
Thermodynamic properties of a Bose gas with tuneable interactions

Robert Lorne Dugald Campbell

Jesus College
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

November 2011



*To my parents Lorne and Lucy,
and my brother Neill.*

I declare that this thesis is my own work and is not substantially the same as any that I have submitted or am currently submitting for a degree, diploma or any other qualification at any other university. No part of this thesis has already been or is being concurrently submitted for any such degree, diploma or any other qualification. This thesis does not exceed the word limit of sixty thousand words, including tables, footnotes, bibliography and appendices, set out by the Faculty of Physics and Chemistry.

Abstract

Thermodynamic properties of a Bose gas with tuneable interactions

Robert Lorne Dugald Campbell

This thesis describes the design, construction and optimisation of a new set of apparatus for the creation of ^{39}K Bose-Einstein condensates. This system is used to investigate the effects of repulsive inter-atomic interactions on the saturation of excited states and the shift in the critical temperature of harmonically trapped Bose gases.

Bose-Einstein condensates of over 4×10^5 atoms have been produced in a crossed optical dipole trap after sympathetic cooling with ^{87}Rb in a quadrupole-Ioffe-configuration magnetic trap. Condensation is made possible by tuning the scattering length of ^{39}K to positive values by means of a broad magnetically-tuneable Feshbach resonance centred at 402.5 G.

The ability to tune the scattering length in this system is used to scrutinise the concept of saturation of the thermal component of a Bose gas. We observe strong deviations from the ideal theory over a wide range of interaction strengths. Extrapolation to the non-interacting limit shows that the saturation picture is recovered for an ideal gas.

High-precision measurements of the shift in critical temperature of a harmonically trapped Bose gas are performed, using a novel differential measurement method. In the weakly-interacting limit the data show excellent agreement with mean-field theory. For stronger interactions an additional positive shift is observed, characteristic of beyond-mean-field effects due to critical correlations. Non-equilibrium behaviour is observed in the limits of very-strong and very-weak interactions.

Acknowledgements

Firstly, I would like to thank my supervisor, Zoran Hadzibabic, for taking me on as a founding member of our now-flourishing group. The task of building up the lab from scratch has at times been a daunting one, but has been greatly eased by the never-ending stream of ideas and optimism from above. I would also like to thank Jean Dalibard for his inspiring presence during his brief sabbatical. I couldn't have asked for more illustrious or pleasant company in the office.

The presented work could never have happened without the joint efforts of all of the past and current members of the group. I would especially like to thank Naaman and Rob. I am most fortunate to have spent my entire time here with Naaman, who has been the perfect companion through all the highs and lows of experimental research. I'm glad that we shared so many interests (our mutual appreciation of Bob Dylan saw us through many a long day's work), and to have gained so many new ones from him. Rob has been around for almost as long, and has been an invaluable asset to the group. I cannot thank him enough for his encyclopaedic knowledge and willingness to answer questions (however daft) - a winning combination. Life in the lab wouldn't have been half as fun without either of them.

The arrival of Scott and Stuart two years ago came as a welcome addition to the team, not least because of the formation of a full doubles pairing for tennis. Each has been a great help to me, and I now can't imagine the place without them. I must particularly thank Scott for his expert assistance when in the painful throws of rewriting our first paper for the umpteenth time. Stuart, who clearly has an alternative career looming with his surprising blu-tack sculpting skills, must also be praised/censured for recklessly introducing us all to the Empress pub quiz. Long live the lunchtime practice 'quiz of the day'.

The newest members of the group, Tobias, Igor, Alex and Richard, are already making their mark, and I thank them for the assistance they have already provided. Essential contributions to the experiment have also been made by the array of coworkers to have passed through the group. Jun-Young, Phil, Nils, Alexandre, Mehdi, Diana, and Mingjie have all helped to add plenty of colour to life in the lab. I am grateful to Michael Köhl and all of his students, who have provided invaluable support, whether by lending equipment or through helpful discussions.

I would like to thank all of the staff in the Cavendish, especially those working in the main workshop and stores, for all the support they have provided. A special mention is due to Pam Hadder for her impeccable efficiency and organisation. I am much

indebted to Nigel Palfrey in the student workshop for all his patient (and sometimes not-so-patient) assistance. Though it seems a while since I had cause to give him any real bother, I am sure he has had enough of me to last him until his retirement.

Outside the lab, I wish to thank all the G&S/musical theatre crew and my housemates past and present (Lucy, James, Danielle and Sybil) for making life in Cambridge such a pleasant experience. Also the old Univ crowd, for giving me the opportunity to let off some steam in our all-too-infrequent London gatherings (especially Laura and Ana for regular soft surfaces to spend the night). Claire is to blame for the contentious appearance of P. G. Wodehouse throughout this thesis (and many thanks to the man himself for keeping my morale high while writing up!) Finally I express my sincere gratitude to Siân for her magnificent attempts to keep me sane, and even more for her partial success.

Above all else I would like to thank my family for their unfaltering support. My brother Neill has acted as sibling, housemate, fellow cast-member, bank (I owe you £10...) and friend throughout the course of my PhD. My parents Lucy and Lorne have always been there for me with kind words and wise heads. Thank you all for your love and encouragement.

Contents

1	Introduction	1
1.1	Why ^{39}K ?	3
1.2	Ideal Bose gas	4
1.2.1	The semi-classical approximation	5
1.2.2	Ideal Bose gas in a harmonic trap	7
1.3	The Bose-Einstein condensate	8
1.3.1	Ideal condensate	8
1.3.2	Condensate with repulsive interactions	9
1.4	Properties of alkali atoms	12
1.4.1	Hyperfine structure	12
1.4.2	Interactions with magnetic fields	12
1.4.3	Interactions with light	15
1.5	Collisions between ultra-cold atoms	20
1.5.1	Scattering length	20
1.5.2	Feshbach resonances	24
1.5.3	Inelastic processes	26
2	Experimental setup	27
2.1	Vacuum system	28
2.1.1	Differential pumping design	28
2.1.2	Components	31
2.1.3	Assembly and bake-out	37
2.2	Magnetic coils	42
2.2.1	Coils around the MOT cell	43
2.2.2	Coils around the science cell	44
2.3	Laser cooling	44
2.3.1	Optical Molasses	46
2.3.2	Doppler limit and sub-Doppler cooling	47
2.3.3	Magneto-optical trap	49
2.3.4	Laser cooling for real atoms	49
2.4	Laser system	51
2.4.1	Saturated absorption spectroscopy	52
2.4.2	Frequency locking of lasers	54
2.4.3	Required frequencies	56

2.4.4	Laser table summary	57
2.4.5	Fibre-port cluster	60
2.4.6	Optics around the MOT cell	60
2.5	Magnetic transport	60
2.5.1	Coil holder design	62
2.5.2	Translation stage	63
2.5.3	Atom transportation	64
2.5.4	Majorana spin-flips	64
2.6	Quadrupole Ioffe-configuration trap	65
2.7	Evaporative cooling	67
2.7.1	Sympathetic cooling	70
2.7.2	Thermalisation	71
2.7.3	Microwave antenna	72
2.8	Crossed optical dipole trap	72
2.8.1	CDT laser	73
2.8.2	Forming the trap	74
2.8.3	Trap potential	77
2.9	State transfer	79
2.9.1	RF antenna	80
2.10	Absorption Imaging	80
2.10.1	Optical density	81
2.10.2	Camera	81
2.10.3	Imaging light	82
2.10.4	Objective	83
2.10.5	Imaging sequence	84
2.10.6	‘Fudge-factor’	84
2.11	Image Fitting	88
2.11.1	Thermal distribution	88
2.11.2	Thomas-Fermi distribution	89
2.11.3	Fitting routines	90
2.12	Computer control	90
2.12.1	Software	91
2.12.2	Hardware	92
2.13	Field and resonance calibration	92
2.13.1	Field calibration	93
2.13.2	Resonance measurement	93
2.13.3	Zero crossing measurement	95

3	Experimental sequence	97
3.1	Laser cooling	97
3.1.1	Monitoring MOT loading	97
3.1.2	MOT optimisation	98
3.1.3	Dual-species MOT	100
3.1.4	Optical molasses	101
3.2	Magnetic trapping and transport	102
3.3	Sympathetic cooling	102
3.3.1	Choice of trapping potential	102
3.3.2	Transfer to the QUIC trap	103
3.3.3	Cooling in the QUIC trap	103
3.3.4	Sympathetic cooling results	105
3.4	Quantum degeneracy in the CDT	107
3.4.1	Cooling methods in the dipole trap	107
3.4.2	Transfer to the CDT	111
3.4.3	Evaporation in the CDT	112
3.4.4	Number calibration	113
3.4.5	Condensation	115
3.4.6	Summary	116
4	Lack of saturation in a Bose gas	121
4.1	Saturation: another approach	121
4.2	Context	124
4.3	Experimental procedure	126
4.3.1	Image fitting for saturation data	126
4.3.2	Data processing	128
4.3.3	Results	129
4.4	Hartree-Fock Approximation	129
4.5	Interpretation of results	132
4.5.1	Popov approximation	133
4.5.2	Small condensate limit	134
4.5.3	Large condensate limit	135
4.6	Conclusions	136
5	Effect of interactions on the critical temperature	139
5.1	History	139
5.1.1	Uniform system	141
5.1.2	Harmonic confinement	142
5.1.3	Past experiments	143
5.1.4	Aims of the investigation	143
5.2	Mean-field shift in a harmonic trap	144

5.2.1	Analytic solution	144
5.2.2	Numerical solution	147
5.2.3	Introducing critical effects	148
5.3	Experimental procedure	149
5.3.1	Image fitting routine	152
5.3.2	Differential measurement process	154
5.4	Critical temperature shift results	154
5.5	Assessment of errors	156
5.5.1	In-trap momentum distribution	156
5.5.2	Interaction energy	157
5.6	Non-equilibrium effects	161
5.7	Conclusions	163
6	Conclusion	165
6.1	Summary	165
6.2	Outlook	166
6.2.1	Condensed fraction induced by critical correlations	166
6.2.2	Out-of-equilibrium studies	167
6.2.3	Two-dimensional Bose gases	167
6.2.4	Annular Bose-Einstein condensates	168
6.2.5	Power-law traps	168
6.2.6	Superconducting magnetic trap	168
6.2.7	Final remarks	169
A	Useful Physical Quantities	171
B	Small system	175
C	Popov Approximation	181
D	Assessment of absolute errors	185
D.1	Finite imaging time	185
D.2	Feshbach field curvature	187

Chapter 1

Introduction

“ I don’t know if you have had the same experience, but the snag I always come up against when I’m telling a story is this dashed difficult problem of where to begin it. ”

– P. G. Wodehouse (*Right Ho, Jeeves*)

The quantum regime in which a Bose-Einstein condensate (BEC) can exist seems a far cry from day-to-day experiences of the world around us. Classical mechanics provides an excellent description for our interactions with macroscopic objects, which we naturally find to be intuitive. However, the quantum mechanical description takes into account the true wave nature of matter. Yet for a gas at room temperature, the extent of these wave-packets is much smaller than the inter-particle separation. The atoms and molecules can be thought of as hard billiard balls colliding with one another, and the behaviour of the gas is also well-described by a classical statistical treatment.

If, however, the spacing is decreased or the wave-packets spread out such that they start to overlap with one another, a classical description is not sufficient, and quantum mechanical behaviour takes hold. In the case that the particles are identical bosons (which have integer quantum mechanical spin), the remarkable result is that, below a critical finite temperature, they start to macroscopically occupy the single-particle quantum ground state of the system. At this point the particles are completely indistinguishable, forming a single matter-wave; in other words a BEC.

This phase transition was predicted by Einstein [1], following from the work of Bose [2], in 1925. However, it wasn’t until 70 years later that it was first realised in a gas of atoms in the group of Wieman and Cornell [3] and soon after in the group of Ketterle [4]. It took so long to achieve because of the extreme low temperatures required. To avoid the formation of a solid the atoms must be kept extremely dilute, several million times lower than the density of air. The thermal wavelength of the particles must therefore be increased to an extent that can only be achieved by cooling the atoms to near the photon recoil limit (the temperature associated with the momentum ‘kick’ of a single photon). This is typically around 1/10,000,000 of a degree above absolute zero, where atoms move at speeds of millimetres per second rather than at hundreds of




Figure 1.1: Logarithmic temperature scale. Bose-Einstein condensation typically occurs near the recoil limit in a dilute atomic gas.

metres per second (as they do at room temperature). This is put into context in Fig. 1.1 where it may be seen that on a logarithmic scale we live at a temperature closer to that at the centre of the sun than at which a BEC is formed.

The bulk of the cooling of atoms from room temperature to quantum degeneracy is achieved with a combination of the scattering force of laser light (the same force that directs a comet’s tail away from the sun) and evaporation of the most energetic atoms in the latter stages (the process that our bodies employ to keep cool via perspiration). The development of laser cooling and trapping techniques for neutral atoms was recognised with the 1997 Nobel prize in physics, while the 2001 prize was awarded for the attainment of the BEC itself.

The motivation to produce these condensates in the lab has not simply been driven by the desire to verify Einstein’s prediction. Though quantum effects naturally originate at the atomic scale, they also lead to some more familiar macroscopic behaviour, such as superconductivity, superfluidity and semiconductor properties. Such many-body quantum phenomena can be studied using ultracold gases to simulate relevant theoretical models, where the experimentalist has a high degree of control over nearly all of the system parameters. Often the inter-atomic interactions are weak in these systems, and the corresponding theoretical problems are easily solved. The system parameters can then be tuned to a regime in which these theories break down, providing new insight into previously impenetrable processes.

Recent additions to the experimental arsenal include the introduction of optical confinement to reduce the dimensionality of the system and to create strong periodic potentials, and the ability to directly tune the strength of interactions using a Feshbach resonance. These methods allow the study of strongly correlated systems, more typically associated with quantum liquids (such as superfluid helium). The overall picture is that of a system in which we have great control over most parameters, enabling the study of a wide range of phenomena in different regimes. An extensive review of recent experimental and theoretical progress concerning many-body physics with ultracold

gases is provided in [5].

Before we are able to experimentally study these many and various fascinating topics, we first have to build a system that allows us to reach the ultracold regime. This requires the combination of many technological disciplines, including those of ultra-high vacuum, laser light, electromagnetic coils, radio frequency (RF) and microwave radiation, while carefully synchronising the operation of the components by means of computer control. After giving the arguments for the choice of atomic species used in our system the remainder of this chapter is given over to basic theory of the Bose condensed state and interactions of bosons with electromagnetic fields and one another. Chapters 2 and 3 describe the experimental setup and process respectively, while Chapters 4 and 5 report on the first new science to have been studied with the system. Finally a summary of our findings and future directions for research is provided in Chapter 6.

1.1 Why ^{39}K ?

As a brand new group setting out to build a complete system for the study of quantum gases it would be understandable to stick to a ‘safe’ atomic species. Though our field is still young there are now well over a hundred research groups worldwide producing BECs every day. Most of these groups use isotopes of lithium, sodium, rubidium and caesium in their experiments, providing a wealth of information to aid the prospective atomic physicist. In contrast the successful condensation of ^{39}K was first reported only a matter of months before the formation of our group, after a series of pioneering experiments at LENS in Florence [6, 7]. This achievement itself was founded on previous work with the other isotopes of potassium [8, 9]. Potassium was a latecomer to the quantum degenerate regime due to unfavourable atomic properties which hamper its cooling and in the case of ^{39}K condensation itself.

Initially we concentrated on the relatively straight-forward task of condensing ^{87}Rb , whilst preparing for the possibility of cooling ^{40}K to work with fermions in the future. However, it became apparent that the advantages of working with ^{39}K might well outweigh the inherent difficulties. Our overall goal is to study universal phenomena in many-body physics; this should not depend on the particular properties of the atom we are using, only its more generic properties. The beauty of working with ultracold gases to this end is that we have the tools to control the temperature, spin, confinement, motion, rotation and other such properties of the system. The final ‘knob’ we would like control of is the inter-atomic interaction strength. Remarkably we may play with this too, via the phenomenon of Feshbach resonances. These allow us to control the scattering length of our atoms by simply applying a uniform magnetic field of appropriate strength. However, their accessibility and usefulness (for our purposes) vary widely between species.

The behaviour of the scattering length with magnetic field strength near a Fesh-

1.2 Ideal Bose gas

Atom	B_∞ (G)	Δ (G)	a_{bg}/a_0	$a_{\text{bg}}/(a_0\Delta)$ (G^{-1})	Ref.
^7Li	736.8	-192.3	-25	0.13	Pollack <i>et al.</i> , 2009 [10]
^{23}Na	1195	-1.4	62	-44.3	Stenger <i>et al.</i> , 1999 [11]
^{39}K	402.5	-52	-29	0.56	Zaccanti <i>et al.</i> , 2009 [12]
^{41}K	51.4	-0.3	60	-200	Kishimoto <i>et al.</i> , 2009 [13]
^{85}Rb	155.04	10.7	-443	-41.4	Claussen <i>et al.</i> , 2003 [14]
^{87}Rb	1007.4	0.21	100	476	Dürr <i>et al.</i> , 2004 [15]
^{133}Cs	-11.7	28.7	1720	59.9	Lange <i>et al.</i> , 2009 [16]

Table 1.1: Characteristic parameters of selected Feshbach resonances in stable bosonic alkali atoms. These are experimentally verified resonances with the most promising properties for control of interaction strength for each respective atomic species. The broad resonances in ^7Li and ^{39}K stand out as prime candidates. References are to the most recent known experimental measurements. The data was collated in Ref. [17], some values being the result of unpublished work.

bach resonance may usually be characterised by just three parameters¹ (see §1.5.2): a resonance centre B_∞ , the corresponding resonance width Δ and background scattering length a_{bg} . Near the resonance the scattering length diverges, while there is what is known as a zero-crossing (where the scattering length goes to zero) for a field value $B_0 = B_\infty + \Delta$. Near this zero-crossing the variation in scattering length with field is roughly linear, with gradient a_{bg}/Δ . Thus prime resonances for good experimental control of the scattering length have a small magnitude for this figure of merit. As can be seen in Table 1.1, the quoted resonances for ^7Li and ^{39}K have gradients at the zero-crossing about two orders of magnitude smaller than any of the others. By this criterion alone ^7Li would appear to have the best experimental control, and indeed is being used with great success [10]. All things considered, there isn't much to choose between these two species, and with the more convenient lower-field resonance of ^{39}K (and with our experiment already shaping up to cater for cooling rubidium and potassium) it was the obvious choice to concentrate our efforts on.

1.2 Ideal Bose gas

Bose-Einstein condensation is fundamentally different to all other phase transitions between different states of matter in that it occurs even in the absence of interactions between particles, relying on purely statistical processes². In a normal classical gas the thermal energy $k_{\text{B}}T$ is far greater than the level spacings of the system, and the behaviour of the gas can be described with standard Maxwell-Boltzmann statistics, where the bosonic or fermionic nature of the particles is insignificant. However, Einstein ob-

¹This is valid as long as there are no other resonances nearby.

²Although thermal equilibrium is assumed, generally requiring interactions between particles.

served in 1925 [1] that this description is inadequate for thermal energies for which the spreading of the atomic wavefunction is comparable to the inter-particle separation, where the indistinguishable nature of the identical particles becomes important. Einstein generalised Bose's work on the statistics of (massless) photons [2] to consider massive particles, describing what is now referred to as Bose-Einstein statistics.

The essential feature is that the mean occupation of single-particle states i for non-interacting bosons in thermal equilibrium is given by the Bose-Einstein distribution function

$$f(\epsilon_i) = \frac{1}{e^{(\epsilon_i - \mu)/k_B T} - 1}, \quad (1.1)$$

where ϵ_i is the energy of the single-particle state i and μ is the chemical potential, which is set by particle number conservation. From this the total number of particles is easily evaluated as

$$N = \sum_i f(\epsilon_i). \quad (1.2)$$

Equation (1.1) also provides the important constraint $\mu < \epsilon_0$ for the ideal Bose gas, where ϵ_0 is the energy of the lowest single-particle state. As $\mu \rightarrow \epsilon_0$ the occupation of the ground state

$$N_0 \equiv f(\epsilon_0) = \frac{1}{e^{(\epsilon_0 - \mu)/k_B T} - 1} \quad (1.3)$$

increases and becomes comparable to N . In fact, this is the mechanism behind Bose-Einstein condensation, occurring under specific conditions as discussed below.

1.2.1 The semi-classical approximation

We can treat (1.3) as a special case such that the total number of atoms is given by

$$N = N_0 + N_T \quad (1.4)$$

where N_T is the number of atoms in excited states (i.e. the number of thermal atoms) given by

$$N_T = \sum_{i \neq 0} f(\epsilon_i). \quad (1.5)$$

In general it is not straightforward to evaluate such summations. However, as long as the thermal energy $k_B T$ is much larger than the energy level spacing of the system we may convert the sum into an integral by replacing the discrete ϵ_i with the continuous variable $\epsilon(\mathbf{r}, \mathbf{p})$ dependent on position \mathbf{r} and momentum \mathbf{p} , which corresponds to the classical energy associated with the single-particle Hamiltonian for the system. This does not take into account the contribution from $i = 0$. In three dimensions there is on average one quantum state per volume of phase space $(2\pi\hbar)^3$. Integrating over all phase space and dividing by this factor thus yields the total number of occupied excited

1.2 Ideal Bose gas

states,

$$N_T = \frac{1}{(2\pi\hbar)^3} \iint \frac{1}{e^{\epsilon(\mathbf{r}, \mathbf{p}) - \mu}/k_B T} - 1} d\mathbf{p} d\mathbf{r} . \quad (1.6)$$

This transformation to an integral is known as the semi-classical approximation. For a gas in a box, $\epsilon(\mathbf{r}, \mathbf{p})$ is simply $p^2/2m$ and we may separate the integral $\int d\mathbf{r} = V$, where V is the volume of the containing box. We may then use the substitution $x = p^2/2mk_B T$ to calculate the number density as a function of μ

$$n_T \equiv \frac{N_T}{V} = \frac{1}{(2\pi\hbar)^3} \int_0^\infty \frac{1}{e^{(p^2/2m - \mu)/k_B T} - 1} 4\pi p^2 dp = \frac{1}{\lambda_T^3} g_{3/2}(e^{\mu/k_B T}) , \quad (1.7)$$

where

$$\lambda_T = \sqrt{\frac{2\pi\hbar^2}{mk_B T}} \quad (1.8)$$

is the thermal de Broglie wavelength. $g_{3/2}(z)$ is a special case of the set of polylogarithmic functions

$$g_s(z) = \frac{1}{\Gamma(s)} \int_0^\infty \frac{x^{s-1}}{z^{-1}e^x - 1} dx = \sum_{k=1}^\infty \frac{z^k}{k^s} \quad (1.9)$$

where the Gamma function is defined as

$$\Gamma(s) = \int_0^\infty x^{s-1} e^{-x} dx , \quad (1.10)$$

giving $\Gamma(3/2) = \sqrt{\pi}/2$. The value of n_T obtained from (1.7) increases with increasing μ . However, there is a fundamental limit from our constraint $\mu < \epsilon_0 = 0$ for a free particle in a box. This means that as N increases the available excited states become saturated at a critical density

$$n_c \equiv \frac{N_c}{V} = \frac{\zeta(3/2)}{\lambda_T^3} \approx \frac{2.612}{\lambda_T^3} , \quad (1.11)$$

where $\zeta(s) = g_s(1)$ is the Reimann zeta function, with all further atoms forced to occupy the ground state, forming a Bose-Einstein condensate. In this ideal case the number of thermal particles is fixed at $N_T = N_c$ for $N > N_c$. Note that this phenomenon does indeed occur at a point where interparticle separation is comparable to the spread of the atomic wavefunction with $n_c^{-1/3} \sim \lambda_T$. The critical value of the phase space density $n\lambda_T^3$ of 2.612 is a key prediction of Einstein's theory and provided a useful target for the experimental pioneers aiming to realise these extreme conditions in the laboratory. We may reformulate (1.11) using (1.8) to give the corresponding result for the critical temperature,

$$k_B T_c = \frac{2\pi\hbar^2}{m} \left(\frac{n}{\zeta(3/2)} \right)^{2/3} \approx 3.313 \frac{\hbar^2 n^{2/3}}{m} . \quad (1.12)$$




Figure 1.2: Density distribution (solid line) of a harmonically trapped Bose gas at the critical point in terms of phase space density. The peak density is equal to the critical phase space density for a uniform Bose gas. The dashed line shows the form of the harmonic trapping potential.

1.2.2 Ideal Bose gas in a harmonic trap

Most experimental setups produce ultracold atomic samples that are trapped in a harmonic potential. We may describe such a system by now setting

$$\epsilon(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + V(\mathbf{r}) , \quad (1.13)$$

where

$$V(\mathbf{r}) = \frac{1}{2}m (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2) \quad (1.14)$$

describes a generic harmonic potential with trapping frequencies ω_k . Using (1.13) with (1.6) and performing only the integral over momentum space gives us the spatial distribution of thermal atoms in the trap as

$$n_T(\mathbf{r}) = \frac{1}{\lambda_T^3} g_{3/2} \left[e^{(\mu - V(\mathbf{r}))/k_B T} \right] . \quad (1.15)$$

Once more the upper bound for our thermal atom number is acquired in the limit of $\mu \rightarrow \epsilon_0 \approx 0^1$. This gives a critical density distribution (as displayed in Fig. 1.2)

$$n_c(\mathbf{r}) = \frac{1}{\lambda_T^3} g_{3/2} \left\{ \exp \left[-\frac{1}{2} \left(\frac{x^2}{R_{T_x}^2} + \frac{y^2}{R_{T_y}^2} + \frac{z^2}{R_{T_z}^2} \right) \right] \right\} , \quad (1.16)$$

where the characteristic widths of the distribution $R_{T_k} = \sqrt{k_B T / m \omega_k^2}$ are known as the thermal radii. We immediately recover the critical density given in (1.11) at the centre of the trap by setting $x = y = z = 0$. Performing instead the \mathbf{r} integral of (1.6) would give a similar expression for $n_T(\mathbf{p})$. To obtain the critical number we must integrate (1.16) over \mathbf{r} , finding

¹This ignores the zero point energy of the potential.

1.3 The Bose-Einstein condensate

$$N_c = \zeta(3) \left(\frac{k_B T}{\hbar \bar{\omega}} \right)^3 \approx 1.202 \left(\frac{k_B T}{\hbar \bar{\omega}} \right)^3, \quad (1.17)$$

where the geometric mean trapping frequency $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is employed. Once again above this critical atom number additional atoms must occupy the ground state as a Bose-Einstein condensate. Alternatively we can describe the critical point in terms of a critical temperature T_c for a fixed atom number given by

$$k_B T_c = \hbar \bar{\omega} \left(\frac{N}{\zeta(3)} \right)^{1/3} \approx 0.940 \hbar \bar{\omega} N^{1/3}. \quad (1.18)$$

The above expression may be used to more closely scrutinise the applicability of the semi-classical approximation. The level spacings of a quantum harmonic oscillator being $\hbar \omega_k$, the condition $k_B T \gg \hbar \omega_k$ must hold. It is apparent then from (1.18) that the semi-classical approximation is only compatible with the quantum degenerate regime if $N^{1/3} \gg 1$.

