

Contact Measurements on a Strongly Interacting Bose Gas

by

Robert Johannes Wild

B.S. Physics, University of Arizona, 2004

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

2012

This thesis entitled:
Contact Measurements on a Strongly Interacting Bose Gas
written by Robert Johannes Wild
has been approved for the Department of Physics

Dr. Eric A. Cornell

Dr. Deborah S. Jin

Date _____

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

Wild, Robert Johannes (Ph.D., Physics)

Contact Measurements on a Strongly Interacting Bose Gas

Thesis directed by Dr. Eric A. Cornell

A powerful set of universal relations, centered on a quantity called the contact, connects the strength of short-range two-body correlations to the thermodynamics of a many-body system with delta-function interactions. For bosons, the fact that contact spectroscopy can be used to probe the gas on short timescales is potentially useful given the decreasing stability of BECs with increasing interactions. Successfully measuring the contact requires careful control of experimental parameters such as the magnetic field and the RF probe pulse. In this thesis I report on measurements of the contact, using RF spectroscopy, for an ^{85}Rb atomic Bose-Einstein condensate (BEC). The measured contact exhibits beyond-mean-field behavior, the degree of which is dependent on the rates of change of the scattering length. A potential complication is the added possibility, for bosons, of three-body interactions. In investigating this issue, we have located an Efimov resonance for ^{85}Rb atoms with loss measurements and thus determined the three-body interaction parameter.

Acknowledgements

As with most academic undertakings, this thesis is a result of much collaboration. Even though it is just my name on the cover, I could not have produced this work without all the scientific, technical, theoretical, emotional, and moral support that I received over the years. I am deeply grateful for all of it, which has affected not only this thesis but also the person I have become.

First and foremost I'd like to thank my two advisors, Eric Cornell and Debbie Jin. I am constantly amazed by their physical insights and am grateful to have had the opportunity to work with them. From them I learned how to be scientifically rigorous and to always question what I think I know. My improved reasoning skills are the most valuable result of my work with them, and in spite of their scientific prowess, they always managed to be friendly and easy-going. I am glad to have been able to build great personal relationships with both.

I would also like to thank my undergraduate advisor at the University of Arizona, Alex Cronin, who showed me how to work in an “unstructured” lab and instilled confidence in my abilities, without restricting my freedom to participate in my numerous extracurricular activities.

Thanks also to John Obrecht and Scott Papp, who taught me much about ultracold atom experiments in my first few years at CU. A special thanks goes to Juan Pino and Phil Makotyn, with whom I spent the bulk of my time in the lab (in which I spent the bulk of my time in general). Juan, Phil, and I worked together for several years, both in pairs and as a team of three, and we now share that special bond that comes from sharing the frustrations and successes of experimental physics and the many late nights in the lab. Our relationships consisted mainly of making fun of

each other at every opportunity, often with two ganging up against the third, and I'm glad to have had them to help make the work a bit more fun. Cathy Klauss joined just months before I left, but I am grateful for the overlap. She is not only intelligent, but quickly learned to hold her own in the laboratory banter.

The tri/bi-group and all of its members over the years were also extremely helpful. They were always available to answer those "dumb" questions you don't want to ask your advisors, and I am glad to have become friends with so many of them.

I also want to thank the various undergrads with whom I worked. Rachel Umbel for beam profiling work, Sophie Letournel for building a very good laser for the lab, Zach Newman for his work on a crossed optical trap, and Chris Poulton for all the various excellent work he has already done and continues to do in the lab.

Of course, JILA has been a fantastic environment to perform research, and I want to thank all the technical help I have received from the machine shop, the electronics shop (especially Carl Sauer), the computing team, and the rest of the staff that make the place run smoothly and allowed me to focus on the experiment.

A special thanks to Brandon Peden, roommate for five years and great friend, with whom I learned how to brew beer, make cheese, and enjoy scotch, and who philosophized with me over topics ranging from math to women to everything in between, not being afraid to discuss absurd hypothetical situations and all their ramifications.

I would also like to thank my family for all their support. Their unwavering belief in my abilities was an immense help throughout the years of my studies, especially when my own belief was not quite as steadfast.

Finally, I would like to thank Colleen Cummings, who for the last five years has given me so much love and support, especially through those worst of times, that I am forever indebted. I can only hope to amply return what she has given me, and am ever grateful to have her by my side.

Contents

Chapter	
1	Introduction 1
1.1	Universality 1
1.2	Strongly correlated systems 3
1.3	Contact spectroscopy 4
1.4	Technical challenges and thesis contents 7
2	Experimental Basics 9
2.1	^{85}Rb and its Feshbach resonance 9
2.2	The road to BEC 10
3	Magnetic Field 15
3.1	Gradient cancellation 15
3.2	Magnetic-field stabilities 19
3.2.1	60 Hz AC line noise 20
3.2.2	Field monitoring 21
3.2.3	Magnetic-field-induced atomic transitions 24
4	Radiofrequency pulses 25
4.1	Gaussian pulses 25
4.2	Shape checks 27
4.3	Rectification 31

5	Detection	34
5.1	Condensate absorption imaging	34
5.1.1	Imaging corrections	34
5.1.2	BEC collapse and high-field imaging	38
5.1.3	High-intensity imaging	41
5.2	In-trap imaging for contact measurements	45
5.2.1	Transfer to the imaging state	45
5.2.2	Small cloud effects	50
6	Two-body Contact	54
6.1	What is the contact?	54
6.2	The contact signal and its characteristics	58
6.3	Extracting the contact	60
6.4	Contact measurements	65
7	Three-body Contact	72
7.1	The Efimov resonance	74
7.2	Searching for C_3	80
8	Summary and Outlook	85
8.1	In Essence	85
8.2	The next few weeks...	86
	Bibliography	88
	Appendix	
A	Effective range contribution	92

B Detailed procedure for the contact measurements

Figures

Figure

- 1.1 Sketch of RF transition rate Γ . (a) Shows the ideal noninteracting spectrum, where the whole population is resonant at a single frequency. (b) Correlations brought about by interactions give rise to a $\omega^{-3/2}$ tail, shown by the red line, as well as a small shift of the resonance frequency. (c) The experimental broadening of the lineshape obscures this shift. Nearly all the atoms are within the area of the resonant lineshape, which gives us a measure of the single atom Rabi frequency. We use different methods to probe the resonant lineshape and the tail regimes, and the final contact measurement involves a ratio of these excitation rates. The magnitude of the tail has been greatly exaggerated for illustrative purposes. 6
- 2.1 The scattering length in units of the Bohr radius a_0 as a function of magnetic field. We create our cold atom clouds on the high-field side of the resonance to optimize collisional properties. From there, we can easily tune the value of the scattering length to either large positive or negative values. 11
- 2.2 A schematic of our evaporation. The optical trap ensures overlap between the ^{85}Rb (red points) and the ^{87}Rb (blue points), and is centered on the ^{85}Rb magnetic equilibrium position. Lowering its trap depth evaporates ^{87}Rb . The ^{87}Rb cools the ^{85}Rb via collisions until all the ^{87}Rb has fallen out, and only ^{85}Rb remains in a purely magnetic trap. 14

3.1 Our scheme to create a uniform field. We start with a BEC in a purely magnetic trap, whose center is shifted upward by a set of large shim coils. We then quickly ramp down the coils, causing the minimum of the magnetic field to move to the location of the atoms. A small shim coil fine tunes this location. 17

3.2 Measuring the magnetic-field gradient. We create a large thermal cloud (pink circle with dashed outline) and perform RF spectroscopy on a magnetically sensitive transition (in this case the $|2, -2\rangle$ to $|2, -1\rangle$ transition). As long as the spectral width of the pulse is smaller than the cloud, different parts of the cloud are resonant at different fields, and will outcouple at different positions in space. The resulting slope can be converted to a magnetic-field gradient. 18

3.3 Correlating noise to the 60 Hz AC line. Part (a) shows how we correlated the noise to the AC line. (a) Taking data on the side of a transition and using the slope allowed us to convert a scatter in signal to a scatter in B-field. Correlating this scatter to the phase of the AC line at the time we applied the RF pulse shows a modulation in agreement with the 60 Hz line. A histogram (c) of the scatter reveals the characteristic double-peaked shape of sinusoidally distributed values. This structure essentially disappears in (d), where we have synchronized the experiment to the AC line 1 second before the B-field measurement. The bin size and axis range of (c) and (d) are identical. 22

3.4 The B-field resonance giving rise to 2-photon transitions. Part (a) shows Fourier transforms of oscilloscope traces taken with a hand held pickup coil close to the trap. The blue curve shows the problem-causing 40 kHz feature. Changing the gain of the B-field servos allowed us to eliminate the feature, as illustrated by the green curve. A schematic of the 2-photon transition is shown in (b). 24

- 4.1 Some RF pulses in frequency space. In (a) we see the power spectrum of two standard pulses (inset). A square pulse in time becomes a Sinc^2 function, which looks much like a Gaussian down to the first minimum, but then exhibits significant wings far from the center (red curve). A Gaussian in time, however, becomes a Gaussian in power (black curve). The curves have been scaled to give similar peaks and widths. In (b) we see the same curves on a logarithmic scale, which makes the difference clear. An $\omega^{-3/2}$ line has been added (blue), to show that the problematic wings of a Sinc^2 function persist with detuning. 26
- 4.2 The Fourier transform of the RF power at system saturation. The labels correspond to the output power of the synthesizer before all amplification. At 9 dBm output we begin to risk damage to the amplifiers. The dashed black line shows a perfect Gaussian for reference (a parabola on this scale). We see that saturation does not affect the nicely Gaussian shape of the pulse down to at least 60 dB below the peak power. 28
- 4.3 The Fourier transform of truncated Gaussians. The inset shows the shape of the truncated Gaussians, at ± 1 (blue), 2 (green), and 4τ (black). The Fourier transform of the power shows the resulting frequency spectrum. The $\pm 4\tau$ truncation looks unaffected, whereas the $\pm 1\tau$ truncation clearly begins to look similar to a square pulse. 29
- 4.4 Atoms spin-flipped by truncated Gaussians. A truncated pulse is less efficient at transferring atoms between magnetic sublevels. The pulse still exhibits lower efficiency at 2τ , but is fully Gaussian at 4τ . The theory line is scaled to the final value. 30

4.5	Effect of RF rectification. (a) We send out a $\tau = 100 \mu\text{s}$ Gaussian pulse, and probe the resulting field with a $50 \mu\text{s}$ square pulse. The $100 \mu\text{s}$ delay gives the servo time to respond and shift the field. (b) At our maximum allowable power (8 dBm at the synthesizer), we see rectification amounting to 17 mG. At 8 dB lower power (c), the rectification has all but disappeared (d). Our contact measurements required powers no more than -9 dBm.	33
5.1	A measurement of our effective saturation intensity. We measure the OD vs I_i in counts per pixel, and then invert the axes for plotting and fitting via Equation 5.3. The red line shows this fit. The measured OD is highly sensitive to the probe intensity, especially at low intensities, where one might think that the correction due to I/I_{sat} is negligible.	36
5.2	Comparison between measured and predicted BEC size. The different shapes correspond to different ARP efficiencies, giving rise to different ODs. The red line is the modeled BEC size in expansion.	39
5.3	Effect of photon re-absorption on OD. The black circles show the measured OD of a BEC, calculated with Equation 5.6, as a function of I/I_{sat} . For intensities near I_{sat} , the high scattering rate in conjunction with the high OD causes a large fraction of atoms to re-absorb photons emitted from surrounding atoms. This makes them unavailable to absorb from the probe beam, lowering the apparent OD. As the intensity increases, the scattering rate remains roughly constant, translating to a smaller fractional amount of rescattered light present in the cloud. The red squares show the same experiment on thermal clouds with lower OD, where this effect does not occur.	44

- 5.4 State diagram showing transfer to the imaging state. The atoms (shown in blue) start in the $|2, -2\rangle$ state, and the RF pulse for contact spectroscopy moves a small fraction to the $|2, -1\rangle$ state. These are transferred via shaped ARP to the $|3, -2\rangle$ state and RF π pulse to the $|3, -3\rangle$ state (solid green arrows). Once there, the probe beam cycles them to the $F'=4$ manifold (dotted green arrow). The entire procedure, from the $|2, -1\rangle$ state to imaging, takes $< 200 \mu\text{s}$. The transition of the cloud remainder to the $|3, -1\rangle$ state (red dashed arrow) is unwanted and would produce significant false signal. The energy splittings are given in MHz, calculated for a field of 160 G. 47
- 5.5 Shaped ARPs. A standard ARP is shown in part (a), where the RF power and therefore the energies of the dressed states are constant in time. The frequency must start far from resonance and end far from resonance for a well behaved ARP. A shaped ARP (or ShARP) seen in (b) starts at low detuning and low power. The power is adiabatically increased as the frequency ramps through resonance and then decreased the same way. This allows for a quick, efficient ARP, sweeping over a very small frequency range around the resonance. 48
- 5.6 Simulation of pixelation effects. An absorption profile is incident on a 2-dimensional array of pixels, and the resulting pixelated OD, the sum over the pixels Σ_{px} , is compared to the integral, I_G , of the OD that gave rise to the absorption profile. Part (a) shows the two configurations under which the simulation was run. In one configuration, the Gaussian profile was centered on the center of a pixel in both dimensions. The corner configuration formed the other “extreme”. Simulation results can be seen in part (b). When the Thomas-Fermi (TF) radius is roughly the size of a pixel, the effects can become very pronounced, even exhibiting unintuitive structure. The error increases linearly with OD, but stays at a 3% or less correction for our normal operating conditions of peak $\text{OD} < 1$ and a TF radius of 2.4 pixels. 51

- 5.7 Schematic of the pixelsum technique. Shown is a standard image of a BEC with 30% thermal component, taken in the magnetic trap. For the pixelsum analysis, we sum all the values of the inner square, and subtract the background as determined by the “sidewalk” (hatched area). 53
- 6.1 Schematic of the interparticle potential (a) and scattering wavefunction (b). The potential $V(r)$ looks like an attractive potential down to a distance given by the van der Waals potential. Outside of this effective range r_e , which is ignored in the zero-range limit, the wavefunction can be universally described by the scattering length a 55
- 6.2 RF contact spectroscopy for a ^{85}Rb BEC. (a) Schematic of the contact transition. We drive the interacting state into a lower magnetic sublevel, but the interaction-induced energy requires less energy from the RF photon to effect the spin-flip. This results in the transition frequency of the contact to be lower than the single atom transition. (b) Example of spectroscopy signal $S(\omega)$, normalized so that $\int_{-\infty}^{\infty} S(\omega)d\omega = 1 \text{ s}^{-1}$. (c) The same signal as (b), shown on a smaller scale. The solid red line is a fit to the expected frequency dependence from Equation 6.6, while the dotted blue line shows a fit ignoring $\beta(\omega)$. The green line shows the expected Gaussian signal from the resonant lineshape. On the positive side of the transition the signal is consistent with zero. (d) The tail signal multiplied by $|\omega|^{3/2}$. This is how the fits to the tail were performed, and more clearly shows the final-state effects. The density here is $\langle n \rangle = 5.8 \times 10^{12} \text{ cm}^{-3}$. The data from the tail and main lineshape come from different experimental runs, but are shown together for illustration purposes. 59

- 6.3 Expansion of the outcoupled atom cloud. (a) The size of the cloud of outcoupled atoms after the contact pulse with $|\omega| = 2\pi \times 50$ kHz detuning. The black line shows a sketch of a truncated Gaussian RF pulse for reference. The cloud expands with a kinetic energy of $\frac{1}{2}\hbar\omega$, as the excess energy is shared between two atoms. The red line is the predicted size due to the kinetic energy, added in quadrature with the resolution limit of our imaging system. (b) The peak OD of the clouds plotted in part (a). The red line is a fit to a $\frac{1}{\omega^2}$ dependence ($\sim \frac{1}{f^2}$), varying only an overall scaling factor. (c) The energy of the outcoupled cloud as a function of detuning, at 4.5 ms time-of-flight expansion. We calculate the energy from the size of the outcoupled cloud, accounting for the size of the cloud of non-spin-flipped atoms at $\omega = 0$. The solid line is $\frac{1}{2} \frac{|\omega|}{2\pi}$ 61
- 6.4 Deviation from signal linearity. On both the peak (a) and tail (b) of the transition, we measure the outcoupled fraction f_{meas} of atoms vs RF pulse time and/or power, and plot it vs the fraction expected f_{exp} if the dependence were linear. We fit this to the simple saturation model of Equation 6.13 to extract a value for the asymptote f_{ayasm} and calculate the magnitude of the correction to the data. In (b) we combine data from two different scattering lengths for a larger range in signal. Individual fits to the two sets give the same result within the error bars. 64
- 6.5 The contact vs a , measured at $|\omega| = 2\pi \times 40$ kHz. (a) Here I plot $\frac{\alpha(a)}{\beta(\omega)} \frac{C_2}{N_0}$, which is directly proportional to the strength of the measured $\omega^{-3/2}$ RF tail. (b) The contact per particle $\frac{C_2}{N_0}$. The solid lines in (a) and (b) are the mean-field predictions. The final-state effects shift what is a parabola centered about $a = 0$ in (b) to one centered about $a' = -565 a_0$ in (a), which enhances our signal at small a 66

- 6.6 Contact measurements as a function of the LHY energy E_{LHY} . (a) The contact per particle normalized by density^{1/3}. The red curve shows the mean-field prediction for the contact, and the blue curve shows the total contact including the LHY term. In (b) the data are normalized to the mean-field value of the contact to compare to theory more clearly. In the limit of low interaction strength, the data match the theory but cannot distinguish the LHY contribution. At higher values, the data are systematically low. At values of E_{LHY} approaching unity, we do not expect the perturbative LHY theory to be valid. 67
- 6.7 Contact measurements as a function of ramp rates, expressed in terms of the “adiabaticity parameter” \mathcal{R} . Higher values of this parameter correspond to faster ramps, which show a value of the contact closer to the mean-field prediction. Slower ramps result in higher values closer to the LHY prediction. Different shapes correspond to different days on which the data were taken. The small upper plot is a reproduction of Figure 6.6b, to illustrate the values of E_{LHY} shown here. For lower interaction strength, the LHY term is too small compared to the error bars to discern a ramp rate dependence. To calculate the mean-field and LHY energies, we require knowledge of the cloud density. The change in densities due to the ramp in a were calculated using a model for expansion that included only the expected mean-field energy, and vary from 40% for the slowest ramps to 10% for the fastest. This reasoning is somewhat circular, but including the LHY term in this model causes at most an extra 5% decrease in density for the slowest ramps, and the resulting change is shown by the open circles. 70

- 7.1 A sketch of Efimov states. The vertical scale gives the wave number $K = \pm\sqrt{m|E|/\hbar^2}$, and the horizontal gives $1/a$, such that the origin corresponds to unitarity ($a \rightarrow \infty$). The green line shows the state corresponding to a bound dimer plus a free atom, and the blue lines show the bound trimers, for which there is an infinite series approaching unitarity. Note that the scaling factor here is ~ 2 to make the behavior visible and to match the literature on Efimov states. A realistic Efimov state with scaling factor 22.7 would have an energy that is nearly indistinguishable from the dimer energy on the scale of this plot. 73
- 7.2 Example of a loss rate measurement to extract K_3 . These data correspond to a scattering length of $730 a_0$. Heating of the cloud causes its size to increase as seen in (a). The parameters extracted from a linear fit to the volume go into a fit to the atom number via Equation 7.10, as seen in (b), resulting in a value of $K_3 = 5.64 \times 10^{-22} \text{cm}^{-6}/\text{s}$ 76
- 7.3 A three-body loss resonance for ^{85}Rb . We plot the three-body event constant K_3 vs the scattering length a , for clouds with a temperature of roughly 8 nK. From fitting Equation 7.11 to the black points, for which $a < 1/k_{\text{thermal}}$, we extract $a_- = -759(6)a_0$ and $\eta = 0.057(2)$ 77
- 7.4 Measurement of η vs temperature. (a) The loss resonances. For the data at 30 and 140 nK, we did not calibrate the scattering length as carefully as the data at 80 nK, which is the same data as in Figure 7.3 shown over a smaller range of a . Moreover, the 30 nK clouds are likely not in thermal equilibrium, and the ensuing error in the average density will cause a systematic error in the calculated value of K_3 . (b) Measured values of η vs temperature. We do not see a significant change in η over a large range in temperature, suggesting that the experiments are performed in the low-temperature limit. 79

- 7.5 G_{RF} as a function of ω , with $\kappa_*=39(1) \mu\text{m}^{-1}$. We plot G_{RF} for negative ω to ease comparison with the measured RF tail, which occurs at negative detunings in our experiment. G_{RF} has a node at $\omega \simeq 2\pi \times 27 \text{ kHz}$, suggesting that one should look for a C_3 contribution to the RF tail for smaller detunings. 81
- 7.6 Changing the scattering length for increased density. We jump the scattering length a (green curve) to $50 a_0$, causing an inward breathe with a period of roughly 50 ms. Close to the turnaround point we ramp to $982 a_0$ and apply the RF pulse to measure the contact (grey line). The blue curve shows the prediction for the relative density of the PG model. 81
- 7.7 (a) The calculated frequency dependence of $G_{\text{RF}}(\omega)$, shown again on a logarithmic ω scale for reference. (b) The measured frequency dependence of the tail of the RF spectrum for $a = 982 \pm 10 a_0$. The solid red line is a fit to the expected frequency dependence of the two-body contact C_2/N_0 including final-state effects. The dotted blue line corresponds to the same value of C_2/N_0 , but ignores final-state effects. For comparison, the fit plus a trial C_3/N_0 term of $0.1 \mu\text{m}^{-2}$ is shown with the dashed black line. Our measurements are consistent instead with a C_3/N_0 of zero. For these data the density is $\langle n \rangle = 1.0 \times 10^{13} \text{ cm}^{-3}$ 82
- 7.8 Searching for a tail in the RF spectrum at $-760 a_0$. We detected no signal indicative of a tail on the RF spectrum. The error bars give an estimate of our detection limit. The green line shows the prediction for the contribution to the signal S from the two-body contact, which is highly suppressed due to final-state effects and expected to be below our detection limit. A heuristic calculation from loss rates via Equation 7.13 predicts a three-body contribution shown by the blue line. The density here is $\langle n \rangle \simeq 7 \times 10^{11} \text{ cm}^{-3}$ 84

- A.1 Effects of the first-order correction due to a finite effective range. (a) and (b) show the fractional corrections introduced by using $\alpha'(a)$ and $\beta'(\omega)$, respectively. For the calculation of $\alpha'(a)$, a detuning of $\omega = 2\pi \times 40$ kHz was assumed. The magnitude of the effect roughly corresponds to the size of the error bars in our data, putting it just below our detection limit. (c) The combined effect on the total contact is on the order of a few percent. 94
- B.1 Timing schematic of a typical contact measurement. The vertical scale is arbitrary. The widths of the grey RF pulses and the red probe pulse have been exaggerated to show up at this scale. The relative timings and the width of the blue RF pulse are to scale. 97

Chapter 1

Introduction

1.1 Universality

Scientific advancement has historically been reductionist, wherein humanity has aspired to explain things by simpler underlying principles. All the complex matter in the world was found to consist of combinations of only a hundred or so different elements. The elements were then found not to be so different after all, as they are merely made of different combinations of the same three particles. And even some of these particles have been broken down to their constituents. However, although this method of understanding the world has brought us extremely far, knowing everything about a constituent part may tell us very little about the whole. For example, detailed knowledge of a single water molecule does not readily lead to the prediction that a collection of water will freeze at a certain temperature, much less that this solid form will actually be less dense than the liquid. Although determined by the constituent parts, such emergent behavior is too complex to predict only using knowledge of the parts. We cannot predict complex chemical behavior if we know everything about the atoms, just as we cannot predict a person's actions by knowing exactly how a neuron works. Instead, we study the different levels of complexity in their own right, with the hope to sometime connect the emergent behavior with our knowledge of the constituent parts.

An example of this is liquid Helium. In 1908, Helium was finally cooled to such low temperatures as to liquefy it [1] (Helium has the lowest boiling point of all gases). The liquefaction itself was expected, since all other gasses behaved similarly. But one can imagine the astonishment that emerged when its superfluid properties were discovered in 1938 [2, 3]. It behaved like no other

liquid, flowing completely without friction and over the sides of any open container. Similarly, some metals exhibit superconductivity when cooled enough, allowing currents to flow through them without resistance (a superfluid of electrons, if you will). These quintessential examples of unexpected emergent behavior have been subject to intense study, but trying to predict this behavior from detailed knowledge of the individual parts proves impractical at best. Instead, we attempt to explain these systems using only a few “universal” properties.

The concept of universality is a powerful tool that allows one to study analogues of complex and diverse systems in simpler, more controlled conditions, and to make theoretical predictions based on the “most important” parameters. The basic idea behind universality is that many systems can be completely described by a small number of parameters and do not depend strongly on details specific to the constituents. For example, much of astronomical behavior can be explained using the universal concept of gravity, where the universal parameter is the mass of an object. Knowing only the masses of objects involved, we can very accurately predict the motions of heavenly bodies regardless of their size, shape, and constitution. Of course, an object’s shape could affect the path of a very nearby mass, so one must make sure that predictions and experiments stay within relevant limits of the universal regime. An atomic example of a universal parameter is the scattering length.

When two atoms interact via low energy s-wave collisions, the scattering length, a , describes how strongly these atoms interact with each other (where a small magnitude scattering length corresponds to weak interactions). The usefulness comes from the fact that details such as the type of atom or particle, or the specific shape of the interatomic potential are irrelevant, and the system of atoms can be completely described by the scattering length. Liquid Helium is not fully in this regime, since its density is so high that the details of the interparticle potential can not necessarily be ignored completely. Nevertheless, much theory has been developed in the universal framework to approach some of the universal physics involved. To realize these universal systems and test the theories, we need low temperature dilute gasses, the difficulty of which has kept the science on a mostly theoretical level through most of the twentieth century. In 1995, however, Bose-Einstein condensation was achieved in dilute gasses [4, 5, 6], making quantum mechanical

many-body systems described by the universal parameter a experimentally accessible.

1.2 Strongly correlated systems

Weakly interacting dilute gases have been well understood in the framework of mean-field theory, where the interactions between particles are characterized simply by the mean-field energy gn , where n is the density of the atoms and $g = \frac{4\pi\hbar a}{2m}$, m being the atomic mass [7]. This provides a highly accurate description of the system in the dilute limit [8], which assumes $na^3 \rightarrow 0$. For interacting Bose gases, however, the interactions give rise to correlations that modify the energy. Systems with strong quantum correlations, such as liquid Helium, represent a frontier in our understanding of the complex quantum systems found in nature. A first step to understanding these correlated systems was taken by Lee, Huang, and Yang in 1957, who found that the leading order correction to the energy density of a BEC is given by [9, 10]

$$\mathcal{E} = \frac{2\pi\hbar^2 an^2}{m} \left(1 + \frac{128}{15\sqrt{\pi}} \sqrt{na^3} + \dots \right), \quad (1.1)$$

where the term proportional to $\sqrt{na^3}$ is often called the LHY correction.

