E.L. Raab, R.N. Ahmad-Bitar, and D.E. Pritchard. Continuous stopping and trapping of neutral atoms. *Phys. Rev. Lett.*, 58:2194, 1987.

[85] Y.V. Gott, M.S. Ioffe, and V.G. Tel'kovskii. *Nuclear Fusion Supplement*, 3:1045, 1962.

[86] C.J. Myatt, E.A. Burt, R.W. Ghrist, E.A. Cornell, and C.E. Wieman. Production of two overlapping Bose-Einstein condensates by sympathetic cooling. *Phys. Rev. Lett.*, 78:586, 1997.

[87] T. Esslinger, I. Bloch, and T.W. Hänsch. Bose-Einstein condensation in a quadrupole-Ioffe-configuration trap. *Phys. Rev. A*, 58:R2664, 1998.

[88] J. Söding, D. Guéry-Odelin, P. Desbiolles, G. Ferrari, and J. Dalibard. Giant spin relaxation of an ultracold cesium gas. *Phys. Rev. Lett.*, 80:1869, 1998.

[89] L.V. Hau, B.D. Busch, C. Liu, Z. Dutton, M.M. Burns, and J.A. Golovchenko. Near resonant spatial images of confined Bose-Einstein condensates in the '4D' magnetic bottle. *Phys. Rev. A*, 58:R54, 1998.

[90] U. Ernst, A. Marte, F. Schreck, J. Schuster, and G. Rempe. Bose-Einstein condensation in a pure Ioffe-Pritchard field configuration. *Europhys. Lett.*, 41:1, 1998.

[91] L.D. Landau and E.M. Lifshitz. *Fluid Mechanics*. Addison-Wesley, Massachusetts, 1959.

[92] Wolfgang Ketterle and N.J. Van Druten. Evaporative cooling of trapped atoms. In *Advances in Atomic, Molecular, and Optical Physics*, volume 37, pages 181–236. Academic Press, San Diego, 1996.

[93] Charles S. Adams, Heun Jin Lee, Nir Davidson, Mark Kasevich, and Steven Chu. Evaporative cooling in a crossed dipole trap. *Phys. Rev. Lett.*, 74(18):3577–3580, May 1995.

[94] Nir Davidson, Heun Jin Lee, Charles S. Adams, Mark Kasevich, and Steven Chu. Long atomic coherence times in an optical dipole trap. *Phys. Rev. Lett.*, 74(8):1311–1314, February 1995.

[95] S. Inouye, M.R. Andrews, J. Stenger, H.-J. Miesner, D.M. Stamper-Kurn, and W. Ketterle. Observation of Feshbach resonances in a Bose-Einstein condensate. *Nature*, 392:151, 1998.

[96] D.M. Stamper-Kurn, H.-J. Miesner, A.P. Chikkatur, S. Inouye, J. Stenger, , and W. Ketterle. Quantum tunneling across spin domains in a Bose-Einstein condensate. *Phys. Rev. Lett.*, 83(4):661–665, July 1999.

[97] J. Stenger, D.M. Stamper-Kurn, M.R. Andrews, A.P. Chikkatur, S. Inouye, H.-J. Miesner, and W. Ketterle. Optically confined Bose-Einstein condensates. *J. Low Temp. Phys.*, 113:167, 1998.

[98] *Melles Griot Catalog*, 1995/96. (Introduction to Gaussian Beam Optics).

[99] Yu.B. Ovchinnikov, I. Manek, and A.I. Sidorov. Gravito-optical atom trap based on a conical hollow beam. *Europhys. Lett.*, 43(55):510–515, 1998.

[100] B.P. Anderson and M.A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, 282:1686, 1998.

[101] M.R. Andrews. *Bose Condensates and the Atom Laser*. PhD thesis, Massachusetts Institute of Technology, 1998.

[102] Eric L. Bolda and Dan F. Walls. Detection of vorticity in Bose-Einstein condensed gases by matter-wave interference. *Phys. Rev. Lett.*, 81(25):5477–5480, December 1998.

[103] M. Lewenstein and L. You. Quantum phase diffusion of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 77:3489, 1996.

[104] Y. Castin and J. Dalibard. Relative phase of two Bose-Einstein condensates. *Phys. Rev. A*, 55:4330, 1997.