1.3 The Bose-Einstein condensate

The preceding section gives us an adequate description of the distribution of the thermal atoms of an ideal gas below the critical temperature, but we would also like to know about the distribution of the condensate. In order to calculate this we may invoke the Bogoliubov approximation [18] to describe the wavefunction of the condensate as $\Psi_0 = \sqrt{N_0} \varphi_0$, where φ_0 is the single-particle ground state wavefunction of the system, obeying

$$N_0 = \int |\Psi_0(\mathbf{r})|^2 d\mathbf{r}. \quad (1.19)$$

Ψ_0 is also known as the order parameter, being zero above the critical temperature and finite below it¹.

1.3.1 Ideal condensate

For an ideal Bose gas, Ψ_0 solves the familiar Schrödinger equation

$$\left(-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) \right) \Psi_0(\mathbf{r}) = \mu \Psi_0(\mathbf{r}), \quad (1.20)$$

where the chemical potential μ replaces the usual energy term as we keep the number of particles fixed. It is easy to check that for the harmonic potential given in (1.14)

¹The Bogoliubov treatment is justified for a large number of particles occupying the same state, where the non-commutativity of field operators is not important and a classical function may be used instead.




Figure 1.3: Density distribution in terms of phase space density for a non-interacting Bose gas in a harmonic trap, where $N_c = 10^6$ and $N_0 = 10^4$. The solid line gives the full distribution while the dashed line shows that of the thermal gas only. The x -axis is in units of thermal radii while the y -axis is given in terms of phase space density. The sharp peak in density is a signature of the onset of condensation in a harmonically trapped Bose gas. It is clearly visible despite the condensed fraction N_0/N of only $\sim 1\%$.

this is solved by the Gaussian wavefunction

$$\Psi_0(\mathbf{r}) = \frac{\sqrt{N_0}}{\pi^{3/4} \bar{a}_{\text{ho}}^{3/2}} \exp \left[-\frac{1}{2} \left(\frac{x^2}{a_{\text{ho}x}^2} + \frac{y^2}{a_{\text{ho}y}^2} + \frac{z^2}{a_{\text{ho}z}^2} \right) \right], \quad (1.21)$$

where $a_{\text{ho}k} = \sqrt{\hbar/m\omega_k}$ is the oscillator length along the k axis and \bar{a}_{ho} is the geometric mean $(a_{\text{ho}x}a_{\text{ho}y}a_{\text{ho}z})^{1/3}$. This is not surprising, being simply $\sqrt{N_0}$ times the lowest single-particle state wavefunction for a harmonic potential. This directly gives the ideal condensate density distribution

$$n_0(\mathbf{r}) = |\Psi_0(\mathbf{r})|^2 = \frac{N_0}{\pi^{3/2} \bar{a}_{\text{ho}}^3} \exp \left[-\left(\frac{x^2}{a_{\text{ho}x}^2} + \frac{y^2}{a_{\text{ho}y}^2} + \frac{z^2}{a_{\text{ho}z}^2} \right) \right]. \quad (1.22)$$

Note that the characteristic length scales a_{ho} and R_T , for the ideal condensate and the thermal distribution respectively, differ by a factor $[\zeta(3)/N_c]^{1/6} \approx 0.1$ for $N_c = 10^6$. This means that in a harmonic trap a sharp peak in the spacial distribution is a signature of the onset of Bose-Einstein condensation, as shown in Fig. 1.3 for $N_c = 10^6$ and $N_0 = 10^4$.

1.3.2 Condensate with repulsive interactions

The introduction of interactions complicates matters, but we are fortunate to be considering dilute systems in which the inter-particle separation is much larger than the range of the interatomic forces. We may generally then just consider the effects of pairs

1.3 The Bose-Einstein condensate

of particles interacting, and we may also use the asymptotic expression for the wavefunction for the relative motion of the atoms which is fixed by the scattering amplitude. In the ultracold regime the particles have very low momenta such that the scattering amplitude is fully determined by the s -wave scattering length a (this is discussed in a little more detail in §1.5). We may also assume that the variation of $\Psi_0(\mathbf{r})$ is small over the range of the inter-atomic force (i.e. for a harmonic trap $a_{\text{ho}}/a \gg 1$), such that the effects of interactions are fully captured by introducing a self-interaction term $g|\Psi_0(\mathbf{r})|^2$ to the Schrödinger equation, where the interaction coupling constant is defined as

$$g = \frac{4\pi\hbar^2 a}{m} . \quad (1.23)$$

The addition of this non-linear term results in the time-independent Gross-Pitaevskii equation,

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g|\Psi_0(\mathbf{r})|^2 \right) \Psi_0(\mathbf{r}) = \mu\Psi_0(\mathbf{r}) . \quad (1.24)$$

This important equation is a key tool for investigating the behaviour of nonuniform ultracold Bose gases. A thorough derivation can be found in [19]. This provides a starting point to find the condensate density in a self-consistent manner.

A remarkable further simplification may yet be made. If we take as a starting point the wave function (1.21) for $g \rightarrow 0$ we can compare the expectation values of the interaction and kinetic energy terms of (1.24). Assuming no extreme anisotropy of the trapping potential, the ratio of the two terms is $\simeq N_0 a / \bar{a}_{\text{ho}}$. This means that the interaction term swamps the kinetic energy provided the condition

$$N_0 \frac{a}{\bar{a}_{\text{ho}}} \gg 1 \quad (1.25)$$

holds. Once more under standard experimental conditions this criterion is met, and we may use the so-called Thomas-Fermi approximation to drop the kinetic energy term of (1.24). This may then be simply arranged in the form

$$n_0(\mathbf{r}) = |\Psi_0(\mathbf{r})|^2 = \frac{\mu - V(\mathbf{r})}{g} \quad (1.26)$$

up to the point in the trap where $\mu = V(\mathbf{r})$, while the trivial solution $|\Psi_0(\mathbf{r})|^2 = 0$ holds otherwise. For a uniform system this gives the result $\mu = gn_0(0)$, being positive for $a > 0$ below the critical temperature, in contradistinction to the ideal case. We can also see that the condensate density is drastically altered from the ideal case, reflecting the shape of the confining potential. Going back to our harmonic potential, the density may be written in the form

$$n_0(\mathbf{r}) = \frac{\mu}{g} \left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2} \right) , \quad (1.27)$$




Figure 1.4: In-trap density distributions for pure interacting ($a = 135 a_0$, solid line) and ideal (dashed line) ^{39}K condensates of 400 000 atoms with $\bar{\omega} = 2\pi \times 70 \text{ Hz}$. Densities are normalised to the central density of the ideal condensate, approximately equivalent to 10^{16} cm^{-3} . Note that the central density in the interacting case is only $\sim 0.8\%$ that of the ideal case, so the displayed distribution has been multiplied by a factor of 50.

where $R_k = \sqrt{2\mu/m\omega_k^2}$ is the Thomas-Fermi radius at which the condensate density drops to zero along the k axis. The relation of N_0 and μ may be found by integrating (1.27) over the ellipsoid with semi-axes R_k . Doing so gives

$$N_0 = \frac{8\pi}{15} \frac{\mu}{g} \bar{R}^3, \quad (1.28)$$

\bar{R} representing the geometric mean $(R_x R_y R_z)^{1/3}$. This may be rearranged to give the Thomas-Fermi chemical potential in the form

$$\mu = \frac{1}{2} \hbar \bar{\omega} \left(\frac{15 N_0 a}{\bar{a}_{\text{ho}}} \right)^{2/5}, \quad (1.29)$$

which we shall make good use of in Chapter 4. We can calculate $n_0(0)$ by dividing (1.29) by g , while the Thomas-Fermi radius is given in terms of N_0 as

$$R = a_{\text{ho}} \left(\frac{15 N_0 a}{a_{\text{ho}}} \right)^{1/5}, \quad (1.30)$$

showing how significantly the condensate spreads out in comparison to the ideal case, where the size is independent of N_0 . Fig. 1.4 compares the in-trap distributions for pure condensates of ^{39}K in the ideal case ($a = 0$) and when $a = 135 a_0$. Interactions suppress the central density by a factor of ≈ 125 with some typical experimental parameters.

1.4 Properties of alkali atoms

The production of BECs is only possible due to the interactions of atoms with electromagnetic fields, so it is best to have an idea of the properties of the atoms themselves before getting into the practical applications of such interactions.

1.4.1 Hyperfine structure

The electronic ground state configurations of the alkali atoms (group IA of the periodic table) are characterised by a single electron in their outer shell, for example:

$$\begin{aligned} \text{K} & 1s^2 2s^2 2p^6 3s^2 3p^6 4s, \\ \text{Rb} & 1s^2 2s^2 2p^6 3s^2 3p^6 3d^{10} 4s^2 4p^6 5s \end{aligned}$$

for potassium and rubidium. It is enough simply to refer to the outer electron to describe electronic states for our purposes as it is the only one that will be affected on the energy scales used (the inner electrons are much more tightly bound). Thus the potassium and rubidium ground states are referred to as $4^2S_{1/2}$ and $5^2S_{1/2}$ respectively. Here the superscript 2 refers to the number of spin components of the angular momentum S (equal to $2S + 1$), the ‘S’ refers to the orbital angular momentum L (S, P, D... for 0, 1, 2...) and the subscript 1/2 refers to total electron angular momentum J .

We concern ourselves with this ground state and the lowest excited nP states. The interaction between the spin and orbital components of the angular momentum (the origins of which are discussed in [20] as well as many other texts) create a fine splitting of the excited state into a doublet with $J = 1/2$ and $3/2$. The two frequencies associated with the energy gap between the ground state and these two excited states are referred to respectively as the D_1 and D_2 lines.

Each of these states is split further by the interaction between the nuclear spin I and J , described by the Hamiltonian

$$\hat{H}_{\text{hf}} = A_{\text{hf}} \mathbf{I} \cdot \mathbf{J}, \quad (1.31)$$

forming the hyperfine structure. Assuming $J \leq I$, this produces $2J+1$ levels labelled by the total atomic angular momentum F , taking the values $F = I - J, I - J + 1, \dots, I + J$. For our experiments we shall mainly be concerned with the D_2 line and the subsequent hyperfine splitting, which is shown in Fig. 1.5 for ^{39}K and ^{87}Rb . Conveniently, both species have $I = 3/2$, allowing us to consider just one configuration of S and I .

1.4.2 Interactions with magnetic fields

The degeneracy of the hyperfine levels may be lifted by applying a magnetic field. Each is split into $2F + 1$ sublevels, giving a total of $(2J + 1)(2I + 1)$ for each fine




Figure 1.5: Hyperfine structure of the ³⁹K and ⁸⁷Rb D₂ lines. Both species have $I = 3/2$ resulting in similar structures, with the frequency splitting between hyperfine levels given in MHz. It should be noted however that the scale of the entire ⁸⁷Rb structure is scaled down by a factor of ten with respect to that for ³⁹K, while the scale of the lower hyperfine splitting for both species is scaled down by a factor of ten with respect to that of their upper scales. Potassium data is taken from [21, 22], while the rubidium data is collated in [23].

1.4 Properties of alkali atoms

structure configuration. For low magnetic fields (where the interaction with the external field is weaker than the hyperfine interaction) each sublevel may be labelled by the good quantum numbers F and its projection along the quantisation axis m_F . In the strong field regime however we must use instead the quantum numbers J, I and their projections m_J and m_I . Using $m = m_F$ for low fields and $m = m_J + m_I$ for high fields provides a consistent parameter to label states over the full range of magnetic field strength.

It is possible to find expressions to describe the behaviour of the hyperfine sublevels in the weak- and strong-field regimes by considering the magnetic interaction as a perturbation of the hyperfine interaction eigenstates and vice versa, but to examine the behaviour in the intermediate region one must diagonalise the full hyperfine-magnetic Hamiltonian

$$\hat{H} = \hat{H}_{\text{hf}} + \hat{H}_B = A_{\text{hf}} \mathbf{I} \cdot \mathbf{J} - \boldsymbol{\mu}_J \cdot \mathbf{B} - \boldsymbol{\mu}_I \cdot \mathbf{B} , \quad (1.32)$$

where $\boldsymbol{\mu}_J = -g_J \mu_B \mathbf{J} / \hbar$ and $\boldsymbol{\mu}_I = -g_I \mu_B \mathbf{I} / \hbar$ are the angular momentum and nuclear magnetic moments respectively, where g_J and g_I are Landé g-factors and μ_B is the Bohr magneton. For $J = 1/2$, this results in the Breit-Rabi formula [24],

$$E_{\text{BR}}(F_{\pm} = I \pm 1/2, m) = -\frac{\Delta E_{\text{hf}}}{2(2I+1)} - g_I \mu_B m B \pm \frac{1}{2} \sqrt{\Delta E_{\text{hf}}^2 + \frac{4m}{2I+1} (g_J - g_I) \mu_B B \Delta E_{\text{hf}} + (g_J - g_I)^2 \mu_B^2} , \quad (1.33)$$

where ΔE_{hf} is the hyperfine splitting for $B = 0$. Turning to dimensionless variables

$$\tilde{B} = \frac{B}{B_0} = B \frac{(g_J - g_I) \mu_B}{\Delta E_{\text{hf}}} \text{ and } \epsilon(F, m) = \frac{E_{\text{BR}}(F, m)}{\Delta E_{\text{hf}}} \quad (1.34)$$

this may be written in the dimensionless form

$$\epsilon(F_{\pm}, m) = -\frac{1}{2(2I+1)} - \frac{m}{g_J/g_I - 1} \tilde{B} \pm \frac{1}{2} \sqrt{1 + \frac{4m}{2I+1} \tilde{B} + \tilde{B}^2} . \quad (1.35)$$

This is plotted in Fig. 1.6 for $I = 3/2$, representing the ^{39}K and ^{87}Rb electronic ground states. In the weak-field ($\tilde{B} \ll 1$) limit, and neglecting $g_I \sim g_J \times 10^{-3}$, we recover the familiar Zeeman effect,

$$E_Z = E_{\text{BR}}(F, m_F, B) - E_{\text{BR}}(F, m_F, B = 0) = \mu_F B , \quad (1.36)$$

where $\mu_F = g_F \mu_B m_F$. The factor g_F arises from the projection of \mathbf{J} onto \mathbf{F} , which for $J = 1/2$ is simply given by

$$g_{F_{\pm}} = \pm \frac{1}{2I+1} g_J . \quad (1.37)$$

Magnetic trapping of atoms therefore depends on creating an inhomogeneous magnetic field such that a potential minimum is created. Given that Maxwell's equations




Figure 1.6: Energies of the magnetic sublevels of the $n^2S_{1/2}$ fine structure ground state of an alkali atom with nuclear spin $I = 3/2$. In the low-field regime the states are grouped by F while in the high-field regime they are grouped by m_J . Note that B_0 is 2438.8 G for ^{87}Rb while it is 164.7 G for ^{39}K .

prohibit the formation of a magnetic maximum in space, only the atoms in states where magnetic potential increases with field strength are magnetically trappable, i.e. those for which $g_F m_F$ is positive. Note that magnetic potentials are proportional to $|B|$, as the direction of the field defines the precession axis of the system, such that there will always be a potential minimum where $B = 0$.

1.4.3 Interactions with light

The development of techniques to cool and trap atoms with laser light was the key to opening up the world of ultracold atomic gases, being recognised with the bestowal of the 1997 Nobel Prize for Physics on Steven Chu, Claude Cohen-Tannoudji and William Phillips. That light has an effect on an atom stems from the fact that a dipole \mathbf{d} may be induced in an atom in the presence of an external electromagnetic radiation field \mathbf{E} , in general given by the expectation value

$$\langle \Psi | \mathbf{d} | \Psi \rangle = \alpha \mathbf{E} , \quad (1.38)$$

where α is the complex polarisability of the atom. The effects of light on atoms can be split into two categories: reactive effects and absorptive effects, the former being more

1.4 Properties of alkali atoms

important when the frequency of the external field is far from resonance, the latter when the frequency is near resonance. These two effects are directly related to the real and imaginary parts of α . Much information may therefore be garnered from finding α , for which we need to know the electronic wavefunction Ψ . This may be approximated by considering the interaction as a perturbation of the isolated atom system, with the total Hamiltonian

$$\hat{H} = \hat{H}_0 + \hat{H}_{\text{dip}} , \quad (1.39)$$

where \hat{H}_0 is the Hamiltonian describing the atom in the absence of the field, and the Hamiltonian representing the perturbing dipole interaction is given by

$$\hat{H}_{\text{dip}} = -\mathbf{d} \cdot \mathbf{E} , \quad (1.40)$$

noting that $\mathbf{d} = -e\mathbf{r}$, where we are considering the coordinate system where the atomic nucleus is located at $\mathbf{r} = 0$. This shows that the dipole moment has odd parity which is used in the following discussion.

As is standard in perturbation theory we set about finding the solution to the time-dependent Schrödinger equation

$$\hat{H}\Psi(\mathbf{r}, t) = i\hbar \frac{d}{dt}\Psi(\mathbf{r}, t) \quad (1.41)$$

by considering the wavefunction to be a superposition of the eigenstates of the unperturbed system:

$$\Psi(\mathbf{r}, t) = \rho_1(t)\Psi_1(\mathbf{r}, t) + \rho_2(t)\Psi_2(\mathbf{r}, t) , \quad (1.42)$$

with

$$\Psi_i(\mathbf{r}, t) = e^{-\frac{i}{\hbar}E_i t}\varphi_i(\mathbf{r}) . \quad (1.43)$$

Here we have simplified the full description to a two-level system (lower state 1, upper state 2) where the time dependent coefficients $\rho_i(t)$ are normalised such that the sum of the state populations $|\rho_1|^2 + |\rho_2|^2 = 1$. The quantities E_i are the energy eigenvalues of the unperturbed Hamiltonian:

$$\hat{H}_0\Psi_i = E_i\Psi_i . \quad (1.44)$$

Using (1.41), (1.42), (1.43) and (1.44) we obtain

$$\hat{H}_{\text{dip}}(\rho_1\Psi_1 + \rho_2\Psi_2) = i\hbar(\Psi_1\dot{\rho}_1 + \Psi_2\dot{\rho}_2) , \quad (1.45)$$

where $\dot{\rho}_i \equiv d\rho_i/dt$. If we assume that the applied radiation field oscillates harmonically in time with frequency $\omega/2\pi$ we can write

$$\mathbf{E}(t) = \mathbf{E}_0 \cos(\omega t) , \quad (1.46)$$

where it is also assumed that any spatial variation of the field amplitude is insignificant over the length scale of the atom. To eliminate the spatial dependence of (1.45) we can multiply through by the complex conjugate wavefunction Ψ_i^* and integrate over space. Doing this for both $i = 1$ and $i = 2$ yields

$$\dot{\rho}_1(t) = \frac{i}{\hbar} \rho_2(t) \mathbf{d}_{12} \cdot \mathbf{E}_0 \cos(\omega t) e^{-\frac{i}{\hbar}(E_2 - E_1)t} \quad (1.47)$$

$$\dot{\rho}_2(t) = \frac{i}{\hbar} \rho_1(t) \mathbf{d}_{21} \cdot \mathbf{E}_0 \cos(\omega t) e^{\frac{i}{\hbar}(E_2 - E_1)t} . \quad (1.48)$$

Here we have introduced the dipole matrix elements

$$\mathbf{d}_{ij} = \langle i | \mathbf{d} | j \rangle = \int \varphi_i^*(\mathbf{r}) \mathbf{d}(\mathbf{r}) \varphi_j(\mathbf{r}) dV \quad (1.49)$$

coupling the two states. For convenience we may choose the phases of the wavefunctions such that the dipole matrix elements are real. We can then define the Rabi frequency Ω (a measure of the strength of the coupling between the light and the applied field) as

$$\Omega = \frac{|\mathbf{d}_{12} \cdot \mathbf{E}_0|}{\hbar} = \frac{|\mathbf{d}_{21} \cdot \mathbf{E}_0|}{\hbar} . \quad (1.50)$$

This is not quite the full story as we have not included the effect of spontaneous emission. This is easily rectified by introducing a damping term $(\Gamma_\omega/2)\rho_2(t)$. Also defining the resonant transition frequency $\omega_0 = (E_2 - E_1)/\hbar$ we may rewrite (1.47) and (1.48) as

$$\dot{\rho}_1(t) - \frac{\Gamma_\omega}{2} \rho_2(t) = i \rho_2(t) \Omega \cos(\omega t) e^{-i\omega_0 t} \quad (1.51)$$

$$\dot{\rho}_2(t) + \frac{\Gamma_\omega}{2} \rho_2(t) = i \rho_1(t) \Omega \cos(\omega t) e^{i\omega_0 t} . \quad (1.52)$$

$\Gamma_\omega = (\omega^2/\omega_0^2)\Gamma$ represents the off-resonant damping coefficient [25] which reduces to Γ close to resonance (or indeed in the absence of an applied field). Note that setting $\mathbf{E}_0 = 0$ (i.e. $\Omega = 0$) in (1.52) gives the result that, in the absence of external radiation, the upper state population decays at a rate Γ , and

$$|\rho_2(t)|^2 = |\rho_2(0)|^2 e^{-\Gamma t} , \quad (1.53)$$

as expected. However, we are more interested in the steady-state solutions in the presence of an external field, i.e. the case where the state populations are constant:

$$\frac{d}{dt} |\rho_i(t)|^2 = 0 . \quad (1.54)$$

For near resonant light we may make the rotating-wave approximation, $\delta = \omega - \omega_0 \ll \omega + \omega_0$, such that in the steady-state situation the time-dependence of $\rho_1 \rho_2^*$ is simply

1.4 Properties of alkali atoms

given by the factor $\exp(i\delta t)$. With these conditions, and using

$$\frac{d}{dt}(\rho_i \rho_j^*) = \rho_i \dot{\rho}_j^* + \dot{\rho}_i \rho_j^* , \quad (1.55)$$

we may use (1.51) and (1.52) to obtain the equations

$$\Omega e^{+i\delta t} \rho_2 \rho_1^* - \Omega e^{-i\delta t} \rho_1 \rho_2^* - 2i\Gamma \rho_2 \rho_2^* = 0 \quad (1.56)$$

$$\Omega(\rho_1 \rho_1^* - \rho_2 \rho_2^*) + (2\delta - i\Gamma) e^{-i\delta t} \rho_1 \rho_2^* = 0 . \quad (1.57)$$

We may solve these directly to obtain, for example,

$$\rho_2 \rho_2^* = |\rho_2|^2 = \frac{\Omega^2}{\Gamma^2 + 4\delta^2 + 2\Omega^2} \quad (1.58)$$

for the upper state population, which reaches a maximum of $1/2$ as $\Omega \rightarrow \infty$. We may now find a value for α by considering the expectation value of the dipole moment,

$$\langle \Psi | \mathbf{d} | \Psi \rangle = \int \Psi^* \mathbf{d} \Psi \, dV = \rho_1^* \rho_2 \mathbf{d}_{12} e^{-i\omega_0 t} + \rho_1 \rho_2^* \mathbf{d}_{21} e^{i\omega_0 t} = \alpha \mathbf{E}_0 \cos(\omega t) . \quad (1.59)$$

In the near-resonance case that we have been discussing we may use the steady-state solutions obtained from the above equations to find

$$\alpha_{\text{near}} = \frac{\mathbf{d}_{12} \mathbf{d}_{21}}{\hbar} \frac{-\delta + i\Gamma/2}{\Gamma^2/4 + \delta^2 + \Omega^2/2} . \quad (1.60)$$

In the far-detuned case we can no longer make the rotating-wave approximation, but instead we may assume $\Gamma^2 \ll \delta^2$ and that the coupling Ω is weak. Similar arguments eventually lead to an expression for the polarisability in this regime,

$$\alpha_{\text{far}} = \frac{\mathbf{d}_{12} \mathbf{d}_{21}}{\hbar} \frac{2\omega_0 [\omega_0^2 - \omega^2 - i(\omega^3/\omega_0^2)\Gamma]}{(\omega_0^2 - \omega^2)^2} . \quad (1.61)$$

The reactive dipole force may be found from

$$\mathbf{F}_{\text{dip}} = -\nabla U_{\text{dip}} \quad (1.62)$$

where

$$U_{\text{dip}} = -\frac{1}{2} \langle \mathbf{d} \cdot \mathbf{E} \rangle_t = -\frac{1}{2\epsilon_0 c} \text{Re}(\alpha) I(\mathbf{r}) . \quad (1.63)$$

$\langle \dots \rangle_t$ represents the time average and the factor of $1/2$ appears as the dipole is induced rather than permanent. We have also used $I(r) = (1/2)\epsilon_0 c |\mathbf{E}_0(r)|^2$ and the fact that the time-averaged value of $\cos^2(\omega t)$ is $1/2$. The \mathbf{r} dependence is now on the macroscopic scale, not the atomic scale, and we see a force that is proportional to the gradient of the intensity of the external field. As we are taking the real part of the polarisability we can see that this force is vanishing in the near-resonant limit.

The absorptive scattering force is found from

$$F_{\text{sc}} = \hbar \mathbf{k} \Gamma_{\text{sc}} , \quad (1.64)$$

where \mathbf{k} is the wave-vector of a photon of the laser field (so its direction is that of the beam propagation), and Γ_{sc} is the scattering rate of the atom. The force is then considered to arise from the atom receiving a kick of momentum $\hbar \mathbf{k}$ each time it absorbs a photon, while we consider the momentum kicks from re-emission to average to zero. This picture will be considered later in §2.3 when discussing the limits of laser cooling methods. Γ_{sc} can be found by considering the power absorbed by an atom and dividing by the energy attributed to an incident photon, $\hbar\omega$:

$$\Gamma_{\text{sc}} = \frac{P_{\text{abs}}}{\hbar\omega} = \frac{1}{\hbar\omega} \langle \dot{\mathbf{d}} \cdot \mathbf{E} \rangle_t = \frac{1}{\hbar\epsilon_0 c} \text{Im}(\alpha) I(\mathbf{r}) . \quad (1.65)$$

Near-resonant light

Taking all of the above into account we can introduce the saturation intensity I_{sat} where

$$\frac{I}{I_{\text{sat}}} = \frac{2\Omega^2}{\Gamma^2} \quad (1.66)$$

to finally obtain the scattering force as

$$F_{\text{sc}} = \hbar \mathbf{k} \frac{\Gamma}{2} \frac{I/I_{\text{sat}}}{1 + (2\delta/\Gamma)^2 + I/I_{\text{sat}}} . \quad (1.67)$$

This force is applied in the early stages of the experiment to trap and cool atoms as is discussed in §2.3. Note that Γ_{sc} could also have been found by multiplying the value of $|\rho_2|^2$ given in (1.58) by Γ , being simply the steady-state population of the upper state multiplied by the decay rate from this state.

Far-detuned light

By using the quantum mechanical result for the spontaneous decay rate of the upper state (see e.g. [26, 27])

$$\Gamma = \frac{\omega_0^3}{3\pi\epsilon_0 \hbar c^3} |\langle 1 | \mathbf{d} | 2 \rangle|^2 \quad (1.68)$$

with (1.61) and (1.63), we may obtain an expression for the dipole potential

$$U_{\text{dip}}(\mathbf{r}) = -\frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) I(\mathbf{r}) \quad (1.69)$$

and using (1.65) the corresponding scattering rate

$$\Gamma_{\text{sc}}(\mathbf{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\omega}{\omega_0} \right)^3 \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right)^2 I(\mathbf{r}) . \quad (1.70)$$

1.5 Collisions between ultra-cold atoms

The dipole force is stronger than the scattering force in the far-detuned regime, but due to the typical temperatures of atoms in traps formed by dipole potentials it is useful to know the scattering rate of the atoms, as only a few scattering events can cause losses from the trap. Assuming for a moment that the detuning is small enough that $1/\delta \gg 1/\omega_0$, we can see that U_{dip} is roughly proportional to I/δ while the scattering rate is roughly proportional to I/δ^2 . It is therefore generally preferable to detune the laser beam far from resonance while compensating for the potential strength by increasing the power in the laser as far as possible, for greater lifetimes and lower heating rates in a dipole potential. Also note from (1.69) that if $\omega < \omega_0$ (red-detuning) a potential minimum can be formed by an intensity maximum, while for $\omega > \omega_0$ (blue-detuning) a potential minimum is formed by an intensity minimum.