However, it has proven difficult to study strongly interacting atomic BEC, and only a few groups have successfully measured beyond-mean-field effects in these system [11, 12, 13]. The main difficulty is a fundamental one, as increasing a in order to increase the interaction strength brings with it a rapidly increased rate of inelastic three-body collisions [14, 15]. In one of these collisions, two atoms bind to form a molecule, and the third atom recoils with a kinetic energy equal to 2/3 of the binding energy of the molecule for a single-species BEC. This results in the loss of all three atoms from the BEC, as well as heating. The heating comes from collisions, as they leave the cloud, of the energetic particles with other BEC atoms, as well as the fact that loss is density dependent, which causes more loss from colder parts of the cloud (anti-evaporation). Moreover, this three-body recombination rate scales as a^4 , so experiments on these systems either must stay at relatively low interactions strength (done successfully by the groups around Salomon [13] and Hadzibabic [12]), or must probe the system very quickly, before the losses become significant. The latter is our strategy;

our measurement technique uses a short radio-frequency (RF) pulse to measure a quantity called the contact. I will introduce the contact below, and a more thorough treatment will be given in Chapter 6. Previously, our group has probed strongly interacting gasses using Bragg spectroscopy, whereby an excitation is created with a two-photon transition [11]. The appeal of this method was the same; it allowed for a fast probe of the system before three-body losses became significant. However, we were trying to resolve the extra energy provided by the LHY term, meaning that our probe pulses had to be long enough to provide the required energy resolution. The contact is not an energy measurement, so this otherwise fundamental trade-off between probe time and energy resolution is not an issue, as explained below.

1.3 Contact spectroscopy

Measuring the contact through RF spectroscopy is a way to quickly probe the BEC, and is based on measuring short-range correlations in the gas. In 2005, Shina Tan introduced a parameter termed the contact, which quantifies the strength of short-range correlations in an ultracold gas, and showed how this quantity connects quantitatively to macroscopic thermodynamic properties of the many-body system via a set of universal relations [16, 17, 18]. Tan’s predictions have been explored theoretically [19, 20, 21, 22, 23, 24] and experimentally [25, 26] for ultracold Fermi gases.

One of these relations pertains to its measurement technique. The contact manifests itself in the tail of an RF spectrum, which can be thought of as due to RF “dissociation” of pairs of atoms that are close to one another. The rate of transferring atoms to a non-interacting state in this tail is given by

$$\lim_{\omega \rightarrow \infty} \Gamma(\omega) = \frac{\Omega^2}{4\pi} \sqrt{\frac{\hbar}{m}} \frac{\alpha(a)}{\beta(\omega)} \frac{C}{\omega^{3/2}}, \quad (1.2)$$

where the integrated RF lineshape is $\int_{-\infty}^{\infty} \Gamma(\omega) d\omega = \pi \Omega^2 N$, ω is the detuning from the single-atom resonance, and Ω is the single atom Rabi frequency [27]. Further discussion of this will follow in Chapter 6, but the point here is that one can measure the outcoupled number of atoms on the tail of an RF lineshape as a measurement of the contact.

A sketch of how we actually measure the contact is shown in Figure 7.1. The frequency spectrum of a single atom (or of a cloud of noninteracting atoms) looks like a delta function (Figure 7.1a), but interactions give rise to the tail in the frequency spectrum (Figure 7.1b), as parameterized by the contact in Equation 1.2. The interactions also cause a slight shift in the resonant frequency, but to resolve this sub-kHz shift would require long interrogation times on the order of milliseconds. In reality, of course, the single atom resonance is widened by things such as inhomogeneous magnetic fields and finite probe times. This results in a spectrum more akin to Figure 7.1c, where the tail has been greatly exaggerated compared to the resonant lineshape. We use different methods to probe the different regimes of this spectrum. We use a short, spectrally broad pulse to probe the resonant lineshape, which tells us the resonant frequency ω_0 and gives us a measure of the single atom Rabi frequency Ω . Longer, spectrally narrow pulses are used to probe the tail of the transition, giving us the excitation rate $\Gamma(\omega)$.

Another powerful relation states that the contact, C , is connected to the derivative of the total energy of the system, E , with respect to a via the adiabatic sweep theorem [28, 29]:

$$\frac{dE}{da} = \frac{\hbar^2}{8\pi ma^2} C. \quad (1.3)$$

With this, we can rewrite the LHY result in terms of the contact for a condensate as

$$C = 16\pi^2 na^2 \left(1 + \frac{5}{2} \frac{128}{15\sqrt{\pi}} \sqrt{na^3} + \dots \right) N_0, \quad (1.4)$$

where N_0 is the number of atoms in the BEC. This enables us, for instance, to look for deviations from the mean-field prediction for the contact, where the lowest order correction comes from the LHY term of Equation 1.1.

Since this method inherently measures an amplitude instead of a frequency, we can use very short RF pulses to probe the system at timescales short compared to three-body loss rates, even though the spectral broadening of these pulses would be much too wide to measure the change in energy (via Equation 1.1) directly. Of course we are not evading the uncertainty principle, and this method does not actually tell us what the energy is. Instead, it tells us how the energy changes

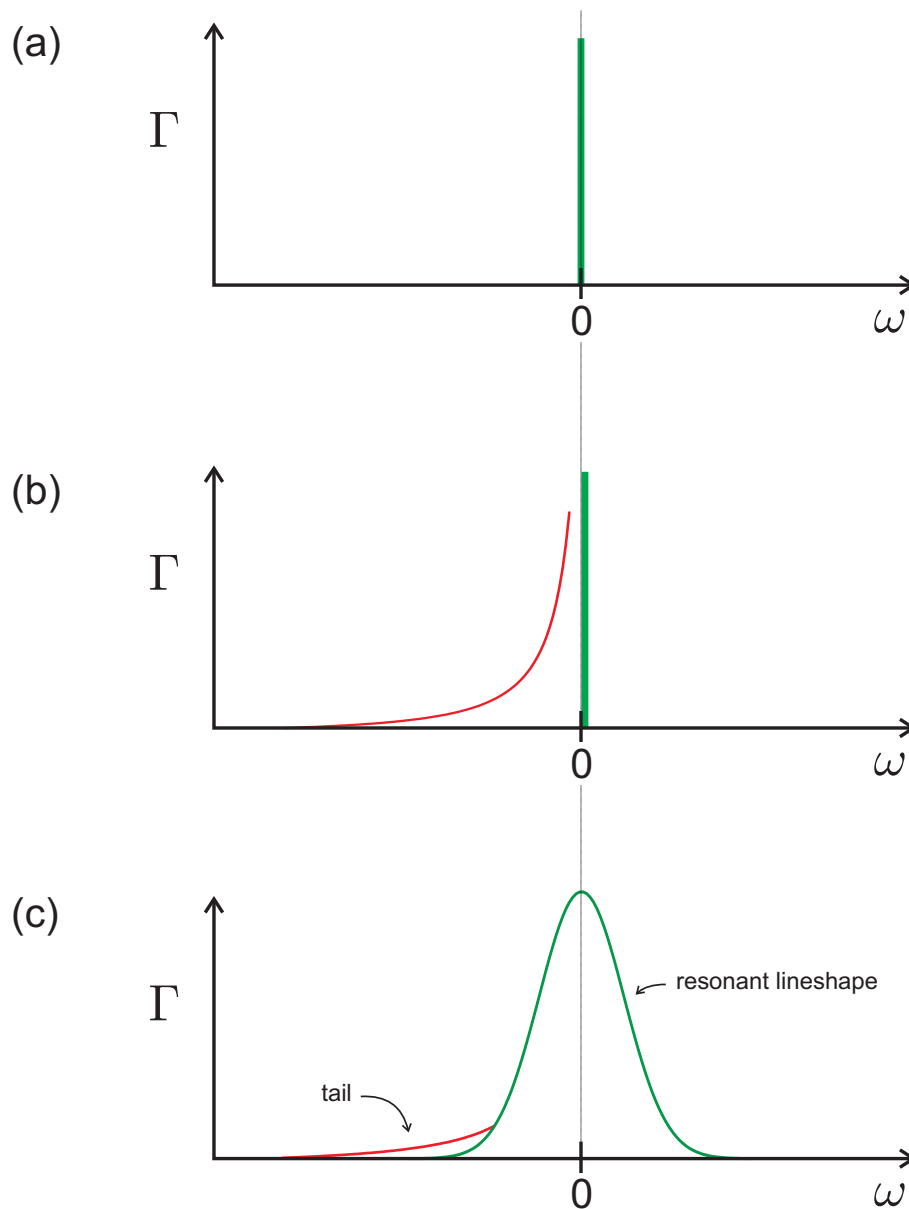


Figure 1.1: Sketch of RF transition rate Γ . (a) Shows the ideal noninteracting spectrum, where the whole population is resonant at a single frequency. (b) Correlations brought about by interactions give rise to a $\omega^{-3/2}$ tail, shown by the red line, as well as a small shift of the resonance frequency. (c) The experimental broadening of the lineshape obscures this shift. Nearly all the atoms are within the area of the resonant lineshape, which gives us a measure of the single atom Rabi frequency. We use different methods to probe the resonant lineshape and the tail regimes, and the final contact measurement involves a ratio of these excitation rates. The magnitude of the tail has been greatly exaggerated for illustrative purposes.

with scattering length (Equation 6.3). This provides a solution to overcome the major hurdle of short lifetimes caused by inelastic losses. Of course, that does not make the experiment easy, and the technical requirements to make this approach work provide many smaller experimental hurdles. Overcoming these hurdles required much work, and the resulting solutions form a large part of this thesis.

1.4 Technical challenges and thesis contents

What are the experimental capabilities then, required to measure the contact? First, we need to be able to create stable condensates, and we must be able to tune the interaction strength to access the mean-field and LHY regimes. We want to do this quickly to minimize three-body losses over the duration of the ramp. When we increase the scattering length on a timescale shorter than the trap period, the BEC will be out of equilibrium. This starts a breathe mode and the density initially drops very quickly along with our signal. Yet working at sub-millisecond times is difficult due to experimental constraints. For example, coil inductances and current-servo bandwidths often cause magnetic fields to ring down before stabilizing. In order to alleviate that less fundamental but still problematic issue, we need a weak spherically symmetric trap to maximize the density oscillation period. These requirements were already met by our BEC machine before we started thinking about the contact, as explained in detail in the thesis of Dr. Juan Pino [30]. I will merely give an overview in Chapter 2, where I will also give a rundown of our cooling process.

To measure the contact via RF spectroscopy, we need the frequency width of our resonant lineshape to be smaller than the detuning at which we are measuring the tail. This requires us to minimize the large magnetic-field gradients present in our magnetic trap. Also, since the magnitude of the signal on the tail is a strong function of the RF detuning ω , we have strict requirements on our trap magnetic-field stability. These requirements and our solutions are discussed in detail in Chapter 3.

The RF pulse must also meet certain requirements. The transition rate on the tail is much smaller than on the resonant lineshape, so great care must be taken to prevent any spectral “wings”

on the pulse from giving us a false signal. For this reason we require Gaussian shaped pulses, and our RF system must in general be well behaved and not have adverse effects on the rest of the experiment. These issues are illuminated in Chapter 4.

Another important element is that of detection. Our weak, spherical trap has the effect that the condensate does not expand much when the trap is turned off, which causes two major challenges for imaging. Even after 30 ms expansion, the optical depth (OD) is still too large to avoid severe saturation problems with standard absorption imaging techniques. Also, the large, negative background scattering length causes the BEC to collapse once the magnetic fields are turned off, and the low expansion energy is not enough to overcome it. Aside from problems associated with imaging the condensate, we must also be careful about our detection schemes for the tail and for the resonant lineshape, so that in the end we can claim an accuracy on the contact measurement to within better than 10%. The imaging and detection techniques and systematics will be described in Chapter 5.

With these technical issues properly delineated, I will talk about the contact in more detail. I will illustrate its origins and limits, and will show data that exhibit beyond mean-field behavior, as well as interesting time-dependent behavior in Chapter 6.

Recently, there has been made some theoretical advancement predicting a three-body contact for bosons, related to Efimov physics [31]. Chapter 7 will expand on this subject, show data on Efimov effects in our system, and show recent developments on the three-body contact.

Chapter 2

Experimental Basics

In this chapter I will give an overview and general background of our experiment. This experiment hinges mostly on the fact that we can produce a Bose-Einstein Condensate (BEC), and that we can tune the interactions of the constituent atoms, which are characterized by the two-body scattering length a . I will describe the various steps and processes we utilize to give us that ability. First, I will describe how we tune a via a Feshbach resonance. Then, I will describe our sympathetic cooling techniques and various trapping schemes that result in a magnetically trapped, spherically symmetric BEC. This chapter is included mostly for the sake of completeness and many details will be left out, as more complete descriptions can be found in the theses of Dr. Scott Papp [32] and Dr. Juan Pino [30].

2.1 ^{85}Rb and its Feshbach resonance

Bose-Einstein Condensation is a hugely interesting phenomenon that has sparked countless studies of basic quantum mechanics (and since they're so countless, I will refrain from filling this page with references to back up my claim). One of the workhorses of these studies has been ^{87}Rb , and by now its cooling and trapping properties are so well developed and documented that one needs a good reason to work with a different atom. This is especially the case for ^{85}Rb , which has such unfortunate scattering properties that it is difficult to cool to degeneracy [32]. Our reason, then, for working with ^{85}Rb is that it has a Feshbach resonance at an experimentally accessible magnetic-field strength of 155 G.

A Feshbach resonance occurs when a two-body bound state is tuned to become degenerate with the atomic scattering threshold. More in-depth explanations can be found here [33, 34], but the point is that the two-body scattering length behaves as

$$a(B) = a_{bg} \left(1 - \frac{\Delta}{B - B_{peak}} \right) \quad (2.1)$$

near the resonance, where for ^{85}Rb , $a_{bg} = -443(3) a_0$ is the background scattering length in units of the Bohr radius, $\Delta = 10.71(2)$ G is the width of the resonance, and $B_{peak} = 155.041(18)$ G is its location in magnetic field B , as measured in [35]. The shape of this resonance can be seen in Fig. 2.1. Thus, by tuning the magnetic field at the position of the atom cloud, we can access a huge range of interaction strengths, allowing us to study condensates beyond the mean-field limit.

2.2 The road to BEC

By now, the production of BEC in alkali atoms alone no longer makes the headlines, and methods for trapping and cooling are omitted from papers to make room for more interesting scientific results. Nevertheless, there can be interesting variations in the standard methods, and the painful truth is that graduate students often spend disproportionately large amounts of time on the optimization of these methods; this warrants a section devoted to the various trapping and cooling stages in our experiment. Again, this is just a quick overview, and more information can be found in [32, 30].

Most of the cooling on ^{85}Rb happens via collisions with colder ^{87}Rb (this process is referred to as sympathetic cooling), since cooling ^{85}Rb by RF evaporation directly has proven difficult due to its collisional properties. Thus, the ^{85}Rb and ^{87}Rb clouds follow much of the same path through the experiment. We start out with a dual-species magneto-optical trap (MOT) to initially cool ^{85}Rb and ^{87}Rb from a room temperature gas. After the MOT stage, we load the atoms into a magnetic trap and transfer the atoms to our science chamber, where the significantly higher vacuum gives us long lifetimes of ~ 500 seconds. This transfer uses pairs of coils that turn on and off in succession,




Figure 2.1: The scattering length in units of the Bohr radius a_0 as a function of magnetic field. We create our cold atom clouds on the high-field side of the resonance to optimize collisional properties. From there, we can easily tune the value of the scattering length to either large positive or negative values.

as well as a moving coil mounted on a track. In the science chamber we load the atoms into a Ioffe-Pritchard (IP) style magnetic trap, and perform standard RF evaporation on ^{87}Rb , thereby also cooling the ^{85}Rb sympathetically. However, since the ^{85}Rb and ^{87}Rb atom have different magnetic moments (in the $|2, -2\rangle$ and $|1, -1\rangle$ states, respectively), the cloud centers are spatially separated due to gravitational sag, and the sympathetic cooling will only be efficient until the sizes of the clouds are on the order of their spatial separation. This occurs at $\sim 10 \mu\text{K}$ and 1 million ^{85}Rb atoms, which is still far from degeneracy for a trap where $\omega_x = 2\pi \times 13$ and $\omega_y = \omega_z = 2\pi \times 210$ Hz, and our cooling procedure requires a change in strategy. At this point things become more interesting, as we have developed a hybrid optical and magnetic trap to continue the sympathetic cooling process and end with a BEC of ^{85}Rb in a 10 Hz spherical trap.

As mentioned in the Introduction, we require a weak, spherical trap to give us more time to probe our strongly interacting BEC. This trap has the effect of increasing the separation between the species' equilibrium positions ($800 \mu\text{m}$), as well as lowering the intraspecies collision rate by nearly a factor of ten (given constant phase-space density). That would increase our evaporation time accordingly, and slow the experiments down to unacceptable rates. To avoid these issues, we change the currents in our various IP and shim coils to create our 10 Hz magnetic trap, and quickly turn on a single-beam optical trap (OT) with a 230 Hz radial confining potential. The atoms are now confined by the OT, ensuring that the two species are in the same location and that collision rates remain high. The tricky part here is to shim the magnetic trap fields in a way such that the ^{85}Rb magnetically trapped equilibrium position is aligned with the optical trap location. A graphical depiction of this trapping scheme can be seen in Figure 2.2. Then we simply decrease the OT intensity down to zero at a rate slow compared the relevant collision rates. The ^{87}Rb experiences normal evaporation, with high energy atoms falling out of the optical trap (we apply RF to then transfer them to an untrapped state), and continuously cools the ^{85}Rb cloud, which stays centered on the OT/magnetic trap throughout the whole procedure. When the OT is completely off, all of the ^{87}Rb atoms have fallen out of the trap, and the ^{85}Rb exhibits a significant condensate fraction, which we purify with some RF evaporation at the end. This evaporation

scheme regularly results in ^{85}Rb condensates of 60,000 atoms, although we have seen condensates upwards of 100,000 atoms.

This method proved to be an effective and reliable scheme to create ^{85}Rb condensates for several reasons. For one, we can be certain that the atoms in our cloud are only ^{85}Rb and only in the $|2, -2\rangle$ state, as the ^{87}Rb or other spin states are trapped at very different positions. The same can not be said for optical dipole traps, for which contamination from other atomic species or spin states can be an issue. Also, since evaporation is performed almost entirely on ^{87}Rb , we only need to load relatively few ^{85}Rb atoms into the MOT at the very beginning. This means that the ^{85}Rb density remains low up until the end of the cooling cycle, suppressing losses of ^{85}Rb due to three-body recombination. Loading fewer ^{85}Rb atoms also means that it presents a smaller heat load for the coolant (^{87}Rb), which means that the same amount of initial ^{87}Rb can reach lower final temperatures. As long as we load just enough ^{85}Rb into the MOT to suppress ^{87}Rb condensation at the end of the cycle, the overlap between the two species is guaranteed and we reliably make ^{85}Rb condensates.

Of course, all methods have their drawbacks. The weakest link in the evaporation is ensuring that the OT is centered on the ^{85}Rb equilibrium position. If the two traps are misaligned, then the ^{85}Rb will slosh out of the OT, causing heating and ending the sympathetic cooling process. To align the traps, we use external coils to shim the magnetic trap position to coincide with the OT position, which roughly coincides with the center of the IP trap. For the vertical direction, we use a pair of coils to apply a uniform vertical field of about 35 G, shifting the magnetic field minimum upward. For the horizontal direction perpendicular to the OT, we use a set of coils that produce a 5 G field at the atoms along the OT, with a gradient perpendicular to the OT, producing a force on the atoms. We tune the magnetic trap rather than the OT position, since in our setup we can control the currents in our shim coils much more precisely than the mirror mounts of our optical trap. Due to experimental drift, this alignment generally has to be done every few weeks. However, with the traps aligned, we often do not need to optimize our system in the morning to achieve our standard BEC conditions, owing to the general reliability of this cooling scheme.

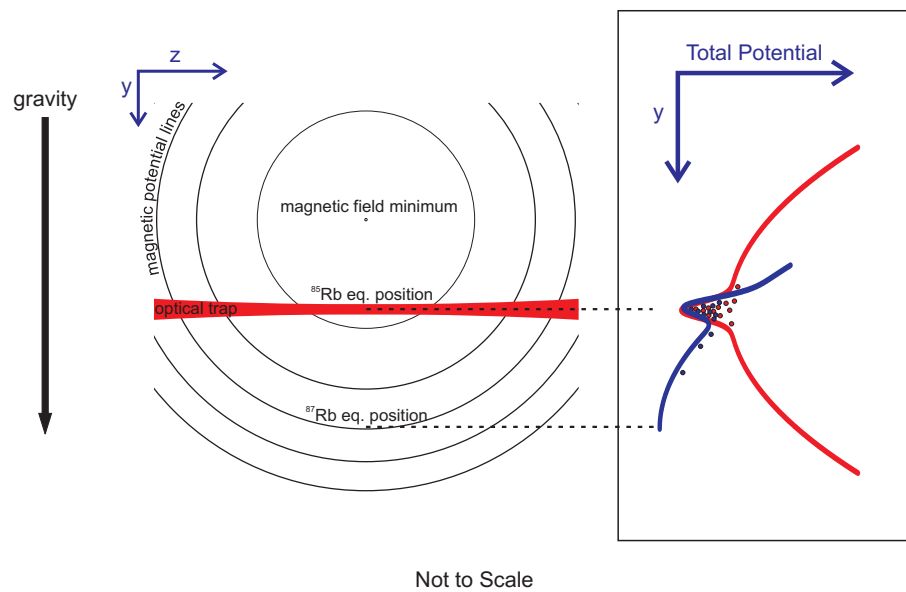


Figure 2.2: A schematic of our evaporation. The optical trap ensures overlap between the ^{85}Rb (red points) and the ^{87}Rb (blue points), and is centered on the ^{85}Rb magnetic equilibrium position. Lowering its trap depth evaporates ^{87}Rb . The ^{87}Rb cools the ^{85}Rb via collisions until all the ^{87}Rb has fallen out, and only ^{85}Rb remains in a purely magnetic trap.

Chapter 3

Magnetic Field

As mentioned in the Introduction, measurement of the contact via RF spectroscopy involves transferring atoms from one magnetic sublevel to another, at frequencies detuned from the center of the resonant lineshape (the single atom resonant frequency) by 10-100 kHz. We require our line width to be much smaller than this detuning, since the transition rate for the tail can be several hundred times smaller than for the resonant lineshape. This puts a significant limit on the width caused by magnetic-field gradients, which cause different parts of the cloud to be resonant at different frequencies. Also, because the contact signal on the tail is a strong function of detuning ($C \sim \omega^{-3/2}$), it is crucial to know the resonant frequency to within < 1 kHz (a 1 kHz uncertainty at 20 kHz detuning results in an 8% uncertainty in the contact). This means that our absolute magnetic-field noise must be less than 2 mG at 155 G, or ~ 12 ppm.

In this chapter I will discuss these problems in two parts. First I will talk about our methods of magnetic gradient cancellation to decrease the RF line width. In the second section I will reveal our techniques for characterizing and stabilizing our magnetic fields, which will be of interest to anyone with magnetic traps and stringent stability requirements.

3.1 Gradient cancellation

One of the problems we were faced with from the beginning was making the width of our resonant lineshape small enough for contact measurements. We needed to probe the tail at low detuning without outcoupling a significant number of atoms due to this width. This problem came

up time and again, and represents one of the main difficulties of RF contact spectroscopy.

One cause of this problem is fundamental to our trap. The fact that we have our BEC in a purely magnetic trap means that we have a strong magnetic-field gradient in the vertical direction to support the atoms against gravity. Specifically, since the atoms are all in the $|F = 2, m_f = -2\rangle$ state, the gradient comes to $\frac{2mg}{3\mu_B} = 22.4 \text{ G/cm}$, where m is the mass and $g = 9.8 \text{ m/s}^2$. For 50,000 atoms in our trap, the $14 \mu\text{m}$ Thomas Fermi radius [36] results roughly in a $14 \mu\text{m} \cdot 22.4 \frac{\text{G}}{\text{cm}} \cdot \frac{\mu_B}{3h} \simeq 15 \text{ kHz}$ rms width on our $|2, -2\rangle$ to $|2, -1\rangle$ transition. We quickly found out that this would not be nearly narrow enough to measure a decent contact signal, so before we even worried about RF pulse widths and field noise and such, we had to find a way to circumvent this “trapping gradient” issue.

Our solution to this problem is in no way unique or even necessarily the best solution. It just turns out that with our system it represented the quickest and easiest way to solve the problem satisfactorily. In Chapter 2 I explained that we have to shim the magnetic trap fields so that the cloud’s sag position overlaps with the optical trap. The optical trap loads from a much tighter 230 Hz magnetic trap, so its vertical position essentially coincides with the magnetic minimum (geometric center) of the Ioffe-Pritchard trap. The exact center, of course, corresponds to a vanishing vertical magnetic gradient. The simple solution, then, is to turn off the large vertical shim coils so that the position of the magnetic minimum once again corresponds to the position of the atoms (see Figure 3.1).

This solution is slightly too simple (as most simple solutions simply are), since the position of the atoms (i.e. the optical trap) does not exactly coincide with the magnetic minimum when the large shims are turned off, leaving us with a 5 G/cm vertical gradient. To cancel this gradient, we then turn on another small vertical shim coil to fine tune the position of the magnetic minimum. One might ask why we do not just use the large coils to do the final shimming, which would be a fine question (and our initial mode of operation). But it turns out that those coils are used to carry high currents and supply strong fields several times in the experiment, and the weak currents and fine adjustments necessary for the final shim were in a regime where the high-current servo fails to

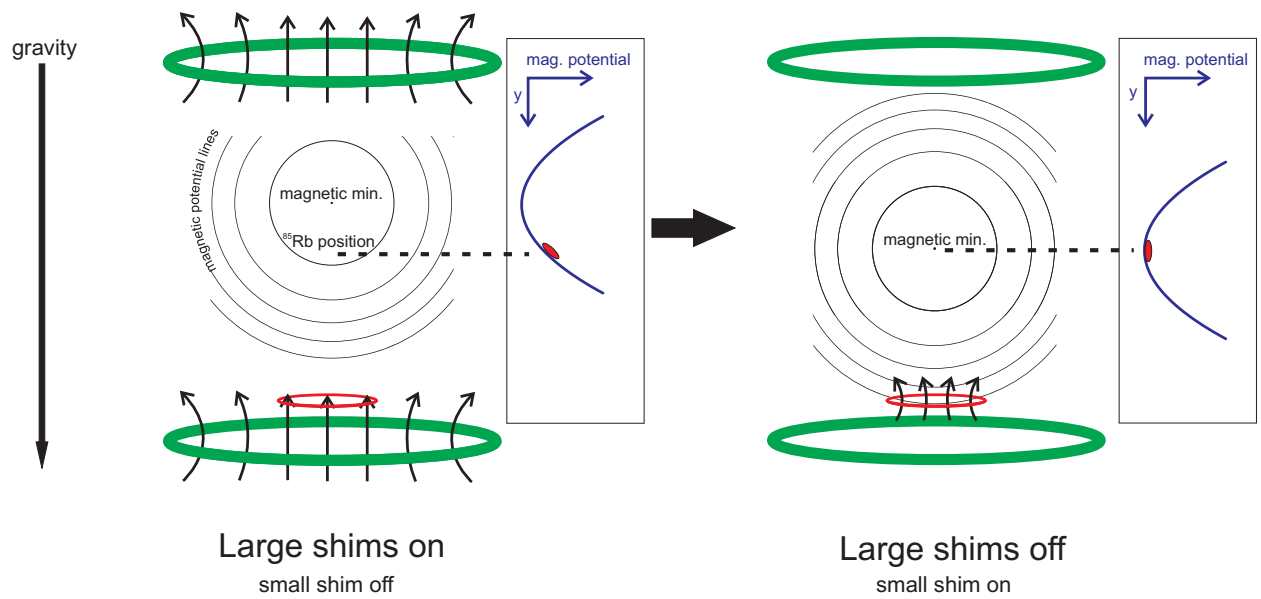


Figure 3.1: Our scheme to create a uniform field. We start with a BEC in a purely magnetic trap, whose center is shifted upward by a set of large shim coils. We then quickly ramp down the coils, causing the minimum of the magnetic field to move to the location of the atoms. A small shim coil fine tunes this location.