[105] E.M. Wright, D.F. Walls, and J.C. Garrison. Collapses and revivals of Bose-Einstein condensates formed in small atomic samples. *Phys. Rev. Lett.*, 77:2158, 1996.

[106] K. Mølmer. Phase collapse and excitations in Bose-Einstein condensates. *Phys. Rev. A*, 58:566, 1998.

[107] J. Javanainen and M. Wilkens. Phase and phase diffusion of a split Bose-Einstein condensate. *Phys. Rev. Lett.*, 78:4675, 1997.

[108] D.S. Hall, M.R. Matthews, C.E. Wieman, and E.A. Cornell. Measurements of relative phase in two-component Bose-Einstein condensates. *Phys. Rev. Lett.*, 81(8):1543–1546, August 1998.

[109] B.P. Anderson and M.A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, 282:1686–1689, November 1998.

[110] H.-J. Miesner and W. Ketterle. Bose-Einstein condensation in dilute atomic gases. *Solid State Comm.*, 107:629, 1998.

[111] David L. Feder, Charles W. Clark, and Barry I. Schneider. Vortex stability of interacting Bose-Einstein condensates confined in anisotropic harmonic traps. *Phys. Rev. Lett.*, 82(25):4956–4959, June 1999.

[112] Tomoya Isoshima and Kazushige Machida. Vortex stabilization in Bose-Einstein condensate of alkali-metal atom gas. *Phys. Rev. A*, 59(3):2203–2212, March 1999.

[113] S. Stringari. Phase diagram of quantized vortices in a trapped Bose-Einstein condensed gas. *Phys. Rev. Lett.*, 82(22):4371–4375, May 1999.

[114] D.S. Jin, J.R. Ensher, M.R. Matthews, C.E. Wieman, and E.A. Cornell. Collective excitations of a Bose-Einstein condensate in a dilute gas. *Phys. Rev. Lett.*, 77(3):420–423, July 1996.

[115] S. Stringari. Collective excitations of a trapped Bose-condensed gas. *Phys. Rev. Lett.*, 77(12):2360–2363, September 1996.

[116] S. Stringari. Clusters of quantum liquids. In G. Scoles, editor, *The Chemical Physics of Atomic and Molecular Clusters*, number CVII, The Chemical Physics of Atomic and Molecular Clusters. Italian Physical Society, North-Holland, 1990.

[117] S.A. Chin and E. Krotschek. Surface excitations of helium droplets. *Phys. Rev. Lett.*, 74(7):1143–1146, February 1995.

[118] F. Dalfovo, S. Giorgini, L.P. Pitaevskii, and S. Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys*, 71:463, 1999.

[119] Francesca Zambelli and Sandro Stringari. Quantized vortices and collective oscillations of a trapped Bose-Einstein condensate. *Phys. Rev. Lett.*, 81(9):1754–1757, August 1998.

[120] C. Raman, M. Köhl, R. Onofrio, D.S. Durfee, C.E. Kuklewicz, Z. Hadzibabic, and W. Ketterle. Evidence for a critical velocity in a Bose-Einstein condensed gas. 1999. (to be published).

[121] L. Landau. The theory of superfluidity of helium II. *J. Phys.*, V(1):71–90, 1941. (U.S.S.R.).

[122] L. Landau. On the theory of superfluidity of helium II. *J. Phys.*, XI(1):91–92, 1947. (U.S.S.R.).

[123] C.G. Kuper. Can a Landau critical velocity be observed? pages 231–240, London and New York, 1966. Academic Press.

[124] R.P. Feynman. Application of quantum mechanics to liquid helium. In C.J. Gorter, editor, *Progress in Low Temperature Physics*, volume 1, Amsterdam, 1955. North-Holland Publishing Company.

[125] Russell J. Donnelly. *Quantized vortices in helium II.* Cambridge University Press, Cambridge, 1991.

[126] D.V. Osborne. The rotation of liquid helium II. *Proc. Roy. Soc.*, A63:909–912, 1950.

[127] David Pines. Richard Feynman and condensed matter physics. *Physics Today*, pages 61–66, February 1989.

[128] L.P. Pitaevskii. Vortex lines in an imperfect Bose gas. *Soviet Phys. JETP*, 13(2):451–454, August 1961.