1.5 Collisions between ultra-cold atoms

Collisions play a central role in both the formation and behaviour of ultra-cold atomic gases. Elastic collisions enable thermalisation of a cloud of atoms, allowing cooling via evaporation in the final step towards quantum degeneracy in many experiments (§2.7), while inelastic collisions limit the effectiveness of such a technique. The modification of the shape of the BEC due to collisional interactions has already been discussed in §1.3.2, while inelastic processes ultimately limit the maximum phase-space density it is possible to achieve in a cooled gas.

1.5.1 Scattering length

In general, collision theory is a complex topic, but the fact that we are dealing with the ultra-cold regime greatly simplifies the description. For now we shall consider the most pertinent case of elastic collisions between just two particles. The easiest way to describe such a system is to ignore the centre-of-mass energy and to consider just the relative position and momentum operators $\hat{\mathbf{r}}$ and $\hat{\mathbf{p}}$:

$$\hat{\mathbf{r}} = \hat{\mathbf{r}}_1 - \hat{\mathbf{r}}_2, \quad \hat{\mathbf{p}} = (\hat{\mathbf{p}}_1 - \hat{\mathbf{p}}_2)/2, \quad (1.71)$$

where the subscripts 1 and 2 refer to the individual atoms, using the notation of [28]. The Hamiltonian for the relative motion of the two particles then naturally follows as

$$\left(\frac{\hat{\mathbf{p}}^2}{2m_r} + V(\mathbf{r}) \right) \psi(\mathbf{k}, \mathbf{r}) = E_k \psi(\mathbf{k}, \mathbf{r}), \quad (1.72)$$

where $V(\mathbf{r})$ represents the inter-atomic potential, \mathbf{k} is the relative wavevector and m_r is the reduced mass, equal to $m/2$ in the case of atoms of the same species. We assume that $V(\mathbf{r}) \rightarrow 0$ as $|\mathbf{r}| = r \rightarrow \infty$, such that the eigenstates of the system have energy

$$E_k = \frac{\hbar^2 k^2}{2m_r}. \quad (1.73)$$

If we assume that the potential is spherically symmetric, the solution to (1.72) in the $r \rightarrow \infty$ limit takes the asymptotic form of a superposition of an incident wave and a scattered wave, given by

$$\psi = e^{ikz} + f(\theta, k) \frac{e^{ikr}}{r} , \quad (1.74)$$

where the z direction has been chosen for the propagation of the incident wave. Here $f(\theta, k)$ is the scattering amplitude which depends on the angle θ between the relative incident and scattered wave-vectors. Due to the axial symmetry with respect to z , the general form of the wavefunction can be expanded as ‘partial waves’ in terms of the Legendre Polynomials $P_l(\cos \theta)$, taking the form

$$\psi = \sum_{l=0}^{\infty} A_l P_l(\cos \theta) R_{kl}(r) . \quad (1.75)$$

Using the usual quantum mechanical result for motion in a centrally symmetric field (see e.g. [26]), the radial wavefunction must satisfy the equation

$$\frac{d^2(rR_{kl}(r))}{dr^2} + \left(k^2 - \frac{l(l+1)}{r^2} - \frac{2m_r V(r)}{\hbar^2} \right) rR_{kl}(r) = 0 . \quad (1.76)$$

For $k \rightarrow 0$, the relative particle energy will be far lower than the barrier formed by the addition of the centrifugal potential $\hbar^2 l(l+1)/2m_r r^2$, such that only the partial wave for $l=0$ will ‘see’ the potential $V(r)$, and all other waves will not contribute to the scattering behaviour¹. The scattering amplitude for $l=0$ is independent of θ , and so is referred to as the s -wave scattering length, which is written as a . We therefore expect in the limit $r \rightarrow \infty$ and $k \rightarrow 0$ to find the solution

$$\psi_0 = 1 + \frac{a}{r} . \quad (1.77)$$

This is simply the result for a hard-core potential of radius a , giving a means of interpreting the scattering length. It can be shown [28, 29] that in the low k limit a is independent of k , so low energy scattering is completely determined by a . This applies to the total scattering cross-section σ , which for identical bosons is given by

$$\sigma = 8\pi a^2 . \quad (1.78)$$

The extra factor of two over the usual expression results from the symmetry of the wavefunction under the exchange of the two sets of particle coordinates for identical bosons. The scattering length also gives the condition for a gas to be considered dilute, being the regime where $an^{1/3} \ll 1$. The quantity $an^{1/3}$ is therefore referred to as the diluteness parameter.

As the behaviour of the atoms depends only on the scattering length, we can choose

¹Technically this depends on the form of r as $r \rightarrow \infty$. When this is given by the Van de Waals potential $\propto r^{-6}$ it is a good approximation. See [26] for details.

1.5 Collisions between ultra-cold atoms

an effective pseudopotential to model the system as long as the scattering length a is retained. The simplest model describing interactions between two particles uses the contact potential

$$V(\mathbf{r}) = g\delta(\mathbf{r}) , \quad (1.79)$$

where $\delta(\mathbf{r})$ is the Dirac delta function centred at $\mathbf{r} = 0$. The proper way to treat the contact interaction is discussed in [30], and it leads to a total s -wave elastic scattering cross-section for polarised Bosons of [28]

$$\sigma(k) = \frac{8\pi a^2}{1 + k^2 a^2} , \quad (1.80)$$

where the contact energy g is found to be the interaction coupling constant given in (1.23). This gives two very different results in the small and large ak limits:

$$\sigma(k) \simeq \begin{cases} 8\pi a^2 & \text{for } ka \ll 1 \\ 8\pi/k^2 & \text{for } ka \gg 1 . \end{cases} \quad (1.81)$$

The first case is the result (1.78) for identical bosons at ultra-low energy, while the second case reveals the upper limit to the cross-section for s -wave scattering. This is known as the unitarity limit, the regime in which the inter-particle spacing is considerably greater than the range of the two-body interaction, but considerably smaller than the magnitude of the scattering length. In this scenario both the form of the interatomic potential and the precise scattering length cease to matter when considering few-body scattering. This leads to intriguing universal properties, such as the shallow three-body bound ‘Efimov states’ in the vicinity of a Feshbach resonance [12, 31]. This limit can be reached in ultra-cold atoms only when the scattering length can take divergent values. For the most-part the former cross-section is appropriate.

Real potentials (unlike hard-core potentials) can result in either positive or negative a . Calculation of the scattering length requires a precise knowledge of the shape of the interaction potential, being very sensitive to small alterations, in particular when the energy of a quasi-bound state is closely matched to the potential depth U_0 . This is illustrated in Fig. 1.7, where the quantity $rR_0(r)$ has been plotted for three different scenarios in the style of [32]. The interatomic interaction has been modelled using a Lennard-Jones potential, which takes the form

$$V(r) = -C_6 r^{-6} + C_{12} r^{-12} \quad (1.82)$$

where C_6 and C_{12} are positive. Fig. 1.7 **a** shows a standard scenario where there is no near-threshold bound state, producing a moderate positive scattering length. If the trap depth is increased, a quasi-bound state with energy slightly greater than U_0 causes the wavefunction to start to ‘turn over’, and the extrapolation from the large r limit designates an increasingly large negative scattering length, as shown in Fig. 1.7 **b**. As




Figure 1.7: Three different cases showing the effect that small changes in the interaction potential can have on the resultant scattering length. The shape of the potential $V(r)$ is in black, while the radial wavefunction $rR_0(r)$ is in red. The well depth U_0 is increased from **a** to **c**. **a** In a simple case where there is no bound state near the dissociation energy limit the scattering length takes a moderate positive value a as indicated. **b** As a bound state approaches resonance the extrapolation of the wavefunction from the large r limit (dotted red line) indicates an increasingly negative scattering length, diverging completely on resonance. **c** Past the resonance position an extra node appears in the wavefunction and a positive scattering length is once again found. These plots are in the style of those found in [32].

the bound state comes into resonance with U_0 the scattering length diverges, until a new bound state exists and an extra node appears in the wavefunction. Just past the resonance the scattering length is therefore large and positive as, depicted in Fig. 1.7 **c**.

In practice, the shape of the interatomic potential is too complicated to directly calculate, and experimental methods such as photoassociation spectroscopy [33], as well as the observation of collisional dynamics in cold atom clouds (as studied in [6] for ^{39}K), are used to determine the scattering length. The scattering properties also depend on the internal states of the atoms. The total electronic spin S in a binary system can be either 0 (singlet) or 1 (triplet). Each case gives rise to a different interatomic potential, resulting in different scattering lengths a_s and a_t respectively. For collisions between maximally polarised atoms (i.e. for $m_F = F$), the single-channel scattering is fully determined by the triplet potential, but the scattering length in other

1.5 Collisions between ultra-cold atoms

states is a mixture of a_s and a_t due to the effects of the hyperfine interaction. However, circumstances can conspire such that a second channel plays a large role in determining the scattering properties, as now described.

1.5.2 Feshbach resonances

The case in which there are two scattering channels to consider is illustrated in Fig. 1.8. In black is the usual open channel, which for large separation corresponds to two free atoms in a gas, while in red is the closed channel. In the ultra-cold regime the atoms have little kinetic energy, and the total atomic energy tends to that of the open channel potential at large separation. Atoms must therefore enter and exit in the open channel as they do not have enough energy to separate in the closed channel. This means there is no coupling at first order between the channels. However, if the closed channel supports a bound state with an energy close to that of the atoms then an extra scattering pathway is available, and the atoms may briefly scatter into the closed channel before emerging in the open channel. In this case even weak coupling can lead to strong mixing between the channels and a large change in the scattering length from the pure single-channel scenario. This second-order coupling scenario is a familiar one, and from perturbation theory we would expect a sum of contributions to the scattering length of the form [29]

$$a \sim \frac{C}{E_0 - E_{\text{res}}} , \quad (1.83)$$

where E_0 is the energy of the particles in the open channel and E_{res} is the energy in the closed channel. For channels with different magnetic moments, as is the case for the singlet and triplet potentials, the application of a magnetic field allows the relative energy of the bound state to be tuned and brought to resonance. This is known as a magnetically-tuneable Feshbach resonance, the benefits of which have already been discussed in §1.1. To reflect the ability to tune with a magnetic field the total scattering length near a particular Feshbach resonance is written in the form

$$a(B) = a_{\text{bg}} \left(1 - \frac{\Delta}{B - B_\infty} \right) , \quad (1.84)$$

as introduced in [34], relating the scattering length in the presence of a magnetic field of strength B to the off-resonant value a_{bg} via a resonance position B_∞ and width Δ . Note that this form holds as long as the resonance is well isolated from other resonances. In this case the sum of the rest of the terms with the form of (1.83) will effectively be constant, and their effect is absorbed into the value a_{bg} . The above equation is derived in a considerably more rigorous fashion in [35] and [17], where it is shown that the width parameter Δ may be written

$$\Delta = \frac{m(2\pi\hbar)^3}{4\pi\hbar^2 a_{\text{bg}} \mu_{\text{res}}} |\langle \psi_{\text{res}} | \hat{H}_{\text{ex}} | \psi_0 \rangle|^2 , \quad (1.85)$$




Figure 1.8: Basic two-channel model for a Feshbach resonance. The phenomenon occurs when atoms incident upon one another in the open channel (black line) have an energy approaching that of a bound state (red dashed line) in the closed channel (red line). Mixing between states leads to a divergent scattering length when the bound state is resonant with the colliding atoms.




Figure 1.9: Characteristic shape of a magnetically tuneable Feshbach resonance. Here the case for the ^{39}K resonance centred at 402.5 G is shown.

very suggestive of the perturbative origins of (1.84). Here \hat{H}_{ex} is the spin exchange coupling Hamiltonian, ψ_0 and ψ_{res} are the open and closed channel energy states respectively, and μ_{res} is the difference between the magnetic moments of the molecular bound state and the asymptotically separated non-interacting atoms.

The resonance profile (1.84) is shown in Fig. 1.9, using the parameters for the ^{39}K resonance in the $|F, m_F\rangle = |1, 1\rangle$ state listed in Table 1.1. Note that both the background scattering length and the resonance width are negative in this case. The zero-crossing point (where the scattering length goes to zero) is given by $B_0 = B_\infty + \Delta$. The existence of such resonances was experimentally confirmed in a BEC of sodium atoms in 1998 [36]. Since then, access to such resonances has become an extremely desirable tool for many experimentalists, while for the studies described in this thesis it is essential.

1.5.3 Inelastic processes

We conclude this chapter with a brief discussion of inelastic collisions which are, for the most part, undesirable but inevitable in ultra-cold atoms experiments. Most obviously, the imperfect vacuum of any experimental system leads to collisions between the trapped gas and the background vapour. This vapour is at room temperature such that the energies of the constituent particles are far greater than that of the trapped atoms. This causes exponential decay of the number of trapped atoms, at a rate determined by the background pressure in the vacuum chamber. This is usually the dominant loss mechanism for clouds at low density.

At higher densities collisions between three particles become more frequent, increasing the probability of three-body recombination, in which two particles form a molecule while the third carries away the excess energy. This is often particularly undesirable, as not only are all three atoms lost, but the coldest atoms are preferentially affected. This is often the dominant loss mechanism for condensates due to their high densities. The process has also been shown to scale as a^4 [37], so can be greatly enhanced near a Feshbach resonance.

Finally, dipole-dipole interactions between polarised atoms in a magnetic field can lead to two-body spin relaxation [38], where the spin state of one or both of the colliding atoms may be flipped. Due to the Zeeman splitting of spin states, this is energetically favourable. When a cloud is magnetically confined this process can lead to atom loss by the scattering of atoms into untrappable states. These three processes all contribute to the overall heating and atom loss which act to oppose the desired increase in phase-space density required for the formation of BECs, so the cooling mechanisms used must work at a higher rate to overcome them.

Chapter 2

Experimental setup

“ The only way out that I can think of is to ask the old gang to let their attention wander for a bit - there are heaps of things they can be doing; washing the car, solving the crossword puzzle, taking the dog for a run - while I place the facts before the newcomers. ”

- P. G. Wodehouse (Much Obligated, Jeeves)

The construction of a working experimental setup to reliably create and detect quantum degenerate gases is no mean feat in itself, and typically still takes a new group a period on the order of years to achieve. With the increasing global knowledge of these systems start-up times are coming down, but the practical details are still inevitably numerous. This chapter provides the particulars of the experimental methods and setup in some detail, with the hope that it will prove useful to future students setting out on a similar road to ours. For those better acquainted with cold atoms research it may be convenient to use as a reference, with the more concise details of our experimental sequence to produce ^{39}K BECs being given in Chapter 3.

The system in which we have condensed and experimented with ^{39}K is, in fact, the second BEC machine in our lab. A trial single-chamber system was first created, the main purpose of which was to familiarise ourselves with the basic experimental methods required for the laser-cooling of atoms by creating a magneto-optical trap (see §2.3). The setup exceeded our expectations though, and was pushed further to produce ^{87}Rb BECs of $\sim 10^5$ atoms in an optically-plugged quadrupole trap [4]. Working with a single chamber creates inherent limits on what is experimentally achievable (e.g. the relatively short lifetimes of trapped atoms due to the pressure required for loading a magneto-optical trap), but it is still hoped that this system will prove fruitful in the future with the aid of planned improvements. To go into detail here would be to stray a little too far from the main story, but a brief description of the trial system may be found in Appendix B. The experimental details given in the remainder of this chapter apply only to the second system.

2.1 Vacuum system




Figure 2.1: Simplified representation of our experimental setup showing a horizontal cross-section through the vacuum system. On the right is the MOT cell, at a relatively high pressure ($\sim 10^{-9}$ mbar) for efficient collection of atoms, while on the left is the ‘science cell’ where evaporative cooling occurs, and the pressure is $\sim 10^{-11}$ mbar. Once the requisite number of atoms has been loaded into the MOT, they are transported to the science cell by the transport coils.

2.1 Vacuum system

In contrast to our trial setup the main system has a two-chamber arrangement as shown in Fig. 2.1. The essential stages in trapping and cooling atoms are a laser-cooling stage in a magneto-optical trap (MOT), and an evaporative cooling stage. The former requires a relatively high pressure to collect a large number of atoms in the first place, while the latter requires as low a pressure as possible to avoid undesired atom loss. The first chamber is therefore referred to from here on in as the MOT cell, where atoms are trapped and laser-cooled from a background vapour. The second chamber is referred to as the science cell, which is held at a lower pressure than the MOT cell, providing the extended lifetimes required for our experiments¹. The pressures required throughout the system are in what is known as the ultra-high vacuum (UHV) regime, typically defined as $\sim 10^{-12}$ to 10^{-9} mbar. The difference in pressure between the two regions is maintained by the continual operation of two ion pumps separated by low-conductance tubes. The background vapour of alkali atoms in the MOT cell is sustained by regular deposition on the chamber walls from atom sources located in an adjoining vacuum tube.

2.1.1 Differential pumping design

The basic design of a two-chamber vacuum system depends on the desired ratio of pressures between the two regions. Applying the old adage that it is better to be safe than sorry, we decided to aim for a drop of around three orders of magnitude across

¹The system is displayed in figures such that the MOT cell is on the right while the science cell is on the left. This is due to its orientation in the laboratory - the familiarity of several years’ work is hard to shake.




Figure 2.2: Representation of our differential pumping system (a simplified representation of Fig. 2.1). The pressure p_0 is assumed to be constant due to the vapour pressure from the rubidium and potassium on the walls of the MOT chamber which is regularly replenished by the atom sources. The low conductances c_1 and c_2 allow a large stable pressure drop to the science cell at p_2 , maintained by the relatively high pumping speeds s_1 and s_2 to the ion pumps.

the system. This could be achieved with a single low-conductance tube and a single pump, but the tube would either have to be excessively long or narrow, which would be detrimental to the transportation of atoms along the system. Instead we chose a two-stage differential-pumping setup, a simplified depiction of which is shown in Fig. 2.2. Pressures are denoted p , conductances c and pumping speeds as s . p_0 is the pressure in the MOT cell, p_2 in the science cell. In reality the system is made up of many components in series, but we may simply calculate total conductances c_T using

$$\frac{1}{c_T} = \sum_i \frac{1}{c_i} . \quad (2.1)$$

As we use identical pumps with similar conductances at their inlets we make the simplification $s_1 = s_2 = s$. By considering the system to have reached a steady state it is possible to equate total pumping speeds to and from each vacuum region, to give the pressure ratio from MOT cell to science cell as

$$\frac{p_2}{p_0} = \frac{c_1 c_2}{c_1 c_2 + c_1 s + 2c_2 s + s^2} . \quad (2.2)$$

Generally $c \ll s$, such that (2.2) may be approximated as $p_2/p_0 \approx c_1 c_2 / s^2$. This may be compared with the fraction c/s for a single pumping stage. The total pressure drop is not the only consideration for the design of the differential pumping system however. Once we have collected the desired number of atoms in the MOT chamber, we want to transport as many of them as possible to the science cell at as low a temperature as possible, so any collisions between trapped atoms and the background gas in the vacuum chamber are undesirable. The main initial concern therefore is dropping the pressure *quickly*. We could of course do this by means of a very small aperture (and hence low conductance) between the MOT cell and the first pump, but there is a limit to how small an aperture we can safely transport the atoms through. In the molecular flow regime¹ the conductance of a long circular tube is given (in litres per second) by

¹The molecular flow regime is roughly defined to be where the mean free path of particles in a gas is greater than the size of the confining system, such that generally the particles travel from chamber wall to chamber wall without encountering other particles. The system is well within this regime.

2.1 Vacuum system




Figure 2.3: Pressure drop along a theoretical vacuum system as described in the text, where the final pressure reached is fixed to be 0.1% of p_0 . l_1 varies from $0.25l_0$ to $2.0l_0$ in steps of $0.25l_0$. The black line represents the case where $l_1 = l_2 = l_0$, while red and blue lines represent $l_1 < l_0$ and $l_1 > l_0$ respectively. There is a clear trade-off between rate of pressure drop-off and the pressure reached after the first stage of differential pumping. This general scaling informed the choice of vacuum components in our system. Inset is the same information on a linear-linear scale, illustrating the linear scaling of the pressure with distance along a tube of uniform cross-section.

[39]

$$c \text{ (1s}^{-1}\text{)} \approx 12.4 \left(\frac{d}{1 \text{ cm}} \right)^3 \left(\frac{1 \text{ cm}}{l} \right) \left(\frac{T}{295 \text{ K}} \times \frac{m_{\text{N}_2}}{m} \right)^{1/2}, \quad (2.3)$$

where d is the diameter and l the length of the tube, m is the particle mass and T the temperature. Fixing d to be the minimum diameter for atom transportation results in conductances inversely proportional to l . The pressure drop for the first stage of pumping is given by

$$\frac{p_1}{p_0} = \frac{c_1(c_2 + s)}{c_1c_2 + c_1s + 2c_2s + s^2}, \quad (2.4)$$

while the pressure drop along a tube is simply linear. The result of all this is that if we fix the total pressure drop p_2/p_0 then the pressure gradient is greater initially for shorter length l_1 of the first tube (corresponding to conductance $c_1 \propto 1/l_1$). This is illustrated in Fig. 2.3 on a logarithmic pressure scale for $p_2/p_0 = 1/1000$, where the black line represents the case where $l_1 = l_2 = l_0$, while red and blue lines illustrate cases where $l_1 < l_0$ and $l_1 > l_0$ respectively. The plot shows the pressure drop along the first tube immediately followed by the pressure drop along the second ignoring any interim space, $\tilde{x} = x/l_0$ representing distance along the system. Here it is clear that the full pressure drop is achieved in the shortest distance for $l_1 \approx l_2$, but faster initial drops are achieved for $l_1 < l_2$. Of course if the first tube is too short then the pressure p_1 won't

be much lower than p_0 , defeating the purpose of using a small l_1 . A happy medium is achieved when p_1 is acceptably low. Our final design has a theoretical $p_1/p_0 \approx 1/27$ and $p_2/p_0 \approx 1/1200$. It can be seen on the cross-section shown in Fig. 2.1 that we have adhered to the spirit of the above arguments as far as possible, with the customised ‘reducing cross’ piece on the right having a wider flange towards the MOT cell so that the pressure is reduced more quickly (corresponding to an effective low l_1), with the narrower flange making up part of the longer second conductance tube.

The simple model we have been considering does not hold so well once we are working at extremely low pressures (in the 10^{-11} mbar regime at the science cell) due to outgassing from the walls of the vacuum chamber becoming a limiting factor. This is reduced as far as possible as discussed in §2.1.3, but nonetheless in reality the pressure at the science cell is closer to only two orders of magnitude lower than at the MOT cell. The effect is mainly on p_2 , and as stated earlier it is worth leaving some leeway in the design. The lifetimes of several minutes achieved for trapped gases in the science cell are extremely satisfactory.

2.1.2 Components

The complete vacuum system is shown in Fig. 2.4 with a breakdown of the key component parts. Most of the pieces are standard items bought from various vacuum companies, while others, such as the ‘4-way reducer crosses’, were customised on request for non-standard lengths and tube diameters. All of the metal near the atom-transport axis is 316LN stainless steel as opposed to standard 304L, the main advantage being the lower magnetic susceptibility of the material (0.001 instead of 0.03). This is to minimise any possible distortion of the magnetic fields we create in the system with external coils. All components in the system are UHV compatible, connected together with standard UHV conflat (CF) flanges and may be baked to well over 200°C (see §2.1.3). Below the important features of some of the vacuum components are discussed.

Pumps

As has already been alluded to, UHV is maintained by the workhorse pumps of the system - a pair of *Gamma Vacuum* TiTan 75S ion pumps, which have a pumping speed of $\sim 60\text{ls}^{-1}$ each. The pumps operate by ionising gases and using an electric field to accelerate these ions into solid electrodes, trapping the particles by chemisorption. As such, no material physically leaves the system, so the pumps can only be operated below $\sim 10^{-5}$ mbar to have a reasonable lifetime. The effective pumping speed to the vacuum chamber is only around 20ls^{-1} , which is limited by the conductance between the pump and the main chamber, so an increase in pump size would have little effect.

In order to reach a pressure at which the ion pumps can operate we use an *Oerlikon Leybold* Turbovac TW70 H turbomolecular (turbo) pump. The turbo pump itself only

2.1 Vacuum system




Figure 2.4: An isometric projection of the vacuum system. **a** Ion pumps. **b** Titanium sublimation pump. **c** Turbomolecular pump. **d** MOT cell. **e** Science cell. **f** Atom sources. **g** All-metal right-angle valve. **h** Bayard-Alpert pressure gauge. **i** ‘Cube’ with windows.

operates with a maximum forevacuum of ~ 20 mbar, so it is backed by an *Oerlikon Leybold* SC 5 D scroll pump. The main advantage of using a scroll pump as a roughing pump is that it is oil free (unlike most alternatives), so there is no chance of oil feeding back into the system which would prevent the attainment of UHV. The turbo pump can itself be used to reach $\sim 10^{-9}$ mbar, well below the pressure required for ion pump operation, and is operated during baking to reduce the load on the ion pumps while pressures in the system are still relatively high. Once the turbo pump has completed its job we use the *VAT* all-metal angle valve to close off the vacuum chamber and we may detach the turbo pump and scroll pump. The ion pumps are supplemented by occasional use of a *Varian* 916-0050 series titanium sublimation (Ti sub) pump. This operates by forming a reactive coating on the walls of the vacuum chamber, effectively removing molecules such as H_2 which are less readily dealt with by the ion pumps. We apply a current of 47 A for ninety seconds roughly twice a year to replenish the deposited titanium.




Figure 2.5: Annotated diagram of the MOT cell with viewports attached. Some important dimensions are shown (in mm), such as the 92 mm total clearance required for the transport coils.

MOT cell

Vacuum chambers for laser cooling require optical access from several angles, and therefore typically fall into one of two broad categories: an all-glass chamber (as used in the trial setup - see Appendix B), or a steel body with viewports bolted on. We took the latter path for the sake of robustness and the fact that viewports can easily be anti-reflection (AR) coated on both surfaces, while typically all-glass chambers can only be coated on the outer surface (once the manufacturing process is complete). We chose to use a standard vacuum flange size of 70 mm diameter for the kodial viewports, which have 38 mm clear window diameter to allow for large MOT trapping beams.

At the point of designing the MOT cell we were yet to settle on the exact configuration of the complete system, most importantly the method of transportation of the atoms to the science cell. We had decided to implement magnetic transportation, but this may be done in two ways: Either with a series of coil pairs along the transportation axis, ramping up each pair in turn such that the magnetic minimum travels smoothly along the system [40], or by using a single coil pair mounted on a translation stage such that the magnetic minimum travels with the trapping coils as they are physically moved [41]. The design therefore had to cater to either eventuality. In either case a low vertical profile is advantageous, so the cell was designed to have a main body height of 44 mm as shown in Fig. 2.5, just slightly larger than the viewable diameter of the windows. This allowed enough margin for tubes to be welded to the main body, extending out to CF flanges for attaching the horizontal viewports. Added to this 44 mm

2.1 Vacuum system




Figure 2.6: Constructed MOT cell with the atom sources in the background. The cell is displayed with some of the supports used to mount the chamber during construction which were subsequently replaced by a single, less obstructive support to allow room for the translation stage and transport coils.

was the minimum flange height to accommodate threaded holes to securely attach the vertical windows, resulting in a total chamber height of ~ 92 mm once the height of the securing bolt heads are taken into account. This limits the trapping coil separation if the translation method is used, while the possibility of a series of coils is catered for by allowing the first pair (with inner diameter greater than 70 mm) to nestle against the MOT cell body with only 44 mm separation.