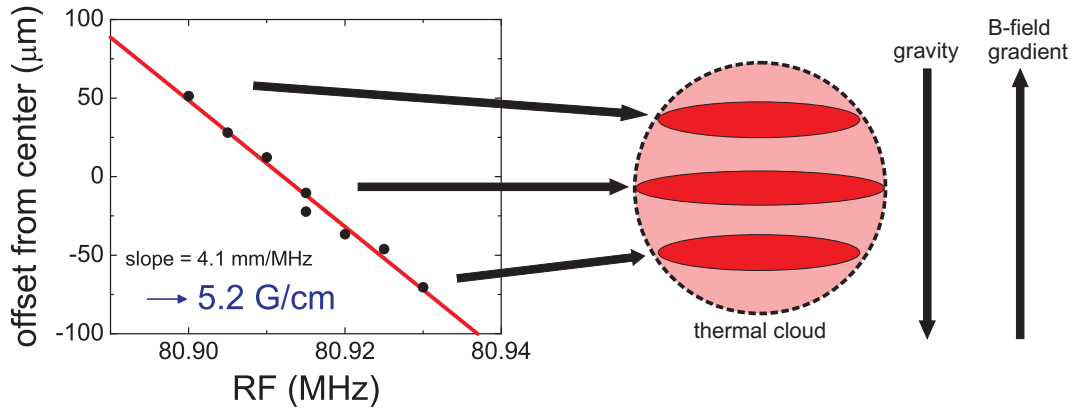


Figure 3.2: Measuring the magnetic-field gradient. We create a large thermal cloud (pink circle with dashed outline) and perform RF spectroscopy on a magnetically sensitive transition (in this case the $|2, -2\rangle$ to $|2, -1\rangle$ transition). As long as the spectral width of the pulse is smaller than the cloud, different parts of the cloud are resonant at different fields, and will outcouple at different positions in space. The resulting slope can be converted to a magnetic-field gradient.

reliably operate. Since we already had a small coil in place to handle smaller currents and produce weaker fields (labeled as the “anti-gravity coil” in Scott Papp’s thesis [32]), it was a simpler solution to use it instead.

This small coil is a 60 turn coil with a 7.6 cm diameter, about 5 cm below the atoms, through which we put about 3 A of current just before we ramp down the large shim coils in 0.5 ms. We tune this current by creating large, thermal clouds of ^{85}Rb and taking an RF spectrum of the atomic resonance. In a gradient, different parts of the cloud will be resonant at different frequencies, so plotting the position of the outcoupled cloud vs frequency gives us a measure of the magnetic-field gradient (see Figure 3.2).

Of course, when we turn off the gradient several things happen. First of all, the cloud begins to fall under gravity, but since the trap is still on and the curvature unchanged, it does not expand like it would if we simply turned the trap off. The falling of the cloud does not pose a problem, since it occurs relatively slowly. In 5 ms of falling, the cloud moves only $120\ \mu\text{m}$ (20 pixels on our camera), and the magnetic field stays constant to within our detection sensitivities. As I explain in Appendix B, we perform all of our experiments within 2-3 ms of turning the gradient off.

Another issue is that we actually change the magnitude of the magnetic field when we ramp off the large shim coils. Those coils create a 26 G magnetic field perpendicular to the trap bias field, which corresponds to an overall decrease of $\sqrt{(155 \text{ G})^2 + (26 \text{ G})^2} - 155 \text{ G} = 2 \text{ G}$ when we turn them off. Since we make BEC near a Feshbach resonance (see Equation 2.1), this will cause the scattering length to jump from $100 a_0$ (where we generally make BEC) to $260 a_0$. For our data in Chapter 6 we simply worked with this, since we then continued ramping the field to higher values anyway. However, for experimental procedures where this ramp is unwanted, we have developed a simultaneous ramp in the bias field that compensates the shim ramp and keeps the field magnitude constant.

This scheme to temporarily create a spatially uniform field works well and rarely requires optimization, but limits us to short times. Perhaps a better and more robust solution would be to work with an optical trap in the first place. To reach similar trapping frequencies and aspect ratios, the obvious setup would be a crossed beam optical trap with a large waist and high power. If this is achieved, then the magnetic fields can be made very uniform with few coils, for as long as is required. For a time we investigated using this setup ourselves, but problems with optical access convinced us to pursue the magnetic route.

3.2 Magnetic-field stabilities

For us to perform trustworthy measurements, the magnetic field (i.e. the resonant transition frequency) must be uniform not only in space but also in time. A typical measurement of the contact (more on that in Chapter 6) involves measuring the resonant transition frequency, and then taking enough data on the tail of the transition (the contact) and the peak of the resonant lineshape (for the Rabi frequency) to achieve fractional error bars no larger than 10%. Any drift in the resonant frequency during the measurements would create a systematic error in the measured contact. So we had two goals for stabilizing our fields. On the one hand, we needed to minimize the shot-to-shot variation, so that we can find the resonance and measure the contact with as few points as possible. On the other hand, we needed to ensure that the field does not drift significantly

over the total time this requires. Our experimental cycle time is about 90 seconds, which sets the former time scale, and we tended to take one half to one hour to complete the measurement, setting the latter.

3.2.1 60 Hz AC line noise

Luckily, we did not have to start from scratch. Scott Papp designed the magnetic trap and controlling servos very well [32], and we already had a stability of about 4 mG shot-to-shot to work with. To improve on this, we first set out to measure the noise from the AC power line. As any table top experimentalist knows, the 60 Hz (50 Hz in Europe) frequency noise is difficult to shield completely, and makes it onto most electronics if you look closely enough. To investigate how much this moved our field, we measured an RF spectrum and took statistics on the side of the gaussian line shape. Assuming the scatter was purely due to the field moving, we could convert the scatter around the expected line to a value of magnetic field for every point. Taking a histogram of these points clearly shows a double-peaked structure, which one would expect for a random sample of a sine wave. We also measured the phase of the AC line for each point, and the oscillation roughly agrees with the expected 60 Hz signal (see Figure 3.3a).

The brute force solution to this problem would be to try to track down the path by which the 60 Hz noise reaches the magnetic trap. However, given that we separately control 5 magnetic field coils with as many servos, and that it would be difficult to suppress the noise without adversely affecting the servo bandwidth, we opted for a more elegant solution. We decided to synchronize the experiment to the AC line such that the RF pulse would always occur at the same phase of the AC line. Our experiment is controlled by a DIO-128 board by Viewpoint Systems with 64 programmable digital output lines, which we load with our timings and output values at the beginning of the experiment, and then trigger to run the experiment off an internal clock. Our change was to set the trigger to wait for the rising edge of the AC line to start its output cycle. However, we found that simply synchronizing the beginning of the experiment to the AC line was insufficient because the phase randomizes too quickly. After just 10 seconds, we measured the phase

noise to be 12% of a full 16.7 ms cycle. Therefore, we synchronize 1 second before the sensitive measurements, where the phase jitter is at a mere 1% of a cycle. This resulted in a shot-to-shot standard deviation of ~ 2 mG with a more Gaussian distribution (Figure 3.3b).

However, there are important issues with this solution, which have to do with experiment timing. Mainly, the extra wait for an AC line trigger introduces timing uncertainties between the times before and after the synchronization. This includes up to 17 ms of hold time from the AC phase, and several hundred milliseconds due to the fact that the DIO board has to load new values for the final part of the experiment. The latter might be circumvented with creative programming and different hardware, but the former is fundamental to the synchronization. This timing uncertainty is a problem when measuring things like cloud oscillations in the trap, which are affected (and usually caused) by various trap changes 10-20 seconds earlier. For this reason we only synchronize the experiment to the AC line when we perform field-sensitive measurements.

3.2.2 Field monitoring

With the synchronization we reached our 2 mG field stability requirement on the 90 second timescale. However, as we were taking numerous frequency spectra throughout the day, we noticed occasional random jumps in the field on the order of 10 mG (5 kHz for the contact transition). If such a jump occurred in the time between measuring the resonant transition frequency and measuring the contact signal on the tail, it would result in a large systematic error in the final value for the contact (20% error for 40 kHz detuning). Therefore we had to start monitoring the fields in order to track down the cause of the jumps.

One monitoring system we use is a $6\frac{1}{2}$ digit multimeter (Agilent 34401A) to record the voltages from the Danfysik current probes that we use to servo the fields. To achieve the necessary precision (40 μ V on the probe correspond to 6 mG in the Bias field), it was required to use the “slow 5 digit” multimeter setting, which averages for 167 ms. Thus, we could only take measurements of the steady-state B-field before our field ramps. However, we did notice occasional jumps in the voltages on our Bias coils, which we could correlate with similar measurements of the servo control

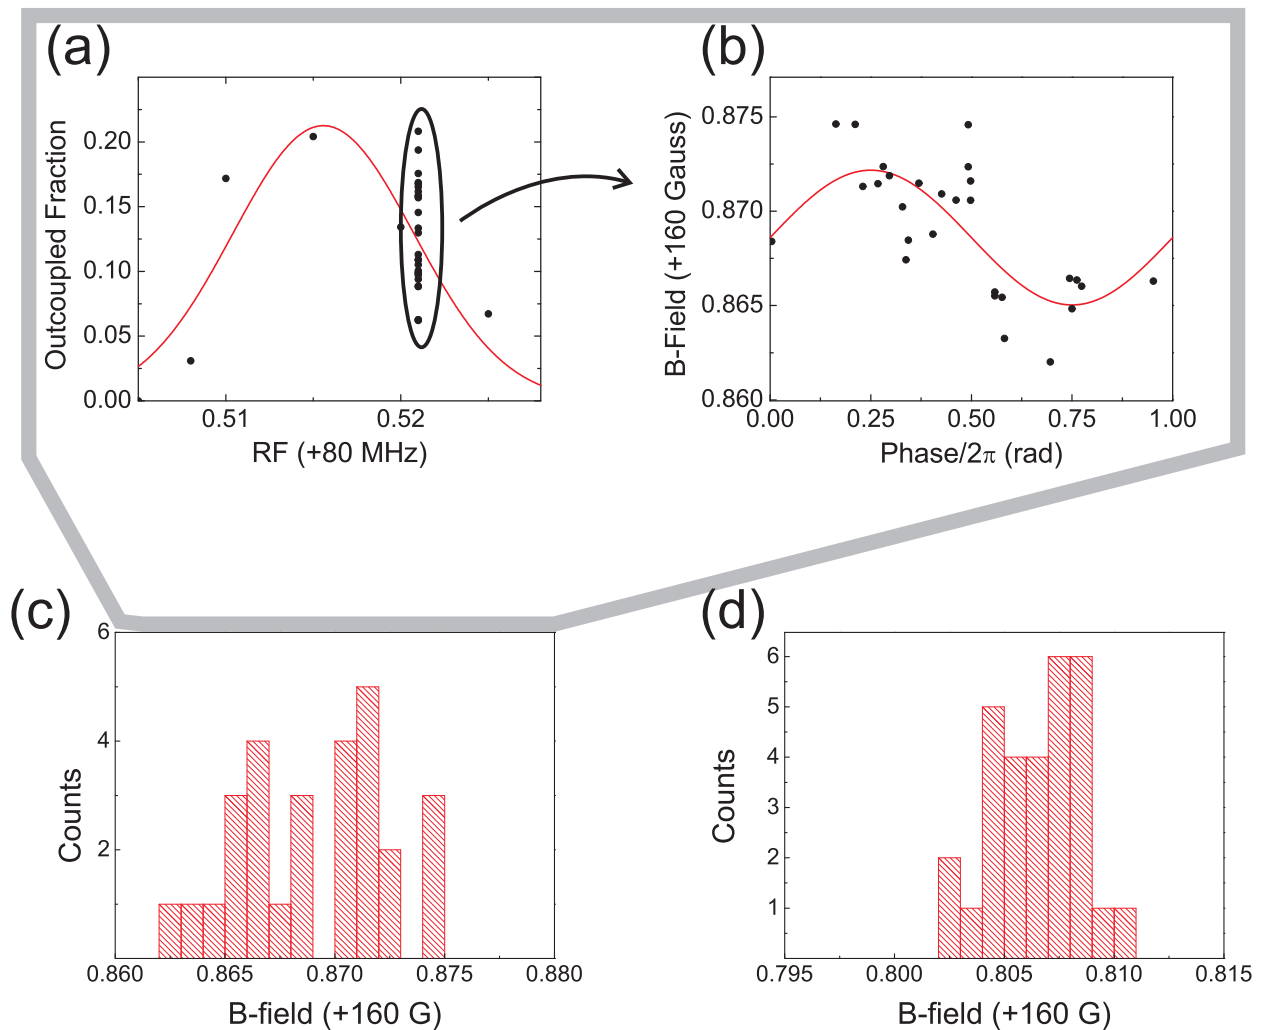


Figure 3.3: Correlating noise to the 60 Hz AC line. Part (a) shows how we correlated the noise to the AC line. (a) Taking data on the side of a transition and using the slope allowed us to convert a scatter in signal to a scatter in B-field. Correlating this scatter to the phase of the AC line at the time we applied the RF pulse shows a modulation in agreement with the 60 Hz line. A histogram (c) of the scatter reveals the characteristic double-peaked shape of sinusoidally distributed values. This structure essentially disappears in (d), where we have synchronized the experiment to the AC line 1 second before the B-field measurement. The bin size and axis range of (c) and (d) are identical.

voltage, which corresponded in magnitude and direction to jumps that we saw in the magnetic field. It turns out that a $100\ \mu\text{V}$ jump in the control voltage caused a $4\ \text{mG}$ jump in field. We found that these tiny voltage jumps were caused mostly by faulty BNC cable connections between our National Instruments PCI-6733 analog output board and the servo control box. Reducing the number of connections, as well as cleaning them with off-the-shelf cleaner for electrical connections essentially solved this problem, and our current probe readings became stable to a B-field equivalent $0.5\ \text{mG}$.

We also monitored external fields, with the worry that there might be changes in the magnetic field environment. To measure external fields, we set up a fluxgate magnetometer roughly 1.5 meters from the position of the magnetic trap, and set it to take and record a measurement of the ambient magnetic field in between each cycle of the experiment, i.e. when all of our magnetic trap coils were turned off. Surprisingly, we noticed and tracked down several effects that could change the magnetic field at the atoms.

The first of these was the freight elevator of the physics building. The elevator is located roughly 13 meters from our experiment and is used numerous times per day mostly for janitorial equipment. We noticed a total change of $1.5\ \text{mG}$, which happened to be parallel and therefore add linearly to our bias field, when we moved the elevator from the basement to the third floor (our lab is on the second floor). It is not enough to cause us problems, but I include it for the interested reader. If the experiment were adjacent to the elevator, the change in field could have been large enough to adversely affect our experiments.

The more relevant findings were related to magnetized metal objects within our laboratory. We found that several items were magnetized enough that moving them toward or away from the 2 meter vicinity of the science cell cause field fluctuations upwards of $5\ \text{mG}$. These included metal stools, trash cans, stepladders, and carts that carried equipment such as oscilloscopes and spectrum analyzers. Also, small magnets attached to large pieces of metal (specifically equipment racks) created significant field variations. Our solution is to avoid placing portable pieces of metal on the side of the room with our science cell, and to simply not move anything in the lab while we take sensitive data.




Figure 3.4: The B-field resonance giving rise to 2-photon transitions. Part (a) shows Fourier transforms of oscilloscope traces taken with a hand held pickup coil close to the trap. The blue curve shows the problem-causing 40 kHz feature. Changing the gain of the B-field servos allowed us to eliminate the feature, as illustrated by the green curve. A schematic of the 2-photon transition is shown in (b).

3.2.3 Magnetic-field-induced atomic transitions

In this last section I will mention an admittedly peculiar problem we encountered in the hopes that it might be interesting (or perhaps even useful) to the reader. Contact spectroscopy relies on detecting low RF transition rates on one side of an atomic resonance at detunings where one would expect negligible rates on the other side. However, for a while we were detecting spurious transitions on both sides, which the widths of our RF pulses could not account for (more on that in Chapter 4). We finally tracked it down to a problem with our magnetic fields. It turns out that the trapping fields were ringing slightly at around 40 kHz, which we detected by placing a hand held pickup coil next to the trap and taking the Fourier transform of the resulting scope trace, seen in Figure 3.4a. The atoms could then absorb a 40 kHz photon from the trap along with an 80 MHz photon from our RF pulse to undergo a 2-photon transition to the final state (Figure 3.4b). Analogously, we can think of the oscillating field as modulating the transition frequency, resulting in sidebands at ± 40 kHz with respect to resonance, and it was the blue sideband photon, along with the carrier photon, that was causing the spin flips. Changing the servo parameters of the system eventually allowed us to eliminate the 40 kHz peak, along with the spurious transition rate.

Chapter 4

Radiofrequency pulses

A recurring theme of this thesis is that the contact measurement on the tail via RF spectroscopy gives a small signal, constantly in danger of being swamped by the nearby resonant atomic transition. To avoid this, we must be very careful to shape our RF probe pulse. This will be the main focus of this chapter, with a note on RF rectification in our system at the end.

4.1 Gaussian pulses

Ideally, one would be able to perform RF spectroscopy using a perfectly narrow function in frequency space. However, since that is experimentally not feasible, we are required to do the best we can with a wider feature. From undergraduate physics courses we know that the shape of a pulse in frequency space is simply the Fourier transform of that shape in time [37]:

$$F(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(t)e^{-i\omega t} dt. \quad (4.1)$$

This applies to the fields, so squaring the result gives the pulse shape in power. The standard and easiest RF pulse shape is a square pulse, usually created by a simple RF switch. However, its power spectrum is a Sinc^2 function, which has oscillatory wings with an appreciable amplitude out to many times the full-width-half-max (FWHM) of the main line (see Figure 4.1a). This would be unworkable, since the spectral component has to decrease much faster with detuning than the $\omega^{-3/2}$ tail from the contact. The logical choice, then, is to use a Gaussian shaped pulse, because its Fourier transform is also a Gaussian and falls off extremely rapidly with detuning, as seen in Figure 4.1b.




Figure 4.1: Some RF pulses in frequency space. In (a) we see the power spectrum of two standard pulses (inset). A square pulse in time becomes a Sinc^2 function, which looks much like a Gaussian down to the first minimum, but then exhibits significant wings far from the center (red curve). A Gaussian in time, however, becomes a Gaussian in power (black curve). The curves have been scaled to give similar peaks and widths. In (b) we see the same curves on a logarithmic scale, which makes the difference clear. An $\omega^{-3/2}$ line has been added (blue), to show that the problematic wings of a Sinc^2 function persist with detuning.

We create the actual Gaussian shape using an Agilent 33220A programmable arbitrary function generator with 800 points per Gaussian. Initially, we tried to create our pulses simply by connecting the output of the function generator to the amplitude modulation (AM) port of our Agilent E4420B frequency generator. However, we found that the frequency generator's 10 kHz bandwidth began to distort the shape of the Gaussian for $\tau < 25\mu\text{s}$, where τ is the rms width. Instead we opted for the LMH6503 linear variable gain amplifier, which has a 100 MHz gain control bandwidth and 70 dB adjustment range. It takes the unmodulated frequency from the synthesizer and the Gaussian shaped voltage pulse from the function generator as inputs, and outputs a Gaussian RF pulse. The power and frequency is controlled by the synthesizer, and the rms width of the pulse by the function generator. It is important to note that the amplifier linearly modulates the *voltage* of the rf signal with the Gaussian

$$V(t) = e^{-\frac{t^2}{2\tau^2}}, \quad (4.2)$$

which is squared to get the power as a function of time. As already mentioned, to get the rms width in frequency, one takes the Fourier transform of the *magnetic field* (proportional to the voltage), and then squares the result to get the power. This results in a final frequency width of $\sigma_\omega = \frac{1}{2\tau}$.

4.2 Shape checks

Of course, this is all nice and tidy in theory, but just because we tell our pulse to be a well behaved Gaussian does not mean that it is one. In fact, several issues could conspire to change the shape of our pulse and create non-Gaussian spectral components at nonzero detuning (wings) on our spectral function.

One of these is the intrinsic nonlinearity of our RF system. The variable gain amplifier will begin to saturate above 0 dBm (1 mW) of peak output power. Later down the line we have another 5 W amplifier (Minicircuits ZHL-5W-1), with a 1 dB compression rating at 5W output. The main result of these nonlinearities are that the peak of the Gaussian is slightly suppressed. We measure this suppression and ensure that the change in overall power is always below a few percent. But that




Figure 4.2: The Fourier transform of the RF power at system saturation. The labels correspond to the output power of the synthesizer before all amplification. At 9 dBm output we begin to risk damage to the amplifiers. The dashed black line shows a perfect Gaussian for reference (a parabola on this scale). We see that saturation does not affect the nicely Gaussian shape of the pulse down to at least 60 dB below the peak power.

does not answer the question of the wings. To measure the spectral function, we used a pickup coil to record the RF pulse on a fast oscilloscope (Tektronix DPO4034), and took its Fourier transform (similar to Section 3.2.3). The results can be seen for several different powers in Figure 4.2. The saturation corrections for the three powers were 2.6% for -3 dBm synthesizer power, 5% for 0dBm, and 11% for 3 dBm (i.e. higher powers than we actually use in the experiment), yet we do not see any broadening above the noise floor.

Another issue is the fact that a true Gaussian extends to positive and negative infinity, yet we have to choose a time to open and close our RF switch, which effectively truncates the Gaussian. We want to give the Gaussian enough of its shape to be well behaved, yet truncate it enough to allow for maximum experimental timing flexibility. We can explore this by creating a Gaussian pulse and truncating it with an RF switch, and then once again calculating the Fourier transform. Figure 4.3 shows the effects of truncating the Gaussian at different multiples of its width. One




Figure 4.3: The Fourier transform of truncated Gaussians. The inset shows the shape of the truncated Gaussians, at ± 1 (blue), 2 (green), and 4τ (black). The Fourier transform of the power shows the resulting frequency spectrum. The $\pm 4\tau$ truncation looks unaffected, whereas the $\pm 1\tau$ truncation clearly begins to look similar to a square pulse.




Figure 4.4: Atoms spin-flipped by truncated Gaussians. A truncated pulse is less efficient at transferring atoms between magnetic sublevels. The pulse still exhibits lower efficiency at 2τ , but is fully Gaussian at 4τ . The theory line is scaled to the final value.

can clearly see how the spectral function begins to look more and more like that of a square pulse as the truncations become tighter. At $\pm 4\tau$, we no longer notice any wings above the noise floor, which is at an appreciable 60 dB below the peak power.

Truncating the Gaussians also decreases the total power in the pulse. Calculations of the Rabi frequency tend to be much easier if we can just assume an ideal Gaussian, so we need to ensure that the integrated power in the truncated pulse is essentially equal to the full pulse. To test this, we resonantly drove our favorite ^{85}Rb atomic transition ($|2, -2\rangle$ to $|2, -1\rangle$), in the limit of low excitation fraction, with various truncated Gaussians. The results can be seen in Figure 4.4. The blue line is the integrated power of a truncated Gaussian, divided by the integral of a full Gaussian. It is then scaled to the maximum transferred fraction. We see reasonable qualitative agreement between the measured outcoupled fraction and the expected shape from our calculations, where we have only adjusted an overall scaling factor.

Given the results from the Fourier transform and the power tests, we eventually decided on truncating our Gaussians at 4τ (8τ total pulse length). A 3τ truncation should be usable as well, but we decided to be conservative in that respect. It essentially guaranteed a highly Gaussian

shape, and timing issues turned out not to be a problem. Typical Gaussian pulses for the contact have a $\tau = 100\mu\text{s}$. We have found this RF system to be very robust and repeatable. For pulses longer than $\tau = 25\mu\text{s}$, the option of using the AM port on the frequency synthesizer seemed to also work very well. Another option is to use a voltage variable attenuator such as ZX73-2500 from Minicircuits, which has a much higher frequency range. However, its gain curve is highly nonlinear, which would have to be carefully accounted for in order to produce clean Gaussian pulses.

4.3 Rectification

One final important issue related to the RF system is that of RF rectification. The contact, as mentioned numerous times already, gives rise to a low transition rate on the tail of the resonance. To be able to measure something, then, we must be able to either increase the pulse length or the pulse power enough for the signal to become detectable. What originally attracted us to measuring the contact is the fact that the system can be probed quickly with respect to three-body loss rates, so increasing the pulse length by many orders of magnitude is somewhat counterproductive. Thus, it is generally the RF power that we increase to bring the contact up to detectable levels. However, although it is relatively simple to buy very powerful RF amplifiers to keep throwing more power at the problem, eventually one is limited by RF rectification.

Rectification, in the sense that we care about, means that the RF that is intended solely for the atoms is transmitted throughout the room and makes it onto the servos that control our magnetic fields. Inside the servos are op-amps, which ideally would reject noise outside their bandwidth (in the hundreds of kHz range). However, internal interferences actually cause the op-amp components to produce a DC offset on the “out” pin [38], which manifests itself in a change in the magnetic field. Great care can be taken with electronics design to avoid this problem, but it is nearly impossible to eliminate completely.

To measure the effect of rectification, we produce two RF pulses: a strong Gaussian pulse to modify the DC magnetic field via rectification, and a weak square pulse to probe the magnetic field using an atomic transition. The strong pulse is detuned from resonance by 1.0 MHz. We can

then compare to the magnetic field without the strong pulse to see if it caused a shift. The results of such a test can be seen in Figure 4.5. We see significant rectification for the strongest pulses our system allows, as seen by the 17 mG magnetic field shift. However, rectification becomes negligible at powers 9 dB higher than the maximum powers needed to probe the contact. We can therefore rule out rectification as a possible mechanism that would give us a systematic shift in the measured contact.