[129] W.F. Vinen. The detection of single quanta of circulation in liquid helium II. *Proc. Roy. Soc.*, A260:218–236, March 1961. (London).

[130] J.B. Mehl and W. Zimmermann. Gyroscopic detection of persistent flow of superfluid liquid helium. *Phys. Rev. Lett.*, 14(20):815–818, May 1965.

[131] G.W. Rayfield and F. Reif. Evidence for the creation and motion of quantized vortex rings in superfluid helium. *Phys. Rev. Lett.*, 11(7):305–308, October 1963.

[132] G.W. Rayfield and F. Reif. Quantized vortex rings in superfluid helium. *Phys. Rev.*, 136(5A):A1194–A1208, November 1964.

[133] D. R. Allum, P.V.E. McClintock, A. Phillips, and R.M. Bowley. The breakdown of superfluidity in liquid ^{4}He: An experimental test of Landau's theory. *Phil. Trans. R. Soc. A*, 284:179–224, 1977.

[134] T. Ellis, C.I. Jewell, and P.V.E. McClintock. Measurement of the Landau velocity in He II. *Phys. Lett.*, 78A(4):358–360, August 1980.

[135] T. Frisch, Y. Pomeau, and S. Rica. Transition to dissipation in a model of superflow. *Phys. Rev. Lett.*, 69(11):1644–1647, September 1992.

[136] Emil Lundh, C.J. Pethick, and H. Smith. Zero-temperature properties of a trapped Bose-condensed gas: Beyond the Thomas-Fermi approximation. *Phys. Rev. A*, 55(3):2126–2131, March 1997.

[137] R.E. Drullinger, D.J. Claze, and D.S. Sullivan. A recirculating oven for atomic beam frequency standards. In *Proceedings of the 39th Annual Frequency Symposium*, May 1985.

[138] E.A. Cornell, C. Monroe, and C.E. Wieman. Multiply loaded, ac magnetic trap for neutral atoms. *Phys. Rev. Lett.*, 67:2439, 1991.

[139] R.J.C. Spreeuw, C. Gerz, L.S. Goldner, W.D. Phillips, S.L. Rolston, C.I. Westbrook, M.W. Reynolds, and I.F. Silvera. Demonstration of neutral atom trapping with microwaves. *Phys. Rev. Lett.*, 72:3162, 1994.

[140] Y.B. Ovchinnikov, I. Manek, and R. Grimm. Surface trap for cs atoms based on evanescent-wave cooling. *Phys. Rev. Lett.*, 79:2225, 1997.

[141] H.J. Lee and S. Chu. Atomic spin polarization by Raman cooling. *Phys. Rev. A*, 57:2905, 1998.

[142] Y.B. Ovchinnikov, I. Manek, A.I. Sidorov, G. Wasik, and R. Grimm. Gravito-optical atom trap based on a conical hollow beam. *Europhys. Lett.*, 43:510, 1998.

[143] D. Boiron, A. Michaud, J.M. Fournier, L. Simard, M. Sprenger, G. Grynberg, and C. Salomon. Cold and dense cesium clouds in far-detuned dipole traps. *Phys. Rev. A*, 57:R4106, 1998.

[144] T. Takekoshi and R.J. Knize. CO_2 laser trap for cesium atoms. *Optics Lett.*, 21:77, 1995.

[145] S. Friebel, C. D'Andrea, J. Walz, M. Weitz, and T.W. Hänsch. CO_2 optical lattice with cold rubidium atoms. *Phys. Rev. A*, 57:R20, 1998.

[146] T.A. Savard, K.M. O'Hara, and J.E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Phys. Rev. A*, 56(2):R1095–R1098, August 1997.

[147] C.C. Agosta, I.F. Silvera, H.T.C. Stoof, and B.J. Verhaar. Trapping of neutral atoms with resonant microwave radiation. *Phys. Rev. Lett.*, 62:2361, 1989.

[148] R.V.E. Lovelace, C. Mahanian, T.J. Tommila, and D.M. Lee. Magnetic confinement of a neutral gas. *Nature*, 318:30, 1985.

[149] E. Riis and S.M. Barnett. A dynamic electric trap for ground-state atoms. *Europhys. Lett.*, 21:533, 1993.