In the end we opted for the translation method for atom transportation as described in §2.5. The finished constructed MOT cell is shown in Fig. 2.6. It allows access for the three orthogonal pairs of beams in a standard MOT configuration with a couple of spare viewports for monitoring fluorescence from the cell, a connection to the atom sources and a narrow connection along the transportation axis to reduce conductance to the rest of the system.

Science cell

At the science cell the requirements are slightly different as we want to be able to take high resolution images of the atoms requiring close-proximity lenses, while also allowing for access for further lasers (e.g. the optical dipole trap) and surrounding magnetic coils. Due to the low profile and optical access that this requires, we decided to buy a high-quality glass cell from *Starna Scientific* as shown in Fig. 2.7. The cell has a 30×30 mm square outer cross-section, which allows a reasonable numerical aperture to the lens systems for both horizontal and vertical imaging, and is 110 mm long with 5 mm thick quartz walls which are AR coated on the outside (as mentioned above, the




Figure 2.7: The science cell shown with its connection to the rest of the vacuum system.

manufacturing technique precludes AR coating of the inner surface). A glass-metal connection extends to a small (34 mm) CF flange for connection to the rest of the vacuum system.

Atom sources

While we are trying to create such low-pressure conditions in the vacuum system, we also need a supply of our intended atomic species to study. As we have to bake the system to a temperature well above the boiling points of both rubidium and potassium to achieve UHV we cannot just leave a lump of material in the system. One solution is to intern a sample in a glass ampule, so that it is contained during baking, then to break the ampule once we are back at room temperature. The release of atoms into the system can then be controlled to some degree by heating or cooling the sample. A more controllable method is to use commercially available vapour sources which rely on a reaction which is activated at high temperatures ($> 300^{\circ}\text{C}$), above the required bake-out temperature. We chose to use sources from *Alvatec*, a company now specifically catering to cold atom experimentalists. These particular sources are based on an alloy of the alkali metal and have an indium seal which allows exposure to the atmosphere during construction of the vacuum system without detrimental effects. The indium seal melts at $\sim 160^{\circ}\text{C}$, after which the sources may be operated by passing a current through the material causing only the alkali metal to sublime. In order to be able to pass the required current of around 4-5 A through the sources we spot weld them to contacts on a UHV electrical feedthrough, allowing connection to an external power supply. When a current is applied the temperature rises in the sources, eventually releasing atoms to coat the walls of the vacuum chamber.

2.1 Vacuum system

Isotope	Natural abundance (%)
^{39}K	93.26
^{40}K	0.012
^{41}K	6.73
^{85}Rb	72.17
^{87}Rb	27.83

Table 2.1: Natural abundances of the isotopes of potassium and rubidium.

A major advantage of using rubidium and potassium for cold atoms experiments is that we may rely on the background pressure from the resultant coating to directly load a MOT. This is due to their relatively high vapour pressures at room temperature of around 4×10^{-7} mbar for rubidium, 2×10^{-8} mbar for potassium, compared with, for example, sodium for which it is about 4×10^{-11} mbar [42]. To reach an acceptable pressure to load a MOT, sodium needs substantial heating to temperatures where the Boltzmann distribution of velocities requires that only a small fraction of atoms will be travelling slowly enough to be captured. This then demands an extra atom cooling stage¹, increasing the complexity of the experiment.

It is usually enough for day-to-day operation in the lab simply to activate an atom source once every twenty-four hours to keep the pressure up, but for more critical experiments we can use shorter pulses each experimental cycle to keep the atom supply as steady as possible over the course of a day. We actually have two sets² of four different sources in the experiment - natural abundance rubidium and potassium sources and enriched ^{87}Rb and ^{40}K versions. The natural abundances are shown in Table 2.1, illustrating another advantage of studying ^{39}K (93.26% abundant). It is clear that it is essential to use enriched sources for studying ^{40}K , those in our system having a $> 10\%$ abundance in preparation for the possible future work with fermions. It is not essential to enrich rubidium, but it is relatively cheap in terms of extra expenditure to potential gain in system performance given that the enriched sources have a 98% abundance of ^{87}Rb , and that MOT performance scales with the ratio of partial pressures of the atom we are trapping to residual gases.

With these four different sources we can in principle have reasonable control over the relative pressures of each species in the MOT chamber. It is not quite straight forward as the sources coat each other, meaning that an increase in pressure is seen

¹Typically a Zeeman slower is used, which relies upon the combination of the Doppler and Zeeman effects to keep a beam of decelerating atoms in resonance with a laser acting at a constant frequency. This generally adds a meter of vacuum chamber and coils to the system, along with an oven to heat the atoms.

²A word of caution - we mounted one set such that the sources are pointing downwards without noting the recommendation that this not be done in the *Alvatec* information manual. This is the likely cause of faulty operation of some of our sources.

for all species whenever a single source is fired, but repeated firing of a single source does result in the expected alteration in composition of gas in the MOT cell. Each source contains around 50 mg of its respective alkali, enough to coat the MOT cell in a monolayer several thousand times over, but the lifetime of a source depends on the pumping rate from the MOT cell. We are yet to exhaust a source after two years' operation.

Pressure Gauge

To monitor the pressure while pumping down the system we have included a Bayard-Alpert style ion gauge from *Vacom*. This works very similarly to a very low speed ion pump, monitoring the ion current to gauge pressure. The current is highly species dependent however, and so the gauge is only useful as a rough guide, though it is quoted to work down to 3×10^{-11} mbar. Due to our chosen positioning for the gauge (see Fig. 2.4) we can only operate it while the system is not in use due to the obstruction caused by connecting the controller cable. The rest of the time we can use the ion pump controllers to give us a rough estimate of the pressure in the system.

Cube

The main curiosity-piece of our design is the central 'cube' with connected windows (see Fig. 2.4i). This was included mainly to enable the imaging of atoms part-way down the system for diagnostic purposes. It was also considered that we might initiate evaporative cooling here with additional magnetic coils to increase the trap gradient before continuing to the science cell. In the end we have not really taken advantage of it, though this owes something to the fortunate lack of major problems with the vacuum system. If we were to make a new system based on this design it would be tempting to leave the cube out to simplify the alignment of the system along the transportation axis.

2.1.3 Assembly and bake-out

Our prior knowledge of UHV systems was thin on the ground at the outset of the project and so much of the procedure for the cleaning, assembly and baking of the chamber was acquired from discussion with members of the extended AMOP group, PhD theses and in fact anyone who would care to give advice in the department at large. A particularly thorough source worth mentioning is the PhD thesis of Kevin Birnbaum [43] and additional associated information available online from the Quantum Optics Group at Caltech. We also experienced a steep learning curve with the construction of the trial system before setting to work on the main event.

2.1 Vacuum system

Cleaning

Briefly then, most vacuum parts were cleaned prior to construction of the vacuum chamber, either by hand with acetone soaked lens paper or by immersion in acetone in an ultrasonic bath where possible, handling components with clean nitrile gloves at all times. These items were then rinsed with isopropanol to remove any acetone residue and wrapped in clean aluminium foil once dry to await incorporation into the vacuum system. Exceptions were the right-angled valve and ion pumps which were left as we received them, and the viewports and science cell which were cleaned with methanol and isopropanol to avoid damage to the AR coatings.

Assembly

Smooth assembly of the system was aided by laying out all necessary parts prior to construction, i.e. nuts, washers, bolts, spanners, copper gaskets, vacuum components, and additional lens paper and methanol for last-minute cleaning of the CF flanges. The sealing mechanism involves a sharp knife edge on each of the pair of flanges that are being joined together. A copper gasket is placed in between, and as the bolts are tightened the knife edges bite into the copper in a complete ring to form a leak-tight seal. A final clean and close inspection of the knife-edges is essential given that this is where the seal is made.

In order to allow deconstruction of the system in the future it is important to lubricate the bolts to avoid seizure after baking. This can be a messy business if using a liquid lubricant which could accidentally end up inside the system, so we have used self-lubricating silver-coated bolts throughout to avoid the hassle. When forming a joint each bolt was made finger-tight before using spanners to gradually tighten the bolts in a star-shaped pattern to ensure an even seal. Once complete the flange faces should not quite be flush, allowing possible adjustment of the seal if a leak becomes apparent. The assembled vacuum system is shown in Fig. 2.8.

Bake-out

Even with the most stringent solvent-based cleaning of vacuum parts it is impossible to reach UHV without baking. This is because in the UHV regime the majority of particles in the system are adsorbed to the walls of the vacuum chamber rather than travelling between walls. The contaminants on the walls desorb at a high enough rate to limit the pressure in the chamber but too slowly to be effectively pumped away at room temperature. A dynamic model of desorption rate [44] based on adsorbed molecules requiring an activation energy to leave the surface of the vacuum chamber (analogous to vapour pressure as described by the Clausius-Calpeyron equation) predicts an adsorption lifetime proportional to $\exp(H_A/RT)$, where H_A is the enthalpy of adsorption and R is the gas constant. This is inversely proportional to pressure in the




Figure 2.8: Assembled system before bake-out. The turbo pump is removed once the bake-out is over.

molecular-flow regime, allowing us to relate temperature and pressure via the equation

$$\ln(P) = \left(-\frac{\Delta H_A}{R}\right) \left(\frac{1}{T}\right) + c. \quad (2.5)$$

If there is a fixed amount of material in the system we may increase the temperature to increase the pressure, and with each order of magnitude increase in pressure the pumping rate effectively increases by an order of magnitude (most pumps have a relatively constant pumping speed in terms of *volume*).

Prior to baking the scroll and turbo pumps are turned on for the initial evacuation of the system. To prepare for the baking (Fig. 2.9) we first cover the science cell with a protective aluminium tube and the viewports with aluminium foil ensuring there is no direct contact with the metal. This is to avoid direct local heating of glass which could cause damage to the AR coatings or possibly even a leak. We then attach thermocouples at key points around the system to monitor temperatures as it is important to avoid large temperature gradients (thirteen thermocouples were used for our bake-out). To allow independent control of heating rates for different sections of the system we wrap several heating tapes (PTFE coated resistance wires) around the system and attach each to a variable AC autotransformer (variac) to control the voltage applied to the wires. The heater tapes used were mainly *Isopad* TeMS 6 tapes, supplying a maximum heating power of 150 W m^{-1} . The ion pumps came with their own heating elements which can be bolted directly to their casings. Insulating the system well is vital to reach the desired baking temperature, so we first wrap the system in glass fibre tapes and blankets before adding several layers of aluminium foil.

Once the system has been fully prepared we can start to turn up the voltages on

2.1 Vacuum system




Figure 2.9: Preparing the system for baking. **a** The science cell is sheathed by an aluminium tube and the viewports are wrapped carefully in aluminium foil to avoid direct heating of glass components. Thermocouples are attached at strategic points to monitor the temperature throughout the system. **b** Several heating tapes are wrapped around the system enabling independent control of heating rates in different regions. **c** Insulation tape and blankets are wrapped around the system. **d** Final additional insulation is provided by layered aluminium foil.

the variacs while monitoring temperatures throughout the system, aiming to keeping the rate of increase in temperature below the recommended limit of $1\text{ }^{\circ}\text{C}$ per minute. It should be noted that the turbo pump should not exceed $\sim 80\text{ }^{\circ}\text{C}$, so we try to keep as gentle a gradient between it and the rest of the system as possible while adhering to this criterion. We do not turn the ion pumps on until we have passed the point at which the indium seals of the atom sources have melted at $\sim 160\text{ }^{\circ}\text{C}$ due to the sudden increase in pressure in the system at this point. Once we are sure we are past this point and the pressure has reached a reasonable level once more below 1×10^{-6} mbar we turn the ion pumps on to increase the overall pumping speed.

The first two bake-out attempts of the system were prematurely halted by broken heater tapes, most likely due to excessive local heating breaking a connection somewhere along the tape. As this would lead to a local cold spot we were forced to gradually bring the temperature in the rest of the system back down to room temperature before stripping away the insulation to replace the faulty part. Our third attempt went far more smoothly, and we brought the temperature of the system up to a steady temperature of $220\text{ }^{\circ}\text{C}$. All the components in the system were quoted as being bakeable to $250\text{ }^{\circ}\text{C}$, so we went as far as we were prepared to risk whilst allowing a safety margin. The baking period was determined by observing a plateau of the system pressure when held at the constant temperature. During the bake-out period we occasionally fired the Ti sub pump to see if it would help matters, and ran a low current through the atom sources to help degas them and make sure they were at a higher temperature than




Figure 2.10: Plot of how the pressure recorded by the Bayard-Alpert gauge varied during the final bake-out of the system. The black points represent the ramping up of temperature in the system (here measured by the thermocouple closest to the pressure gauge) and holding the temperature, while the red points represent the cooling period after baking. Spikes in pressure coincide with application of the Ti sub pump which causes temporary local outgassing and a consequent increase in pressure near the gauge. There are clear steps down in pressure where the temperature was held constant for a long period signifying successful baking. The gradient of the red points gives a ΔH_A of around 40 kJ mol^{-1} .

the surrounding system so as not to localise dirt on the sources. Once the recorded pressure had not changed for a couple of days we began to ease the system back to room temperature. The resultant plot of this final bake-out is shown in Fig. 2.10. It is not quite as dramatic as perhaps it might be given that during our two aborted baking attempts we had already pumped a great deal of material out of the system, so we appear to return to a similar pressure as we started, but this is around the limit of the pressure gauge so cannot fully be trusted. The suggestion is though that at the gauge we have reached a pressure in the 10^{-11} mbar range. The relation of equation (2.5) is observed to hold reasonably well while decreasing the temperature as expected, with the measured gradient giving a ΔH_A of roughly 40 kJ mol^{-1} which is fairly typical of such systems [45]. Once a fully operational UHV chamber with atom sources has been created, we need means of adding light and magnetic fields in order to manipulate the atoms, as will be discussed in the following sections.

2.2 Magnetic coils

It is probably already clear that magnetic fields play a vital role in the cooling and trapping of atoms. They are required for laser cooling in the MOT, transport to the science cell and for trapping the atoms during the first stages of evaporative cooling, as well as for accessing Feshbach resonances, providing guide fields for imaging and optical pumping of atoms and simply to cancel out the effects of the background magnetic field in the lab. The simplest way to produce magnetic fields in a controllable manner is to pass a current through a wire, and generally the easiest shape to work with (both for designing and winding magnetic coils) is a circle. The vast majority of coils in the system consist of sets of circular loops, the key aspects of which are discussed below.

Magnetic field from a circular loop

The magnetic field \mathbf{B} produced at a position \mathbf{r} by passing a steady current I through an element of wire $d\mathbf{l}$ may be calculated from the Biot-Savart law

$$d\mathbf{B} = \frac{\mu_0 I}{4\pi r^3} d\mathbf{l} \times \mathbf{r} , \quad (2.6)$$

where μ_0 is the permeability of free space. For a circular loop of radius a we may write

$$d\mathbf{l} = (-a \sin \theta, a \cos \theta, 0) d\theta \quad (2.7)$$

$$\mathbf{r} = (a \cos \theta - x, a \sin \theta - y, z) \quad (2.8)$$

where we have defined the origin $\mathbf{r} = 0$ to be at the centre of the loop and the axis of the circle is in the z direction. \mathbf{B} is then found by integrating (2.6) around the full loop. This may be found analytically in the on-axis case where, by symmetry, the x and y components of the magnetic field B_x and B_y are zero, while the z component is found to be

$$B_z = \frac{\mu_0 I}{2} \frac{a^2}{(a^2 + z^2)^{3/2}} . \quad (2.9)$$

We can also infer some information about B_y and B_z close to the axis from Gauss' law for magnetism, $\nabla \cdot \mathbf{B} = 0$, and the fact that by symmetry dB_x/dx must be equal to dB_y/dy on axis, such that

$$2 \frac{dB_x}{dx} = 2 \frac{dB_y}{dy} = - \frac{dB_z}{dz} = B' . \quad (2.10)$$

Bias fields

The guide, compensation and Feshbach fields mentioned above should ideally be uniform over the experimental region. A single coil cannot easily provide this, but the field gradients (and in fact all odd terms in a Taylor expansion for the field) at the centre-point of two identical co-axial coils with the same current circulation will cancel.

This creates a field that is uniform below second order, which is also true along x and y via (2.10). The bias field produced at the centre of a coil pair separated by a distance d is simply twice (2.9) with $z = d/2$. It is convenient to work with B in gauss, I in amps and distances in centimetres, giving

$$B_{\text{bias}} = \frac{4\pi I}{10} \frac{a^2}{(a^2 + d^2/4)^{3/2}} \text{ G} . \quad (2.11)$$

The curvature of the field can be found from differentiating (2.9) twice with respect to z and setting $z = d/2$. It may be shown that it disappears for $d = a$ such that the field is uniform to below fourth order expansion terms. This is known as the Helmholtz configuration, which we have tried to use for our bias coils where possible.

Quadrupole fields

Reversing the direction of current flow in one of the pair of coils will instead produce a quadrupole field as required by the MOT and for simple magnetic trapping, where $B = 0$ at the centre (as now all the even expansion terms cancel). B' from (2.10) and (2.9) evaluated at $z = d/2$ is then the total gradient along the ‘weak’ x and y axes (twice the x and y gradient of a single coil), while $-2B'$ gives the total gradient along the ‘strong’ z axis. When again converted to gauss with distances in centimetres this gives

$$B' = \frac{3\pi I}{10} \frac{a^2 d}{(a^2 + d^2/4)^{5/2}} \text{ G cm}^{-1} . \quad (2.12)$$

Analogously to the Helmholtz configuration above, for a separation $d = \sqrt{3}a$ the third order term vanishes at the coil-pair centre, giving the most linear field gradient. This arrangement is known as the anti-Helmholtz configuration. Also useful for coil design is the fact that B' is greatest for $d = a$ for fixed a and for $a = d/\sqrt{6}$ for fixed d .

In practice we use coils made up of several turns each, so one may sum the above expressions over various a and d to estimate the resultant fields. Further considerations of coil design include the coil inductance, power dissipation (which manifests itself as thermal energy) and the consequent water-cooling required.

2.2.1 Coils around the MOT cell

The properties of the coils located around the MOT cell are summarised in Table 2.2. The requirements are fairly minimal as only laser cooling is carried out at this end of the system. Three pairs of compensation coils (providing tuneable bias fields along three axes) are used to cancel out the background magnetic field of the order of 1 G, while the transport coils provide the small field gradient required for magneto-optical trapping. The compensation coils are essential to minimise the magnetic field at the atoms during the optical molasses stage (see §2.3.1). One of the horizontal pairs is also used to provide a guide field for optical pumping (§2.4.3).

2.3 Laser cooling

Coil	Wire (mm)	Turns	Dimensions (cm)	Separation (cm)	B_0/I (G/A)	B'/I (G/(cm A))
Transport	∅4.0	10×4	∅5.1	9.8		0.40
Comp. (Vert.)	∅1.0	3×3	∅9.0	4.4	1.73	
Comp. (Horiz.)	∅1.0	5×8	∅7.2	20.0	0.55	

Table 2.2: Parameters of the coils used around the MOT cell. The number of turns are given in the form $N_{\text{radial}} \times N_{\text{axial}}$, while the given dimensions and separations refer to clear inner distances. The compensation fields have not been verified by direct measurement.

A current of up to 200 A can be supplied to the transport coils by a *Delta Elektronika* power supply for magnetic trapping, producing a maximum field gradient along the weak radial direction of $B' = 80 \text{ G cm}^{-1}$. Such high currents naturally generate a lot of heat, so the coils are made from hollow tubing to allow water cooling. Fail-safes are included to stop the coils over-heating; a flow-meter detects whether water is flowing through the water-cooling tubes and thermistors detect if the coils are over-heating. The flow-meter is connected to the power-supply itself, setting the output to zero if the water flow drops below a certain level. Voltage comparators are used in conjunction with the thermistors to provide a signal to break the electrical connection of the coils via an insulated gate bipolar transistor (IGBT) if the coils reach above $\sim 60^\circ\text{C}$. The IGBT is also used for switching of the coils during the experimental sequence. Similar safety precautions are taken with all the other coils in which high currents can flow.

2.2.2 Coils around the science cell

The properties of the coils located around the science cell are summarised in Table 2.3. The transport coils also operate in this region, while another three pairs of compensation coils are used to cancel the local background fields at this end of the system. In addition, two pairs of bias coils are wound on top of the (rectangular) horizontal compensation coils. The coil pair for providing large bias fields to utilise the Feshbach resonance at 402.5 G is wound from the same hollow copper tubing as the transport coils to allow water cooling. The coils produce a bias field of 528 G when a current of 200 A is flowing through them. Finally, a small Ioffe coil is used in the formation of the QUIC trap described in §2.6. It is mounted on a water-cooled copper block. The layout of the coils around the science cell is shown in Fig. 2.11.

2.3 Laser cooling

The cooling of atoms by lasers relies on atoms absorbing and re-emitting photons. A simple illustration of this concept is given in Fig. 2.12 for an atom travelling towards a laser source. Each time the atom absorbs a photon its momentum is reduced. The subsequent emission of photons occurs in a regular dipole pattern such that these

Coil	Wire (mm)	Turns	Dimensions (cm)	Separation (cm)	B_0/I (G/A)	B'/I (G/(cm A))
Transport	$\varnothing 4.0$	10×4	$\varnothing 5.1$	9.8		0.40
Feshbach	$\varnothing 4.0$	7×2	$\varnothing 7.0$	3.6	2.64	
Bias (Axial)	$\varnothing 1.0$	10	7.4×14.9	5.6	1.43	
Bias (Trans.)	$\varnothing 1.0$	10	6.2×15.2	5.6	1.46	
Comp. (Vert.)	$\varnothing 1.0$	7×3	$\varnothing 19.0$	7.0	2.28	
Comp. (Axial)	$\varnothing 0.7$	8	7.4×14.9	5.6	1.14	
Comp. (Trans.)	$\varnothing 0.7$	8	6.2×15.2	5.6	1.14	
Ioffe	$\varnothing 1.0$	6×3	$\varnothing 1.0$	1.7*		

Table 2.3: Parameters of the coils used around the science cell. The number of turns are either given as a total or in the form $N_{\text{radial}} \times N_{\text{axial}}$. Dimensions (as a diameter for circular coils and as side lengths for rectangular coils) and separations refer to clear inner distances. *‘Separation’ here refers to the distance from the face of the Ioffe coil to the centre of the quadrupole field.




Figure 2.11: Layout of the magnetic coils around the science cell. In grey are the transport coils, black the Feshbach coils and in white the Ioffe coil. The copper coloured coils are the bias and compensation coils (the rectangular coils each contain two sets of coils, wound around the same groove of the coil mount).

2.3 Laser cooling

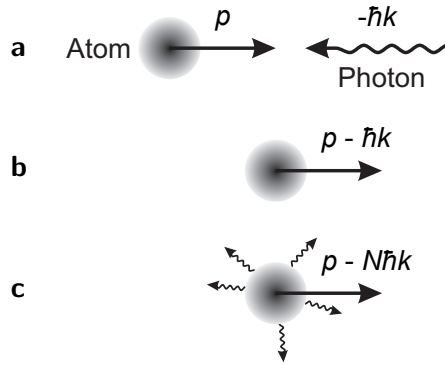


Figure 2.12: Simplistic view of laser cooling. **a** A photon with momentum $-\hbar k$ is incident upon an atom travelling with momentum p . **b** If the photon is absorbed the momentum of the atom is reduced. **c** After absorbing N such photons the momentum of the atom is reduced to $p - N\hbar k$ as the momentum kicks from photon re-emission cancel.




Figure 2.13: The Doppler effect in the case of counter-propagating beams. **a** The frequency of the light in the lab frame is ω . **b** In the rest frame of the atom the frequencies are shifted by $\pm kv$.

momentum kicks tend to cancel one another out. Over many cycles the atom is slowed down, which may be used to reduce the spread of atom velocities for a group of atoms, cooling them down. This is the effect of the scattering force from near-resonant light introduced in §1.4.3.

The key point is that to cool an atomic cloud each atom must mainly absorb photons with momenta in opposition to its direction of motion, but the velocities of atoms in a gas are randomly distributed. A method of coping with this problem was proposed by Hänsch and Schawlow in 1975 [46], now known as optical molasses.

2.3.1 Optical Molasses

Initially considering the one-dimensional case, a counter-propagating beam must be introduced to allow the possibility of atoms receiving momentum kicks in the opposite direction. Some mechanism is then required by which atoms will preferentially absorb photons from the beam that opposes its motion. This is achieved by the neat trick of detuning the beam frequency slightly to the red of resonance. The Doppler shift then acts to bring the deceleration beam back towards resonance while shifting the acceleration beam further off resonance, providing a net damping force. Atoms close to zero velocity are equally off-resonant with both beams and therefore feel no further force. Here we are once more considering the ideal two-level picture. We may describe

this in terms of the scattering force F_{sc} of equation (1.67) by replacing δ with $\delta \pm kv = \omega - \omega_0 \pm kv$ to obtain

$$F_{\text{mol}} = F_{\text{sc}}(\omega - \omega_0 - kv) - F_{\text{sc}}(\omega - \omega_0 + kv) \quad (2.13)$$

$$\simeq -2 \frac{\partial F_{\text{sc}}}{\partial \omega} kv \quad (2.14)$$

$$\simeq -\alpha v \quad (2.15)$$

for $kv \lesssim \Gamma$, where the damping coefficient α is given by

$$\alpha = -4\hbar k^2 \frac{(I/I_s)(2\delta/\Gamma)}{[1 + 2I/I_s + (2\delta/\Gamma)^2]^2}. \quad (2.16)$$

The extra factor of 2 in intensity in the denominator of α arises because saturation must be considered in terms of both beams. This is valid for $v \lesssim 4 \text{ m s}^{-1}$ for potassium and rubidium. α is positive for red-detuned beams so the resultant force does indeed always oppose the motion of the atom. This is equivalent to motion in a viscous fluid (such as molasses).

2.3.2 Doppler limit and sub-Doppler cooling

This picture may be extended by the addition of two further sets of counter-propagating beams along orthogonal axes to provide strong damping in all three dimensions. This might be expected to lead directly to temperatures approaching absolute zero. However, the idea that the scattering via spontaneous emission in Fig. 2.12 **c** cancels to zero is not true for individual atoms, and instead they are heated by a random walk in velocity space where each step is a single recoil velocity $v_r = \hbar k/m$. The atoms also undergo a random walk via absorption (along the paths of the beams) once they are at low velocity. The balance of cooling and heating leads to a limit on the temperature that it is possible to reach in this picture of optical molasses, which is minimised for a frequency detuning $\delta = \Gamma/2$. This is known as the Doppler limit [47] and is given by

$$k_B T_D = \frac{\hbar \Gamma}{2}. \quad (2.17)$$

T_D is $144 \mu\text{K}$ for our atomic species.