The net results of all these systematic RF pulse checks are that we can be confidently measure the tail of an RF transition. The clean Gaussian spectral line allows us to rule out that we are inadvertently exciting atoms via the resonant transition.

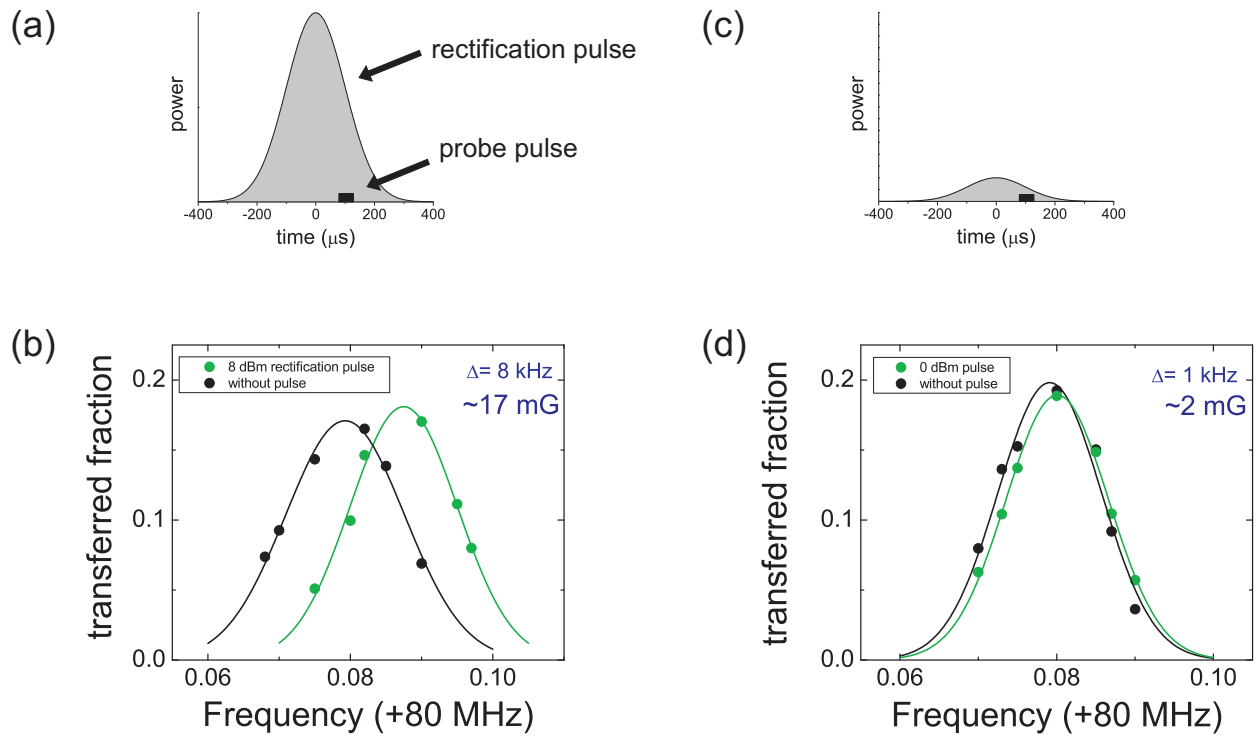


Figure 4.5: Effect of RF rectification. (a) We send out a $\tau = 100 \mu\text{s}$ Gaussian pulse, and probe the resulting field with a $50 \mu\text{s}$ square pulse. The $100 \mu\text{s}$ delay gives the servo time to respond and shift the field. (b) At our maximum allowable power (8 dBm at the synthesizer), we see rectification amounting to 17 mG. At 8 dB lower power (c), the rectification has all but disappeared (d). Our contact measurements required powers no more than -9 dBm.

Chapter 5

Detection

Detection is one of the most important parts of any experiment. It connects us to the object of our studies as an extension of our natural senses, and what it tells us forms the basis of everything we can say about it. However, it is often seen as a black box that “just works” by anyone new in the lab, and not understanding its complexities and subtleties can cause serious problems for the unaware. More specifically, it turns out that accurately calculating atom number from absorption images is not at all straightforward, and many systematics can cause errors in the measured signal. This applies to simply measuring condensate number in expansion, as well as to the in-trap imaging we perform when making measurements of the contact.

5.1 Condensate absorption imaging

5.1.1 Imaging corrections

One of the effects of our 10 Hz trap is a slowly expanding BEC when the trap is turned off, which is a problem when it comes to imaging. In “standard” time-of-flight (TOF) absorption imaging, the trap is released, and the cloud expands until the optical density (OD) is low enough to be imaged. However, the slow timescales our trap causes means that our BEC would have to expand for 70 ms or more before the OD drops to measurable levels. Such long expansion times are not feasible, since the cloud would hit the floor of our science cell in about 30 ms. Our initial solution to this problem (see Section 5.1.3 for our current solution) was to transfer a small fraction (20-30%) of the atoms from our $|F = 2, m_f = -2\rangle$ trapping state to our $|3, -3\rangle$ imaging state via

a deliberately inefficient adiabatic rapid passage (ARP). This reduces the OD of our imaging state by a factor of 3-5, which we then simply multiply up by a correction factor which we calibrate using thermal clouds. But even for clouds with low OD, much care must be taken to correctly calculate the atom number.

The reason that a low OD is required in the first place is twofold. One problem is the limited dynamic range of the camera, which will be addressed in Section 5.1.3. The other is that images usually saturate at some value of OD with standard absorption imaging. A picture of our cloud is generated using three images: the first image uses resonant light to probe the atoms, which casts a shadow onto the camera (I_{shadow}). The second image contains the probe pulse again, but without the shadow of the atoms (I_{light}). The final image triggers the camera and opens all the same shutters, but does not include the probe beam (I_{dark}). This dark frame is then subtracted off the other two frames to account for room lights and a camera offset. I will redefine the resulting frames using the convention in [39]: $I_f = I_{\text{shadow}} - I_{\text{dark}}$ and $I_i = I_{\text{light}} - I_{\text{dark}}$. This is done for every pixel, and the simply measured OD_{meas} is then

$$OD_{\text{meas}} = \ln \frac{I_i}{I_f}. \quad (5.1)$$

However, if any light in the probe beam cannot be absorbed by the atoms due to being off-resonant (what we call “bad light”) or simply by scattering around the atoms, then I_f cannot become arbitrarily small, and the OD_{meas} will saturate to some value OD_{sat} , regardless of how optically thick the cloud actually is. Also, if the dark frame does not effectively subtract light that is *not* resonant, the OD will saturate at lower values. For example, vibrations of the optics or non-repeatable shutter timings could cause less ambient light to enter the dark frame than the shadow frame. For our experiment, OD_{sat} is typically around 3.5, and its finite value can be corrected for to get a modified OD using [40]

$$OD_{\text{mod}} = \ln \frac{1 - e^{-OD_{\text{sat}}}}{e^{-OD_{\text{meas}}} - e^{-OD_{\text{sat}}}}. \quad (5.2)$$

Also, the probe intensity can saturate the atomic transition, which causes a significant change




Figure 5.1: A measurement of our effective saturation intensity. We measure the OD vs I_i in counts per pixel, and then invert the axes for plotting and fitting via Equation 5.3. The red line shows this fit. The measured OD is highly sensitive to the probe intensity, especially at low intensities, where one might think that the correction due to I/I_{sat} is negligible.

in the measured OD even for low intensities. This effect can be corrected for using

$$OD_{\text{actual}} = OD_{\text{mod}} + (1 - e^{-OD_{\text{mod}}}) \frac{I}{I_{\text{sat}}}, \quad (5.3)$$

where I is the intensity of light at the atoms, and I_{sat} is a constant, which is 1.16 mW/cm² for Rb on the cycling transition. To make this correction, then, one can try to carefully measure the intensity of the probe beam at the atoms, taking into account losses at optical surfaces. It is much simpler, however, to measure an effective saturation intensity $I_{\text{sat}}^{\text{eff}}$ in the same units as one measures I_i , in counts per camera pixel. One can then replace I/I_{sat} with the equivalent $I_i/I_{\text{sat}}^{\text{eff}}$. To do this, we create thermal clouds and image them using different probe intensities. We can then use Equation 5.3 as a fitting function with OD_{actual} and $I_{\text{sat}}^{\text{eff}}$ as the fit parameters. An example of this can be seen in Figure 5.1.

The resulting corrections to our total BEC numbers come to about 5% for the OD_{sat} correction and about 10% for the I/I_{sat} correction, given a measured OD of 1 and an $I/I_{\text{sat}} = 0.1$. But since imaging can be so important, we wanted an independent confirmation of our methods. To get an independent measurement, we can look at the sizes of the clouds (which depend on BEC number) and compare to expected sizes. The expected values come from a variational model developed by Pérez-García *et al.* [41] to simulate BEC dynamics using a Gaussian cloud ansatz. This will be referred to as the PG model. Equations 10a-c in [41] reduce to

$$\ddot{w} + \nu^2 w = \frac{\hbar^2}{m^2 w^3} + \sqrt{\frac{2}{\pi}} \frac{a \hbar^2 N}{m^2 w^4} \quad (5.4)$$

for a spherical trap, where $w/\sqrt{2}$ is the rms width of the Gaussian cloud, ν is the trap frequency, m is the mass, a is the scattering length, and N the total BEC number. Setting \ddot{w} to 0, we first solve for the equilibrium size w_0 of this Gaussian condensate given our initial number and scattering length. We then plug in an $a(t)$ that corresponds to the scattering length vs time in our experiment, and numerically integrate Equation 5.4 to get the size as a function of time $w(t)$. Then we must convert the size of the modeled cloud w to our fit size σ_{fit} . The Thomas-Fermi radius R_{TF} of a BEC is calculated to be 1.78 times larger than the equilibrium width w_0 independent of the

parameters. To go from the Thomas-Fermi radius R_{TF} to our fits, we simulate a three dimensional spherical Thomas-Fermi distribution and integrate it through one dimension (as is the effect of imaging). We then fit the resulting distribution to a two-dimensional Gaussian, resulting in a ratio $\sigma_{\text{fit}}/R_{\text{TF}} = 0.44$, which is also independent of the specific initial parameters. Thus, to convert the PG results to the imaging fits, we use $\sigma_{\text{fit}}/w = 1.78 \cdot 0.44 = 0.78$. Using it, we predict the size of the BEC vs number after 28 ms expansion and compare it to our measured sizes, seen in Figure 5.2. For this test, we expanded at $500 a_0$ and subtracted our $8 \mu\text{m}$ resolution limit σ_{res} in quadrature via $\sigma = \sqrt{\sigma_{\text{fit}}^2 - \sigma_{\text{res}}^2}$. The good agreement between the data and the theory provides additional confirmation of the validity of the OD corrections. Unfortunately, the dependence of BEC size on the number is too weak to allow spatial extent to be our primary measure of condensate number.

This method of calculating OD via Equation 5.1 and then applying ARP, OD_{sat} , and I_{sat} corrections works well enough, but can be quite tedious. Also, the corrections can change with experimental drift. OD_{sat} , for example, can change drastically with a small change in probe beam alignment or frequency components of the probe laser, and the ARP calibration has to be measured every day. Carefully performing these necessary corrections is certainly manageable, but in the end we decided to switch to a more robust high-intensity probing scheme described in Section 5.1.3.

5.1.2 BEC collapse and high-field imaging

Another imaging difficulty related to our system arises from the collisional properties of ^{85}Rb . As explained in Chapter 2, ^{85}Rb has a background scattering length a_{bg} of $-443 a_0$, so the condensate begins to collapse as soon as the magnetic fields are turned off (the low expansion energy is not enough to overcome this). We allow the condensate to expand by simultaneously increasing the bias field as we turn the trap off, such that the magnitude of the magnetic field stays constant. The BEC then expands at a positive scattering length for ~ 30 ms until the size has increased to about twice the resolution limit. At this point, the original mode of operation was to turn off the fields as fast as we can and image the BEC once the field stabilizes. Even so, the condensate has about 2 ms time at its background scattering length to ponder the merits of collapsing and act on




Figure 5.2: Comparison between measured and predicted BEC size. The different shapes correspond to different ARP efficiencies, giving rise to different ODs. The red line is the modeled BEC size in expansion.

the result. In order to predict the outcome we return once again to the Pérez-García (PG) model.

If we take a ^{85}Rb BEC of 50,000 atoms at $100 a_0$ at rest in our trap and then instantly turn off all the trapping fields, the PG model predicts that the BEC as a whole will collapse in about 27 ms. But one of the limitations of this model is that the cloud shape is always conserved, meaning that the cloud shrinks nicely as a whole. In reality however, smaller clumps of the cloud can collapse locally on much faster timescales. To get an idea of this timescale, we use the fact that small condensates can be stable at negative scattering length. Assuming no trapping potential ($\nu = 0$), we again solve for the stability condition by setting \ddot{w} in Equation 5.4 to 0:

$$w_{\text{stable}} = -aN\sqrt{\frac{2}{\pi}} \quad (5.5)$$

(This illustrates another reason for making spherical clouds...to make the math easier). Here a is the background scattering length. If we constrain the peak density of this stable cloud to be equal to the initial peak density of our BEC, then we can solve for the width and atom number N_p that will form the smallest element of collapse in our BEC. Then we simply increase the number by $\sqrt{N_p}$, the statistical number fluctuation in that packet, and use the PG model to predict the collapse time of the packet, which comes out to be about 2 ms. This is certainly just a rough calculation, as it does not take into account the expansion energy the cloud already has, but it tells us that we cannot assume that 2 ms is too short to see losses from local implosions.

Our previous way to avoid collapse was simply to increase the scattering length during expansion to provide more kinetic energy and lower the density, but increasing a also greatly increases the rate of three-body losses. And when varying the scattering length in expansion, we did not see a regime that clearly avoids both collapse and three-body loss problems. Thus we decided to instead image at high field. We still need to transfer the atoms to the $|3, -3\rangle$ imaging state, but we can do that using a $7 \mu\text{s}$ RF π pulse, drastically reducing the time at negative a .

To image at high field on the cycling transition, we need to shift the probe laser's frequency by 252 MHz from the zero field cycling transition. As explained in detail in [32], our lasers are controlled via an offset lock. The probe laser is overlapped with a master laser, and the resulting beat note

mixed down with a voltage-controlled oscillator (VCO) to produce a signal that eventually feeds back to control the frequency of the probe laser. The VCO frequency is changed by our computer controller, but does not have enough tunability for both low and high-field applications. To increase its range, we ramp the VCO to its maximum detuning, and then switch over to a different VCO that works at higher frequencies to continue the ramp. We have not had any problems with switching between VCOs, and that method of increasing the laser tuning range has proven to be quite robust for offset lock lasers.

Unfortunately, our imaging axis does not coincide with the bias field direction, meaning that we cannot probe with pure σ_- polarization. Instead, we change the polarization such that it is linear and perpendicular to the quantization axis. This is achieved by a waveplate mounted on a flipper, which we manually put into or out of the beam depending on which imaging scheme we want to use. The polarization geometry results in a superposition of σ_+ and σ_- light (a nice table of different geometries and their resulting transitions can be found in [40]). Since we image in a ~ 160 Gauss field, the σ_+ and σ_- transitions are separated by many linewidths, giving us a pure cycling transition. In fact, the only result of the superposition is that it reduces the effective transition strength by a factor of 2.

The ability to image at high fields enables us much more flexibility than before, as we do not have to wait for fields to turn off to probe our system. The downside is that it is somewhat of a nuisance to have to change the probe frequency every time we decide to image at a different field. However, the 6 MHz natural linewidth of the transition is much larger than the magnetic field fluctuations, so that we can just dial in the correct probe frequency once we know the field without recalibrating every time.

5.1.3 High-intensity imaging

In Section 5.1.1, I described a workaround to deal with the fact that our condensates have very high OD, which involved performing inefficient ARPs to transfer only a small fraction of the cloud to the imaging state. However, the calibration of that fraction is cumbersome and possibly introduces

systematic errors. Also, the move to imaging at high fields means that the RF frequency for that transfer depends on the field, which increases the calibration requirements and the possibility of error. But there is another solution. Fundamentally, a high OD leads to problems because of the limited dynamic range of the camera CCD. Our camera has a 16-bit output, meaning a single pixel can read $2^{16} \approx 65,000$ different values. Thus, all else being perfect, the maximum OD the camera can display is $\ln(2^{16}) = 11.1$. However, as there is noise in the dark frame of the camera and shot noise in the beam, a workable limit to the maximum OD is much lower. We measure 12 counts/pixel rms noise in the dark frame, so with $I/I_{\text{sat}} = 0.1$, an OD signal-to-noise ratio of 2 will limit the maximum workable OD to 5.3. We numerically calculate this based on error propagation through Equation 5.1. To beat this problem we clearly need more light to make it through the atom cloud. I mentioned in Section 5.1.1 that high probe intensity causes problems, because saturation decreases the amount of light that gets absorbed. However, we understand how this happens and can account for it, and if we significantly increase the intensity that gets through, the noise becomes fractionally insignificant. Thus, imaging with intensities much higher than I_{sat} can allow for accurate quantitative measurements of optically thick clouds [39].

Instead of using the simple definition of OD (Equation 5.1) to create the image and then correcting for the saturation intensity, we use the full OD definition given by

$$OD_{\text{new}} = \ln \frac{I_i}{I_f} + \frac{I_i - I_f}{I_{\text{sat}}^{\text{eff}}}, \quad (5.6)$$

which is just a rewritten form of Equation 5.3, assuming that the effect from bad light is negligible. Our imaging software uses this formula, with an empirically supplied $I_{\text{sat}}^{\text{eff}}$ from data such as shown in Figure 5.1, to calculate the OD_{new} for each pixel. Using this method, we have been able to image clouds with an OD on the order of 8 with negligible noise, and higher ODs should be possible.

But with all new and fancy techniques come new issues and systematics. One of these is the problem of photon re-absorption. The calculations for OD assume that photons emitted from the atoms leave the system. However, there is a finite chance that another atom will absorb this photon, which will make it unavailable to absorb from the probe beam, and therefore decrease the

apparent OD. The higher the optical depth, the stronger this effect. But for very high intensities, the scattering rate eventually saturates to 1/2 the excited state decay rate [42]. At that point, the intensity from surrounding atoms stays constant as the probe intensity is increased, and an atom will be increasingly likely to absorb from the probe beam. In other words, once you commit to high intensity imaging, you have to really let loose. An illustration of this effect can be seen in Figure 5.3.

As we could not simply turn up the power of our probe laser to reach intensities of $15 \times I/I_{\text{sat}}$, we mounted a 1 m focal length lens roughly 1 m upstream of the atoms to reach higher intensities by decreasing the beam waist. The lens is also mounted on a flipper mount, which allows us to easily switch between imaging schemes. Such high intensities saturate the camera pixels, so we also use a neutral density filter in front of the camera to decrease the intensity by a factor of 2. One might naively think that this is counterproductive, but the important thing is to have a high intensity incident on the atoms. The neutral density filter does not change the fraction of light absorbed, and all we have to do is to measure our $I_{\text{sat}}^{\text{eff}}$ with the filter in place to account for it.

But one must be careful with high-intensity beams. If too many photons are absorbed by the atoms during the pulse, the resulting momentum transfer will result in a Doppler shift, and the atoms are pushed out of resonance. This again decreases the apparent OD, leading to another systematic. To ensure this does not happen, the probe pulse times t must be kept very short. The scattering rate Γ , the doppler shift δ_{D} , the photons absorbed N_{ph} , and the velocity v of an atom after absorbing photons are related via the coupled equations:

$$\begin{aligned}\Gamma &= \frac{\frac{I}{I_{\text{sat}}} \frac{\gamma}{2}}{1 + \frac{I}{I_{\text{sat}}} + \left(\frac{4\pi\delta_{\text{D}}}{\gamma}\right)^2} \\ N_{\text{ph}} &= \Gamma \cdot t \\ v &= \frac{N_{\text{ph}}h}{\lambda m} \\ \delta_{\text{D}} &= \frac{c}{\lambda} \left(1 - \frac{1}{1 + \frac{v}{c}}\right),\end{aligned}\tag{5.7}$$

where $\gamma = 2\pi \times 5.98$ MHz is the natural linewidth of the transition, $\lambda = 780$ nm is the probe

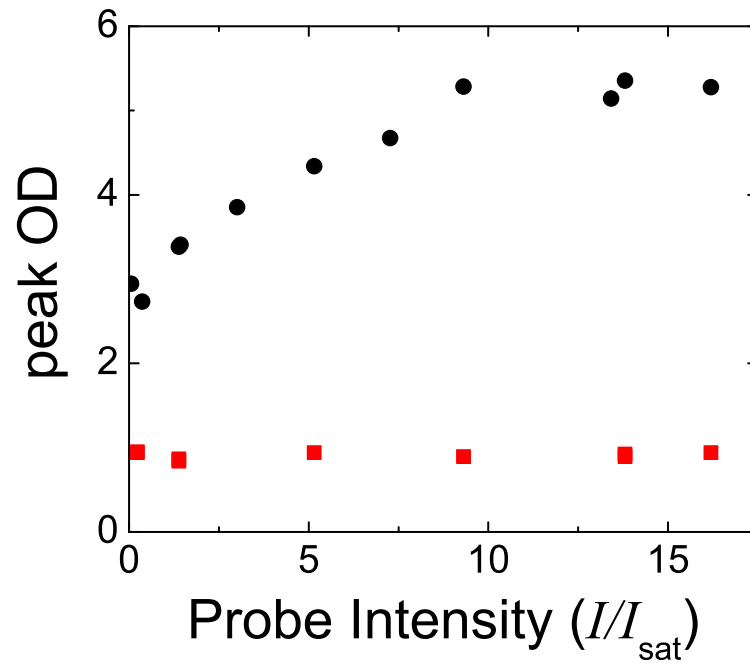


Figure 5.3: Effect of photon re-absorption on OD. The black circles show the measured OD of a BEC, calculated with Equation 5.6, as a function of I/I_{sat} . For intensities near I_{sat} , the high scattering rate in conjunction with the high OD causes a large fraction of atoms to re-absorb photons emitted from surrounding atoms. This makes them unavailable to absorb from the probe beam, lowering the apparent OD. As the intensity increases, the scattering rate remains roughly constant, translating to a smaller fractional amount of rescattered light present in the cloud. The red squares show the same experiment on thermal clouds with lower OD, where this effect does not occur.

wavelength, m is the mass of the atom, and c is the speed of light. We numerically solve these equations vs time to calculate the expected error due to this effect and ensure that it is small. We use $5 \mu s$ probe times, which for an $I/I_{\text{sat}} = 10$ and an OD of 6 corresponds to an error of 0.5%.

Finally, if the OD is low, imaging with high intensities results in unwanted noise. The light and shadow frames will both have high intensity light, but the fractional difference between the two will be small. If we assume shot noise (a lower limit) on the beam and an $I/I_{\text{sat}} = 10$, then an OD of 0.1 results in a signal-to-noise ratio of 1.5 on the OD. Luckily, this is a per pixel noise figure, and low OD regions usually spread over many pixels. The spatial extent of thermal clouds therefore help average down this noise to workable levels. However, because of this noise issue we use high intensity imaging only when there is a high-OD feature we are interested in. For the lower OD pure thermal clouds we switch back to low-intensity imaging.

Overall, the high-intensity imaging technique works very well for us. The value of $I_{\text{sat}}^{\text{eff}}$ is dependent mostly on the reflectivity of the optics between the atoms and the camera, as well as the camera's quantum efficiency. Therefore it is generally not subject to experimental drift and does not have to be recalibrated unless the optics change (such as with the addition of the neutral density filter). We have found that the best measurements of $I_{\text{sat}}^{\text{eff}}$ (see Figure 5.1) are performed with thermal clouds of low OD. Higher ODs can exhibit photon reabsorption, as well as cause problems in the region of low intensity, where imperfections in the dark frame subtraction can cause significant deviations from the saturation curve in Figure 5.1.

5.2 In-trap imaging for contact measurements

5.2.1 Transfer to the imaging state

We still use absorption imaging for our contact measurements, but we require several different additional techniques for our purposes. Mainly, when we perform RF spectroscopy to measure the contact, we transfer 1-2% of the atom cloud from the $|2, -2\rangle$ to the $|2, -1\rangle$ state. The (difficult) trick is to image and count that small fraction in the $|2, -1\rangle$ state, without imaging any of the large

cloud left in the $|2, -2\rangle$ state, which would cause significant systematics. So the imaging process must be highly efficient and highly state-specific.

One important step toward enhancing imaging efficiency, or rather the signal-to-noise ratio, is to wait for as little time as possible between the RF contact pulse and the probe pulse. The atoms that we outcouple on the tail have high momentum, and their expansion energy is $\frac{1}{2}\hbar\omega$, where ω is the detuning from resonance (see Chapter 6). For a typical $2\pi \times 40$ kHz detuning, the atoms move at a velocity of $14 \mu\text{m}/\text{ms}$. Given that the starting size of our BEC is $14 \mu\text{m}$, the OD will drop significantly if given more than a millisecond to expand. This is another reason that we do not turn off the trap, but instead image at high field.

To perform state-specific imaging, we need to efficiently transfer the $|2, -1\rangle$ atoms to our $|3, -3\rangle$ imaging state without putting any of the $|2, -2\rangle$ atoms up there as well. We do this via a two-step RF and microwave (μ -wave) process, as seen in Figure 5.4. A 2.8 GHz μ -wave photon transfers them to the $|3, -2\rangle$ state, followed by a ~ 85 MHz RF photon to get to the $|3, -3\rangle$ imaging state. Since we already have to find the resonant frequency of the $|2, -2\rangle$ to $|2, -1\rangle$ transition for every contact measurement, we can exactly calculate the other transition frequencies using the Breit-Rabi formula [43]. Initially, we used a short π pulse for both frequencies to transfer the atoms as quickly as possible ($22 \mu\text{s}$ for the μ -wave and $16 \mu\text{s}$ for the RF), but we found that this caused noise in our signal. The μ -wave transition has a large field sensitivity of 1.4 MHz/Gauss, and the shot-to-shot field fluctuations were large enough to make this transfer highly unstable. The obvious solution would be to sweep the frequency and transfer them using an adiabatic rapid passage (ARP), which tends to be insensitive to magnetic-field noise. However, the $|2, -2\rangle$ to $|3, -1\rangle$ transition (shown by the red dashed line in Figure 5.4) is only about 140 kHz away, and a standard ARP would transfer the 99% of atoms remaining in the $|2, -2\rangle$ state into the $|3, -1\rangle$ state with very high efficiency. Unfortunately we found that, although we do not send out resonant RF purposefully, some of the many atoms in the $|3, -1\rangle$ state always make it into the $|3, -2\rangle$ state and then get pumped into the $|3, -3\rangle$ state and get imaged. We solved this issue by shaping the ARPs to avoid the unwanted transition.