[150] F. Shimizu and M. Morinaga. Electric trapping of neutral atoms. *Jpn. J. Appl. Phys.*, 31:L1721, 1992.

[151] N.R. Newbury and C.E. Wieman. Resource letter TNA-1: Trapping of neutral atoms. *Am. J. Phys.*, 64:18, 1996.

[152] A.L. Migdall, J.V. Prodan, W.D. Phillips, T.H. Bergeman, and H.J. Metcalf. First observation of magnetically trapped neutral atoms. *Phys. Rev. Lett.*, 54:2596, 1985.

[153] H. Hess, G.P. Kochanski, J.M. Doyle, N. Masuhara, D. Kleppner, and T.J. Greytak. Magnetic trapping of spin-polarized atomic hydrogen. *Phys. Rev. Lett.*, 59:672, 1987.

[154] R. van Roijen, J.J. Berkhout, S. Jaakkola, and J.T.M. Walraven. Experiments with atomic hydrogen in a magnetic trapping field. *Phys. Rev. Lett.*, 61:931, 1988.

[155] W.D. Phillips, J.V. Prodan, and H.J. Metcalf. Laser cooling and electromagnetic trapping of neutral atoms. *J. Opt. Soc. Am. B*, 2:1751, 1985.

[156] J.T.M. Walraven. *Quantum Dynamics of Simple Systems*, page 315. Institute of Physics Publ., London, 1996.

[157] K.-J. Kügler, K. Moritz, W. Paul, and U. Trinks. Nestor - a magnetic storage ring for slow neutrons. *Nuclear Instruments and Methods in Physics Research*, 228:240, 1985.

[158] K.B. Davis, M.O. Mewes, M.A. Joffe, and W. Ketterle. In *The Fourteenth International Conference on Atomic Physics (book of abstracts)*. 1M-3.

[159] W. Petrich, M.H. Anderson, J.R. Ensher, and E.A. Cornell. In *The Fourteenth International Conference on Atomic Physics (book of abstracts)*. 1M-7.

[160] J. Schwinger. On nonadiabatic processes in inhomogeneous fields. *Phys. Rev.*, 51:648, 1937.

[161] T.H. Bergeman, P. McNicholl, J. Kycia, H. Metcalf, and N.L. Balazs. Quantized motion of atoms in a quadrupole magnetostatic trap. *J. Opt. Soc. Am. B*, 6:2249, 1989.

[162] K.B. Davis, M.-O. Mewes, M.A. Joffe, M.R. Andrews, and W. Ketterle. Evaporative cooling of sodium atoms. *Phys. Rev. Lett.*, 74:5202, 1995.

[163] W. Petrich, M.H. Anderson, J.R. Ensher, and E.A. Cornell. A stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, 74:3352, 1995.

[164] V.E. Shapiro. Rotating magnetic quadrupole field traps for neutral atoms. *Phys. Rev. A*, 54:R1018, 1996.

[165] J.J. Tollet, C.C. Bradley, C.A. Sackett, and R.G. Hulet. Permanent magnet trap for cold atoms. *Phys. Rev. A*, 51:R22, 1995.

[166] B. Desruelle, V. Boyer, P. Bouyer, G. Birkl, M. Lecrivain, F. Alves, C.I. Westbrook, and A. Aspect. Trapping cold neutral atoms with an iron-core electromagnet. *European Physical Journal D*, 1:255, 1998.

[167] P.W.H. Pinkse, A. Mosk, M. Weidemüller, M.W. Reynolds, T.W. Hijmans, and J.T.M. Walraven. Adiabatically changing the phase-space density of a trapped Bose gas. *Phys. Rev. Lett.*, 78:990, 1997.

[168] C.V. Sukumar and D.M. Brink. Spin-flip transitions in a magnetic trap. *Phys. Rev. A*, 56:2451, 1997.

[169] H.F. Hess. Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen. *Phys. Rev. B*, 34:3476, 1986.

[170] N. Masuhara, J.M. Doyle, J.C. Sandberg, D. Kleppner, T.J. Greytak, H.F. Hess, and G.P. Kochanski. Evaporative cooling of spin-polarized atomic hydrogen. *Phys. Rev. Lett.*, 61:935, 1988.