This technique was first successfully implemented in 1985 by Chu *et al.* [48], but curiously it was observed not long afterwards that for larger frequency detunings the temperatures reached were well below the Doppler limit [49]. This remarkable phenomenon was explained by Dalibard and Cohen-Tannoudji [50] and relies on the multi-level nature of the atoms. The polarisation of the electric field produced by counter-propagating beams with opposite circular polarisations varies with period $\lambda/2$. The Zeeman sub-levels of the ground state are perturbed by the electric field by differing amounts that depend on the polarisation of the field, resulting in energy levels that oscillate in space

2.3 Laser cooling




Figure 2.14: Simplified depiction of the mechanism behind sub-Doppler Sisyphus cooling. **a** Length scale in terms of the wavelength λ of the optical molasses beams. **b** The variation of polarisation through space produced by two counter-propagating beams. **c** Relevant energy levels, including the splitting of the degenerate ground state which is polarisation dependent. **d** The energy E_a of an atom as it travels from left to right. The arrows in **c** show how an atom in one of the ground states loses kinetic energy by travelling to a position of higher potential energy, at which point it is pumped to a ground state at lower energy.

with a phase offset between them (see Fig. 2.14). As an atom moves in space it loses kinetic energy as it travels up the potential hill. Near the peak the circular polarisation of the light near this point pumps the atom into a sublevel of lower energy. Due to the short excited-state lifetime the position of the atom is essentially constant during this process. The energy loss is accounted for by the difference in energy between the absorbed and emitted photon. Now in a different substate, the process is repeated with sublevels and circular polarisation reversed at a position $\lambda/4$ further along. This process has become known by the term ‘Sisyphus cooling’ due to the resemblance to the Greek myth of the king punished to forever push a boulder uphill only to watch it roll back down again.

The procedure is repeated until the energy of each atom is comparable to the light shift amplitude of the energy levels¹, resulting in temperatures of tens of μK , around

¹A more rigorous treatment given in [51] suggests the limit is around $40 E_r$, when the amplitude of the lower-state light shift is $\approx 100 E_r$, where $E_r = mv_r^2/2 = \hbar^2 k^2/2m$ is the recoil energy (it is not possible to reach arbitrarily low temperatures by simply reducing the laser intensity).

an order of magnitude lower than the Doppler temperature. The success of this process was a major step towards cooling atoms to quantum degeneracy. The only downside is that this technique alone does not actually trap atoms due to the lack of a restoring force. This was rectified by the development of the magneto-optical trap - a concept suggested by Dalibard [52] and first implemented at Bell Laboratories, USA [53].

2.3.3 Magneto-optical trap

A weak magnetic field may be used to split the atomic energy states via the Zeeman effect of equation (1.36) to create an imbalance in the radiation force from counter-propagating beams that is now dependent on position. The beams must have opposite circular polarisations so that they interact with different electronic transitions such that an increase in magnetic field strength takes one beam closer to resonance, the other further away. If a quadrupole field is used the resultant linear energy level splitting allows us to use a very similar treatment to that of Doppler cooling described above. This principle is illustrated in Fig. 2.15. Replacing δ in equation (1.67) now by $\delta \pm (kv + \beta z) = \omega - \omega_0 \pm (kv + \beta z)$, where $\beta\hbar z = (g'm' - gm)\mu_B B'z$ is the energy shift of a transition resonance, directly leads to the one-dimensional scenario where

$$F_{\text{MOT}} \simeq -\alpha v - \frac{\alpha\beta}{k}z, \quad (2.18)$$

valid for $\beta z \lesssim \Gamma$. This can once again be extended to three dimensions by careful choice of the beam polarisations of the three beam pairs as shown in Fig. 2.16. The MOT is a very robust trap and can capture atoms with significantly higher velocity than are slowed by the optical molasses technique, allowing us to load the trap directly from a room temperature vapour for potassium and rubidium.

2.3.4 Laser cooling for real atoms

This outline of the mechanisms involved in laser cooling condenses many years of dedicated research into a few pages, and it is worth noting the huge significance of these major advances. This was recognised by Steven Chu, Bill Phillips and Claude Cohen-Tannoudji being awarded the 1997 Nobel prize “for development of methods to cool and trap atoms with laser light”. A MOT in some form is used in nearly every cold atoms experiment in the world.

The above arguments are most accurate for a two-level system, but in reality we have all of the hyperfine levels to deal with. However, atoms excited to $|F'=3\rangle$ decay only to $|F=2\rangle$, so using light near the $|F=2\rangle \rightarrow |F'=3\rangle$ transition causes the atoms to cycle between these two states. If an atom happens to fall out of this cycling (or ‘cooling’) transition we also require light near the $|F=1\rangle \rightarrow |F'=2\rangle$ ‘repumping’ transition to eject the atom from $|F=1\rangle$.

^{87}Rb nicely adheres to this picture with its large hyperfine splittings, where the

2.3 Laser cooling




Figure 2.15: Simple scheme for a magneto-optical trap for a singlet ground state and triplet excited state. As the atoms move away from the centre the Zeeman shift tunes them towards resonance with the restoring beam and out of resonance with the opposite beam producing a net restoring force. σ^\pm polarised light works on the $\Delta m = \pm 1$ transition.




Figure 2.16: Extension of the MOT principle to three dimensions. The reversal of the field direction from the radial direction to the axial direction (with respect to the quadrupole coils) requires that the beam polarisations also be reversed. The beams are red-detuned from the cycling transition in order to both cool and trap the atoms. In a real system some repumping light must also be present to avoid pumping the atoms into a dark state.

force on the atoms is provided by the cooling light, while just a little repumping light is needed to compensate for the low decay rate to the lower ground state. As such the repumping light can just be sent in as a single separate low-intensity beam while the cooling light must have the signature six-beam geometry.

Conversely, the small splitting of the ^{39}K F states means that if we detune the cooling beam as required for laser cooling we actually bring it closer to resonance with the non-cycling transitions. Pumping into the $|F = 1\rangle$ state occurs at a rapid rate. The repumping light must therefore be about as intense as the cooling light, be detuned from the repumping transition and must be counter-propagating as it now provides a considerable amount of the force. ‘Cooling’ and ‘repumping’ in this scenario start to lose meaning, but will continue to be used for descriptive convenience. In fact it has been neatly demonstrated that a ^{39}K MOT may be formed by using only cooling light on the horizontal axes and only repumping light on the vertical axis [54].

A more detailed discussion of the properties of the MOT setup as used in our experiments will be given in Chapter 3. Now we know what laser frequencies are required for laser cooling we just need a way to reliably produce them.

2.4 Laser system

It is advantageous that today the laser frequencies required for the laser cooling of rubidium and potassium are accessible by diode lasers which are simple to maintain and operate. The basic principle relies on a high-band-width gain medium (laser diode) in combination with an external diffraction grating. The combination of the medium gain, the diode internal cavity, the external cavity between the grating and the back of the diode and the grating profile itself produces high gain at a single frequency with a line-width smaller than the natural line-widths of alkali atoms (Fig. 2.17). The particular lasing mode is selected by tuning the feedback from the grating via its angle, while we may scan the frequency continuously by modifying the length of the external cavity via a piezoelectric transducer (piezo). There is a limit to the tuning range as the laser will eventually ‘mode-hop’ to the adjacent external cavity mode. This mode-hop-free tuning range is typically up to around 10 GHz (roughly restricted by the mode spacing of the external cavity). The temperature and laser current are used to tune where the mode hops occur, affecting the internal cavity modes. Once the laser is tuned near its designated frequency the piezo voltage is modulated to scan the laser frequency (used for locking the laser frequency as discussed below).

Standard diode lasers at our required frequencies produce up to ~ 100 mW of power. Given the losses inherent in AOMs, fibre-coupling, reflection off mirrors etc. this is below what we require to have enough cooling light at the MOT cell, so tapered amplifiers (TAs) are used to boost the power. These contain an amplifying chip with a tapered gain region. The internal modes are controlled by the coupling of a low-power ‘seed’

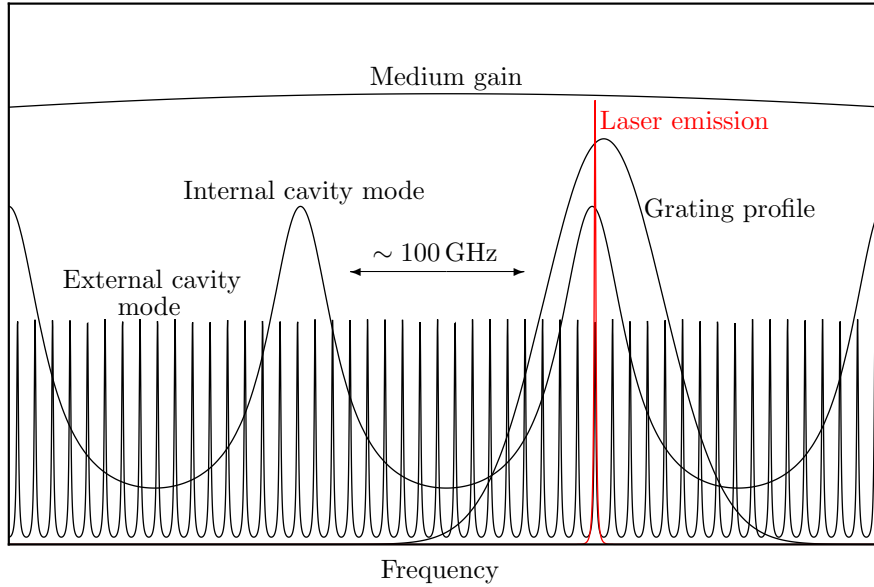


Figure 2.17: Guide to the frequency scales of the gain profiles of the features that make up a diode laser. The grating profile can be shifted for coarse tuning of the emitted laser frequency, while the external and internal modes can be modified for fine tuning as discussed in the text.

laser (i.e. ~ 10 mW beam from a diode laser) so the fine line-width is maintained while allowing amplification to over 1 W of laser power.

The lasers in our system were bought with all the necessary current and temperature regulators and piezo controls, plus the electronics for frequency locking in a feedback loop (§2.4.2). Optical isolators are fitted in every laser which, via their main components of a Faraday rotator and a polariser, vastly suppress unwanted feedback into the optical cavity.

2.4.1 Saturated absorption spectroscopy

The laser frequencies required for laser cooling, optical pumping and imaging of atoms must be accurate to the MHz level, and are generally in the vicinity of the hyperfine transitions of Fig. 1.5. A natural course of action therefore is to somehow lock the laser frequencies directly to the hyperfine features of the relevant atomic spectra. We can observe the absorption of light by an atomic vapour, but in order to discern the hyperfine features the thermal broadening due to the Doppler effect¹ must be suppressed. Typically this is achieved via saturated absorption spectroscopy.

A typical saturated absorption setup is shown in Fig. 2.18. The transmission of a weak ‘probe’ beam is recorded by a photodiode and produces the expected broad absorption signal as the frequency of the laser is scanned. When a strong counter-propagating ‘pump’ beam is introduced at the same frequency there is generally no effect, as atoms travelling at a general velocity v (along the beam axis) will not be on

¹The thermal broadening produces a feature ~ 1 GHz in width, making it impossible to observe the underlying peaks of ~ 6 MHz line-width.




Figure 2.18: Example of a saturated absorption spectroscopy setup. In this case a half-wave ($\lambda/2$) plate is used to adjust the relative strength of the transmitted and reflected components from a polarising beam-splitter (PBS) cube, forming the probe and pump beams respectively. The transmission of the probe beam through the vapour cell is recorded by a photodiode.

resonance with both beams simultaneously. However, atoms travelling at $v = 0$ will be resonant with both beams when the laser frequency matches an atomic transition. In this case the stronger pump beam promotes many of these atoms to the excited state and indeed may cause atoms to be pumped into a ‘dark’ ground state. The consequent reduction in the absorption of the probe beam is observed as a sharp peak in transmission to the photodiode.

In the case of ^{87}Rb both the ground and excited hyperfine states are split well enough to be easily resolved. Considering just one of the ground states, this mechanism would lead to the observation of three peaks for, say, the $|F = 2\rangle \rightarrow |F' = 1, 2, 3\rangle$ transitions. However, when the laser frequency is exactly half way between two transitions ω_1 and ω_2 , atoms travelling at $v = \pm(\omega_1 - \omega_2)/2k$ will again be resonant with both beams simultaneously, so three strong ‘cross-over’ peaks are seen bisecting the three direct transition peaks. This is shown in Fig 2.19 **a**, where the black line shows the saturated absorption spectrum while the red line shows the Doppler broadened spectrum observed when the pump beam is blocked. The linewidth of the six observed peaks is around 16 MHz, not far off the natural lifetime broadening of ~ 6 MHz. The mechanisms producing the spectrum are shown schematically in Fig. 2.20.

The small hyperfine splitting in ^{39}K leads to a different picture. The lower state splitting is small enough to be spanned by the Doppler broadening, while the upper state splitting is so small that even with saturation spectroscopy it is difficult to resolve the structure. The six peaks of the $|F = 2\rangle \rightarrow |F' = 1, 2, 3\rangle$ transition that are so clear in ^{87}Rb are condensed into a single peak in ^{39}K , peak i of Fig. 2.19 **b**. Another merged peak (iv) is seen for the $|F = 1\rangle \rightarrow |F' = 0, 1, 2\rangle$ transitions, while we also see a deep dip in transmission between the two. This results from a similar argument to the cross-over peaks, where now we consider the case where we have two ground states and a single upper state. In this scenario, when the laser frequency bisects the two transition frequencies there is again a velocity class of atoms for which both the probe and the pump beams are resonant. However, this time the pump beam acts to pump atoms back into the complementary ground state (much like the repumping beam in the




Figure 2.19: Saturated absorption spectra for the D₂ transitions of ⁸⁷Rb and ³⁹K. The saturated spectra are shown in black, Doppler broadened versions in red and the difference between the two in blue. **a** Transitions from the $F = 2$ ground state of ⁸⁷Rb. Peaks i, iii and vi are direct transitions to the $F' = 1, 2$ and 3 excited states respectively, the rest being the corresponding cross-over peaks. The x -scale shows the frequency difference in MHz from the $F = 2$ to unperturbed upper state transition. **b** Full spectrum for the D₂ line of ³⁹K. Peaks i and iv are the conglomerated transitions from the $F = 2$ and $F = 1$ ground states respectively. The crossover dips ii and iii are to the $F' = 1$ and $F' = 2$ upper states respectively. To accentuate the features the spectrum was taken using a vapour cell that was heated to 30 °C, but we have not found it necessary to heat the cell when locking our laser. The small residual narrow features in the Doppler broadened profile are likely due to reflection of the probe beam in the system. The x -scale shows the frequency difference in MHz from the unperturbed D₂ transition.

MOT discussed later), increasing the absorption of the probe beam. We can actually just make out two inverted peaks in Fig. 2.19 **b**, which correspond to the $F' = 1$ and $F' = 2$ upper states, being the only two states accessible to both lower states by the standard dipole transition selection rules.

2.4.2 Frequency locking of lasers

The photodiode in the saturation absorption spectroscopy setup provides us with a signal that is either at a local maximum or minimum when we are resonant with one of the hyperfine transitions or one of the cross-over transitions. It is then a relatively straight-forward task to implement a feedback mechanism to lock the laser to one of these frequencies. We use cross-over resonances as they provide the strongest signal.

The principle relies on applying some low amplitude jitter to the laser frequency and comparing the photodiode signal either side of the central frequency. For example, if we are locking the laser to a peak the signal just either side of resonance will be lower than the central signal. When multiplied by the wave-form of the jitter this will




Figure 2.20: Mechanisms producing the features of saturated absorption spectroscopy in the case of a single ground state and two excited states, as seen for ^{87}Rb . **a** For general atom velocity v the pump has no effect on the probe beam and the Doppler broadened absorption is seen. **b** For $v = 0$ both beams are resonant with the same transition when the laser frequency ω is equal to one of the transition frequencies ω_1 and ω_2 . Atoms absorb light from the pump beam that would otherwise absorb light from the probe beam and the transmission of the probe beam is consequently increased. **c** For $v = \pm(\omega_1 - \omega_2)/2k$ each beam is resonant with a different transition when $\omega = (\omega_1 + \omega_2)/2$. As they share the same ground state a peak in probe beam transmission is again observed.




Figure 2.21: Mechanisms producing the features of saturated absorption spectroscopy in the case of a single excited state and two ground states, as seen for ^{39}K . For the cases **a** and **b** the situation is much the same as in Fig. 2.20 **a** and **b**. **c** For $v = \pm(\omega_1 - \omega_2)/2k$ each beam is resonant with a different transition when $\omega = (\omega_1 + \omega_2)/2$. In this scenario the pump beam replenishes the ground state of the probe beam transition, causing an increase in absorption such that a dip in the transmission is observed.

2.4 Laser system

average to zero over a period. If, however, we are slightly above resonance the signal will be lower for positive jitter than for negative jitter, providing a negative signal from the periodic average once multiplied again by the jitter wave-form. This gives a signal proportional to the gradient of the spectroscopy signal. This signal is fed back to the laser and the external cavity is modified to return the frequency to resonance.

In our system we apply the high frequency jitter to the laser itself by oscillating the cavity length. As long as the amplitude is below the required line-width this doesn't cause a problem. A lock-in regulator extracts the slope of the signal and we can adjust the proportional/integral/derivative (PID) components to optimise the stability of the lock.

2.4.3 Required frequencies

The need for cooling and repumping light for the MOT in the vicinity of the $|F=2\rangle \rightarrow |F'=3\rangle$ and $|F=1\rangle \rightarrow |F'=2\rangle$ transitions, respectively, has already been discussed in §2.3.4. Elsewhere in the system we need resonant cycling light for imaging the atoms, and pumping light resonant with the $|F=2\rangle \rightarrow |F'=2\rangle$ transition for when we want to transfer all the atoms into the $|F=2, m_F=2\rangle$ state for magnetic trapping. 'Resonant' here includes a small Zeeman shift due to the presence of a uniform magnetic guide field operated to provide a quantisation axis during imaging and pumping. The pumping is achieved by using σ^+ polarised light so that once the atoms have entered the $|F=2, m_F=2\rangle$ state they will remain there.

Acousto-optical modulators

Having locked our lasers to the cross-over transitions of saturated absorption spectroscopy, we must shift the emitted light by tens and hundreds of MHz to reach the desired final frequencies, which may be done using acousto-optical modulators (AOMs). These consist of a crystal in which a sound wave is generated by a piezoelectric transducer. Incident laser light is Bragg diffracted by an angle θ given by $\sin \theta = m\lambda_l/2\lambda_s$, where λ_l and λ_s are the wavelengths of light and sound respectively and m is the diffraction order. The frequency of the sound ν_s also Doppler shifts the diffracted beam, altering the frequency of the light to $\nu'_l = \nu_l + m\nu_s$. The diffraction efficiency into the first order can be varied from almost 0 to upwards of 90% by varying the power of the acoustic waves. This makes AOMs additionally useful as fast switches; the diffracted light can be turned off in the time it takes a sound wave to travel the width of the laser beam, typically under 100 ns.

As the angle of beam deflection depends on the frequency shift of the AOM, we need to use a double-pass configuration in order to dynamically change the laser frequency without losing beam alignment. This involves placing an AOM at the focal point between two collimating lenses and reflecting the diffracted beam back through the AOM as shown in Fig. 2.22. The beam is diffracted a second time, acquiring a total




Figure 2.22: AOM in double-pass configuration. Identical lenses a focal length either side of the AOM ensure that as ν_s varies the return beam retains its alignment. Retro-reflection through a quarter-wave plate rotates the polarisation so that the frequency-shifted beam may be separated from the incident beam path by a PBS cube. Irises are used to block unwanted diffraction orders

frequency shift of $2m\nu_s$ and returning along the original beam path. This beam can be separated from the original beam at a PBS cube by inserting a quarter-wave plate before the retro-reflection mirror, so that the polarisation of the two beams are orthogonal. In general the diffraction efficiency is peaked at a particular frequency that depends on the AOM in question, such that shifting the frequency alters the power of the resultant beam. However, using the double-pass configuration as part of the laser-lock setup allows frequency control with no power loss from the main beam, so this method is used for most lasers in our system.

The AOMs in our system are from *Crystal Technology* and *AA Opto-electronic* and are controlled by drivers designed and built in-house. The AOM drivers centre on a voltage-controlled oscillator (VCO) which outputs a voltage-dependent frequency, the power being adjusted separately by means of a voltage-variable attenuator. The design allows frequency and power control both externally from a supplied voltage and internally via control of chassis-mounted potentiometers. Fast-switching with a TTL signal is also available. The radio-frequency components in the drivers are from *Mini-circuits*.

Table 2.4 summarises the laser lock points (the saturated absorption features used for frequency stabilisation) and AOM frequency shifts used to get to the final frequencies.

2.4.4 Laser table summary

The entire experiment is carried out on two separate optical tables, designated the laser and vacuum tables. All of the lasers and AOMs are laid out on the laser table to be isolated from the vibrations caused by the translation stage. Only three actual master lasers are used, one for potassium (where the ground-state hyperfine splitting may be bridged using AOMs) and two for rubidium. Specifically, we use a *Toptica photonics* DL Pro laser for potassium, while a DL 100 is used for the rubidium repumping light and a TA 100 (actually a combination diode laser and TA in one box) for the rest. The master lasers are based on the Littrow configuration.

Both the cooling and repumping light for potassium need to end up at high in-

2.4 Laser system

	Function	Master	Locking line	AOM shift (MHz)	Resultant light
³⁹ K	Cooling	DL Pro	$ F=1-2\rangle \rightarrow F'=1\rangle$	$2 \times (70.5 \rightarrow 95.5)$ $-2 \times (185 \rightarrow 215)$	$ F=2\rangle \rightarrow F'=3\rangle$ $-89 \rightarrow +21$ MHz
	Repumping	DL Pro	$ F=1-2\rangle \rightarrow F'=1\rangle$	$2 \times (70.5 \rightarrow 95.5)$ $+40$	$ F=1\rangle \rightarrow F'=2\rangle$ $-50 \rightarrow 0$ MHz
	Imaging	DL Pro	$ F=1-2\rangle \rightarrow F'=1\rangle$	$2 \times 95 - 2 \times 195$	$ F=2\rangle \rightarrow F'=3\rangle$
	Pumping	DL Pro	$ F=1-2\rangle \rightarrow F'=1\rangle$	$2 \times 95 - 2 \times 206$	$ F=2\rangle \rightarrow F'=2\rangle$
	Cooling	TA 100	$ F=2\rangle \rightarrow F'=2-3\rangle$	$2 \times (83 \rightarrow 108)$ -83	$ F=2\rangle \rightarrow F'=3\rangle$ $-50 \rightarrow 0$ MHz
⁸⁷ Rb	Repumping	DL 100	$ F=1\rangle \rightarrow F'=1-2\rangle$	$+78$	$ F=1\rangle \rightarrow F'=2\rangle$
	Imaging	TA 100	$ F=2\rangle \rightarrow F'=2-3\rangle$	$2 \times 100 - 67$	$ F=2\rangle \rightarrow F'=3\rangle$
	Pumping	TA 100	$ F=2\rangle \rightarrow F'=2-3\rangle$	$2 \times 75 - 83 - 200$	$ F=2\rangle \rightarrow F'=2\rangle$

Table 2.4: Summary of the laser-lock transitions and AOM shifts used to provide the laser frequencies used for laser cooling, imaging and optical pumping. The form “ $|F=1-2\rangle \rightarrow |F'=1\rangle$ ” denotes locking to a ground-state crossover feature. Use of an AOM in double-pass configuration results in an overall frequency shift of $2 \times \nu_s$ (all AOM shifts are first order). Note that where an AOM is used in the spectroscopy setup the light that exits the laser is shifted in the opposite direction (i.e. a negative frequency shift in the locking setup is counted as a positive frequency shift).

tensity with the same polarisations, so after the master laser light has been split and frequency shifted it is recombined on a PBS cube before being rotated and split in two at another PBS cube such that we end up with two polarised beams each containing both frequencies. One beam is used for imaging and pumping (requiring relatively little light) while the other (~ 10 mW) beam is used to seed a *Sacher Lasertechnik* TEC-400 tapered amplifier. Dual seeding means we can’t directly measure the relative powers at the two frequencies, so a Fabry-Pérot interferometer is used to provide a rough gauge. This is important as the amplification depends on how well coupled each beam is, and small imbalances in the coupling leads to larger imbalances in relative output power. A total power of up to ~ 1.5 W leaves the TA.

The rubidium repumping light is produced at a fixed frequency by the DL 100 with a single-pass AOM in the lock setup, the 80 mW it produces proving to be plenty. The TA 100 has two output apertures, one for a pre-amplification 10 mW beam, the other for the amplified beam. The weak beam is used for locking and imaging, while the amplified beam of around 600 mW is split and shifted for cooling and pumping light.

The light is passed to the vacuum table by means of polarisation-maintaining optical fibres. In addition to the saturated-absorption spectroscopy components and AOMs there are several other key components which are briefly discussed below.

Anamorphic Prism pairs

Typically diode lasers produce elliptical beams ($\sim 3 \text{ mm} \times 1 \text{ mm}$ for the *Toptica* master lasers) which are unsuitable for efficient fibre coupling. Anamorphic prism pairs are used at the outputs of the diode lasers which are used to magnify or shrink the beam in one dimension resulting in a near-circular profile. The exact magnification can be adjusted by the prism angles so can conveniently be adjusted to suit each particular laser.

Mechanical shutters

Although fast switching of light is available using AOMs there is always a small fraction of light diffracted away from the zeroth order. Sometimes it is necessary that absolutely no near-resonant light reaches the atoms and mechanical shutters must be used. These act more slowly than AOMs (depending on the shutter speed and once again on the beam width) and must be triggered in advance so that the shutter reaches the beam at the desired time if timing is critical. However, they have the advantage of blocking all light. The efficiency of an AOM is unstable when turned on from ‘cold’, so it is best to keep them on for the majority of the time. A mechanical shutter can be used to block the light while the AOM is ‘kept warm’, allowing for higher power stability when light is needed once more. The shutters are controlled via TTL signals and mounted on rubber bases to suppress mechanical vibrations that can affect the laser frequency lock.

Dichroic wave plate

The two-species MOT requires beams containing light for both species with the same polarisation. To achieve this we combine orthogonally polarised beams at a PBS cube before using a dichroic wave-plate to rotate the polarisation of just the potassium light. This wave plate, manufactured by *LENS-Optics*, is precisely machined such that it acts as a $27 + 1/2$ wave plate for light at 766.7 nm and a 27 wave plate for 780.2 nm light. Subsequently only zero-order wave plates are used to keep the rubidium and potassium polarisation matched.

The entire laser table optical configuration is displayed in Fig. 2.23, tying in all the above component pieces discussed above. All of the optics are anti-reflection coated to minimise power loss through the system. For clarity, the central AOM frequencies are quoted because the operating frequencies change throughout the experimental sequence (whether the shift is positive or negative is indicated).

2.5 Magnetic transport

2.4.5 Fibre-port cluster

In order to split the cooling and repumping light into the six beams required for the MOT we use a two- to six-way fibre-port cluster from *Schäfter+Kirchhoff*. This is an enclosed cage-mounted system, providing excellent alignment stability, as required for the coupling between the input and output optical fibres. The combined rubidium and potassium light and the rubidium repumping light are collimated at two input couplers. PBS cubes and rotatable zero-order half-wave plates are used to first combine and then divide up the two beams as shown in Fig. 2.24. PBS cubes produce near-pure polarised light in transmission, but the reflected light contains a small but significant fraction of the incorrect polarisation. Linear polarisers are therefore used to clean up the polarisation of the three beams which undergo reflection at the last cube they encounter.