Figure 5.4: State diagram showing transfer to the imaging state. The atoms (shown in blue) start in the $|2, -2\rangle$ state, and the RF pulse for contact spectroscopy moves a small fraction to the $|2, -1\rangle$ state. These are transferred via shaped ARP to the $|3, -2\rangle$ state and RF π pulse to the $|3, -3\rangle$ state (solid green arrows). Once there, the probe beam cycles them to the $F'=4$ manifold (dotted green arrow). The entire procedure, from the $|2, -1\rangle$ state to imaging, takes $< 200 \mu\text{s}$. The transition of the cloud remainder to the $|3, -1\rangle$ state (red dashed arrow) is unwanted and would produce significant false signal. The energy splittings are given in MHz, calculated for a field of 160 G.

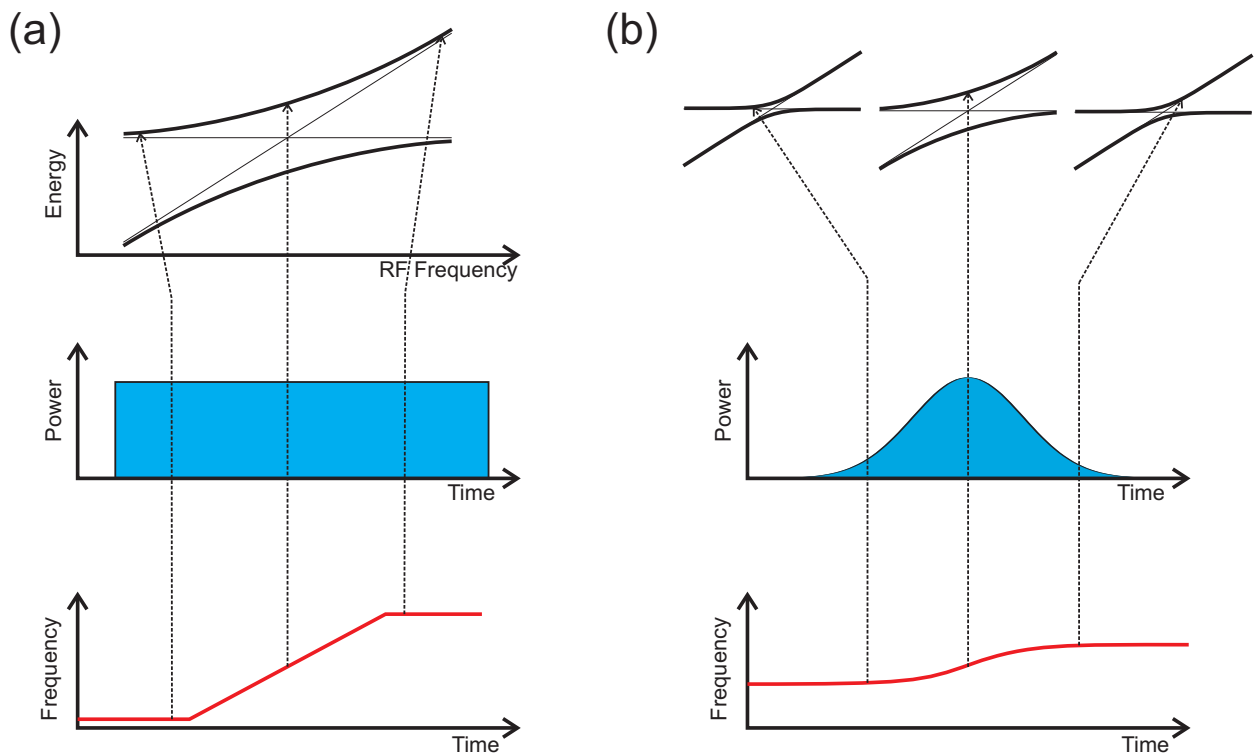


Figure 5.5: Shaped ARPs. A standard ARP is shown in part (a), where the RF power and therefore the energies of the dressed states are constant in time. The frequency must start far from resonance and end far from resonance for a well behaved ARP. A shaped ARP (or ShARP) seen in (b) starts at low detuning and low power. The power is adiabatically increased as the frequency ramps through resonance and then decreased the same way. This allows for a quick, efficient ARP, sweeping over a very small frequency range around the resonance.

In the dressed-state picture, the resonant energy splitting between two states is given by \hbar times the Rabi frequency Ω . To complete a very efficient ARP, one must start and end the frequency sweep far away from resonance to be completely in a bare atom state at the beginning and end of the sweep. What constitutes “far away” is, not surprisingly, relative to Ω . A “standard” ARP, with constant power and a linear frequency ramp, is shown in Figure 5.5a. To be able to perform an efficient ARP quickly (one of our requirements), one needs a high Rabi frequency. But the only way to perform a standard efficient ARP and not have to ramp through the nearby resonance would be to use low power and sweep too slowly for our needs. We found a way around this by shaping both the power and the ramp rate of the ARP as shown in Figure 5.5b. By creating a Gaussian envelope on the power, we start at low detuning and low power to avoid the nearby unwanted transition. As we ramp through the “good” resonance, the power adiabatically increases, which also lets us increase the ramp rate. Using this method, we can ARP $\sim 80\%$ of the atoms from the $|2, -1\rangle$ to the $|3, -2\rangle$ state in $100 \mu\text{s}$, without transferring any detectable atoms from the $|2, -2\rangle$ to the $|3, -1\rangle$ state. The total sweep covers 100 kHz, making this ARP very insensitive to the shot-to-shot fluctuations of the magnetic field.

Once the atoms are in the $|3, -2\rangle$ state, we pulse a $16 \mu\text{s}$ RF π pulse to transfer to the $|3, -3\rangle$ imaging state with 97% efficiency. This transition only has 0.47 MHz/Gauss sensitivity (same as the transition for the contact), so magnetic-field noise is not an issue.

Unfortunately, as mentioned in Section 5.1.2, the fact that we image at high field, with a beam perpendicular to the quantization axis, means that we automatically lose a factor 2 in OD signal. But since we can image the outcoupled atoms $< 200 \mu\text{s}$ after the end of the contact pulse, we probe much higher ODs than we would by waiting for the fields to turn off. When we first started contact measurements, we imaged the atoms at low field about 4 ms after the RF pulse for the contact. The downside of this was that the expanded clouds were not easily discernible on the CCD image due to the low OD, and we had to trust the fitting program to average out the noise and fit a cloud. But imaging after some expansion time also brings a great benefit. Since the expansion energy can be very large for the contact signal, the expanded size can serve as a clear

confirmation that the atoms really are from the tail of the RF spectrum, rather than a systematic unrelated to the contact.

5.2.2 Small cloud effects

With in-trap imaging also comes the problem of resolution. The Thomas-Fermi radius of our BEC is generally $14 \mu\text{m}$, corresponding to 2.4 pixels. With a cloud so small, one might wonder if there are systematics associated with the pixelation. For example, the calculated OD is a nonlinear (logarithmic) function of the intensity, yet for each pixel we get simply a sum of the photons that hit it. Thus, a significant variation of the cloud OD over a length scale corresponding to a single pixel can systematically lead to an error. To get an idea of the size of this error, we simulated a Gaussian absorption profile divided onto a 2-dimensional array of pixels, used the fraction of missing light in each pixel to calculate that pixel's OD value via the simple OD formula (Equation 5.1), and then summed up all the OD values from the pixels. Comparing the integral of the Gaussian OD that produced the shadow profile I_G to the summed pixels Σ_{px} gives an idea of the error associated with the pixelation. Some results of this simulation can be seen in Figure 5.6. This simulation does not account for the diffraction-limited resolution, which blurs the image and decreases the intensity variation on a pixel, so the simulation shown in Figure 5.6 represents a worst-case scenario. Our total resolution limit is about 1.4 pixels, or $8 \mu\text{m}$ (representing the rms width of the Gaussian that fits the smallest clouds we can observe with our imaging system). Note that the effect is dependent on the position of the Gaussian with respect to the pixels, and increases with OD. For the ODs in trap and a Thomas-Fermi radius of ~ 2.4 pixels, we expect that this pixelation effect adds an error $< 3\%$ to the atom number.

When measuring the contact, we alternate taking a measurement on the tail of the transition and the peak of the transition, with the final value for the contact involving a ratio of these two measurements. Since the tail signal is dominated by the high-density BEC in the center, we use a single, two-dimensional Gaussian fit to find the outcoupled number. However, on the peak of the resonant transition the rate is independent of density, and we outcouple equal fractions from the

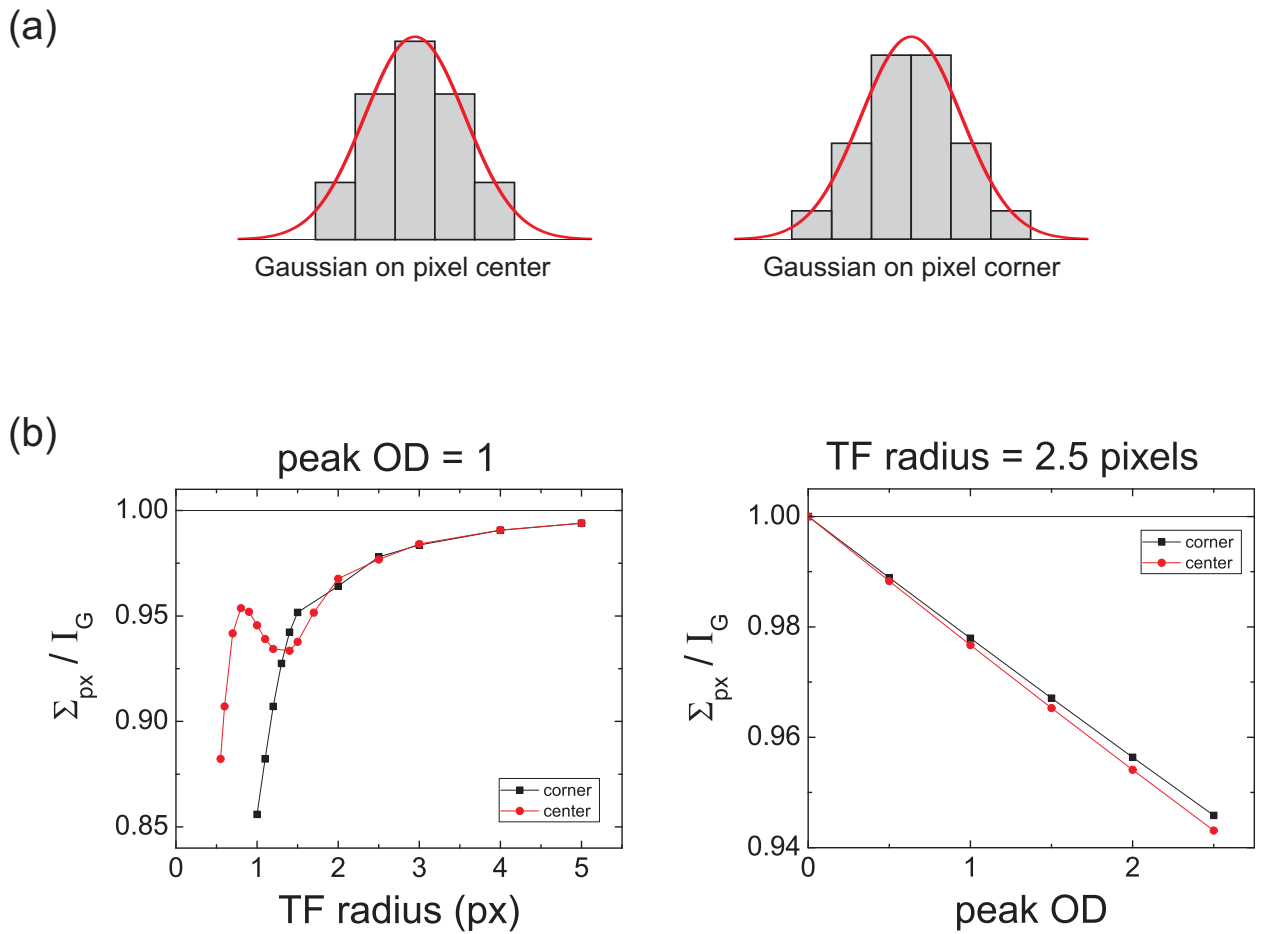


Figure 5.6: Simulation of pixelation effects. An absorption profile is incident on a 2-dimensional array of pixels, and the resulting pixelated OD, the sum over the pixels Σ_{px} , is compared to the integral, I_G , of the OD that gave rise to the absorption profile. Part (a) shows the two configurations under which the simulation was run. In one configuration, the Gaussian profile was centered on the center of a pixel in both dimensions. The corner configuration formed the other “extreme”. Simulation results can be seen in part (b). When the Thomas-Fermi (TF) radius is roughly the size of a pixel, the effects can become very pronounced, even exhibiting unintuitive structure. The error increases linearly with OD, but stays at a 3% or less correction for our normal operating conditions of peak OD < 1 and a TF radius of 2.4 pixels.

BEC and the thermal components. To measure the total number accurately, we require a double-Gaussian fit to account for each component. But as both BEC and thermal are near the resolution limit and similar in size, the fits often cannot clearly distinguish the two components. Also, in terms of general experimental wisdom, using different techniques to measure the tail and the peak allows for unexpected measurement systematics. To ensure that this does not happen, we came up with a technique based on summing over pixels. We choose a small region centered on the cloud and simply calculate the sum of all the pixel values. Then we create another region, a “sidewalk” around the first, and use it to calculate the average value of the background signal (Figure 5.7). We subtract this average value from the inner region, leaving us with the signal from our atom cloud, independent of shape. To decrease the noise, we want the inner box to be as small as possible while still enclosing all of the signal. For the “sidewalk”, a smaller size allows less chance of a systematic offset from different parts of the CCD, but it must be large enough to average out local noise. We analyzed numerous data sets, varying the sizes of the two regions, and find the optimum sizes to be 15 pixels on a side for the inner box and 35 for the outer (see Figure 5.7).

The main result from the analysis was that the method of summing over pixels agrees with the fitted Gaussian method as long as the subtracted background is small. This independently verifies that using the Gaussian fits does not produce a systematic error due to the difficulty of fitting very small clouds. We do not, however, use the pixelsum method for taking our data. When the signal was relatively strong (as in Figure 5.7), we find that the pixelsum method works very well, often even giving smaller errors than the Gaussian fits. However, for small peak ODs, when the signal was comparable in magnitude to the subtracted background, we find that this pixelsum method results in very large fluctuations compared to the Gaussian fits.




Figure 5.7: Schematic of the pixelsum technique. Shown is a standard image of a BEC with 30% thermal component, taken in the magnetic trap. For the pixelsum analysis, we sum all the values of the inner square, and subtract the background as determined by the “sidewalk” (hatched area).

Chapter 6

Two-body Contact

In this chapter I will explain what the contact is and try to elucidate its physical meaning. I will also present data of contact measurements versus interaction strength, as well as some interesting dependences on timescales. In much of the literature up to the writing of this thesis, the contact is presented for fermions. I will only deal with the contact for bosons which, for the most part, merely changes some prefactors in the equations by a factor of 2. However, the possibility of three-body interactions among bosons gives rise to a three-body contact, which I will describe in Chapter 7. For this reason, I will refer to the two-body contact as C_2 and the three-body contact as C_3 . In this thesis, as well as in papers regarding fermions or ignoring three-body interactions, the contact refers to the two-body contact, C_2 , unless specified otherwise.

6.1 What is the contact?

There are numerous working definitions of what the contact is, but I will begin with what I think may be the most intuitive. The contact, C_2 , is essentially the non-classical probability of finding two particles close to one another. Specifically, this probability due to the contact is higher than in a classical, non-interacting gas. With atoms that interact only via s-wave collisions, the presence of a two-body potential (Figure 6.1a) means that the scattering wavefunction will look something like Figure 6.1b. In the zero-range limit we ignore the complicated and species-dependent short-range potential that gives rise to an effective range r_e , and assume that the smooth, longer-wavelength part of the wavefunction extends all the way to zero. Another way of representing the

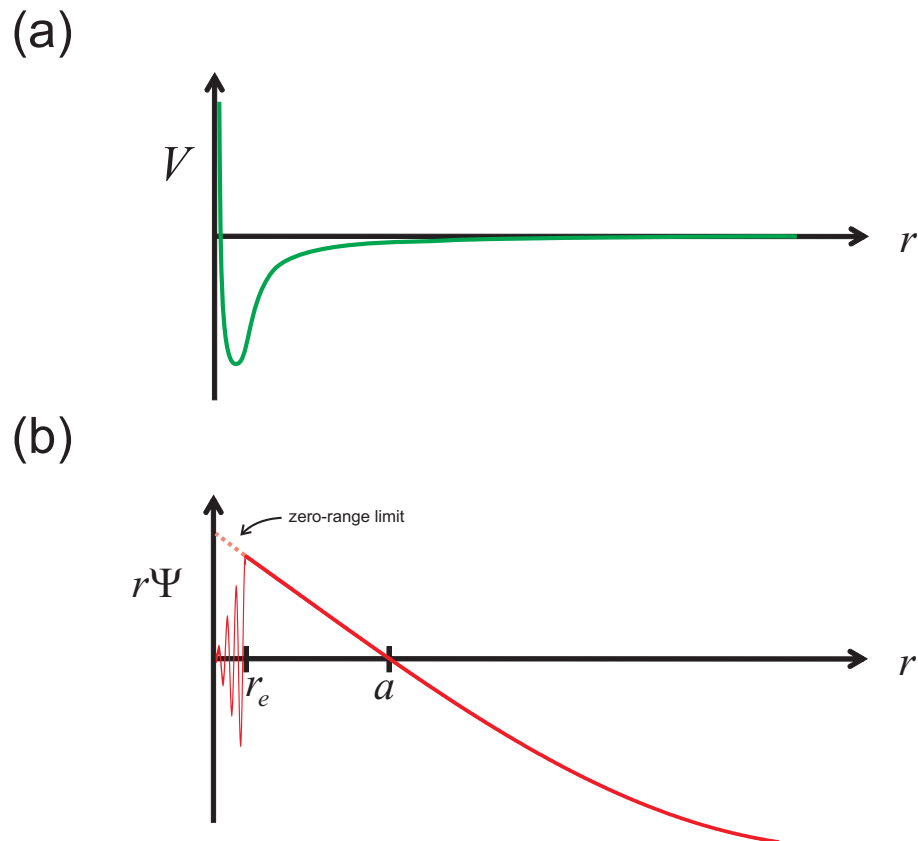


Figure 6.1: Schematic of the interparticle potential (a) and scattering wavefunction (b). The potential $V(r)$ looks like an attractive potential down to a distance given by the van der Waals potential. Outside of this effective range r_e , which is ignored in the zero-range limit, the wavefunction can be universally described by the scattering length a .

two-particle wavefunction $\Psi(r)$ in the zero-range limit is to say

$$\lim_{r \rightarrow 0} r\Psi(r) = \left(1 - \frac{r}{a}\right) A, \quad (6.1)$$

where r is the interparticle distance, a is the scattering length, and A is a constant that depends on the problem [44]. Given this wavefunction, we can calculate the probability of finding a particle near another particle at $r = 0$, given by $4\pi \int_0^{r'} \Psi(r)^2 r^2 dr$, which is proportional to r' in the limit of $r' \ll a$. The probability of finding the particle at 0 just goes as the volume $\frac{4}{3}\pi r'^3$, so the probability of finding a pair in a spherical volume of radius r' goes as r'^4 . Indeed, the coefficient of the pair correlation for small volume is given by the contact density, $\mathcal{C}_2(\mathcal{R})$, and the number of pairs of atoms in a sphere of radius r' is

$$N_{\text{pair}} = \frac{r'^4}{4} \mathcal{C}_2(\mathcal{R}). \quad (6.2)$$

The contact is obtained by integrating $\mathcal{C}_2(\mathcal{R})$ over all space. If the contact is zero, there is no enhanced correlation as $r \rightarrow 0$, and you would expect the probability of finding pairs of particles to go as r'^6 (volume²) as in a non-interacting gas. This classical probability is unrelated to the contact.

With only short-range interactions, the number of pairs close to one another, as parameterized by the contact, has far-reaching consequences to the system as a whole. Indeed, if we change the scattering length, the resulting energy change of the system depends on the contact. This dependence was first derived by Shina Tan in 2005 (published in 2008) [16] and is referred to as the adiabatic sweep theorem:

$$\left. \frac{dE}{d(1/a)} \right|_s = -\frac{\hbar^2}{8\pi m} \mathcal{C}_2, \quad (6.3)$$

where E is the total energy of the gas, and m is the mass of one atom. Combining the adiabatic sweep theorem (Equation 6.3) with the LHY result (Equation 1.1) gives us

$$\mathcal{C}_2 = 16\pi^2 n a^2 \left(1 + \frac{5}{2} \frac{128}{15\sqrt{\pi}} \sqrt{na^3} + \dots \right) N_0, \quad (6.4)$$

where N_0 is the number of atoms in the BEC. Not only can the LHY result be recast in terms of the contact, the fractional contribution of the LHY part of the contact has the added benefit of a

factor 5/2 compared to the fractional contribution of the LHY energy, aiding its measurement.

Perhaps one of the best aspects of the contact is the way by which we can measure it. If we look again to Figure 6.1b, we can ask ourselves what other effects are due to this shape of the wavefunction. To get the momentum distribution n_k of the system, one takes the Fourier transform of the two-particle wavefunction. A noninteracting state looks just like a sine wave, giving a delta function momentum distribution. However, the interacting wavefunction at low a results in a high-momentum tail in the momentum distribution that falls off as $1/k^4$ and is also proportional to the contact:

$$\lim_{k \rightarrow \infty} n_k \rightarrow \frac{C_2}{k^4}. \quad (6.5)$$

Indeed, this relation is sometimes used as the definition of the contact. The momentum distribution can be (and has been) used to measure the contact in fermions [26], but we need a much faster probe than time-of-flight (TOF) imaging to avoid the timescale restrictions due to bosonic three-body losses. If we project the atoms into a noninteracting state via RF spectroscopy, the tail in the momentum distribution gives rise to a tail in the RF spectrum with an $\omega^{-3/2}$ dependence [45, 46]. This is a key result, since we can use very short RF pulses to probe the RF spectrum before significant three-body losses can occur. It turns out that atoms in the final spin state *do* interact with atoms in the initial state, modifying the scattering rate to [27]:

$$\lim_{\omega \rightarrow \infty} \Gamma(\omega) = \frac{\Omega^2}{4\pi} \sqrt{\frac{\hbar}{m}} \frac{\alpha(a)}{\beta(\omega)} \frac{C_2}{\omega^{3/2}}, \quad (6.6)$$

where the integrated RF lineshape is

$$\int_{-\infty}^{\infty} \Gamma(\omega) d\omega = \pi \Omega^2 N \quad (6.7)$$

and Ω is the single atom Rabi frequency. In Equation 6.6, $\alpha(a)/\beta(\omega)$ describes the final-state effects; the a -dependent part is $\alpha(a) = (a'/a - 1)^2$, where a' is the scattering length for interactions between atoms in the final spin state and atoms in the initial spin state, while the frequency-dependent part is $\beta(\omega) = 1 + \hbar|\omega|/E'$, where $E' = \hbar^2/m a'^2$. For our system, the final-state effects are characterized by $a' = -565 a_0$ [47] and $E'/h = 133$ kHz.

Of course, there are limits of applicability of this analysis. In Equation 6.1, the wavefunction was assumed to continue in the given form all the way to $r = 0$. However, once the distances are comparable to the effective range r_e , the wavefunction can no longer be described by the single scattering parameter a , and the details of the interatomic potential must be taken into account. Since these details vary between specific species of atoms, we no longer consider it the universal regime. Hence the above derivation is only valid for $k \ll 1/r_e$, or $\omega \ll \frac{\hbar}{2mr_e^2}$. The effective range is given by [48]

$$r_e = \frac{\Gamma\left(\frac{1}{4}\right) \Gamma\left(\frac{3}{4}\right)}{6\pi^2\sqrt{8} \Gamma\left(\frac{5}{4}\right)} \left(\frac{mC_6}{\hbar^2}\right)^{1/4} \quad (6.8)$$

where C_6 is the coefficient of the van der Waals potential, giving the limit $\omega \ll 2\pi \times 3.5$ MHz. We can include a first-order correction due to the effective range, which turns out to be small for the ranges of a and ω that we probe. A derivation and the effect of this correction can be found in Appendix A. On the other side of the scale, the limit is set by the interparticle spacing. We only consider the two-body wavefunction in the limit $r \rightarrow 0$, and the above equations are valid only for $k \gg n^{1/3}$, or $\omega \gg \frac{\hbar n^{2/3}}{2m}$, where n is the density, resulting in the limit $\omega \gg 2\pi \times 300$ Hz for $n = 10^{13} \text{ cm}^{-3}$. Experimentally we stay within these limits by about a factor of 10 on each side, so the contact formulation is valid for all our experiments.

6.2 The contact signal and its characteristics

For our RF contact spectroscopy, the RF drives a transition to a lower energy spin state and one expects the $|\omega|^{-3/2}$ interaction-induced tail on the low frequency side of the lineshape. This is schematically illustrated in Figure 6.2a. We drive the interacting $|2, -2\rangle$ state into the free-particle-like $|2, -1\rangle$ state with extra kinetic energy. Since the absolute energy of the final state is lower, but with added kinetic energy, the transition frequency is lower than the bare transition. Consistent with this expectation, we observe a tail for large negative detunings, while for similar detunings on the positive side we find that the signal is consistent with zero.

An example of RF contact spectroscopy at $a = 1016 \pm 10 a_0$ is shown in Figure 6.2b. We




Figure 6.2: RF contact spectroscopy for a ^{85}Rb BEC. (a) Schematic of the contact transition. We drive the interacting state into a lower magnetic sublevel, but the interaction-induced energy requires less energy from the RF photon to effect the spin-flip. This results in the transition frequency of the contact to be lower than the single atom transition. (b) Example of spectroscopy signal $S(\omega)$, normalized so that $\int_{-\infty}^{\infty} S(\omega)d\omega = 1 \text{ s}^{-1}$. (c) The same signal as (b), shown on a smaller scale. The solid red line is a fit to the expected frequency dependence from Equation 6.6, while the dotted blue line shows a fit ignoring $\beta(\omega)$. The green line shows the expected Gaussian signal from the resonant lineshape. On the positive side of the transition the signal is consistent with zero. (d) The tail signal multiplied by $|\omega|^{3/2}$. This is how the fits to the tail were performed, and more clearly shows the final-state effects. The density here is $\langle n \rangle = 5.8 \times 10^{12} \text{ cm}^{-3}$. The data from the tail and main lineshape come from different experimental runs, but are shown together for illustration purposes.

define our signal, $S(\omega)$, as $\Gamma(\omega)$ normalized by the integrated lineshape. A detailed explanation of the experimental procedure and timings is found in Appendix B. The solid line in Figure 6.2b shows a fit to the expected frequency dependence from Equation 6.6, while the dotted line shows a fit to $|\omega|^{-3/2}$. One can clearly see the data deviate from the $|\omega|^{-3/2}$ line, owing to the final-state effects manifested in $\beta(\omega)$. The corrections due to $\alpha(a)$ will be illustrated in Section 6.4.