[171] J.M. Doyle, J.C. Sandberg, I.A. Yu, C.L. Cesar, D. Kleppner, and T.J. Greytak. Hydrogen in the submillikelvin regime: Sticking probability on superfluid ^{4}He. *Phys. Rev. Lett.*, 67:603, 1991.

[172] D.E. Pritchard, K. Helmerson, and A.G. Martin. Atomic physics. volume 11, page 179, Singapore, 1989. World Scientific.

[173] T.W. Hijmans, O.J. Luiten, I.D. Setija, and J.T.M. Walraven. Optical cooling of atomic hydrogen in a magnetic trap. *J. Opt. Soc. Am. B*, 6:2235, 1989.

[174] W. Ketterle, K.B. Davis, M.A. Joffe, A. Martin, and D.E. Pritchard, 1993. OSA Annual Meeting, Toronto, Canada, October 3-8.

[175] C.C. Bradley, C.A. Sackett, J.J. Tollet, and R.G. Hulet. Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions. *Phys. Rev. Lett.*, 75:1687, 1995.

[176] H. Wu and C.J. Foot. Direct simulation of evaporative cooling. *J. Phys. B*, 29:L321, 1996.

[177] M. Holland, J. Williams, K. Coakley, and J. Cooper. Trajectory simulation of kinetic equations for classical systems. *Quantum and Semiclassical Optics*, 8:571, 1996.

[178] H. Wu, E. Arimondo, and C.J. Foot. Dynamics of evaporative cooling for Bose-Einstein condensation. *Phys. Rev. A*, 56:560, 1997.

[179] C.A. Sackett, C.C. Bradley, and R.G. Hulet. Optimization of evaporative cooling. *Phys. Rev. A*, 55:3797, 1997.

[180] E. Arimondo, E. Cerboneschi, and H. Wu. (to be published). In *Bose-Einstein Condensation* [186].

[181] P.W.H. Pinkse, A. Mosk, W. Weidmuller, M.W. Reynolds, and T.W. Hijmans. One-dimensional evaporative cooling of magnetically trapped atomic hydrogen. *Phys. Rev. A*, 57:4747, 1998.

[182] B.P. Anderson and M.A. Kasevich. Spatial observation of Bose-Einstein condensation of ^{87}Rb in a confining potential. *Phys. Rev. A*, 59(2):R938–R941, February 1999.

[183] W.D. Phillips. (to be published). In *Bose-Einstein Condensation* [186].

[184] D.J. Han, R.H. Wynar, P. Courteille, and D.J. Heinzen. Bose-Einstein condensation of large numbers of atoms in a magnetic time-averaged orbiting potential trap. *Phys. Rev. A*, 57:R4114, 1998.

[185] D.M. Stamper-Kurn, H.-J. Miesner, A.P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle. Reversible formation of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 81:2194, 1998.

[186] Italian Physical Society. *Proceedings of the International School of Physics "Enrico Fermi"*, Bose-Einstein Condensation 1999 (in press).

[187] C.A. Sackett, C.C. Bradley, M. Welling, and R.G. Hulet. Bose-Einstein condensation of lithium. *Bra. J. Phys.*, 27:154, 1997.

[188] H.D. Politzer. Bose-stimulated scattering off a cold atom trap. *Phys. Rev. A*, 55:1140, 1997.

[189] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-Photon Interactions*. Wiley, New York, 1992.

[190] O. Morice, Y. Castin, and J. Dalibard. Refractive index of a dilute Bose gas. *Phys. Rev. A*, 51:3896, 1995.

[191] L. Allen and J.H. Eberly. *Optical Resonance and Two-Level Atoms*. Dover Publications, New York, 1975.

[192] M.A. Joffe, W. Ketterle, A. Martin, and D.E. Pritchard. Transverse cooling and deflection of an atomic beam inside a Zeeman slower. *J. Opt. Soc. Am. B*, 10:2257, 1993.

[193] O.J. Luiten, H.G.C. Werij, I.D. Setija, M.W. Reynolds, T.W. Hijmans, and J.T.M. Walraven. Lyman-a spectroscopy of magnetically trapped atomic hydrogen. *Phys. Rev. Lett.*, 70:544, 1993.