This imperfection of the PBS cubes means it is also impossible to balance the output beam powers with respect to both inputs, but it may be done for just one or the other. The force-exerting light can therefore be split into three balanced pairs, while various amounts of rubidium repumping light ends up in each beam. The light is coupled into six fibres which transport the light to the required positions on the vacuum table.

2.4.6 Optics around the MOT cell

Each MOT beam (now containing all four frequencies) is allowed to freely expand from the end of an optical fibre before being collimated by a lens, producing beams 3 cm in diameter. To produce pure circular polarisation each beam then passes through linear polarising sheet and a zero-order quarter-wave plate. All of these components are cage mounted for stability and to be sure that the expanding beam and lens share a common axis. Two 2'' steering mirrors are used to allow adjustment of the angle and position of each beam at the MOT cell, finally reproducing the picture of Fig. 2.16.

2.5 Magnetic transport

In order to transport the atoms to the science cell we use the tight magnetic confinement of the quadrupole field produced by the pair of transport coils that were used at low current in the MOT. As mentioned before the coils are mounted on a *Parker* translation stage to allow computer-controlled motion along the vacuum system axis. The coils are clamped tightly in place and held on a set of coil mounts that were manufactured on site to our design.

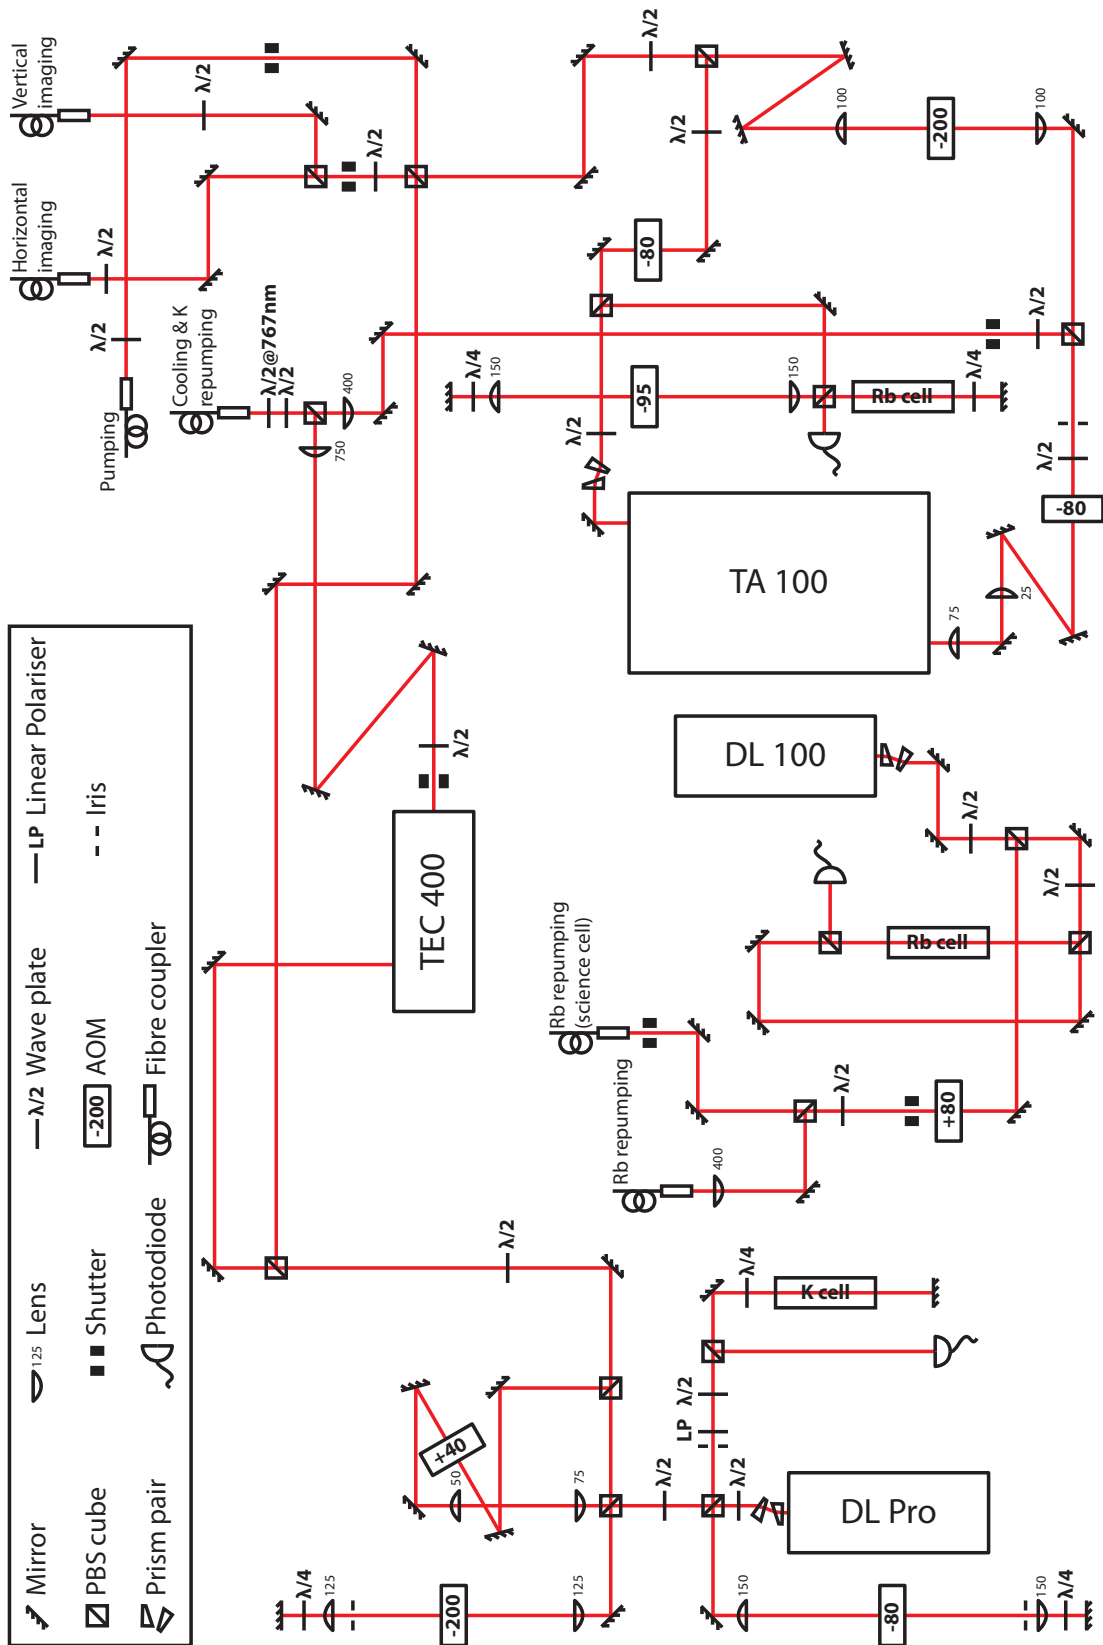


Figure 2.23: Overview of the laser system. The DL Pro is the master laser for potassium, seeding the TEC 400 tapered amplifier. The TA 100 contains both a master laser and tapered amplifier as described in the text, providing all but the repumping light for rubidium. This is catered for by the DL 100 diode laser. The central frequencies of the AOMs are given, and the destinations of the optical fibres are indicated.

2.5 Magnetic transport

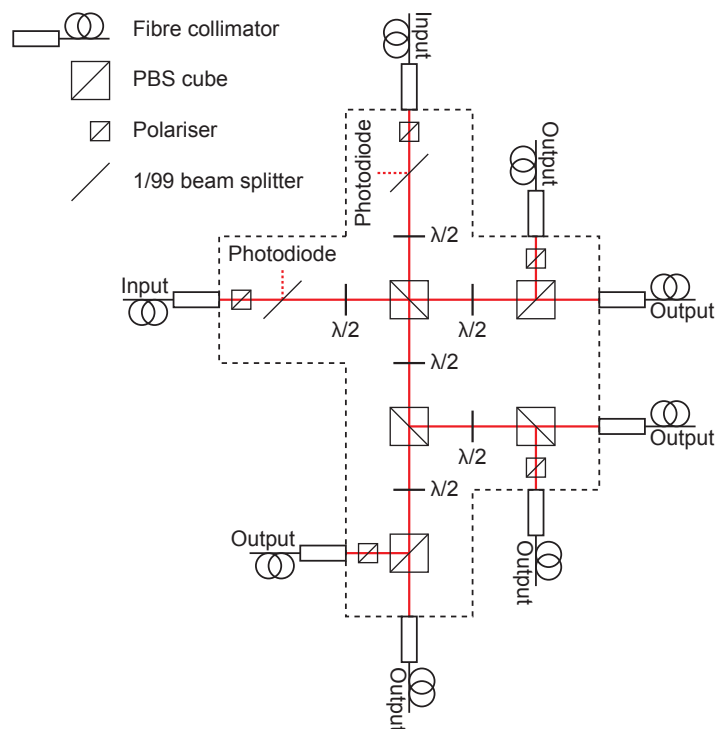


Figure 2.24: Representation of the *Schäfter+Kirchhoff* two- to six-way fibre-port cluster. 1% of each input beam is split off to monitor input power at a photodiode. Though the system is enclosed it is possible to rotate the wave plates in order to obtain the required power balance at the outputs.

2.5.1 Coil holder design

The coil holders are formed from thirteen separately machined aluminium pieces which are bolted together with M6 screws as displayed in Fig. 2.25. The base is screwed directly to the translation stage which comes with an array of threaded holes for this purpose. The arms are the largest pieces and are 25 mm thick for high rigidity, as the coils are held far away from the translation stage base. This is to allow plenty of room for mounting optical elements around the science cell. The vertical struts are far enough away that when the coils are at the MOT cell the 45° horizontal beam may pass by.

The coils themselves are held in clamps comprising four pieces each. The two plates on the inner side are fixed solidly to the arms, fixing the coil separation. The two outer plates clamp onto the inner plates to hold the coils in place. There is a gap in each clamp to reduce eddy currents in the mounts. Grooves in the support arms provide channels for the copper tubing to travel down to the water cooling and power connections.

As the tubes for water-cooling and the high-current power cables need to travel up and down the track with the mount an ‘energy chain’ (resembling caterpillar track) from *Igus* is used to keep everything tidy; serious damage could be caused by the cables snagging on anything as they are pulled along by the powerful translation stage. The power cables and tubes for water-cooling must be suitably flexible to cope with




Figure 2.25: Transport coil holders. The holders are made from thirteen separate pieces of aluminium, held together by screws and bolted directly to the translation stage. The coil separation is fixed to just clear the vacuum system during transport, with the outer plates clamping down to securely hold the coils. The water and electrical connections are mounted off the side of the structure. The arrow shows the direction of motion when the translation stage is operated.

the small bending radius and repeated flexing during every experimental run. The assembled transportation system is shown in Fig. 2.26.

2.5.2 Translation stage

The *Parker* 406XR series translation stage consists of a carriage (to which the coil mount is attached) mounted on a ballscrew drive which is driven by an in-line synchronous servo motor. The lead of the ballscrew thread is 25 mm while the motor rotates at a maximum rate of 6000 rpm, giving a maximum speed of 2.5 m s^{-1} . The track position is monitored by a *Renishaw* RGH24 series digital linear non-contact optical encoder with $0.5 \mu\text{m}$ resolution. Manual limit sensors are included to stop the track in the event of a control error. The stage allows for 900 mm of travel and is very robust, taking a load of up to $\sim 6000 \text{ N}$ and a torque of up to $\sim 270 \text{ N m}$, far more than we actually subject it to. The repeatability of the track end point is given as $\pm 3.0 \mu\text{m}$.

The track is controlled by a *Parker* Compax3 servo drive, computer-operated via RS 232 connectors. We found that it was important to position the servo drive far from the other electrical components of the system, due to the emission of radio-frequency noise that affected the production of condensates. This noise has been noted elsewhere [41], though in this case no detrimental effects were seen. Similarly we have found it essential to de-energise the motor while the atoms are in the science cell. These




Figure 2.26: The transport coils in position around the MOT cell. The coil holders are mounted to the translation stage and travel towards the observer to reach the science cell. The *Igus* ‘energy chain’ used to shepherd the cables can be seen in the bottom right-hand corner of the picture.

precautions vastly increased the reproducibility of condensates in the system.

2.5.3 Atom transportation

The process of moving the track back and forth during each experimental run is accomplished by fixing the track motion parameters (i.e. maximum speed, acceleration and jerk) and setting the two target end points. The track can then be operated by using the ‘ \pm jog’ commands by applying a voltage to the correct pin of the servo drive, prompting the track to move in the designated direction. The coils travel a distance of 740 mm with a maximum speed of 0.5 m s^{-1} , acceleration of 5 m s^{-2} and jerk of 50 m s^{-3} . This takes a little under two seconds end-to-end.

The maximum available current of 200 A is passed through the transport coils during translation to confine the atoms as tightly as possible. The ‘trap acceleration’ $g_F m_F \mu_B B' / m$ is $\sim 50 \text{ m s}^{-2}$ for rubidium ($\sim 110 \text{ m s}^{-2}$ for potassium) at this current, much larger than the maximum acceleration of the track of 5 m s^{-2} , so heating due to the macroscopic trap motion is not an issue. It is possible for the more energetic atoms to reach the sides of the vacuum chamber as they are transported through the minimum aperture of 10 mm, causing cooling and atom loss, which will be quantified in the following chapter.

2.5.4 Majorana spin-flips

Another possible source of heating comes from what are known as Majorana spin-flips. The problem with the linear quadrupole trap created by the transport coils is that there

is a magnetic zero at the centre. As described in §1.4.2, only atoms with $g_F m_F > 0$ can be trapped at low magnetic fields, and will remain trapped as long as they move adiabatically. However, at the centre of the trap the magnetic field direction rapidly changes and the spin of an atom moving through the centre will have a good chance of becoming mis-aligned with the field direction (undergoing a ‘spin-flip’). Though the region this occurs in is relatively small, the colder atoms spend more time near the centre and so are preferentially lost, causing heating as well as atom loss. This loss rate is estimated in [55] as

$$\Gamma_M \approx 7.4 \frac{\hbar}{m} \left(\frac{\mu_F B'}{k_B T} \right)^2, \quad (2.19)$$

where m is the atomic mass. $\Gamma_M \approx 0.15 \text{ s}^{-1}$ for rubidium in our trap at $100 \mu\text{K}$, so for the short time spent in the pure quadrupole field this process is not a major issue. As the atoms get colder this becomes an overwhelming problem though, preventing the continued cooling to the quantum degenerate regime in such a trap.

2.6 Quadrupole Ioffe-configuration trap

The quadrupole Ioffe-configuration (QUIC) trap [56] gets around the problem of Majorana spin-flips by eliminating the magnetic zero altogether, having a magnetic minimum at some finite field strength. This comes at the expense of the sharp trapping confinement and simplicity of a quadrupole trap, but is still relatively straightforward to implement. In fact it brings together the three basic concepts discussed in §2.2; the field from a single coil, the quadrupole field and a bias field. The QUIC trap geometry with the directions of current flow is illustrated in Fig. 2.27, where the origin of the coordinate system is located at the centre of the quadrupole coils. The quadrupole coils ideally produce a purely linear field

$$\mathbf{B}_{\text{quad}} = B' \begin{pmatrix} x \\ y \\ -2z \end{pmatrix} \quad (2.20)$$

as previously discussed, where B' is now the total field gradient along the weak axes. The so-called Ioffe coil is located on the x -axis and produces an axial field that opposes the quadrupole field. For a non-zero trap minimum to exist somewhere along the x -axis there must be a point x_{min} at which

$$\frac{dB_x^{\text{Ioffe}}}{dx} = -B' \quad (2.21)$$

and hence

$$\frac{dB_y^{\text{Ioffe}}}{dy} = \frac{dB_z^{\text{Ioffe}}}{dz} = \frac{1}{2} B'. \quad (2.22)$$

2.6 Quadrupole Ioffe-configuration trap




Figure 2.27: Formation of a QUIC trap. The vertical coils with opposing senses of current flow (indicated by the arrows) provide the quadrupole field while the single horizontal Ioffe coil provides the requisite field curvature. The rectangular bias coils can be used to independently tune the field offset. The resultant trap minimum lies somewhere along the x -axis, about which the trap displays rotational symmetry.

It immediately follows that the total field gradients in the y and z directions at this point are $(3/2)B'$ and $-(3/2)B'$ respectively. Defining the field offset at this minimum to be $-B_0$ and the field curvature produced by the Ioffe coil to be $-(1/2)B''$ (where B_0 and B'' are positive to comply with the chosen coordinate system), the total field near the minimum may be written as

$$\mathbf{B} = \begin{pmatrix} -B_0 - \frac{1}{2}B''(x - x_{\min})^2 \\ \frac{3}{2}B'y \\ -\frac{3}{2}B'z \end{pmatrix}. \quad (2.23)$$

The bias field can be applied to decrease B_0 , the relevance of which will shortly become apparent. The field *amplitudes* along x , y and z are then simply

$$|B|_x = B_0 + \frac{1}{2}B''(x - x_{\min})^2 \quad (2.24)$$

$$|B|_y = B_0 \left[1 + \left(\frac{3}{2} \frac{B'}{B_0} y \right)^2 \right]^{1/2} \approx B_0 + \frac{1}{2} \frac{\left(\frac{3}{2} B' \right)^2}{B_0} y^2 \quad (2.25)$$

$$|B|_z = B_0 \left[1 + \left(\frac{3}{2} \frac{B'}{B_0} z \right)^2 \right]^{1/2} \approx B_0 + \frac{1}{2} \frac{\left(\frac{3}{2} B' \right)^2}{B_0} z^2. \quad (2.26)$$

The identical forms of (2.25) and (2.26) gives the trap its fundamental circular symmetry about the axis of the Ioffe coil as illustrated in Fig. 2.27. Using the expression for the Zeeman shift (1.36) we can obtain expressions for the axial and perpendicular

trapping frequencies

$$\omega_{\text{ax}} = \sqrt{\frac{\mu_F B''}{m}} \quad (2.27)$$

$$\omega_{\perp} = \sqrt{\frac{\mu_F \left(\frac{3}{2}B'\right)^2}{mB_0}}. \quad (2.28)$$

The trapping potential is actually modified slightly by the Earth's gravitational field, but this description provides a good estimate of the true values. The perpendicular trapping frequencies diverge as $B_0 \rightarrow 0$ which is why it is useful to be able to directly tune this value with the bias coils. B'' scales as $I_{\text{Ioffe}}/a_{\text{Ioffe}}^3$ [56], hence the tight winding of the Ioffe coil.

We are able to measure B_0 by finding the microwave frequency that removes all atoms from the trap (see section §2.7.1), while B' for the transport coils has been calibrated using a gauss meter. During normal operation of the QUIC trap they are 5.3 G and 64 G cm^{-1} respectively. B'' is estimated to be 120 G cm^{-2} from a suitable model of the Ioffe coil¹ (which is constrained by our knowledge of B_0 and B'). This predicts trapping frequencies $\omega_{\text{ax}} = 2\pi \times 21(14) \text{ Hz}$ and $\omega_{\perp} = 2\pi \times 79(53) \text{ Hz}$ for potassium (rubidium). A rudimentary check of the vertical trapping frequencies may be provided by measuring the difference in ‘sag’ due to gravity Δz (the distance between the trap minima) for the two species. Assuming perfect harmonic potentials, the trap minimum is displaced by a distance g/ω_z^2 under gravity, where in our case $\omega_z = \omega_{\perp}$. This gives the result

$$\omega_{z\text{K}}^2 = \left(\frac{m_{\text{Rb}}}{m_{\text{K}}}\right) \omega_{z\text{Rb}}^2 = \frac{g}{\Delta z} \left(\frac{m_{\text{Rb}}}{m_{\text{K}}} - 1\right) = \frac{g}{\Delta z} \left(\frac{48}{39}\right), \quad (2.29)$$

which agrees reasonably well with the predicted values. (2.29) typically slightly underestimates the actual trapping frequencies due to the real anharmonicity of the magnetic field.

2.7 Evaporative cooling

Though the temperatures achieved by laser cooling are very impressive it has not yet been possible to cool far enough at sufficient atomic density to reach quantum degeneracy by these processes alone. The suggestion of evaporative cooling as a pathway to higher phase-space densities was first suggested by Hess in 1985 [57], focussing on atomic hydrogen. This was applied to alkali atoms by groups at MIT [58] and JILA [59] in 1994 and indeed subsequently led to the first realisation of alkali atomic BECs in 1995 [3, 4].

The concept is strikingly simple. Removing atoms with an energy greater than the atomic average leaves the remaining atoms with a lower average energy. Upon

¹The coil may be modelled by assuming all eighteen windings have radius 8.5 mm and are situated a distance 18 mm from the centre of the quadrupole field.

2.7 Evaporative cooling




Figure 2.28: Illustration of the process of evaporative cooling in a harmonic trap. The trap depth is gradually lowered from **a** to **d** such that the most energetic atoms escape the trap and the atomic cloud rethermalises at a lower temperature.

rethermalisation the remaining atoms are at a lower temperature. This process can be repeated to obtain temperatures far below the recoil limit (Fig. 2.28). The obvious drawback is that atom loss is an integral part of the process. However, evaporation can be extremely efficient, with phase-space-density growth by three orders of magnitude commonplace for a drop in number by only a single order of magnitude.

In a simple model as described in [60] the process can be viewed as repeated removal of the high-energy tail of the Maxwell-Boltzmann energy distribution¹ and rethermalisation. The removal of atoms is achieved by restricting the depth of the trapping potential to some finite value. This is characterised by the truncation parameter η , giving trap depth $\eta k_B T$. As long as η is large the treatment that the system is close to equilibrium is valid. The atoms that are lost from the trap take away an average energy $(\eta + \kappa)k_B T$, where κ is a small number usually somewhere between 0 and 1. Thus for large η a rough approximation can be made that κ is negligible. For the temperature to decrease then we must have $\eta k_B T$ greater than the mean atomic energy $(3/2 + \delta)k_B T$, where the mean kinetic energy is $(3/2)k_B T$ in three dimensions and the mean potential energy is $\delta k_B T$, and it is easily shown that $\delta = 3/2$ for a harmonic trap. The rate of energy loss is therefore simply given by

$$\frac{dE}{dN} = \eta k_B T, \quad (2.30)$$

while the total energy of the harmonically trapped sample is given by

$$E = 3Nk_B T. \quad (2.31)$$

Equating (2.30) with the derivative of (2.31) with respect to N leads us to the temperature variation with N ,

$$T \propto N^{(\eta-3)/3}, \quad (2.32)$$

showing that for temperature reduction we indeed require $\eta > 3$. The volume of a harmonically trapped sample V is proportional to $T^{3/2}$, which then directly gives us

¹For the most part evaporation is carried out in the regime where this is still a reasonable approximation.

Quantity ($\propto N^x$)	Exponent, x
Atom number, N	1
Temperature, T	$(\eta - 3)/3$
Volume, V	$(\eta - 3)/2$
Density, n	$(5 - \eta)/2$
Phase-space density, $D = n\lambda_T^3$	$4 - \eta$
Elastic scattering rate, $\Gamma_{\text{el}} = n\sigma v$	$(6 - \eta)/3$

Table 2.5: Scaling laws for physical quantities (in the spirit of [60]) in terms of truncation parameter η for a harmonic trap in the simple model described in the text. Within this model it is clear that for evaporative cooling to be successful (i.e. for phase-space density to increase) η must be greater than 4, while to be in the ‘run-away’ regime (i.e. for the elastic scattering rate to increase during evaporation) η must be greater than 6.

the number dependence of all other relevant physical quantities in the system as given in Table 2.5.

More important than the temperature scaling is that of the phase-space density, and also that of the elastic scattering rate to maintain the evaporation process. The above model suggests that for $\eta > 6$ the elastic scattering rate increases as N drops, in what is known as ‘run-away’ evaporation. Of course these values of η are not really very large, so the model is only used as a guide to the key concepts. A more sophisticated model [61] gives $\eta > 4.6$ for the run-away regime (in the absence of any additional loss processes). Most cooling procedures operate by keeping η relatively constant, forcing the evaporation as the temperature drops by reducing the trap depth correspondingly. This can be achieved in a magnetic trap by applying time-varying RF or microwave frequency radiation (discussed later), and in optical traps by reducing the optical power.

The above model suggests that the efficiency of evaporation has no upper bound, as in the commonly cited example of waiting for a single particle to have the entire energy of the system. But there is a trade-off between cooling efficiency and evaporation timescale, while inherent loss mechanisms (such as collisions with the background gas) limit the time we may take over evaporation. The key requirement in a real system then is that the elastic collision rate Γ_{el} be much higher than the loss rate Γ_{loss} of the sample. The value of η available experimentally is constrained by the ratio $\alpha = \Gamma_{\text{el}}/\Gamma_{\text{loss}}$. If α is small we can no longer go to high η as the trap losses become overwhelming. The quantity α may be added into evaporation models to include trap losses as described in [60], where it is seen that the minimum value of α for run-away evaporation¹ is ~ 100 for $\eta = 6$, rising to ~ 300 for $\eta = 8$. Note that the value of Γ_{loss} will change during evaporation; for example, the three-body inelastic collision rate scales strongly with

¹This minimum value is a factor $2\sqrt{2}$ lower than given in [60] as there the elastic collision rate was defined for the trap centre rather than the cloud as a whole.

2.7 Evaporative cooling

density.

Considering the direct evaporation of rubidium alone in the QUIC trap in our experimental setup, α is large (> 1000) owing to the long trap lifetime and the large scattering cross-section, allowing highly efficient evaporation at high η . To quantify the efficiency of the evaporation process we introduce the local quantity

$$\gamma = -\frac{d(\ln D)}{d(\ln N)} \quad (2.33)$$

where $D = n\lambda_T^3$ is the phase-space density. So in the simple lossless model $\gamma = \eta - 4$. Alternatively, a global value γ_{tot} is given by taking the difference between final and initial conditions rather than the derivative. Typical experimental values of γ_{tot} for direct evaporation are in the range 1 to 3 [13, 60, 62, 63]. This is fairly species dependent; for example, evaporation of ^{87}Rb usually gives a value closer to 3 due to its favourable scattering length. This is a good benchmark for comparison with our results.

It is also shown in [60] that the rate of atom loss due to evaporation in the limit of large η is

$$\Gamma_{\text{ev}} = \frac{\eta e^{-\eta}}{\sqrt{2}} \Gamma_{\text{el}}. \quad (2.34)$$

If we keep η and Γ_{el} constant (for forced evaporation on the cusp of the run-away regime) then, in the absence of losses, the atom number decays exponentially during evaporation with time constant $1/\Gamma_{\text{ev}}$. As $T \propto N^x$ the sweep of the trap depth should also be exponential to maintain constant η with a time constant $\sim 1/(x\Gamma_{\text{ev}})$. Again this may be used only as a rough guide, though most experiments (including ours) use an exponential ramp of the trap depth with an experimentally determined optimum time constant.