The atoms outcoupled on the $|\omega|^{-3/2}$ tail have kinetic energy from the interactions, and the kinetic energy released is given by the detuning from resonance. This excess energy is shared between two pairwise interacting atoms, meaning that each atom will have a kinetic energy equal to $\frac{1}{2}\hbar\omega$. The absolute direction of this motion is random, so on average the energy is divided equally into the three degrees of freedom. Thus we expect the radius w of the outcoupled cloud to expand in time t due to the kinetic energy as $w_{\text{KE}}(t) = \sqrt{\frac{\hbar\omega}{3m}}t$. Adding our imaging resolution $w_{\text{res}} = 7.5 \mu\text{m}$ in quadrature gives us the expected size versus time, $w(t) = \sqrt{w_{\text{KE}}(t)^2 + w_{\text{res}}^2}$. We see precisely this expansion in Figure 6.3a. The Gaussian RF pulse, in this case cut off at $\pm 2\tau$, is shown for reference. Here the cloud was still confined in a trapping potential, but as the breathe period is 50 ms, we expect the initial behavior to resemble that of an untrapped cloud. Also shown is the measured peak OD of the cloud in Figure 6.3b, which falls off as $1/w^2$, and illustrates the benefit of imaging as soon as possible after the RF pulse to maximize the signal (as described in Chapter 5). Conversely, if we look at the energy of the expanded cloud versus detuning, we again see the expected kinetic energy, shown by the line in Figure 6.3c. In this case the measurement was taken in expansion at $t = 4.5$ ms. We calculate the energy from the width of the expanded cloud using $E = \frac{3}{2}m\frac{w^2 - w_0^2}{t^2}$, where w_0 is the size of the expanded cloud measured at $\omega = 0$.

6.3 Extracting the contact

The derivations in the literature assume an infinitely narrow probe in frequency that measures the RF spectrum, along with an easily measured Rabi frequency. But in reality, we must limit ourselves to finite probe times, which necessarily broadens the RF probe pulse in frequency. In Chapter 4 I explained that we utilize Gaussian-shaped RF pulses for our contact measurements.

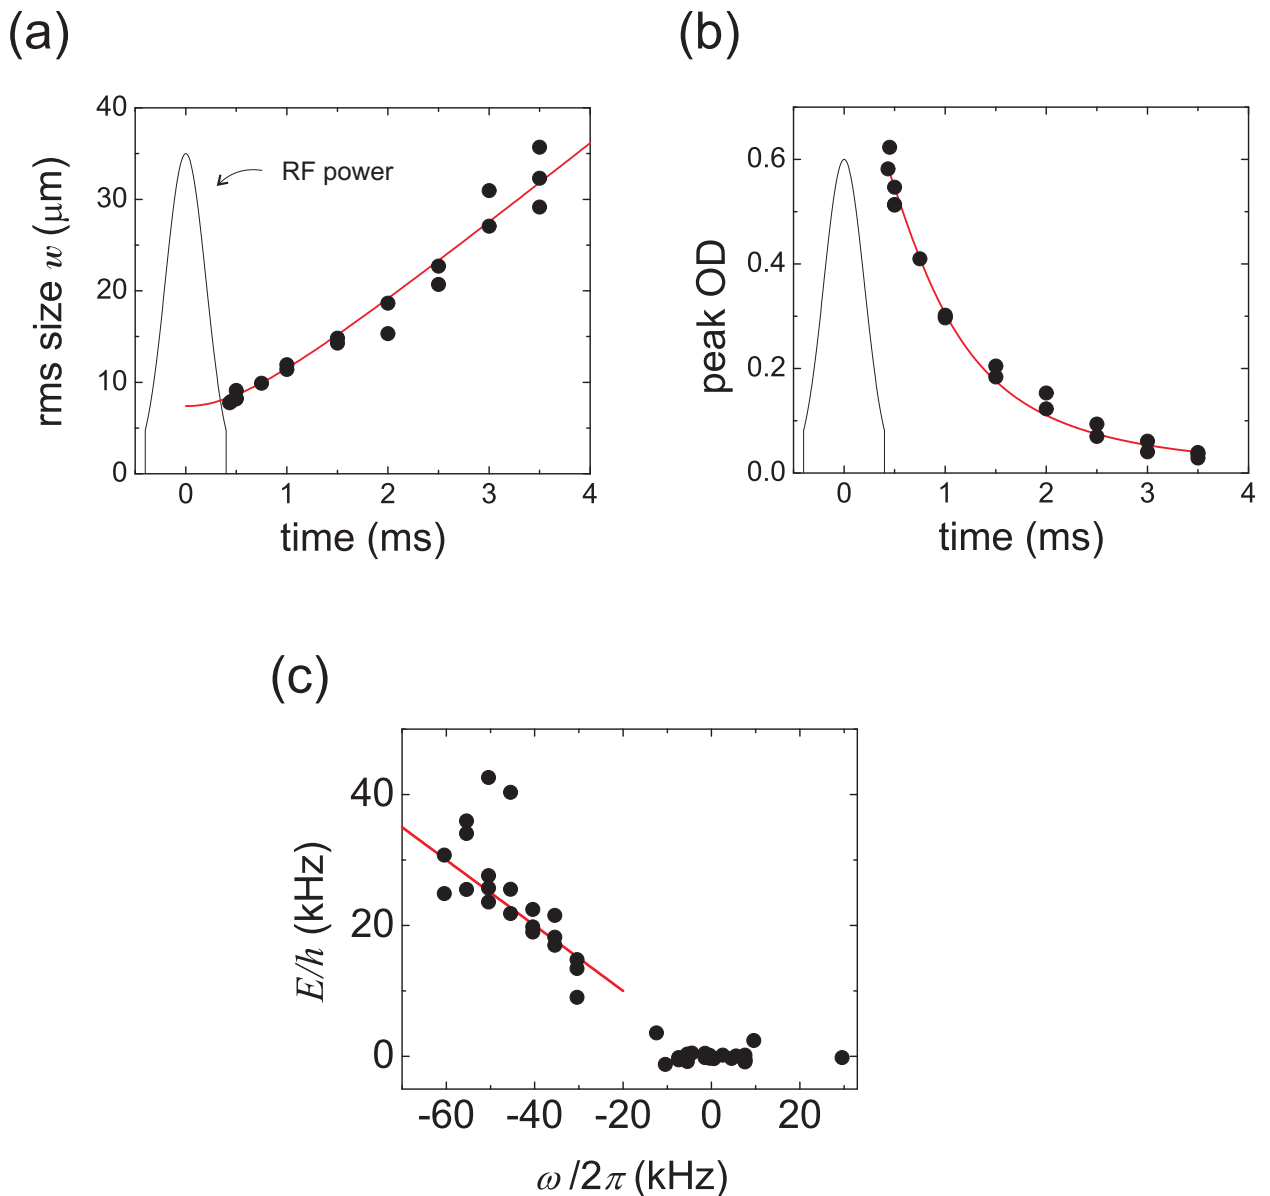


Figure 6.3: Expansion of the outcoupled atom cloud. (a) The size of the cloud of outcoupled atoms after the contact pulse with $|\omega| = 2\pi \times 50$ kHz detuning. The black line shows a sketch of a truncated Gaussian RF pulse for reference. The cloud expands with a kinetic energy of $\frac{1}{2}\hbar\omega$, as the excess energy is shared between two atoms. The red line is the predicted size due to the kinetic energy, added in quadrature with the resolution limit of our imaging system. (b) The peak OD of the clouds plotted in part (a). The red line is a fit to a $\frac{1}{w^2}$ dependence ($\sim \frac{1}{t^2}$), varying only an overall scaling factor. (c) The energy of the outcoupled cloud as a function of detuning, at 4.5 ms time-of-flight expansion. We calculate the energy from the size of the outcoupled cloud, accounting for the size of the cloud of non-spin-flipped atoms at $\omega = 0$. The solid line is $\frac{1}{2} \frac{|\omega|}{2\pi}$.

This section describes in greater detail how we rewrite the equations for the contact in terms of finite Gaussian pulses, and how we extract a value of the contact from our measured quantities. In Equation 6.6, the single atom Rabi frequency Ω is really a Gaussian function of time given by $\Omega(t)^2 = \Omega_p^2 e^{-\frac{t^2}{\tau^2}}$, where Ω_p is the peak Rabi frequency of the Gaussian. If we integrate over all time, we get a slightly modified version of Equation 6.6:

$$\int_{-\infty}^{\infty} \lim_{\omega \rightarrow \infty} \Gamma(\omega, t) dt = N_{\text{tail}} = \frac{\tau \Omega_p^2}{4\sqrt{\pi}} \sqrt{\frac{\hbar}{m}} \frac{\alpha(a)}{\beta(\omega)} \frac{C_2}{\omega^{3/2}}, \quad (6.9)$$

where N_{tail} is the number of atoms we measure on the tail for a single pulse at ω . We do not directly measure the Rabi frequency, but instead use the integral of the main lineshape to normalize the data. This ensures that issues such as imaging efficiencies and efficiencies of the transfer to the imaging state are common-mode and cancel out. As above, the Rabi frequency in Equation 6.7 is also a Gaussian in time, so

$$\int_{-\infty}^{\infty} \Gamma(\omega, t) d\omega = \pi \Omega_p^2 e^{-\frac{t^2}{\tau^2}} N. \quad (6.10)$$

Integrating this over both time and frequency, and solving for Ω_p^2 , we get

$$\Omega_p^2 = \frac{\sqrt{2} N_{\text{peak}} \sigma_{\text{LS}}}{\pi N_{\text{tot}} \tau}, \quad (6.11)$$

where N_{peak} is the number of atoms we measure on the peak of the transition, N_{tot} is the total number of atoms, and σ_{LS} is the Gaussian frequency width of the lineshape. If we assume that we have a condensate fraction f and only BEC number N_0 contributes to the contact, then we can combine Equations 6.9 and 6.11 to solve for the contact per particle:

$$\frac{C_2}{N_0} = \frac{4\pi^2}{I \cdot f} \sqrt{\frac{m}{\hbar}} \alpha(a) \left(N_{\text{tail}} \beta(\omega) |\omega|^{3/2} \right), \quad (6.12)$$

where $I = \sqrt{2\pi} N_{\text{peak}} \sigma_{\text{LS}}$ is the integrated lineshape. We use different powers and pulse lengths for measuring N_{peak} and N_{tail} to outcouple a small fraction of the cloud, so we have to scale the numbers accordingly. We check that for both cases the outcoupled numbers are linear in time and power, and correct for a small deviation from linearity. We have found that the deviation seems to be dependent only on the outcoupled fraction, as long as we stay far below the saturation point of

our RF amplifier. In Figure 6.4 we show the deviation from linearity of both the peak (a) and the tail (b) as a function of expected outcoupled fraction f_{exp} , where the expected fraction is calculated such that it asymptotes to the measured fraction f_{meas} for low values. The fit is a simple saturation model given by

$$f_{\text{meas}} = \frac{1}{\frac{1}{f_{\text{exp}}} + \frac{1}{f_{\text{aysm}}}}, \quad (6.13)$$

where $\frac{f_{\text{meas}}}{f_{\text{aysm}} - f_{\text{meas}}}$ is the deviation from linearity. We outcouple 1-2% of the atoms on the tail of the transition, and 2-3% on the peak, resulting in a $\sim 5\%$ correction for both. These mostly cancel, so the final correction due to nonlinearities is on the order of 1%.


We also know that non-condensed (thermal) atoms contribute to the contact signal as well, so we have to account for them. In order to know what our thermal fraction is, we make sure to leave enough thermal component in the cloud to be able to clearly measure two separate components, while keeping the condensate fraction as high as possible. This results in a condensate fraction of $\sim 70\%$ for most of our data. We then model the contribution of the thermal atoms to the contact based on the energy density and the spatial overlap of the two components. The total interaction energy is given by (see Equation 13.15 in [36]),

$$E_{\text{int}} = \int \left[\frac{g}{2} \left(n_B(\mathbf{r})^2 + \frac{128}{15\sqrt{\pi}} \sqrt{n_B(\mathbf{r})^5 a^3} \right) + 2gn_B(\mathbf{r})n_T(\mathbf{r}) + gn_T(\mathbf{r})^2 \right] d\mathbf{r}, \quad (6.14)$$

where n_B is the density of BEC atoms and n_T the density of thermal (non-condensed) atoms, and we have included the LHY correction to the energy in the first term. We then use this result in conjunction with the adiabatic sweep theorem (Equation 6.3) to calculate the total contact including the thermal contribution. We repeat the calculation for the BEC atoms (just the first term in Equation 6.14), and compare the two to get a fractional contribution from the thermal atoms, with which we adjust our data. The calculated correction due to thermal atoms is 5-12% for our range of data. We check that the contact on a purely thermal cloud is consistent with the calculated value of

$$\frac{C_{\text{therm}}}{N} = 32\pi^2 na^2. \quad (6.15)$$

(a) Peak saturation



(b) Tail saturation

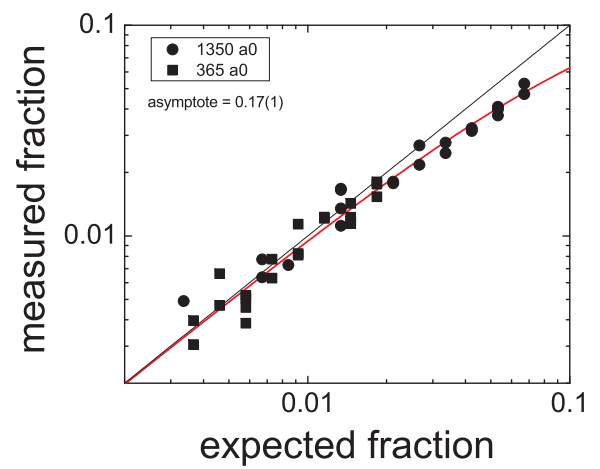


Figure 6.4: Deviation from signal linearity. On both the peak (a) and tail (b) of the transition, we measure the outcoupled fraction f_{meas} of atoms vs RF pulse time and/or power, and plot it vs the fraction expected f_{exp} if the dependence were linear. We fit this to the simple saturation model of Equation 6.13 to extract a value for the asymptote f_{asym} and calculate the magnitude of the correction to the data. In (b) we combine data from two different scattering lengths for a larger range in signal. Individual fits to the two sets give the same result within the error bars.

The very low signal (due to the low densities) currently prevents us from measuring it with high precision, but future technical improvements should be able to allow the experimenters to further explore the thermal contact.

Finally, to compare the contact to the theory prediction given by Equation 6.4, we need to know both the scattering length and the density of the cloud during the RF probe pulse. Precise measurements of the Feshbach resonance, which are used to determine the scattering length as a function of magnetic field, were performed in [35]. Since we determine the resonant RF frequency for every contact measurement, we can exactly calculate the magnetic field, which gives us the scattering length a . The error bars from the Feshbach resonance measurements, along with the uncertainty in the magnetic field, result in final errors on the scattering length of roughly 1% in the range that we probe (300 to 3500 a_0). We determine the density of the BEC using a numerical model based on mean-field interactions (the PG model described in Section 5.1.1), calculating the space- and time-averaged density of the cloud after ramping to various values of a . Since we have a weak spherical trap and a fast probing technique, our average density does not drop more than 20% for fast ramps in a over our full range of scattering lengths.

6.4 Contact measurements

To investigate LHY physics represented by Equation 6.4, we measure the contact for different values of the scattering length. Here is where the final-state effects due to $\alpha(a)$ manifest themselves. The contact signal comes from the fact that we project the wavefunction of the interacting particles onto a state with different interactions. If the interactions of that final state were exactly the same as the initial state, we would not expect to outcouple any atoms. Thus the size of the signal is highly dependent on the difference between the scattering lengths of the final and initial state. The effect of $\alpha(a)$ then, is to shift what would be a parabola centered about $a = 0$ to one centered about $a' = -565 a_0$, which enhances our signal at small a . This effect can be seen in Figure 6.5. The solid line is the mean-field prediction including final-state effects, whereas the dotted line ignores them. For these data, the space-averaged density has been averaged over the different points to




Figure 6.5: The contact vs a , measured at $|\omega| = 2\pi \times 40$ kHz. (a) Here I plot $\frac{\alpha(a)}{\beta(\omega)} \frac{C_2}{N_0}$, which is directly proportional to the strength of the measured $\omega^{-3/2}$ RF tail. (b) The contact per particle $\frac{C_2}{N_0}$. The solid lines in (a) and (b) are the mean-field predictions. The final-state effects shift what is a parabola centered about $a = 0$ in (b) to one centered about $a' = -565 a_0$ in (a), which enhances our signal at small a .

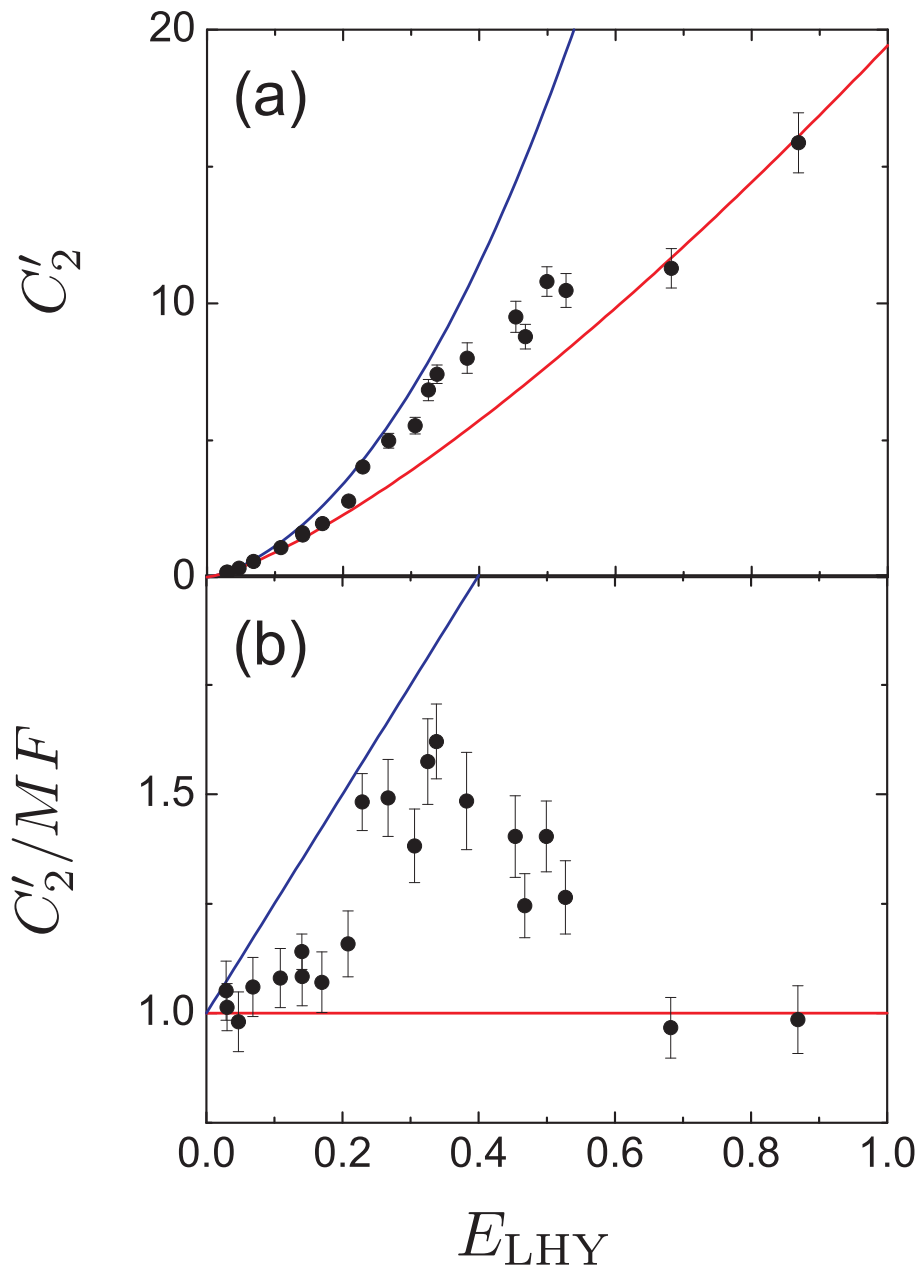


Figure 6.6: Contact measurements as a function of the LHY energy E_{LHY} . (a) The contact per particle normalized by density^{1/3}. The red curve shows the mean-field prediction for the contact, and the blue curve shows the total contact including the LHY term. In (b) the data are normalized to the mean-field value of the contact to compare to theory more clearly. In the limit of low interaction strength, the data match the theory but cannot distinguish the LHY contribution. At higher values, the data are systematically low. At values of E_{LHY} approaching unity, we do not expect the perturbative LHY theory to be valid.

be $\langle n \rangle = 5.8 \times 10^{12} \text{ cm}^{-3}$. Since the data generally have varying densities, depending on the scattering length that we ramp to, we have to normalize out the density. If we normalize C_2 to the interparticle spacing $n^{1/3}$ and BEC number N_0 , we can rewrite Equation 6.4 to give the contact in its dimensionless form C'_2 :

$$C'_2 = \frac{C_2}{N_0 n^{1/3}} = 16\pi^2 \left(\sqrt{na^3} \right)^{4/3} \left(1 + \frac{5}{2} \frac{128}{15\sqrt{\pi}} \sqrt{na^3} + \dots \right). \quad (6.16)$$

In Figure 6.6a we plot C'_2 versus E_{LHY} , where $E_{\text{LHY}} = \frac{128}{15\sqrt{\pi}} \sqrt{na^3}$ is the value of the fractional LHY correction to the energy. The red line is the mean-field prediction, and the blue line includes the contribution from the LHY term. The advantage of our measurement method is clearly evident in the range of the x-axis. We can measure the contact up to such high interaction strengths that E_{LHY} can no longer be said to be a perturbation on the mean-field energy, and we no longer expect the LHY result to hold.

In order to see the data more clearly, we then normalize the contact per particle by the mean-field prediction to get

$$\frac{C_2}{N_0 16\pi^2 na^2} = \left(1 + \frac{5}{2} E_{\text{LHY}} + \dots \right), \quad (6.17)$$

shown in Figure 6.6b. The mean-field value is now 1, and the LHY term is a line with a slope of 5/2. For the lowest values of E_{LHY} , where the separation between the mean-field and the LHY prediction is on the order of the error bars, the data show excellent agreement with the predicted value. At values of E_{LHY} approaching unity, we would not expect the perturbative LHY theory to be valid, so the mismatch between data and theory comes as no great surprise. However, we do expect the LHY theory to be valid in the intermediate regimes, so we must take a closer look at our methods.

When we change the scattering length, we ramp a as quickly as possible to avoid the destructive three-body loss rates. The resulting increase of interaction energy brings the system out of global equilibrium and starts a breathe oscillation of the cloud. However, an ongoing assumption is that the system is in equilibrium locally, because we change a on a timescale that is adiabatic

with respect to the binding energy of a dimer, with \dot{a}/a never reaching more than $0.01\hbar/(ma^2)$ (\dot{a} being the time derivative of a). One might then wonder if we are possibly ramping too quickly for the physics underlying the LHY correction to evolve. Going one step further, we can ask what determines the relevant timescale. One (possibly naive) guess would be to say that the relevant time scale is given by the mean-field chemical potential $\mu = \frac{4\pi\hbar^2}{m}na$. We then look at how large the fractional rate of change $\dot{\mu}/\mu$ of this energy is, compared to the characteristic time given by that energy, μ/\hbar . This gives us a unitless “adiabaticity parameter” (\mathcal{R}) for our ramps,

$$\mathcal{R} = \frac{\dot{\mu}/\mu}{\mu/\hbar}. \quad (6.18)$$

For the data in Figure 6.6, \mathcal{R} ranges between 2 and 3. When we ramp to our final scattering lengths, we ramp in such a way to keep \mathcal{R} roughly constant throughout the ramp.

In Figure 6.7 we show the results of changing our ramp rates at two values of E_{LHY} . Here the data are normalized to represent where, fractionally, they lie between the mean-field and LHY predictions. The red line at 0 represents the mean-field contact, and the blue line at 1 is the full contact with the LHY term. We can clearly see a dependence on the ramp rate, suggesting that we are indeed ramping too quickly for the LHY energy to completely manifest itself. Also, the fact that the value of the contact is in between the mean-field and the LHY values right around when our \mathcal{R} is equal to 1, suggests that the chemical potential is the relevant energy scale. For lower values of E_{LHY} , the LHY term is too small compared to the error bars to discern a ramp rate dependence.