2.7.1 Sympathetic cooling

When direct cooling of a particular species is difficult (as for ^{39}K) we may make use of a second to act as a coolant [64]. Both species are loaded into a trap but only one is selectively evaporated, while interspecies collisions allow the other to be sympathetically cooled (Fig. 2.29). This process breaks down once the atom numbers of the two species are comparable, so efficient sympathetic cooling requires a large initial ratio of coolant to load. Ideally, no atoms are lost from the load species while the phase-space density is greatly increased, so it is possible to create large condensates despite the relatively low initial numbers in the trap. Naturally, losses do occur in real systems, and the sympathetic cooling efficiency can be characterised by γ_{tot} for the load species. The larger the initial load the lower the cooling efficiency will be [63].

Considering for a moment the situation at the start of sympathetic cooling in the QUIC trap, we can estimate the elastic scattering rate for potassium atoms colliding




Figure 2.29: Sympathetic cooling in a harmonic trap. The trap depth is lowered as in Fig. 2.28 for the coolant species (black), while the trap depth remains far greater than the temperature of the load species (red). As the temperature of the coolant is reduced by evaporation the load is cooled by inter-species elastic collisions, while no load atoms are intentionally removed. The load is thus cooled without significant reduction in atom number.

with rubidium atoms to be given by

$$\Gamma_{\text{K-Rb}} = \bar{n}_{\text{Rb}} \sigma_{\text{K-Rb}} \bar{v}_{\text{K-Rb}} . \quad (2.35)$$

This quantity is important, as it gives the limiting rate of thermalisation of potassium during sympathetic cooling. Here $\bar{n} = n(0)/(2\sqrt{2})$ is the average density in a harmonic trap, $\sigma_{\text{K-Rb}} = 8\pi a_{\text{K-Rb}}^2$ is the inter-species scattering cross-section and

$$\bar{v}_{\text{K-Rb}} = \sqrt{\frac{8k_{\text{B}}}{\pi}} \sqrt{\frac{T_{\text{K}}}{m_{\text{K}}} + \frac{T_{\text{Rb}}}{m_{\text{Rb}}}} \quad (2.36)$$

is the average relative velocity between atoms of different species. At the start of our evaporation $T_{\text{K}} > T_{\text{Rb}}$ and in fact T_{K} always lags a little throughout the evaporation process as the elastic scattering rate for collisions between rubidium atoms $\Gamma_{\text{Rb}} \approx 5\Gamma_{\text{K-Rb}}$ when $T_{\text{K}} \approx T_{\text{Rb}}$. For our experimental parameters we find that initially $\Gamma_{\text{K-Rb}} \sim 4\text{ s}^{-1}$ owing to the weak trapping in the QUIC trap and the low inter-species scattering length $a_{\text{K-Rb}} = 36 a_0$ [65]. However, our vacuum is good enough that $\Gamma_{\text{loss}} \sim 1/300\text{ s}^{-1}$, allowing suitably long evaporation sweeps for T_{K} to stay close to T_{Rb} (important for maintaining the spatial overlap between the two species for cross-thermalisation). A more detailed discussion of thermalisation in mixtures of ultracold gases may be found in [66].

2.7.2 Thermalisation

The constant unavoidable loss of atoms from the trap means that the system is never fully in equilibrium. However, as long as $\alpha \gg 1$ the system is in what is termed quasi-equilibrium, where the thermal distribution is a very good approximation. Thermalisation of the cloud is usually taken to require ~ 3 elastic scattering events per particle [67]. Extreme scattering cross-sections can lead to conditions where thermalisation is no longer able to keep up with atom loss, and the system is deemed to be out of equilibrium. Small cross-sections reduce thermalisation below the collision rate

2.8 Crossed optical dipole trap

with the background vapour, while large cross-sections drive the three-body loss rate above the elastic scattering rate. Both of these regimes are accessible with a Feshbach resonance, so it is well to be aware of the contingency.

2.7.3 Microwave antenna

When evaporating a single species it is common to use resonant RF radiation to transfer high energy atoms into untrappable states (the process is similar to the state transfer described in §2.9). Signal generators and amplifiers spanning the kHz to tens of MHz range required for RF evaporation from standard magnetic traps are readily available. However, the Zeeman substate splitting for potassium and rubidium are almost identical, so it is impossible to use RF evaporation for sympathetic cooling as forced evaporation would also occur for potassium. The large hyperfine splitting of 6.8 GHz for the rubidium ground states, though, makes the use of microwaves a viable option, as frequencies around this point are far off resonance for potassium. In our experiments, rubidium atoms are transferred from the $|2, 2\rangle$ state to the untrapped $|1, 1\rangle$ state during microwave evaporation.

The design of the microwave antenna was undertaken by Igor Gotlibovych [68]. The relatively short wavelength required (~ 4 cm) is comparable to the antenna distance from the atoms, allowing consideration of directional radiation intensity. A free-space helix antenna was created to maximise the field strength at the atoms. The circumference of the helix matches the wavelength at 6.8 GHz and consists of ten loops with a pitch angle of 13° . The signal is provided by either of two microwave signal generators: an *HP* HP83622A and an *Agilent* N5183A MXG. A suitable amplifier for 6.7–7.0 GHz was built to specification by *Microwave Amplifiers Ltd.*, having a saturated output power of 39.5 dBm (~ 9 W).

2.8 Crossed optical dipole trap

In order to make use of Feshbach resonances the atoms must be trapped in a purely optical potential. This may be done by making use of the dipole potential (1.69) as discussed in §1.4.3. Trapping is achieved by focussing a gaussian shaped red-detuned beam to a $1/e^2$ waist size of w_0 to create an intensity maximum in space, forming an optical dipole trap. The Rayleigh range given by [69]

$$z_R = \frac{\pi w_0^2}{\lambda} \quad (2.37)$$

is generally far greater than the beam waist, so for tight confinement in three dimensions a second beam is added, crossing the first to form a crossed dipole trap (CDT) as illustrated in Fig. 2.30.

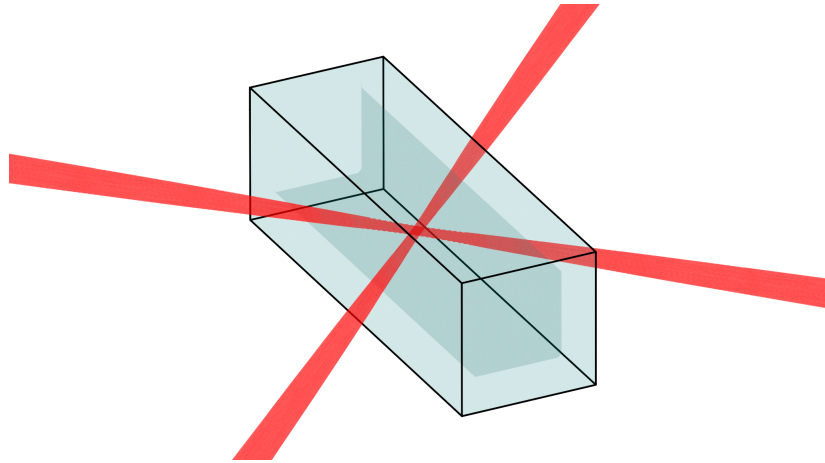


Figure 2.30: Basic two-beam trapping scheme, forming the CDT in our science cell. Crossing two beams produces tight confinement in three dimensions, as opposed to the single-beam case where trapping along the axis of the beam is usually relatively weak.

2.8.1 CDT laser

As suggested in §1.4.3 we must try to fulfil the requirements of a low scattering rate and a large trap depth by balancing the laser detuning with the laser power required. These parameters also depend on the beam waist of course, but at the time of purchasing the laser we based our decision on the dipole trap parameters of other research groups. We settled on a 10 W linearly polarised ytterbium fibre laser from *IPG Photonics* (YLM-10-LP-SC) at 1070 nm. The laser is simple to operate given it produces light at a fixed frequency with a large linewidth of 1.2 nm. The light is collimated at the end of the 10 m long output fibre with a beam diameter of ~ 5 mm.

Power stability

The beam power stability is important as large long-term drifts alter the trap depth over time and high frequency modulations cause extra heating in the trap. We therefore implement a negative feedback loop in order to stabilise the power. We control the power by passing the light through an AOM and adjusting the acoustic amplitude in the crystal to vary the power in the first-order diffracted beam. The power in this beam is monitored by picking off a fraction of the trapping light once it has passed through the system, deflecting it into a fast photodiode. The voltage signal obtained from the photodiode is compared with a desired voltage (set by a computer-controlled analogue channel) by a PID controller [70]. The PID increases or decreases, as necessary, its output voltage, which is sent to the voltage variable attenuator (VVA) in the AOM driver. This directly effects the RF power sent to the AOM, completing the loop. The trapping beam power may therefore be controlled by adjusting the voltage signal sent to the PID.

The disadvantage of such a scheme is that a fraction of beam power is lost in the

2.8 Crossed optical dipole trap

zeroth order of the AOM which is sent straight to a high-power beam dump. This fraction is minimised by several measures. Firstly, the beam diameter at the AOM is closely matched with the active aperture of the AOM ($\varnothing 1$ mm and 3 mm \times 1.25 mm respectively), as the rise time may be traded off for beam diffraction efficiency by increasing the beam diameter. This still gives a rise time of only ~ 300 ns. Secondly, the voltage controlled oscillator (VCO) in the AOM driver is sent a fixed voltage that has been tuned (via adjustment of a trimmer resistor) such that it outputs a frequency of exactly 110 MHz, the central frequency of the chosen AOM where the highest diffraction efficiency occurs. Thirdly, the output of the laser is oriented such that polarisation is 90° to the AOM mounting plate as required for optimal diffraction efficiency. Finally, the AOM angle was carefully tuned to be Bragg matched for first order diffraction, the AOM being mounted on a modified mirror mount to allow simple angular adjustments.

The AOM used is product number 3110-197 from *Crystal Technology* and was chosen specifically for the possibility of high diffraction efficiency. The photodiode is a *Thorlabs* DET36A/M model which has a large (13 mm²) detection area and 14 ns rise time. The key components of the AOM driver are shown in Table 2.6, and the PID device was designed and built by Philip Richman in the lab. The losses in the system mean that the trapping beam has a maximum power of 7.5 W at the atoms, though we usually limit this to 7 W to ensure the PID can function properly to provide a stable beam power.

2.8.2 Forming the trap

An overview of the CDT setup is shown in Fig. 2.31. The laser output coupler is held securely in a cage mount along with two lenses to form a telescope to obtain the desired beam diameter for efficient diffraction by the AOM. The zeroth order light is directed straight into a high-power beam dump, while the first order beam is reflected off two mirrors onto the optical axis of a second cage system. This contains a PBS cube to make sure the polarisation is pure, deflecting the orthogonal component of the

Product number	Description	Supply voltage (V)	Control Voltage (V)
ZX95-200-S+	VCO	+12/Gnd.	~ 4 (for 110 MHz)
ZX73-2500-S+	VVA	+5/Gnd.	0 \rightarrow 12
ZXSWA-2-50DR	Switch	± 5 /Gnd.	0/5

Table 2.6: Key components of the AOM driver used to control the CDT power. All components are from *Mini-circuits*. The VCO has a fixed control voltage as the AOM is only used to control power. The VVA is controlled from an external voltage input during normal operation, but may be controlled by a chassis-mounted potentiometer during testing. The switch is included so that the beam may be turned off rapidly via a TTL trigger. An external RF amplifier (*Delta RF* LA2-1-525-30) is used to allow all the components to be run off a $\pm 5/12$ V power supply in the driver box.

light downwards into another beam dump. The two lenses in this cage mount first expand then focus the beam as is discussed below. A periscope with adjustable mirrors elevates the beam to the height of the science cell (185 mm above the surface of the optical table) and allows the beam to be directed into the science cell at 45° .

To increase the effectiveness of our available optical power we form the CDT in the so-called ‘bow-tie’ configuration by re-collimating the trapping beam after its first pass through the science cell, reflecting it off a couple of mirrors and then re-focussing the beam such that the second focal point coincides with the first. In this way almost the full beam power is recycled. In our system we cross the beams at 90° . The transmission through the science cell is maximised by having the beam p-polarised with respect to the cell surface for both passes. With this polarisation the transmission is $\sim 96\%$ per window (recalling that only the outer surface has been anti-reflection coated). There is no problem with beam interference as the linewidth is sufficiently large that the path length of the bow-tie section is far greater than the coherence length.

Specialised optics are not required to form the optical dipole trap, owing to the large beam waists desired (the system is effectively diffraction limited by design). All of the lenses and mirrors used are broadband anti-reflection coated 1” pieces from *Thorlabs*.

Tuning the beam waist

When setting up the CDT we wanted the ability to alter the beam waist with minimal realignment time. We also chose to have the focussing lenses mounted in a cage system for mechanical stability and to allow simple translation along the optical axis. We were also constrained by the coil mount around the science cell. The result is that the back focal length, d_{BFL} , of the lens system must be greater than ~ 450 mm. This means that the beam must first be expanded from the 0.5 mm waist required at the AOM, otherwise the waist at the atoms would end up being $> 270 \mu\text{m}$.

In order to calculate the required lenses and separations in the system we can make the paraxial approximation and use simple ray transfer matrices to find the angle and off-axis distance of a ray at the beam waist (ignoring the diffraction limit). The position and angle of this ray are recorded in the form

$$\mathbf{x} = \begin{pmatrix} x \\ \theta \end{pmatrix}, \quad (2.38)$$

and we sequentially apply to the initial conditions \mathbf{x}_i the transfer matrices \mathbf{D} for traversing a certain distance d along the optical path and \mathbf{L} for the advent of a lens of focal length f , which take the form

$$\mathbf{D} = \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \text{ and } \mathbf{L} = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}. \quad (2.39)$$

For the two-lens system illustrated in Fig. 2.32 the final conditions are given by

2.8 Crossed optical dipole trap




Figure 2.31: CDT setup overview. The output of the fibre laser passes through a telescope resulting in a collimated beam that is size-matched to the active aperture of the 110 MHz AOM. Zeroth order light is sent directly to a beam dump. The polarisation of the first order light is purified using a PBS cube, with light of the incorrect polarisation sent to a beam dump below the cube. This beam is expanded and then focussed by the next lens pair before being raised to the height of the science cell by a periscope. The focal point lies at the centre of the science cell. After the first pass the beam is collimated and refocussed for the second pass, crossing itself perpendicularly. A small fraction of the beam is then split off and sent to two photo-detectors. One is used in the feedback loop for stabilising the beam power, the other for independent monitoring of the beam power. The rest of the light is sent to another beam dump.

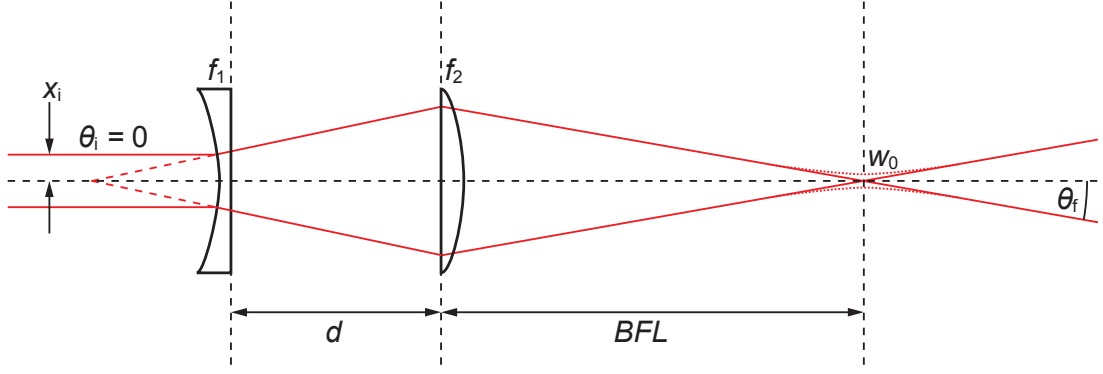


Figure 2.32: Two-lens system for producing the desired focal beam-waist w_0 for optical trapping.

$$\mathbf{x}_f = \begin{pmatrix} x_f \\ \theta_f \end{pmatrix} = \mathbf{D}_{\text{BFL}} \mathbf{L}_2 \mathbf{D} \mathbf{L}_1 \mathbf{x}_i . \quad (2.40)$$

The angle of convergence allows us to find the true diffraction limited waist from

$$w_0 = \left| \frac{\lambda}{\pi \theta_f} \right| , \quad (2.41)$$

while setting $x_f = 0$ lets us find the back focal length. This results in expressions for w_0 and d_{BFL} dependent on the initial beam waist x_i , the lens focal lengths f_1 and f_2 and their separation d :

$$d_{\text{BFL}} = \frac{f_2(d - f_1)}{d - f_1 - f_2} , \quad (2.42)$$

$$w_0 = \left| \frac{\lambda}{\pi x_i} \frac{f_1 f_2}{(d - f_1 - f_2)} \right| . \quad (2.43)$$

This enables us to find combinations of lenses that produce the desired beam waist within the constraints of our setup. In our finalised system we use $f_1 = -50$ mm, $f_2 = 100$ mm and $d = 74$ mm, giving $w_0 = 140$ μm and $d_{\text{BFL}} = 511$ mm. We have found that using the cage system it is a simple task to modify the beam waist by translation of the lenses and almost as easy to simply switch out a lens to change the accessible range of beam waists. The above lenses allow tuning of w_0 between ~ 120 and 150 μm . The focal beam waist was verified by diverting the beam directly into a charge-coupled device (CCD) camera and extracting the waist at several points along the optical axis.

2.8.3 Trap potential

The Rayleigh range given by (2.37) for the above trap parameters is ~ 6 cm and gives a beam waist at the science cell walls of only $\sim 1.03 w_0$. A close approximation to the true trapping potential near the focal point is that given by collimated gaussian beams of waist w_0 . With x and y the beam axes and z the vertical axis the total potential

2.8 Crossed optical dipole trap

may be written as

$$U(x, y, z) = mgz - \frac{A}{2}e^{-2(x^2+z^2)/w_0^2} - \frac{A}{2}e^{-2(y^2+z^2)/w_0^2}, \quad (2.44)$$

where the gravitational potential is included. Here g is the acceleration due to gravity, m is the atomic mass of a trapped atom and $-A$ is the peak optical potential shift at the centre of the two overlapped beams. The value of the latter may be found from (1.69);

$$A = -U_{\text{dip}}(\mathbf{r} = 0) = \frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) \frac{2P}{\pi w_0^2}, \quad (2.45)$$

where P is the total power of the trapping beams.

A trap minimum is created at $x = y = 0$, $z = z_0$, where gravity displaces the minimum vertically with respect to the centre of the dipole trap beams. The trapping frequencies can be calculated from (2.44) and are given by

$$\omega_x^2 = \omega_y^2 = \frac{2A}{mw_0^2} e^{-2z_0^2/w_0^2} \quad (2.46)$$

$$\omega_z^2 = 2 \left(1 - 4 \frac{z_0^2}{w_0^2} \right) \omega_x^2, \quad (2.47)$$

giving a cylindrically symmetric trap around z . In the high-optical-power limit where $A \gg mgw_0$, the quantity $z_0/w_0 \rightarrow 0$ such that ω_x tends to $\sqrt{2A/mw_0^2}$ and ω_z tends to $\sqrt{2}\omega_x$. From (2.47) it is clear that in the low-optical-power limit the trap is destroyed if the trap minimum displacement reaches half a beam waist. The ratio between the vertical and horizontal trapping frequencies is shown for the possible values of z_0 in Fig. 2.33. When $z_0 = -w_0/2\sqrt{2}$ the trap is spherically symmetric.




Figure 2.33: Dependence of the ratio of the vertical to horizontal trapping frequencies of the CDT on the sag z_0 of the trap minimum with respect to the beam waist w_0 . Displacement z_0 is always negative using the convention that gravitational potential increases in the positive z direction.

2.9 State transfer

To be able to work with the broad Feshbach resonance in the ^{39}K $|F, m_F\rangle = |1, 1\rangle$ state we need to transfer the atoms from the $|2, 2\rangle$ state. This is possible to achieve directly via a magnetic dipole transition, using radio frequency (RF) radiation of frequency ω in a uniform magnetic field B . The frequency splitting of the $|1, 1\rangle$ and $|2, 2\rangle$ states at low magnetic fields is given by $\omega_{\text{res}} \approx \omega_{\text{hf}} + 3\mu_B B/2\hbar$, where $\omega_{\text{hf}} = 2\pi \times 461.7\text{MHz}$ is the zero-field lower-state hyperfine splitting in ^{39}K and we have taken $g_F \approx \pm 1/2$ for the two hyperfine states. In what is known as the dressed-atom picture we view the lower state as being shifted up by an energy $\hbar\omega$. However, coupling $\hbar\Omega_0/2$ (where Ω_0 is the resonant Rabi frequency) between the two states means that as ω or B is tuned they undergo an avoided crossing, and the eigenstates of the system are split by an energy $\hbar\Omega_0$ where they would otherwise cross (as shown in Fig. 2.34). If the magnetic field is increased adiabatically then all of the atoms initially in the $|2, 2\rangle$ state remain in the adiabatic lower energy state $|\phi_l\rangle$. Removing the RF field once past the crossing point leaves the atoms in the $|1, 1\rangle$ state and state transfer has been successful.

For a finite rate of magnetic field change there is a chance that atoms will undergo Landau-Zener transitions, ending up in the diabatic $|2, 2\rangle$ state. The probability of this occurring is given by the Landau-Zener equation

$$P = e^{-2\pi\Lambda} \quad (2.48)$$




Figure 2.34: The avoided crossing scheme of the $|2, 2\rangle$ and $|1, 1\rangle$ Zeeman levels in the dressed-atom picture in the presence of an RF field. Where the diabatic energy levels (black) cross, the adiabatic energy levels $|\phi_u\rangle$ and $|\phi_l\rangle$ (coloured) are split by $\Omega_0\hbar$. Sweeping the magnetic field strength adiabatically from low to high transfers atoms from one state to the other. The atoms remain in this state once the RF field is removed.

2.10 Absorption Imaging

where

$$\Lambda = \frac{(\Omega_0/2)^2}{\frac{d}{dt}(\omega - \omega_{\text{res}})} . \quad (2.49)$$

In our experiment we chose to perform the state transfer by sweeping the magnetic field in the presence of a steady RF frequency of 469.3 MHz. We sweep the field at a rate of 15.4 G s^{-1} about the 3.6 G resonance position, equivalent to sweeping the detuning $\omega - \omega_{\text{res}}$ at a rate of $2\pi \times 32 \text{ MHz s}^{-1}$. In the Landau-Zener model a transfer efficiency above 99% is achieved for a Rabi frequency $\Omega_0 \gtrsim 2\pi \times 2 \text{ kHz}$.

2.9.1 RF antenna

The antenna used to provide the RF is a single 18 mm radius loop of 1 mm diameter copper wire, situated a few centimetres from the atoms. A single loop was chosen to minimise the inductance of the loop which is the main contributor to the impedance of the antenna. In order to maximise transfer efficiency, Ω_0 must be maximised which is proportional to the amplitude of the RF field. This in turn depends on the power output from the antenna. To maximise this a tuneable capacitor is used to impedance match the imaginary part of the impedance of the antenna due to its inductance. The tuning of the capacitor was achieved by maximising transfer efficiency at low power. Once the full available power is applied we do not observe any atoms left in the $|2, 2\rangle$ state.

2.10 Absorption Imaging

Absorption imaging is the standard method of recording the atomic distribution after some time-of-flight in cold atoms experiments. The atoms are released from the trapping potential and allowed to freely expand until a resonant or near-resonant pulse of light is applied. The transmission along the direction of propagation of the light pulse depends on the number of atoms along the path. The shadow produced is then imaged onto a camera for processing as discussed below. Analysis of the resultant images provides us with details of the imaged cloud such as temperature and atom number as discussed in §2.11. The process is unfortunately destructive, so the full experimental sequence must be run through each time a new image is produced.

2.10.1 Optical density

As a beam of light travels through a cloud of atoms, a fraction of the photons are absorbed dependent on the total atom number per unit area integrated along the beam path (referred to as the column density, n_{2D}) and the atomic absorption cross-section σ . It is simple to show that the intensity is attenuated according to Beer's law,

$$I = I_0 e^{-\sigma n_{2D}} , \quad (2.50)$$

where I_0 and I are the respective incident and emergent intensities of the imaging beam. The quantity $OD = \sigma n_{2D}$ is known as the optical density. The standard method of determining OD requires a set of three images: one where the atoms were present, I_{atoms} , one with the unattenuated imaging light (where no atoms were present), I_{light} , and one with no imaging light for normalisation purposes, I_{dark} . The images here are rectangular arrays of electron counts at each pixel of the imaging camera's CCD. The counts of I_{dark} are produced by any stray non-imaging light that makes its way to the camera, so ideally the optical density is given by

$$OD = -\ln\left(\frac{I_{\text{atoms}} - I_{\text{dark}}}{I_{\text{light}} - I_{\text{dark}}}\right). \quad (2.51)$$

2.10.2 Camera

For cold atom experiments the choice of CCD camera often hinges on a trade-off between price and quantum efficiency (QE), which is the conversion factor of incident photons to stored electrons. High quantum efficiency cameras have several advantages, e.g. allowing shorter/lower-intensity imaging pulses and offering a better overall signal-to-noise ratio (SNR). This is most important in studies of such areas as density fluctuations in ultracold gases [71, 72] where real fluctuations in atom number must be deduced by removing systematic noise in the image. A high QE is always advantageous, but there are other features in cheaper cameras with relatively low QE which can make them more desirable when QE is not critical.

One such feature is the ‘double-shutter’ mode available in some interlining CCD cameras. In these cameras the CCD is subdivided into alternating columns of sensor and storage wells. The storage region is masked, so at the end of the exposure the electrons that have accumulated in the sensor wells are transferred to the storage wells in a matter of microseconds. The main advantage for our imaging is that a second image can be taken almost immediately after the first. A common problem in absorption imaging is that there are fringes in the light intensity (due to the coherence of the laser) which shift in time due to mechanical vibrations in the system. If the time between the two illuminated images is small compared to the period of the vibrations the fringes will overlap and a clear fringeless optical density picture is acquired when dividing one by the other. If, however, the fringes shift between pictures they will also appear in the final optical-density image. This is problematic for image analysis, especially when trying to detect small condensates. It is possible to remove fringes in post-processing by Fourier filtering methods [73] or by matching the I_{atoms} image with a linear superposition of real I_{light} images [74], but it is simplest to avoid the issue in the first place.

Our chosen camera, the *PCO Pixelfly* QE Double Shutter camera, can take two images less than $10 \mu\text{s}$ apart, far shorter than standard full-frame CCD cameras (where the inter-frame time is roughly equal to the digital read-out time, 88 ms for our camera).

2.10 Absorption Imaging

This is because the second image can be taken as soon as the first has been shifted to the masked region of the interline CCD. The second exposure must, however, be at least as long as the read-out time because the first image must be registered before the second can be shifted to the masked region. This leaves additional time after the imaging light pulse for stray light to be acquired in the image. This means that in our case two dark images must be taken for normalisation, one short to match I_{atoms} , one long to match I_{light} , such that optical density is given by

$$OD = -\ln \left(\frac{I_{\text{atoms}}^s - I_{\text{dark}}^s}{I_{\text{light}}^l - I_{\text{dark}}^l} \right). \quad (2.52)$$

Occasionally images do nonetheless exhibit fringes. Just as with normal shuttering where it is possible to have a smooth image, there are times when the oscillations in the system are unfortunately caught at peak velocity and even double-shuttering produces a fringed image. The extent and regularity of the problem are, however, much reduced, making it a worthwhile technique.

Despite the relatively low quantum efficiency of the camera of $\sim 28\%$ for $\lambda = 767\text{ nm}$, this is higher than the QE of standard interline cameras as the CCD is overlaid with an array of microlenses. These act to direct the light towards the sensor regions so that the light that would otherwise strike the masked storage regions doesn't go to waste, hence the 'QE' in the camera name.