A more convincing argument would be to show that the contact actually saturates to the LHY and mean-field values for very slow and very fast ramps, respectively. Unfortunately, achieving these rates present a nontrivial experimental challenge. The fastest ramps we are currently able to achieve are limited by the gain of our servos and the inductance of our magnetic-field coils. This limit could be surpassed by installing a separate low-inductance coil pair driven by a capacitor bank, which could change the fields by the necessary amount as fast as $\sim 5 \mu\text{s}$ [49], giving us another factor of 20 in ramp rate. This is quite feasible, and it is likely that the experiment will

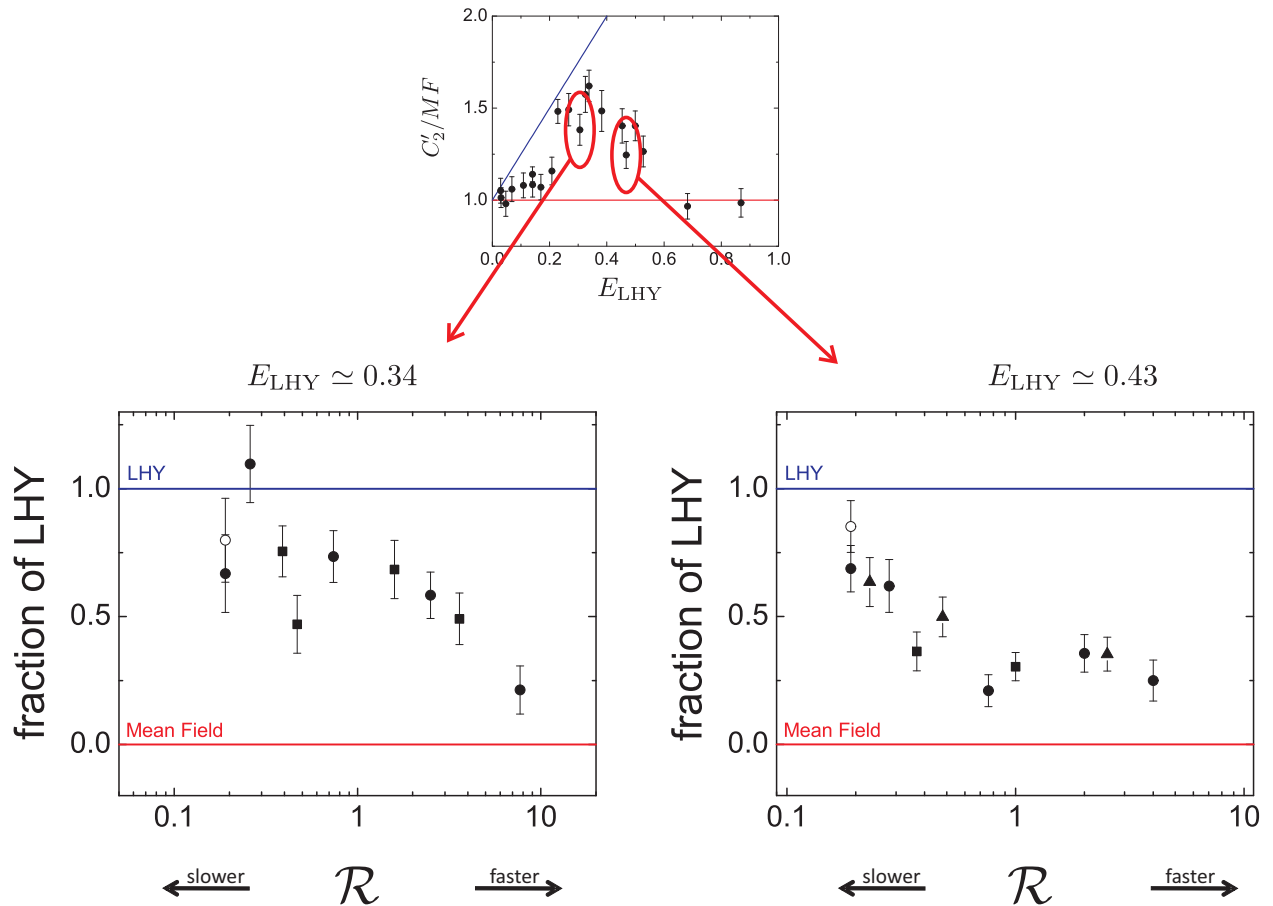


Figure 6.7: Contact measurements as a function of ramp rates, expressed in terms of the “adiabaticity parameter” \mathcal{R} . Higher values of this parameter correspond to faster ramps, which show a value of the contact closer to the mean-field prediction. Slower ramps result in higher values closer to the LHY prediction. Different shapes correspond to different days on which the data were taken. The small upper plot is a reproduction of Figure 6.6b, to illustrate the values of E_{LHY} shown here. For lower interaction strength, the LHY term is too small compared to the error bars to discern a ramp rate dependence. To calculate the mean-field and LHY energies, we require knowledge of the cloud density. The change in densities due to the ramp in a were calculated using a model for expansion that included only the expected mean-field energy, and vary from 40% for the slowest ramps to 10% for the fastest. This reasoning is somewhat circular, but including the LHY term in this model causes at most an extra 5% decrease in density for the slowest ramps, and the resulting change is shown by the open circles.

move in this direction in the not-too-distant future. The other extreme for the ramp rate is more problematic.

Achieving very slow ramp rates is technologically trivial, but two major issues prevent us from easily obtaining trustworthy measurements. One of these is the change in density. When we change a , the resulting change in interaction energy causes the cloud to expand for the start of a breathe mode. We rely on the PG model to calculate the change in density of our cloud from the interaction-induced breathe and obtain the value of E_{LHY} . But the results of that model depends on the energy of the gas, and whether or not we must include the LHY corrections to the energy. For fast ramps the density drop is on the order of 10%, and the inclusion (or exclusion) of the LHY energy in the PG model will not incur a significant total error. For slow ramps the density will drop more like 30% from the initial density, and exact knowledge of the energy is much more important. For the slowest ramps shown in Figure 6.7, using the LHY rather than the mean-field energy causes a density change of $\sim 5\%$, which moves it closer to the LHY prediction as shown by the open circles. Because of this, measuring an LHY term tends to be circular when large density corrections are made. The other major issue is the ever-recurring problem of three-body losses. RF contact spectroscopy is an appealing measurement technique because it can be performed quickly to avoid the loss issues. Taking a long time to ramp a means we once again have to account for significant losses in the condensate for every measurement. Again, this is a fundamental issue inherent in bosons.

To summarize, we have used RF contact spectroscopy to measure beyond-mean-field effects in strongly interacting BEC. At low interaction strengths, where the LHY term is a $< 5\%$ perturbation on the energy, the measured contact shows excellent agreement with theory. At higher interaction strengths, we see clear deviation from the mean-field contribution to the contact. These beyond-mean-field effects are dependent on ramp rate, and evidence suggests that the relevant time scale is determined by the chemical potential. Nonetheless, RF contact spectroscopy is a powerful tool for probing strongly interacting Bose gasses before three-body losses become significant, even letting us reach a regime where the LHY contribution to the energy is calculated to be of order 1.

Chapter 7

Three-body Contact

Beyond some prefactors in the formulas due to particle statistics, the two-body contact applies equally to bosons and fermions. However, the possibility of three-body interactions in bosonic systems means that one has to take into account the possibility of a three-body contact. The three-body contact, similar to its counterpart, measures the probability for *triples* of identical bosons to be close to one another [31]. Continuing the analogy, just as the two-body universal parameter is the scattering length, a , the universal parameter describing the three-body physics is the Efimov parameter κ_* . An important thing to keep in mind is that the three-body contact also connects a few-body interaction parameter with many-body effects. One fundamental property of Efimov physics is that it predicts an infinite series of successively more weakly bound trimers with discrete scaling factor $e^{\pi/s_0} \approx 22.7$, where $s_0 \approx 1.00624$ for identical bosons [50]. A sketch of this behavior can be seen in Figure 7.1. The parameter κ_* can be defined by the energies of trimers at unitarity, where $a \rightarrow \infty$, as

$$E_{\text{trimer}} = \frac{\hbar^2 \kappa_*^2}{m} \left(e^{-2\pi/s_0} \right)^l, \quad (7.1)$$

where l is an integer, and m is the mass of one atom. In this way, we can understand κ_* to be the characteristic momentum of the most tightly bound trimer, and its inverse gives the approximate size of the trimer molecule [50].

The three-body contact C_3 has an “adiabatic sweep theorem” like Equation 6.3, which defines




Figure 7.1: A sketch of Efimov states. The vertical scale gives the wave number $K = \pm\sqrt{m|E|/\hbar^2}$, and the horizontal gives $1/a$, such that the origin corresponds to unitarity ($a \rightarrow \infty$). The green line shows the state corresponding to a bound dimer plus a free atom, and the blue lines show the bound trimers, for which there is an infinite series approaching unitarity. Note that the scaling factor here is ~ 2 to make the behavior visible and to match the literature on Efimov states. A realistic Efimov state with scaling factor 22.7 would have an energy that is nearly indistinguishable from the dimer energy on the scale of this plot.

C_3 in terms of the derivative of the energy

$$\left(\kappa_* \frac{\partial E}{\partial \kappa_*}\right)_a = -\frac{2\hbar^2}{m} C_3. \quad (7.2)$$

The three-body contact also manifests itself in an additional contribution to the tail of the momentum distribution that goes as

$$\lim_{k \rightarrow \infty} n(k) = \frac{F(k)}{k^5} C_3, \quad (7.3)$$

where the log-periodicity of Efimov physics manifests itself in the function

$$F(k) = A \sin[2s_0 \ln(k/\kappa_*) + 2\phi] \quad (7.4)$$

with numerical constants $A = 89.2626$ and $\phi = -0.669064$ [31]. This results in an additional term to the RF tail at large detunings that should be added to the right-hand side of Equation 6.6:

$$\frac{\hbar\Omega^2}{2m} \frac{G_{\text{RF}}(\omega)}{\omega^2} C_3. \quad (7.5)$$

Here, $G_{\text{RF}}(\omega)$ is a log-periodic function given by

$$G_{\text{RF}}(\omega) = 9.23 - 13.6 \sin \left[s_0 \ln \left(\frac{m\omega}{\hbar\kappa_*^2} \right) + 2.66 \right]. \quad (7.6)$$

As of the writing of this thesis, there is no prediction for final-state effects on the C_3 contribution to the RF tail.

Our aim, then, is to see if we can find a contribution to the interaction-induced RF tail from three-body effects, or to show that possible three-body effects are dominated by two-body correlations and can be ignored in measurements of the contact. Although four-body interactions (or greater) can occur, it is expected that they do not require extra universal parameters, which means that there is no need to define a four-body (or higher) contact. The total contact in the universal regime can be completely described by C_2 and C_3 [51].

7.1 The Efimov resonance

To measure a three-body contact via Equation 7.5, we need to know the value of κ_* . This will tell us the frequencies for which $G_{\text{RF}} = 0$ and the ones for which G_{RF} is high, which can in turn aid us in maximizing the chances of detecting C_3 . One way to measure κ_* , which is a value unique to each atomic species, is to find the value of the scattering length a_- on the negative side of the Feshbach resonance where the energy of a trimer becomes degenerate with the energy of three free atoms (see Figure 7.1). At this value of a we expect that the probability of three atoms close together is resonantly enhanced, manifesting itself in a peak in the three-body recombination rate. Similar measurements of Efimov resonances have been reported for several other ultracold atom systems [52, 53, 54, 55, 56]. For identical bosons, the value of a_- is related to the three-body parameter through $\kappa_* = -1.56(5)/a_-$ [57].

We measure the ^{85}Rb Efimov resonance using non-condensed clouds of 1.5×10^5 atoms at a temperature $T = 80$ nK. We ramp the magnetic field to realize the desired scattering length on the $a < 0$ side of the Feshbach resonance and wait for a variable time t . We then turn off the magnetic fields (except for a small quantization field), ARP all the atoms to the $|3, -3\rangle$ imaging state, and use high-intensity imaging (see Section 5.1.3) to probe the clouds at 3 ms expansion time. This time is small compared to the trap frequency, so the clouds still exhibit position information (we expect the cloud size to have increased by 2% during this time). We measure the number and size

vs time, and extract the three-body event rate constant K_3 for various values of the scattering length.

The three-body event rate constant for a homogeneous gas with a number density n is defined by

$$\frac{d}{dt}n = -3K_3n^3, \quad (7.7)$$

where the factor of 3 in front comes from the fact that, in our case, three atoms are lost per event. But since losses also cause heating, which contributes to a decrease in density, we integrate Equation 7.7 over all space for the number loss rate

$$\frac{d}{dt}N = -3K_3\langle n^2 \rangle N, \quad (7.8)$$

where $\langle n^2 \rangle$ is the density-weighted mean square density $\langle n^2 \rangle = \frac{\int n^3 dr}{\int n dr}$. Our 500 s vacuum-limited lifetime and previous experiments on ^{85}Rb suggest that one- and two-body losses can be ignored for the range of magnetic fields we probe [58]. For a thermal distribution in a harmonic trap, we can rewrite this as

$$\frac{d}{dt}N = -3K_3 \frac{8}{\sqrt{27}} \frac{N^3}{V^2}, \quad (7.9)$$

where the volume $V = 8\pi^{3/2}w^3$ and w is the rms width of the spherical cloud. We perform a linear fit to the volume of the cloud given by $V(t) = A + Bt$. Plugging this into Equation 7.9 and solving for $N(t)$ gives us our fitting function for the rate:

$$N(t) = N_0 \sqrt{\frac{A(A + Bt)}{\frac{48}{\sqrt{27}} N_0^2 K_3 t + A(A + Bt)}}. \quad (7.10)$$

Figure 7.2 shows an example of this procedure. The volume vs time is fit in Figure 7.2a, setting A and B for the fit to the number data in 7.2b via Equation 7.10. Using this procedure, we extract K_3 values for a range of a , plotted in Figure 7.3. We fit the measured K_3 vs a to the expected form for an Efimov resonance for non-condensed atoms [57],

$$K_3 = \frac{4590 \sinh(2\eta)}{\sin^2[s_0 \ln(a/a_-)] + \sinh^2 \eta} \frac{\hbar a^4}{m}. \quad (7.11)$$




Figure 7.2: Example of a loss rate measurement to extract K_3 . These data correspond to a scattering length of $730 a_0$. Heating of the cloud causes its size to increase as seen in (a). The parameters extracted from a linear fit to the volume go into a fit to the atom number via Equation 7.10, as seen in (b), resulting in a value of $K_3 = 5.64 \times 10^{-22} \text{cm}^{-6}/\text{s}$.




Figure 7.3: A three-body loss resonance for ^{85}Rb . We plot the three-body event constant K_3 vs the scattering length a , for clouds with a temperature of roughly 8 nK. From fitting Equation 7.11 to the black points, for which $a < 1/k_{\text{thermal}}$, we extract $a_- = -759(6)a_0$ and $\eta = 0.057(2)$.

Because this expression comes from a $T = 0$ theory, we only fit the data for $a < 1/k_{\text{thermal}}$, where $k_{\text{thermal}} = \sqrt{2mk_B T}/\hbar$ and k_B is Boltzmann's constant. From the fit, we extract $a_- = -759(6) a_0$ and $\eta = 0.057(2)$. The error in a_- is relatively small because of the fact that the Feshbach resonance for ^{85}Rb was measured with very high precision in [35], and the uncertainty from those measurements dominate our errors. Uncertainties of our measurement of the magnetic field or from the fit to Equation 7.11 are negligible. In fitting our observed loss rates to Equation 7.11, we initially included an overall multiplicative scaling factor, to account for systematics in our measurements, but found it to be consistent with one. We do not expect to be able to measure another resonance a factor of 22.7 more negative in scattering length ($-17,000 a_0$) due to the finite temperature of the gas. A factor of 22.7 lower ($-33 a_0$), is smaller in magnitude than the van der Waals length, where universality no longer applies.


Since the η parameter essentially determines the width of the resonance, we can think of it as proportional to the inverse lifetime of the trimers. The trimers can decay into a deeply bound dimer and a free atom. The width can conceivably increase if the temperature of the cloud is too hot, due to the increased thermal de Broglie wavelength moving the system closer to unitarity. To check that we are in the regime where η is independent of temperature, we repeat the measurement of η for 30 and 140 nK. The results are plotted in Figure 7.4. We can see that over a factor of four in temperature, η does not show significant change, suggesting that we are indeed in the low-temperature limit.

The value of a_- is not expected to be universal, since it comes from the details of the short-range potential [15, 59, 57]. However, it gets interesting when expressed in units of the mean scattering length of the van der Waals potential [60]

$$\bar{a} = \frac{1}{\sqrt{8}} \frac{\Gamma(\frac{3}{4})}{\Gamma(\frac{5}{4})} \left(\frac{mC_6}{\hbar^2} \right)^{1/4}, \quad (7.12)$$

where C_6 is the coefficient of the van der Waals potential. For ^{85}Rb , $\bar{a} = 78.5 a_0$, so we find a value for $\frac{a_-}{\bar{a}}$ of $-9.67(7)$. Note that in other works, a_- is sometimes normalized instead by the van der Waals range $R_{\text{vdW}} = \frac{1}{2} \left(\frac{mC_6}{\hbar^2} \right)^{1/4}$ [34]. It turns out that other experiments that have measured

(a)



(b)




Figure 7.4: Measurement of η vs temperature. (a) The loss resonances. For the data at 30 and 140 nK, we did not calibrate the scattering length as carefully as the data at 80 nK, which is the same data as in Figure 7.3 shown over a smaller range of a . Moreover, the 30 nK clouds are likely not in thermal equilibrium, and the ensuing error in the average density will cause a systematic error in the calculated value of K_3 . (b) Measured values of η vs temperature. We do not see a significant change in η over a large range in temperature, suggesting that the experiments are performed in the low-temperature limit.

Efimov resonances have found very similar result in ^{133}Cs [61], ^6Li [62], and ^7Li [55, 56, 63], with values for $\frac{a_-}{a}$ ranging 8-10. This empirical evidence suggests that the three-body parameter depends only on the coefficient of the $1/r^6$ part of the two-body potential and not on the details of a three-body potential at short range [61].

7.2 Searching for C_3

From the measured location of the Efimov resonance, we can extract $\kappa_* = 39(1) \mu\text{m}^{-1}$. This gives us the expected shape of $G_{\text{RF}}(\omega)$, which is plotted in Figure 7.5. Note that $G_{\text{RF}}(\omega)$ has a node at $\omega \simeq 2\pi \times 27 \text{ kHz}$ and a smaller magnitude at larger ω . Since the C_3 term in Equation 7.5 goes as $G_{\text{RF}}(\omega)/\omega^2$, this suggests that the largest contribution from C_3 will be for $\omega < 2\pi \times 27 \text{ kHz}$. Like the C_2 term, the prediction for the C_3 term is valid for $\omega \rightarrow \infty$. For the case of the C_2 term, the RF tail arises from two-body short-range correlations at distances that are small compared to the interparticle spacing, which is always satisfied. However, for the case of C_3 , the prediction for the three-body contribution to the RF tail may have a more limited range of applicability [64]. In particular, the C_3 theory may only be applicable for $\omega > \frac{\hbar}{ma^2}$ [65], where the frequency dependence makes it less likely to contribute significantly to the RF tail.

The results of our search for C_3 can be seen in Figure 7.7, where we examine the frequency dependence of the RF tail for a BEC at $a = 982 \pm 10 a_0$. To maximize our chances of measuring a C_3 contribution, we increase the density of the clouds by 65% by jumping the scattering length from $100 a_0$ to $50 a_0$ and waiting 1/4 of the trap period before ramping to the final value of $982 a_0$. This creates a density oscillation as seen in Figure 7.6, which we predict using the PG model (see Section 5.1.1). We fit the data to the predicted frequency dependence of the C_2 contribution, shown by the solid line. The dotted line is the same fit but shown without including the final-state correction $1/\beta(\omega)$. We can see that our data fit very well to the expected frequency-dependence for the two-body contact with final-state effects, and we do not observe any deviation consistent with a three-body term. Fitting the data to both contributions gives an upper limit for C_3/N_0 of $0.07 \mu\text{m}^{-2}$. Also shown is a trial C_3/N_0 term of $0.1 \mu\text{m}^{-2}$ with the dashed line. For reference,




Figure 7.5: G_{RF} as a function of ω , with $\kappa_* = 39(1) \mu\text{m}^{-1}$. We plot G_{RF} for negative ω to ease comparison with the measured RF tail, which occurs at negative detunings in our experiment. G_{RF} has a node at $\omega \simeq 2\pi \times 27$ kHz, suggesting that one should look for a C_3 contribution to the RF tail for smaller detunings.

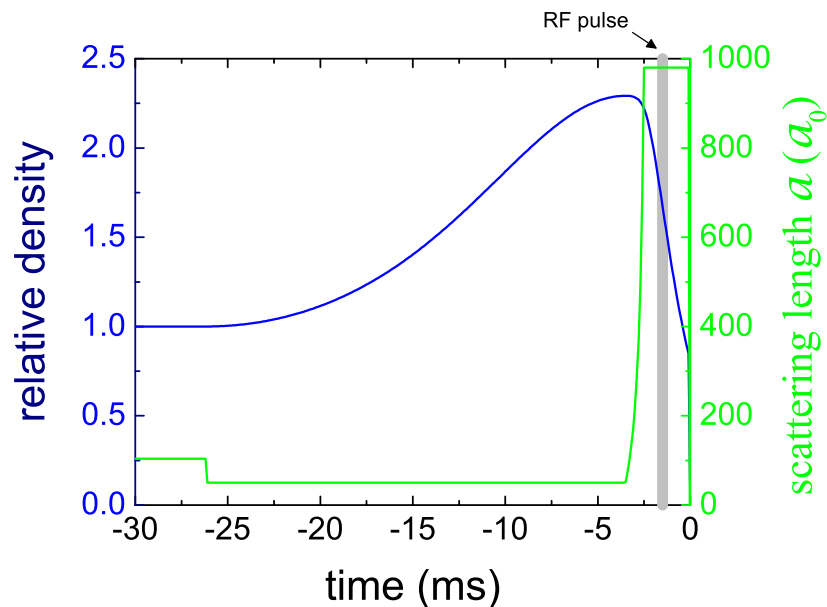


Figure 7.6: Changing the scattering length for increased density. We jump the scattering length a (green curve) to $50 a_0$, causing an inward breathe with a period of roughly 50 ms. Close to the turnaround point we ramp to $982 a_0$ and apply the RF pulse to measure the contact (grey line). The blue curve shows the prediction for the relative density of the PG model.




Figure 7.7: (a) The calculated frequency dependence of $G_{\text{RF}}(\omega)$, shown again on a logarithmic ω scale for reference. (b) The measured frequency dependence of the tail of the RF spectrum for $a = 982 \pm 10 a_0$. The solid red line is a fit to the expected frequency dependence of the two-body contact C_2/N_0 including final-state effects. The dotted blue line corresponds to the same value of C_2/N_0 , but ignores final-state effects. For comparison, the fit plus a trial C_3/N_0 term of $0.1 \mu\text{m}^{-2}$ is shown with the dashed black line. Our measurements are consistent instead with a C_3/N_0 of zero. For these data the density is $\langle n \rangle = 1.0 \times 10^{13} \text{ cm}^{-3}$.

$\frac{\hbar}{ma^2} = 277$ kHz for this scattering length.

Another idea was to look for a C_3 contribution on the peak of the Efimov resonance. The thought was that the final-state effects suppress the two-body contribution to the contact, while the presence of the Efimov resonance might enhance the contribution from the three-body contact. Since a BEC begins to implode at negative scattering length, we first create condensates at $100 a_0$, then we decrease the density by a factor of 5 by jumping a to $400 a_0$ and waiting $1/4$ of a trap period, before jumping the magnetic field to the Efimov resonance. This should give us ~ 4 ms to measure the RF spectrum before the condensate begins to collapse locally (see Section 5.1.2 for an explanation of this prediction). The results of that test are shown in Figure 7.8, where all measurements of the tail are consistent with zero. The error bars roughly represent our detection limit. The two-body contact prediction is shown by the green line, which is below our detection limit due to the aforementioned final-state suppression. The blue line shows a heuristic prediction for the three-body contact from the loss rate, given by [65]

$$\frac{d}{dt}N = -4\eta\frac{\hbar}{s_0m}C_3. \quad (7.13)$$

Given that this prediction lies significantly above our detection limit, the data suggest that the $\omega > \frac{\hbar}{ma^2}$ limit *does* define the range of applicability for the C_3 theory, and/or that final-state effects for the three-body contribution also cause significant suppression of the signal, or that Equation 7.13 is incorrect. Of interest is the prediction of Equation 7.5, that if the scattering length were chosen such that the C_2 term is zero, the transition rate goes *negative*, indicating that the theory is at this point still incomplete.

In short, our investigations for a three-body contribution to the contact have shown that the short-range correlations in the BEC are dominated by two-body effects, as we see no clear signature of three-body effects in the frequency dependence of the interaction-induced tail in RF spectroscopy.




Figure 7.8: Searching for a tail in the RF spectrum at $-760 a_0$. We detected no signal indicative of a tail on the RF spectrum. The error bars give an estimate of our detection limit. The green line shows the prediction for the contribution to the signal S from the two-body contact, which is highly suppressed due to final-state effects and expected to be below our detection limit. A heuristic calculation from loss rates via Equation 7.13 predicts a three-body contribution shown by the blue line. The density here is $\langle n \rangle \simeq 7 \times 10^{11} \text{ cm}^{-3}$.

Chapter 8

Summary and Outlook

8.1 In Essence

In this thesis I presented an experiment on ^{85}Rb BEC near a Feshbach resonance, for which we measured the contact via RF spectroscopy. To fulfill the technical requirements for this experiment we stabilized our trapping magnetic fields to ~ 2 mG rms fluctuations, and minimized the magnetic field variation across the cloud to a ~ 4 kHz rms width on a 0.47 MHz/G transition. We also created Gaussian-shaped RF probing pulses without stray spectral components down to 60 dB below the peak power. To probe our BECs we developed imaging technology that easily spans imaging in magnetic fields of zero to upwards of 200 G, giving us the ability to probe clouds with optical depths greater than 6. We also developed a technique to efficiently measure only the atoms outcoupled from the tail of an RF lineshape, without letting the 99% of atoms in the original state contaminate our images.

With these technical requirements resolved, we measured the two-body contact, experimentally probing the tail of the RF spectrum to high precision. We clearly see a deviation from the $1/|\omega|^{3/2}$ tail consistent with predictions for the final-state effects. From the ratio of the transition rate on the tail to the resonant transition rate, we extracted a value for the two-body contact. We then measured the contact for varying values of the scattering length, ranging from a regime that is well-described by the mean-field approach, through a regime where we expect the correction to the energy to be described by the Lee-Huang-Yang (LHY) formalism, all the way to a regime where na^3 is so high that we no longer expect the LHY theory to be valid. In the low a regime our data match

the mean-field prediction, and at higher interaction strengths we see beyond-mean-field behavior. The data do not match the LHY prediction, likely due to the physics underlying the LHY energy not evolving adiabatically with fast changes in the interaction strength. When we vary the ramp rate, we see a clear trend towards the mean-field prediction for the contact for faster ramps and toward the LHY prediction for slower ramps. This opens up an exciting area of research regarding the dynamics of strongly interacting bose gasses.

We also investigated the possibility of a three-body contact, related to Efimov physics. In order to predict the shape of the three-body contribution to the RF tail of the frequency spectrum, we measured the location of an Efimov resonance on the $a < 0$ side of the Feshbach resonance. We measured loss rates on non-condensed clouds to extract the three-body event rate constant K_3 as a function of the scattering length. This gives us a resonance located at $a_- = -759(6) a_0$ and a width given by $\eta = 0.057(2)$. Evidence suggests this data was taken in the low-temperature regime. From the location of the Efimov resonance we calculate the three-body universality parameter $\kappa_* = 39(1) \mu\text{m}^{-1}$. Our subsequent measurements for a three-body contribution to the RF tail are consistent with zero. In the regime of perturbative interactions, such as assumed in the LHY calculation, one would expect that the short-range correlations in the BEC are dominated by two-body effects. In general, this paves the way for using RF spectroscopy to measure the two-body contact for BECs and thus measure beyond-mean-field physics and probe non-equilibrium many-body dynamics.

8.2 The next few weeks...

Strongly interacting systems have long been a challenge for theorists and experimentalists alike, and the ^{85}Rb experiment is poised to investigate much that has not been explored by either group. The dynamics of the LHY correction that we have already measured lack a theoretical foundation, and the experiments on those dynamics will likely be refined and extended to give clear results in the future. The two-body contact in general has proven such a useful tool that the experiment will probably continue to use it to study strongly-interacting BECs. Specifically, work is currently in progress to measure the contact at unitarity, where the scattering length is no

longer a relevant length scale. Instead, the interparticle spacing becomes the relevant parameter, changing the density dependence of three-body losses. Also, I imagine that the book on the three-body contact away from unitarity is not yet completely closed, as the current team will come up with new ideas and techniques to search for it in different regimes. It will be exciting to see what sorts of new and interesting science the ^{85}Rb experiment (through its capable keepers) will reveal in the future.

Bibliography

- [1] J. Wilks, *The properties of liquid and solid helium* (Clarendon P., Oxford, 1967).
- [2] P. Kapitza, Viscosity of Liquid Helium below the λ -Point *Nature* **141**, 74 (1938).
- [3] J. F. Allen and A. D. Misener, Flow of Liquid Helium II *Nature* **141**, 75 (1938).
- [4] M. H. Anderson, J. R. Ensher, M. R. Matthews, and C. E. W. and E. A. Cornell, Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor *Science* **269**, 198 (1995).
- [5] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions *Phys. Rev. Lett.* **75**, 1687 (1995), *ibid.* **79**, 1170 (1997).
- [6] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. Durfee, D. M. Kurn, and W. Ketterle, Bose-Einstein Condensation in a Gas of Sodium Atoms *Phys. Rev. Lett.* **75**, 3969 (1995).
- [7] N. Bogoliubov, On the theory of superfluidity *J. Phys. USSR* **11**, 23 (1947).
- [8] J. Steinhauer, R. Ozeri, N. Katz, and N. Davidson, Excitation Spectrum of a Bose-Einstein Condensate *Phys. Rev. Lett.* **88**, 120407 (2002).
- [9] T. D. Lee and C. N. Yang, Many-Body Problem in Quantum Mechanics and Quantum Statistical Mechanics *Phys. Rev.* **105**, 1119 (1957).
- [10] T. D. Lee, K. Huang, and C. N. Yang, Eigenvalues and Eigenfunctions of a Bose System of Hard Spheres and Its Low-Temperature Properties *Phys. Rev.* **106**, 1135 (1957).
- [11] S. B. Papp, J. M. Pino, R. J. Wild, S. Ronen, C. E. Wieman, D. S. Jin, and E. A. Cornell, Bragg Spectroscopy of a Strongly Interacting *Rb85* Bose-Einstein Condensate *Phys. Rev. Lett.* **101**, 135301 (2008).
- [12] R. P. Smith, R. L. D. Campbell, N. Tammuz, and Z. Hadzibabic, Effects of Interactions on the Critical Temperature of a Trapped Bose Gas *Phys. Rev. Lett.* **106**, 250403 (2011).
- [13] N. Navon, S. Piatecki, K. Gnter, B. Rem, T. C. Nguyen, F. Chevy, W. Krauth, and C. Salomon, Dynamics and Thermodynamics of the Low-Temperature Strongly Interacting Bose Gas *Phys. Rev. Lett.* **107**, 135301 (2011).

- [14] P. O. Fedichev, Y. Kagan, G. V. Shlyapnikov, and J. T. M. Walraven, Influence of Nearly Resonant Light on the Scattering Length in Low-Temperature Atomic Gases *Phys. Rev. Lett.* **77**, 2913 (1996).
- [15] B. D. Esry, C. H. Greene, and J. P. Burke, Recombination of Three Atoms in the Ultracold Limit *Phys. Rev. Lett.* **83**, 1751 (1999).
- [16] S. Tan, Large momentum part of a strongly correlated Fermi gas *Ann. Phys.* **323**, 2971 (2008).
- [17] S. Tan, Generalized Virial Theorem and Pressure Relation for a strongly correlated Fermi gas *Ann. Phys.* **323**, 2987 (2008).
- [18] S. Tan, Energetics of a strongly correlated Fermi gas *Ann. Phys.* **323**, 2952 (2008).
- [19] E. Braaten and L. Platter, Exact Relations for a Strongly-interacting Fermi Gas from the Operator Product Expansion *Phys. Rev. Lett.* **100**, 205301 (2008).
- [20] E. Braaten, D. Kang, and L. Platter, Universal relations for a strongly interacting Fermi gas near a Feshbach resonance *Phys. Rev. A* **78**, 053606 (2008).
- [21] S. Zhang and A. J. Leggett, Universal properties of the ultracold Fermi gas *Phys. Rev. A* **79**, 023601 (2009).
- [22] R. Haussmann, M. Punk, and W. Zwerger, Spectral functions and rf response of ultracold fermionic atoms *Phys. Rev. A* **80**, 063612 (2009).
- [23] D. Blume and K. M. Daily, Universal relations for a trapped four-fermion system with arbitrary s-wave scattering length *Phys. Rev. A* **80**, 053626 (2009).
- [24] F. Werner, L. Tarruel, and Y. Castin, Number of closed-channel molecules in the BEC-BCS crossover *Eur. Phys. J. B* **68**, 401 (2009).
- [25] E. D. Kuhnle, H. Hu, X.-J. Liu, P. Dyke, M. Mark, P. D. Drummond, P. Hannaford, and C. J. Vale, Universal Behavior of Pair Correlations in a Strongly Interacting Fermi Gas *Phys. Rev. Lett.* **105**, 070402 (2010).
- [26] J. T. Stewart, J. P. Gaebler, T. E. Drake, and D. S. Jin, Verification of Universal Relations in a Strongly Interacting Fermi Gas *Phys. Rev. Lett.* **104**, 235301 (2010).
- [27] E. Braaten, D. Kang, and L. Platter, Short-Time Operator Product Expansion for rf Spectroscopy of a Strongly Interacting Fermi Gas *Phys. Rev. Lett.* **104**, 223004 (2010).
- [28] R. Combescot, F. Alzetto, and X. Leyronas, Particle distribution tail and related energy formula *Phys. Rev. A* **79**, 053640 (2009).
- [29] A. M. J. Schakel, Tan Relations in Dilute Bose Gasses, arXiv:1007.3452v1 (unpublished).
- [30] J. M. Pino, Strongly Interacting Bose-Einstein Condensates: Probes and Techniques, Ph.D. thesis, University of Colorado, 2010.
- [31] E. Braaten, D. Kang, and L. Platter, Universal Relations for Identical Bosons from Three-Body Physics *Phys. Rev. Lett.* **106**, 153005 (2011).

- [32] S. B. Papp, Experiments with a two-species Bose-Einstein condensate utilizing widely tunable interparticle interactions, Ph.D. thesis, University of Colorado, 2007.
- [33] J. L. Roberts, Bose-Einstein condensates with tunable atom-atom interactions: the first experiments with ^{85}Rb BECs, Ph.D. thesis, University of Colorado, 2001.
- [34] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, Feshbach resonances in ultracold gases *Rev. Mod. Phys.* **82**, 1225 (2010).
- [35] N. R. Claussen, S. J. J. M. F. Kokkelmans, S. T. Thompson, E. A. Donley, E. Hodby, and C. E. Wieman, Very-high-precision bound-state spectroscopy near a ^{85}Rb Feshbach resonance *Phys. Rev. A* **67**, 060701 (2003).
- [36] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation* (Oxford University Press, ADDRESS, 2003).
- [37] E. Hecht, *Optics* (Addison-Wesley, ADDRESS, 1998).
- [38] Analog Devices, RFI Rectification Concepts, <http://www.analog.com/static/imported-files/tutorials/MT-096.pdf>.
- [39] G. Reinaudi, T. Lahaye, Z. Wang, and D. Gury-Odelin, Strong saturation absorption imaging of dense clouds of ultracold atoms *Opt. Lett.* **32**, 3143 (2007).
- [40] H. J. Lewandowski, Coherences and correlations in an ultracold Bose gas, Ph.D. thesis, University of Colorado, 2002.
- [41] V. M. Pérez-García, H. Michinel, J. I. Cirac, M. Lewenstein, and P. Zoller, Dynamics of Bose-Einstein condensates: Variational solutions of the Gross-Pitaevskii equations *Phys. Rev. A* **56**, 1424 (1997).
- [42] H. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer, New York, 1999).
- [43] D. Steck, Rubidium 85 D Line Data, <http://steck.us/alkalidata/rubidium85numbers.pdf> (unpublished).
- [44] J. J. Sakurai, in *Modern Quantum Mechanics*, edited by S. F. Tuan (Addison-Wesley Publishing Company, Inc., ADDRESS, 1985).
- [45] W. Schneider, V. B. Shenoy, and M. Randeria, Theory of Radio Frequency Spectroscopy of Polarized Fermi Gases, arXiv:0903.3006v1 (unpublished).
- [46] A. Perali, P. Pieri, and G. C. Strinati, Competition between final-state and pairing gap effects in the radio-frequency spectra of ultracold Fermi atoms. *Phys. Rev. Lett.* **100**, 010402 (2008).
- [47] J. L. Bohn, , private communication (unpublished).
- [48] C. Chin and P. S. Julienne, Radio-frequency transitions on weakly bound ultracold molecules *Phys. Rev. A* **71**, 012713 (2005).
- [49] N. R. Claussen, Dynamics of Bose-Einstein condensates near a Feshbach resonance in ^{85}Rb , Ph.D. thesis, University of Colorado, 2003.

- [50] E. Braaten and H.-W. Hammer, Efimov physics in cold atoms *Annals of Physics* **322**, 120 (2007).
- [51] E. Braaten, Universal Relations for Identical Bosons from 3-Body Physics, DAMOP conference talk (unpublished).
- [52] T. Kraemer *et al.*, Evidence for Efimov quantum states in an ultracold gas of caesium atoms *Nature* **440**, 315 (2006).
- [53] S. Knoop, F. Ferlaino, M. Mark, M. Berninger, H. Schobel, H.-C. Nagerl, and R. Grimm, Observation of an Efimov-like trimer resonance in ultracold atom-dimer scattering *Nat Phys* **5**, 227 (2009).
- [54] M. Zaccanti, B. Deissler, C. D'Errico, M. Fattori, M. Jona-Lasinio, S. Muller, G. Roati, M. Inguscio, and G. Modugno, Observation of an Efimov spectrum in an atomic system *Nat Phys* **5**, 586 (2009).
- [55] S. E. Pollack, D. Dries, and R. G. Hulet, Universality in Three- and Four-Body Bound States of Ultracold Atoms *Science* **326**, 1683 (2009).
- [56] N. Gross, Z. Shotan, S. Kokkelmans, and L. Khaykovich, Observation of Universality in Ultracold ^7Li Three-Body Recombination *Phys. Rev. Lett.* **103**, 163202 (2009).
- [57] E. Braaten and H. Hammer, Universality in few-body systems with large scattering length *Phys. Rep.* **428**, 259 (2006).
- [58] J. L. Roberts, N. R. Claussen, S. L. Cornish, and C. E. Wieman, Magnetic Field Dependence of Ultracold Inelastic Collisions near a Feshbach Resonance *Phys. Rev. Lett.* **85**, 728 (2000).
- [59] J. P. D'Incao, C. H. Greene, and B. D. Esry, The short-range three-body phase and other issues impacting the observation of Efimov physics in ultracold quantum gases, 2009.
- [60] G. F. Gribakin and V. V. Flambaum, Calculation of the scattering length in atomic collisions using the semiclassical approximation *Phys. Rev. A* **48**, 546 (1993).
- [61] M. Berninger, A. Zenesini, B. Huang, W. Harm, H.-C. Ngerl, F. Ferlaino, R. Grimm, P. S. Julienne, and J. M. Hutson, Universality of the Three-Body Parameter for Efimov States in Ultracold Cesium *Phys. Rev. Lett.* **107**, 120401 (2011).
- [62] A. N. Wenz, T. Lompe, T. B. Ottenstein, F. Serwane, G. Zrn, and S. Jochim, Universal trimer in a three-component Fermi gas *Phys. Rev. A* **80**, 040702 (2009).
- [63] N. Gross, Z. Shotan, S. Kokkelmans, and L. Khaykovich, Nuclear-Spin-Independent Short-Range Three-Body Physics in Ultracold Atoms *Phys. Rev. Lett.* **105**, 103203 (2010).
- [64] S. Tan, , private communication (unpublished).
- [65] E. Braaten, , private communication (unpublished).
- [66] C. Chin and P. S. Julienne, Radio-frequency transitions on weakly bound ultracold molecules. *Phys. Rev. A* **71**, 012713 (2005).

Appendix A

Effective range contribution

We calculate the transition rate for a dimer for a first-order approximation to include effects from the short-range potential in the two-body contact. We start with Equation 7 of [66] for the Frank-Condon factor between the initial state with scattering length a and the final state with a' ,

$$F(K) = \frac{2m}{\pi\hbar^2k} (1 + k^2a^2)^{-2} (\sin \delta' + ka \cos \delta')^2, \quad (\text{A.1})$$

where δ' is given by

$$k \cot \delta' = -\frac{1}{a'} + \frac{r'_e}{2} k^2 + \dots, \quad (\text{A.2})$$

and the effective range r'_e is defined in the limit $|a'| \gg \bar{a}$ by

$$r'_e = \frac{\Gamma(\frac{1}{4})^4}{6\pi^2} \bar{a}. \quad (\text{A.3})$$

The mean scattering length of the van der Waals potential \bar{a} is given by

$$\bar{a} = \frac{1}{\sqrt{8}} \frac{\Gamma(\frac{3}{4})}{\Gamma(\frac{5}{4})} \left(\frac{mC_6}{\hbar^2} \right)^{1/4}. \quad (\text{A.4})$$

This results in an effective range for ^{85}Rb of $229 a_0$. Following the derivation in [66], we express $k^2a^2 = K/E_b$ and $k^2a'^2 = K/E'_b$ where, $K = \hbar^2k^2/m$ and $E_b = \hbar^2/ma^2$ is the binding energy, and combine Equations A.1 and A.2. After quite a bit of algebra, this comes out to be

$$F(K) = \frac{2}{\pi} \left(\frac{a'}{a} - (1 - \rho) \right)^2 \frac{K^{1/2} E_b^{1/2} E'_b}{(K + E_b)^2 (K + E'_b (\rho - 1)^2)}, \quad (\text{A.5})$$

where $\rho = \frac{r'_e k^2 a'}{2}$. If we used only the first term in Equation A.2, then $\rho = 0$ and Equation A.5 reduces to Equation 10 in [66]. Now we define the RF offset energy $E = E_b + K$, and take the limit

of $E \gg E_b$ (large detuning) to get

$$F(K) = \frac{2 E_b^{1/2}}{\pi E^{3/2}} \frac{\left(\frac{a'}{a} - (1 - \rho)\right)^2}{\left((1 - \rho)^2 + \frac{E}{E_b}\right)}. \quad (\text{A.6})$$

Now we plug in $E = \hbar\omega$ and $E_b = \hbar^2/ma^2$ to put things into more familiar units:

$$F(K) = \frac{2}{\pi} \frac{1}{a\omega^{3/2}\sqrt{\hbar m}} \frac{\left(\frac{a'}{a} - (1 - \rho)\right)^2}{\left((1 - \rho)^2 + \frac{ma'^2\omega}{\hbar}\right)}. \quad (\text{A.7})$$

We can also substitute in the two-body contact for the dimer, $C_2 = \frac{8\pi}{a}$, using the adiabatic sweep theorem (Equation 6.3). Finally, we use $\Gamma(\omega) = \frac{\hbar\Omega^2}{2} F(\omega)$ from Equation 3 of [66] to get

$$\Gamma(\omega) = \frac{\Omega^2}{4\pi} \sqrt{\frac{\hbar}{m}} \frac{\alpha'(a)}{\beta'(\omega)} \frac{C_2}{\omega^{3/2}}, \quad (\text{A.8})$$

where the only difference between this result and Equation 6.6 is a modified definition for the terms that include final-state effects:

$$\alpha'(a) = \left(\frac{a'}{a} - (1 - \rho)\right)^2 \quad (\text{A.9})$$

$$\beta'(\omega) = \left((1 - \rho)^2 + \frac{ma'^2\omega}{\hbar}\right). \quad (\text{A.10})$$

The effect of including this effective range contribution is seen in Figure A.1. It may be measurable as a slight deviation from the ω -dependence, but the expected deviation is on the order of our error bars and as of yet too small to detect. Since the total contribution of final state effects to the contact go as $\beta(\omega)/\alpha(a)$, the corrections largely cancel, as seen in Figure A.1c.

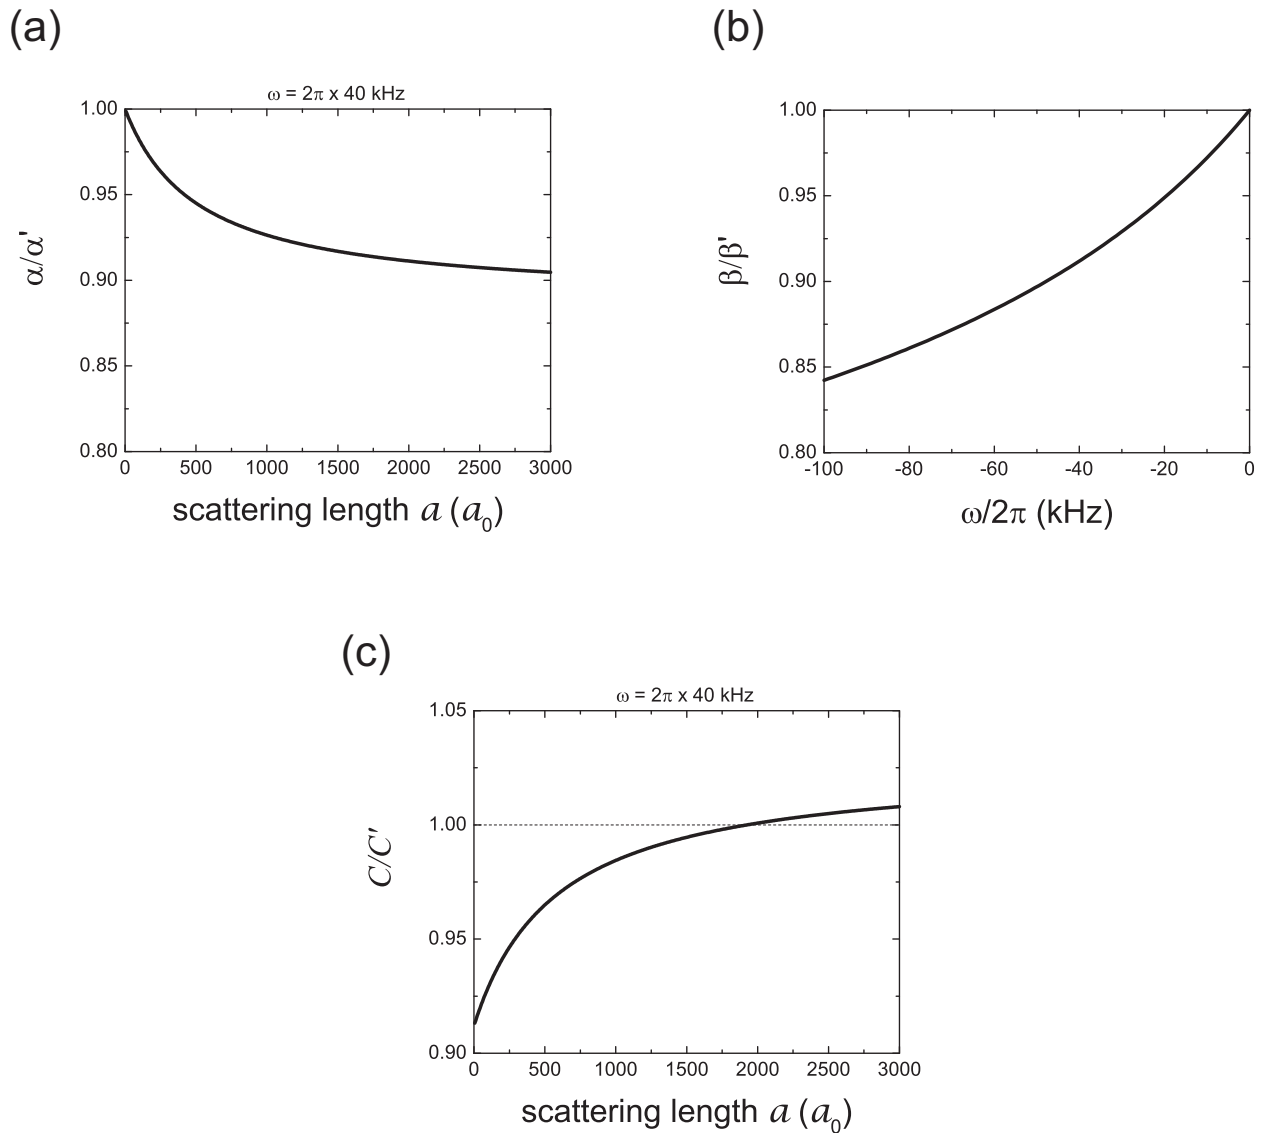


Figure A.1: Effects of the first-order correction due to a finite effective range. (a) and (b) show the fractional corrections introduced by using $\alpha'(a)$ and $\beta'(\omega)$, respectively. For the calculation of $\alpha'(a)$, a detuning of $\omega = 2\pi \times 40$ kHz was assumed. The magnitude of the effect roughly corresponds to the size of the error bars in our data, putting it just below our detection limit. (c) The combined effect on the total contact is on the order of a few percent.

Appendix B

Detailed procedure for the contact measurements

The procedure described below is typical for our experiment, but may not be the exact procedure used for any specific set of data. Before and after a contact measurement, we take measurements of BEC number using high-field and high-intensity imaging as described in Chapter 5. We check that those measurements do not vary greatly, and average them to determine the condensate number during the contact measurements. The contact measurement itself consists of three parts, for which the magnetic-field ramps are identical. For mostly historical reasons, I will define 0 ms to be the time when we turn off the trapping fields. We ramp off the large shim coils to minimize the magnetic-field gradient at -4 ms in 0.5 ms, while simultaneously ramping the bias field to compensate and keep the scattering length constant. At -3.5 ms, we begin a 1 ms ramp in the bias field to reach our final value of a . The Gaussian RF pulse is always centered on -1.5 ms.

We first determine the center frequency of the single-atom transition. For this measurement we make non-condensed clouds, and probe the $|2, -2\rangle$ to $|2, -1\rangle$ transition with Gaussian pulses of an rms width $\tau = 100 \mu\text{s}$, outcoupling 10-20% of the atoms with a Gaussian line width of ~ 4 kHz. We also include the RF pulses that normally transfer the $|2, -1\rangle$ atoms to the $|3, -3\rangle$ imaging state, but detune them by 5 MHz to avoid this transfer. This ensures that any RF rectification will be common-mode between the measurement of the center frequency and the measurement of the RF tail. We then turn off the magnetic fields at 0 ms, and expand for 15 ms, applying a vertical magnetic gradient for 5 ms to physically separate the magnetic sublevels. This imaging is performed like standard absorption imaging, at low field and with low probe intensities. We find

that fitting both clouds and looking at the outcoupled fraction significantly reduces noise due to number fluctuations, enabling us to reach < 1 kHz uncertainties in the center frequency with 5-6 well-picked points on the lineshape.

Once we know the single-atom frequency, we calculate and set the resonant frequencies for the RF transfer to the imaging state, as well as the probe laser frequency. To measure the integrated lineshape of the BEC, we tune the $|2, -2\rangle$ to $|2, -1\rangle$ RF to resonance and probe with a short, $\tau = 5$ μ s Gaussian pulse, outcoupling on the order of 2-3% of the atoms. Such a short pulse is always pulse-width limited, and we only need the peak of the lineshape to determine its integral to within a small correction described in Section 6.3. The RF transfer pulses occur at -1 ms and -0.9 ms to pump the atoms to the imaging state. We then probe at high-field (in-trap) but with low intensity, using a 25 μ s probe pulse.

For the measurement on the tail of the transition, we change only the initial RF pulse length to $\tau = 100$ μ s, and increase the power, always outcoupling $\sim 1 - 2\%$ of the atoms. Since we “truncate” our Gaussians at $\pm 4\tau$, this pulse “begins” at -1.9 ms (600 μ s after the end of the magnetic field ramp) and ends at -1.1 ms (100 μ s before the first RF transfer pulse). In order to avoid systematics caused by experimental drift, we alternate between measuring the peak and the tail of the transition. A sketch of the timings can be seen in Figure B.1.

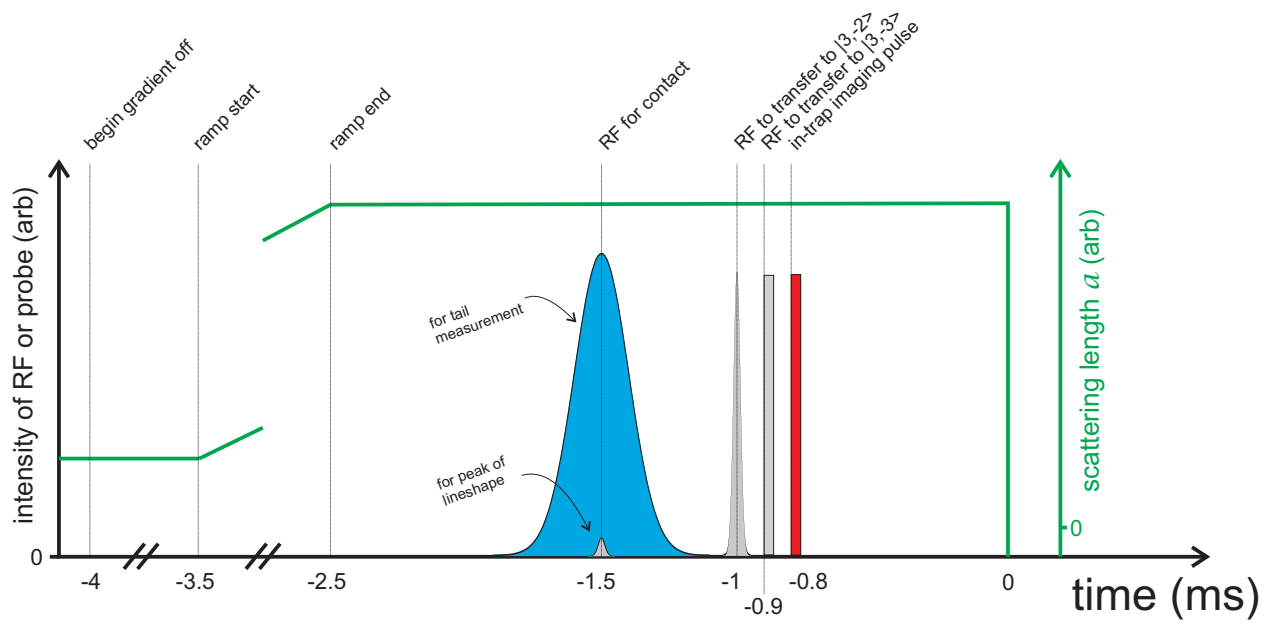


Figure B.1: Timing schematic of a typical contact measurement. The vertical scale is arbitrary. The widths of the grey RF pulses and the red probe pulse have been exaggerated to show up at this scale. The relative timings and the width of the blue RF pulse are to scale.