Another neat trick when double-shuttering is available is the ability to image two species near-simultaneously during time-of-flight. This is done by imaging one set of atoms with the first double-shutter exposure and the other set of atoms with the second. The images with no atoms are taken after a pause during read-out of the initial exposures, and finally the background images. This brings back the issue of fringes in the optical density images, but the technique was very useful for trouble-shooting. For example, observing whether there was a correlation between anomalously hot ^{39}K clouds and low ^{87}Rb atom numbers during sympathetic cooling.

2.10.3 Imaging light

All of our samples are imaged with a beam of σ^+ polarised light that is resonant with the $|2, 2\rangle \rightarrow |3, 3\rangle$ cycling transition. A quantisation axis is provided by a weak (few gauss) guide field provided by the transverse bias coils around the science cell. It is important for the imaging beam to be near-uniform as the absorption cross-section varies with intensity, being given by

$$\sigma = \sigma_0 \left(\frac{1}{1 + (2\delta/\Gamma)^2 + I/I_{\text{sat}}} \right), \quad (2.53)$$

as may be deduced from (1.67). The resonant cross-section σ_0 for σ^+ polarised light on the $|2, 2\rangle \rightarrow |3, 3\rangle$ transition is [23]

$$\sigma_0 = \frac{3\lambda^2}{2\pi} . \quad (2.54)$$

Some intensity variation across the beam is inevitable, so it is best to use low intensities ($I \ll I_{\text{sat}}$) when possible¹. It is also sometimes useful (for dense clouds) to detune the beam to deliberately reduce σ by a determined amount. However, for our critical measurements we tune the imaging beam to resonance as small variations in the detuning here have little effect on σ , while its dependence on δ is stronger on the sloped sides of the Lorentzian lineshape.

The light pulse we use for imaging has a relatively high intensity of around $I_{\text{sat}}/3$ and lasts for $80 \mu\text{s}$. This evolved due to the balancing of several factors. The effective pixel size (discussed below) of our CCD, the quantum efficiency, the conversion factor of electrons to counts and the level of read-out noise from the camera require this much light so that the accessible range of OD is not too severely truncated. The optical density can also be limited by any off-resonant light in the imaging beam, causing saturation at a relatively low value. From observations of OD saturation in our images we choose to trust only optical densities below 3.

2.10.4 Objective

To resolve features on the small length-scales of the atomic cloud an objective was designed in our group by Alexandre Dareau. The aim was to create an objective out of cheap standard optical elements, but designed specifically to take into account the spherical aberrations introduced by the wall of our science cell. The natural aberrations of the optical elements are calculated to balance one another (and those of the science cell wall). The design used for horizontal imaging (from which the vast majority of the data in this thesis is taken) was based on one by Wolfgang Alt [75], designed for similar purposes. A full description of the finalised objective may be found in [76], a representation of which is shown (not to scale) in Fig. 2.35. The design has been shown to have a resolution better than $2 \mu\text{m}$. The full set of lenses provides a magnification factor of 2.78, determined by observing the acceleration of the atoms due to gravity in a sequence of absorption images. This magnification factor gives an effective pixel size for our camera of $2.32 \times 2.32 \mu\text{m}^2$, close to the fundamental resolution limit of the objective.

¹When imaging on resonance ($\delta = 0$), a 20% change in I alters σ by $\sim 10\%$ around $I = I_{\text{sat}}$, but only by $\sim 2\%$ around $I = I_{\text{sat}}/10$.

2.10 Absorption Imaging




Figure 2.35: Structure of the objective lens system used for absorption imaging. The four blue optics form the objective which collimates the light from the atoms while compensating for spherical aberrations introduced at the science cell wall. The red lens is a diffraction limited achromatic doublet used for focussing the light to form an image at the camera. The 1” optics are held in a lens tube with fixed spacers to accurately set the required inter-lens separations. The lenses used to acquire the data in this thesis resulted in a magnification factor of 2.78.

2.10.5 Imaging sequence

Usually the atoms leave the trap in the $|1, 1\rangle$ state, so just before imaging a $100\ \mu\text{s}$ pulse of repumping light transfers the atoms to the $|2, 2\rangle$ state with near-perfect efficiency. The atoms are then imaged (in the presence of a guide field of a few gauss) with an $80\ \mu\text{s}$ pulse of light to produce I_{atoms} . With the double shutter technique the next exposure starts only $10\ \mu\text{s}$ later. However, in this short time the atoms are still in the path of the image beam. To get around this issue a strong magnetic field is ramped up to take the atoms far off resonance with the imaging light. This is done rather than changing the frequency of the light with an AOM to avoid possible changes in the light intensity. The field-ramping time means we must leave a $60\ \mu\text{s}$ gap before taking I_{light} with a second $80\ \mu\text{s}$ imaging pulse. The two I_{dark} images are taken over a second later during which the sequence of events just before the first set of images is mimicked to make sure the same amount of stray light appears in these images as in the first pair. The acquired images are then immediately processed into a single OD image by our MATLAB ‘Pixel Machine’ software (designed by Naaman Tammuz).

2.10.6 ‘Fudge-factor’

Assuming a perfect absorption cross-section σ_0 typically causes underestimation of the atom number [72, 77]. The multiplication factor used to compensate for the difference between the atom number as calculated using σ_0 and the true atom number has colloquially become known as the ‘fudge-factor’, roughly equivalent to σ_0/σ . This factor for references [77] and [72] is 4.0 and 1.4 respectively. The reduction in cross-section is usually mainly attributed to imperfect polarisation of the imaging light. However, in our case a significant shift is produced by using a relatively high intensity beam and a long imaging pulse. This results in the cross-section dropping due to both saturation of the cycling transition and the Doppler shifting of the atoms off resonance due to the acceleration of the atoms by photon recoil.

The combined effect can be estimated from the scattering force (1.67) and the dependence of the cross-section on intensity and detuning given in (2.53). Assuming




Figure 2.36: Time-averaged ^{39}K absorption cross-section $\langle\sigma\rangle_t$ for imaging light that is resonant for a stationary atom. The variation is shown with light intensity I and imaging pulse duration τ . The limit $\tau \rightarrow 0$ shows the standard dependence on intensity, while for $I \rightarrow 0$ the ideal cross-section is recovered for all τ . Imaging with this frequency with our I and τ gives $\langle\sigma\rangle_t \approx 0.5$, but this frequency does not give the maximum value of $\langle\sigma\rangle_t$, as discussed in the text.

that initially the cloud of atoms is at rest and the imaging light is resonant leads to the dependence on intensity and imaging pulse duration τ shown in Fig. 2.36 for ^{39}K atoms. This takes into account the change in cross-section and scattering force over time, resulting in a time-averaged cross-section $\langle\sigma\rangle_t$. In this simplified model the atoms are assumed not to drop out of the cycling transition. In this picture the reduction in effective cross-section for long imaging pulses is significant; for our imaging parameters $\langle\sigma\rangle_t/\sigma_0$ is predicted to be about 0.5.

The effect is exaggerated compared to what is seen experimentally as we tune the imaging frequency to give the largest OD . This results in the imaging frequency being tuned slightly above resonance (for stationary atoms) such that as the atoms are accelerated during the imaging pulse they are first Doppler shifted onto resonance before being shifted away again. This is illustrated in Fig. 2.37 where the dependence of σ with time is shown for a beam that starts on resonance and for a beam that is initially tuned above resonance, becoming resonant half-way through the imaging pulse. The latter case produces the optimal effective absorption cross-section. The resultant effective line-shape for this ideal model is shown in Fig. 2.38 for a range of imaging parameters. This illustrates the fact that not only does the effective cross-section decrease for a finite pulse duration and intensity, but the peak frequency of the profile is shifted, the shape is significantly skewed and the line-width increased.

It is not possible of course to directly compare the line-shape of our chosen setup

2.10 Absorption Imaging




Figure 2.37: Variation of ^{39}K absorption cross-section with time when the imaging light is initially resonant (black line) and when it becomes resonant half-way through the pulse duration (red-line). The total pulse duration is $80\ \mu\text{s}$ and the light intensity is $I_{\text{sat}}/3$. The time-averaged value gives the effective cross-section in the $OD \rightarrow 0$ limit.




Figure 2.38: Imaging beam frequency resonances for ^{39}K with σ^+ polarisation on the cycling transition for the Doppler model discussed in the main text. The grey dashed line shows the ideal lineshape. The black, blue and red lines show the lineshapes where imaging period $\tau \rightarrow 0, 40$ and $80\ \mu\text{s}$ respectively for $I/I_{\text{sat}} = 0.1 \rightarrow 0.505$ in steps of 0.015. For $\tau \rightarrow 0$ only the saturation of the transition affects the cross-section. The blue and red series of lines show how increasing imaging pulse duration and intensity distorts the perfect Lorentzian lineshape and reduces the effective cross-section.

with the ideal form, but we can observe the difference between profiles taken with different sets of parameters. To confirm that the above effects manifest themselves in a real system, two series of images were taken of atom clouds created with identical experimental sequences. Within each series the imaging frequency was varied over a range of 8 MHz in 0.5 MHz steps, one set being taken with 40 μs imaging pulses, the other with 80 μs pulses. The series were interleaved during the experiment and the whole process repeated for a total of six sets to reduce possible errors.

The data, shown in Fig. 2.39, confirms the qualitative trend expected for an increase in imaging-pulse duration, and we can fit the data using a model that assumes this Doppler-based effect is the only one affecting the cross-section. The frequency spacing between datapoints was controlled in the experiment (and the pulse duration set). This left I/I_{sat} , absolute frequency and an absolute cross-section scaling factor as free fitting parameters. The two datasets were fitted simultaneously (producing a single value for each of the stated fitting parameters), keeping relative frequency and effective cross-section fixed. The fitted resonance curves are shown as blue and red solid lines in Fig. 2.39, while the predicted resonance curve for a stationary atom (where only the finite I/I_{sat} affects the ideal lineshape) is given by the black solid line. The ideal Lorentzian case is shown as a grey dashed line.

The value $I/I_{\text{sat}} \approx 0.23$ extracted from the fit is inconsistent with the recorded value $I/I_{\text{sat}} \approx 1/3$ (determined by pixel counts on the imaging camera's CCD). The fudge-factor predicted from the fit of ~ 1.4 is also inconsistent with the value 1.6 calculated by observing the critical point of an atomic cloud with minimal interactions. This is not surprising, as imperfect imaging beam polarisation must also be taken into account. This causes the line-shape to be less skewed than for perfect polarisation for a fixed intensity or effective cross-section. A model that includes this effect is unlikely to produce a more meaningful fit to the data due to the over-abundance of free parameters. However, it does conveniently explain the underestimated I/I_{sat} and fudge-factor values.

A time-varying σ actually leads to non-linearities in the conversion of OD to n_{2D} , as what we directly measure in the image is actually the transmission, proportional to $\langle e^{-n_{2D}\sigma} \rangle_t$, rather than the mean optical density, proportional to $\langle \sigma \rangle_t$. The effect is generally small though, with variations of $< 1\%$ for $OD < 1$ in our scenario. However, it is important when measuring number fluctuations as mentioned in [72], which is why they avoid the problem by using a high QE camera. This allows a short low-intensity light pulse during which the atoms on average absorb only 6 photons as opposed to ~ 300 in our procedure. Future experiments in the group may require a high QE camera, but for the experiments covered in this thesis the smooth imaging produced by double-shuttering with suitable cross-section calibration is more than adequate. It is well to be aware, though, of the warping of the line-shape and shift in resonance produced when altering imaging timings and intensities.

2.11 Image Fitting




Figure 2.39: Imaging beam frequency resonance data for ^{39}K . The grey dashed line again shows the ideal Lorentzian lineshape. The blue and red points are experimental data series taken with $40\ \mu\text{s}$ and $80\ \mu\text{s}$ exposures respectively, with statistical error bars shown based on six readings per point. The solid lines show fits using the model discussed in the text, with the black line representing the $0\ \mu\text{s}$ exposure lineshape based on the fitted value for $I/I_{\text{sat}} \approx 0.23$. The data clearly shows the expected skewing of the lineshape to higher frequencies and reduction in effective cross-section.

2.11 Image Fitting

All of our experimental information is garnered from the images produced by absorption imaging after time-of-flight. Most of the time this means extracting the atom number (both condensed and thermal parts) and the temperature. With an appropriate function the number can be calculated from the integral under the fitted distribution, while the temperature can be found from the widths of the distribution.

2.11.1 Thermal distribution

We have already found the in-trap distribution of an ideal Bose gas, and the expanded distribution can be found in a similar way, as described in [78]. Assuming free expansion, a particle can only end up at final position \mathbf{r} if it has momentum $\mathbf{p} = m(\mathbf{r} - \mathbf{r}_0)/t$, where its initial position was \mathbf{r}_0 . The expanded distribution can therefore be found by multiplying the integrand on the right-hand side of (1.6) by $\delta(\mathbf{r} - \mathbf{r}_0 - \mathbf{p}t/m)$ and integrating over both \mathbf{p} and \mathbf{r}_0 (rather than \mathbf{r}). Doing this for a harmonically trapped gas gives

$$n_{\text{T}}(\mathbf{r}, t) = \frac{1}{\lambda_T^3} \prod_{k=1}^3 \left(\frac{1}{1 + \omega_k^2 t^2} \right) g_{3/2} \left\{ \exp \left[\frac{\mu}{k_{\text{B}}T} - \frac{m}{2k_{\text{B}}T} \sum_{k=1}^3 x_k^2 \left(\frac{\omega_k^2}{1 + \omega_k^2 t^2} \right) \right] \right\}. \quad (2.55)$$

For the critical density distribution ($\mu \rightarrow 0$) this can be rewritten as

$$n_c(\mathbf{r}, t) = \frac{n_c(0, t)}{\zeta(3/2)} g_{3/2} \left\{ \exp \left[-\frac{1}{2} \left(\frac{x^2}{R_{T_x}^2(t)} + \frac{y^2}{R_{T_y}^2(t)} + \frac{z^2}{R_{T_z}^2(t)} \right) \right] \right\}, \quad (2.56)$$

where the thermal radii $R_{T_k}(t)$ have been redefined as

$$R_{T_k}(t) = \left(\frac{1 + \omega_k^2 t^2}{m \omega_k^2} k_B T \right)^{1/2}, \quad (2.57)$$

such that they now also apply during time-of-flight. For $t \rightarrow 0$ the in-trap critical density distribution (1.16) is recovered. Imaging the atoms gives us a two-dimensional optical density map, equivalent to integration of the three-dimensional distribution along y :

$$n_{c\ 2D}(\mathbf{r}_{xz}, t) = \int_{-\infty}^{\infty} n_c(\mathbf{r}, t) dy. \quad (2.58)$$

Here \mathbf{r}_{xz} is the radial vector in circular polar coordinates in the x - z plane. By using the sum representation of $g_s(z)$ given in (1.9) it is straightforward to show that

$$n_{c\ 2D}(\mathbf{r}_{xz}, t) = \frac{n_{c\ 2D}(0, t)}{\zeta(2)} g_2 \left\{ \exp \left[-\frac{1}{2} \left(\frac{x^2}{R_{T_x}^2(t)} + \frac{z^2}{R_{T_z}^2(t)} \right) \right] \right\} \quad (2.59)$$

where

$$n_{c\ 2D}(0, t) = \sqrt{2\pi} R_{T_y}(t) \frac{\zeta(2)}{\zeta(3/2)} n_c(0, t). \quad (2.60)$$

By fitting an image with this g_2 function we can obtain $R_{T_x}(t_{\text{TOF}})$ and $R_{T_z}(t_{\text{TOF}})$ which, along with our trapping frequencies, allow us to calculate the temperature of the cloud using (2.57). It is worth noting that fitting full thermal distributions of atoms near the critical point with a gaussian function produces significant errors in temperature measurement as mentioned in [41]. As long as we know the conversion of optical density to number density the integral under the fitted g_2 function also gives the total thermal atom number as

$$N_T = \frac{\zeta(3)}{\zeta(2)} 2\pi R_{T_x}(t) R_{T_z}(t) n_{c\ 2D}(0, t). \quad (2.61)$$

2.11.2 Thomas-Fermi distribution

We have seen that, even for very weak interactions, if the number of condensed atoms N_0 is large that the ideal gaussian distribution for the condensate is warped into a parabolic shape. If we are in this Thomas-Fermi regime it can be shown (as discussed in e.g. [19]) that this form also holds during time-of-flight, such that the distribution is given by

$$n_0(\mathbf{r}, t) = n_0(0, t) \left(1 - \frac{x^2}{R_x^2(t)} - \frac{y^2}{R_y^2(t)} - \frac{z^2}{R_z^2(t)} \right) \quad (2.62)$$

2.12 Computer control

inside the ellipsoid with semi-axes $R_k(t)$, and it is 0 outside this region. Now the Thomas-Fermi radii $R_k(t)$ are time dependent, where at $t = 0$ they take the in-trap values

$$R_k(0) = \sqrt{\frac{2\mu}{m\omega_k^2}}. \quad (2.63)$$

Here μ is the Thomas-Fermi chemical potential given by (1.29). In our experiments the time dependence of these radii is not important as we are only trying to extract the atom number from the fits to the condensate. Of general interest though is the aspect ratio in time-of-flight. Considering for simplicity an axially symmetric trap $\omega_x = \omega_y = \omega_\perp$, we can denote the in-trap aspect ratio of the condensate $\lambda = R_\perp/R_z$. It can be shown analytically [79] that for $\lambda \ll 1$ (i.e. a ‘cigar-shaped’ trap symmetric around z) the aspect ratio for long time-of-flight goes to $2/\pi\lambda \gg 1$, so there is an inversion of aspect ratio during time-of-flight. In general an aspect ratio inversion will be observed for any asymmetric trap. This of course differs from a thermal cloud where the distribution tends towards spherical symmetry with increasing time, and is another signature of Bose-Einstein condensation.

We must integrate (2.62) over y to yield the two-dimensional distribution

$$n_{02D}(\mathbf{r}_{xz}, t) = n_{02D}(0, t) \left(1 - \frac{x^2}{R_x^2(t)} - \frac{z^2}{R_z^2(t)} \right)^{3/2}, \quad (2.64)$$

where

$$n_{02D}(0, t) = \frac{4}{3} R_y(t) n_0(0, t). \quad (2.65)$$

Integrating over the remaining dimensions yields the total atom number

$$N_0 = \frac{1}{5} 2\pi R_x(t) R_z(t) n_{02D}(0, t). \quad (2.66)$$

2.11.3 Fitting routines

The image fitting is carried out by our ‘Analysis GUI’ program, created in *MATLAB* by Naaman Tammuz. Essentially, it applies a set of non-linear least squares fitting routines to the data, with the specifics of each iteration dependent on the parameter we are trying to extract from the image. These are discussed in more detail in the relevant experimental sections. The essential parameters for processing the images (such as the trapping frequencies, apparent pixel size, effective imaging cross-section etc.) are set by the user. The program returns a database of the fit results and calculated quantities for all the images that were selected for processing.

2.12 Computer control

Typically our experimental sequences last around a hundred seconds, while some steps in the sequence need to be timed with microsecond precision. A computer system must

be used to deterministically trigger and control our experimental equipment, with a suitable user interface for simple construction of complex sequences. With the large number of experimental components to control during each run this is a somewhat daunting task. The implementation of such a system in our lab was made much more straightforward thanks to a set of applications created and maintained by Aviv Keshet in Wolfgang Ketterle's group at MIT, which he has made freely available to the scientific community. This software is designed to be compatible with *National Instruments* analogue and digital output hardware, and allows communication to devices via RS-232 and GPIB protocols.

The system allows digital TTL voltage signals to be sent to components such as mechanical shutters which just require an on/off signal, while analogue voltage signals are used to set power-supply outputs, AOM amplitudes and so forth. Items such as the translation stage and RF generators are controlled via RS-232, while the microwave generator uses the GPIB protocol for communication.

2.12.1 Software

The control software is comprised of Cicero, the client user interface, and Atticus, the server. Sequences are constructed in the intuitive Cicero interface, which are then compiled by Atticus into a command sequence that can be passed on to the output hardware. Cicero allows batches of sequences to be run where parameter values are cycled according to user-defined lists, simplifying the acquisition of experimental data sets. The software is naturally attuned to the operation of cold atoms experiments, with nearly all of the functionality we desire being available. The only major feature missing is the ability to allow feedback from system inputs, which would require major alterations and additions to the software and is generally unnecessary.

A significant feature added by Aviv while we were setting up our experiments was the ability to use a variable time-base. Normally a sequence allowing for microsecond resolution with a hundred output channels would essentially require a set of a hundred million commands to be sent for every second of the duration of the experimental sequence. The amount of data that can be stored in the output buffer generated by Atticus limits the duration of a sequence created in this way to a rather modest length, and the large amount of data to be compiled creates substantial delays between sequences. However, the variable time-base feature vastly reduces the buffer size by updating outputs only when something is changing. Trigger signals to change state are generated separately, so when all outputs are constant in the system instead of sending the same set of commands every microsecond, just one command is sent for the entire period. This allows us to keep the microsecond resolution where needed while permitting the sequence length required for our experiments.

As we have made no modifications ourselves to the software there is no need to discuss its use and structure here. The full functionality is covered in the user manuals

2.13 Field and resonance calibration

Part description	Company	Part number	Quantity
Chassis	National Instruments	NI PXIe-1062Q	1
Digital card	National Instruments	NI PXI-6533	1
Digital card	National Instruments	NI PXI-6534	1
Analogue card	National Instruments	NI PXI-6713	2
RS-232 card	IntaShield	IS-400	1
GPIB card	National Instruments	NI GPIB-USB-HS	1
Analogue connector block	National Instruments	NI BNC-2110	2

Table 2.7: Summary of the hardware used for computer control. All items were chosen to be compatible with the Cicero sequence generator software.

provided by Aviv Keshet [80].

2.12.2 Hardware

Having opted to use the Cicero software the choice of hardware was greatly simplified. A desktop computer was chosen to run the software with 4 GB of RAM to cope with the large output buffers. Most of the rest of the hardware is made by *National Instruments* to be compatible with the software. The parts used in our system are listed in Table 2.7.

The basic hardware setup is shown in Fig. 2.40. The software is run on a desktop PC, with the GPIB card and the 4-port RS-232 card directly connected. The analogue and digital cards are connected to the *National Instruments* chassis. When a sequence has been compiled by the Atticus software the sequence data is sent to the chassis while timing data is sent to a separate FPGA board. The FPGA board synthesises a variable frequency clock to produce the timing triggers which are sent directly to one of the analogue cards and shared between the analogue and digital cards. During a run successive groups of signals are sent from the output cards on receiving each timing signal. Break-out boxes to connect the 32 digital channels of each digital card to BNC connectors were made, with opto-couplers installed for isolation purposes. Commercial connector boxes were used for the analogue channels. The output signals are sent on from here to the experiment. As the RS-232 and GPIB cards do not share the timing signal of the digital and analogue cards, there is a lack of complete synchronisation, but precise timing of the RS-232 and GPIB commands are not critical in our setup.

2.13 Field and resonance calibration

The purpose of our experiments is not to investigate the Feshbach resonance itself, rather simply to use it for tuning interactions. However, it was still necessary to

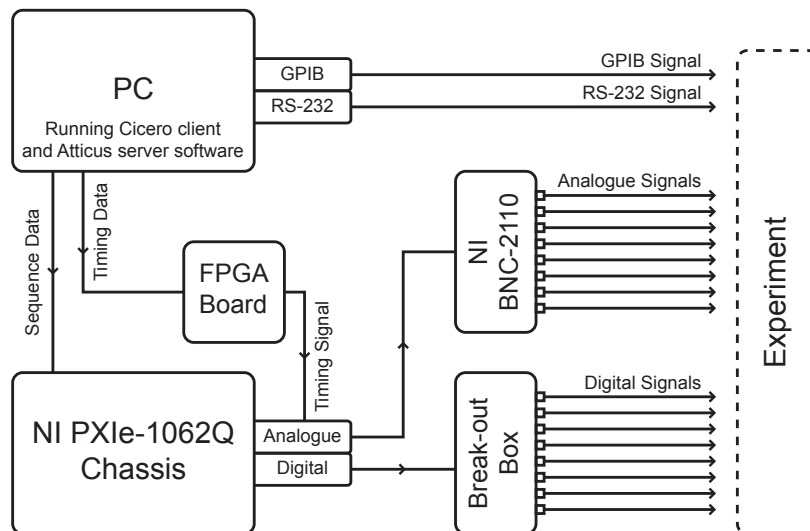


Figure 2.40: Simple schematic of the hardware setup for computer control. In our system there are two analogue and two digital cards connected to the chassis. The analogue cards have eight channels each while the digital cards have thirty-two.

calibrate the field produced by the Feshbach coils for a given electric current, and to check that our own resonance position and width calibrations agreed within errors with the values found during the more stringent calibrations carried out at LENS [6, 12, 81].

2.13.1 Field calibration

In order to calibrate the field produced by our Feshbach coils we looked at state transfer of rubidium in the presence of an RF field. The most precise determination of the ^{39}K $|1, 1\rangle$ resonance position B_∞ to date was carried out in [12], giving a value of 402.50(3) G. We chose this as a natural point to calibrate our current-to-field conversion. Returning to the Breit-Rabi equation (1.33), the splitting between the $|2, 2\rangle$ and $|2, 1\rangle$ states in ^{87}Rb at this field value is 249.444 MHz. An RF field at this frequency was applied to a thermal cloud of rubidium held in the optical trap while the current in the Feshbach coils was swept over a 160 mA range in 200 ms. This was repeated while incrementing the central current in steps of 80 mA. The expectation is that atoms are transferred to the $|2, 1\rangle$ state if 402.50 G lies within the range of the magnetic field sweep. To distinguish between atoms in the $|2, 2\rangle$ and $|2, 1\rangle$ states we spatially separated them by applying a Stern-Gerlach field gradient during time-of-flight. This produced a clear signal, allowing us to determine the current to produce a field of 402.50 G at the atoms to be 152.46(8) A. This gives the conversion factor given in Table 2.3.

2.13.2 Resonance measurement

Our initial calibration of the resonance position relied on the enhancement of three-body losses in the vicinity of a Feshbach resonance. ^{39}K atoms were prepared in the optical trap using a standard sequence to provide a fixed starting number and temperature.

2.13 Field and resonance calibration

For the data-set shown in Fig. 2.41 **a** the field was first swept to 403.4 G (near to the resonance), before sweeping from there to the final value in 20 ms and holding at this field for 100 ms¹.

As expected, there is a dip in final atom number near the expected resonance position, but it is asymmetric and looks to be centred to the attractive side of the resonance position. This was also found in similar experiments carried out by Massimo Inguscio's group at LENS when the first studies of the resonance were made [81]. In their case, fitting a single gaussian profile to the data gave a resonance position of 403.4(7) G, similar to the apparent shift in our data. The asymmetry in the profile can be attributed to the enhancement of three-body losses on the attractive side of the resonance, due to the mechanism described in [82]. The losses are greatest on resonance but decrease more sharply on the repulsive side of the resonance, as can readily be seen in Fig. 1 of [12]. This is compounded by the increase in average density for a cloud with attractive interactions which also enhances three-body recombination. The minimum remaining atom number actually closely coincides with the accepted resonance position, marked in Fig. 2.41 by the dashed line.

We decided to attempt a more accurate determination of the resonance position following a similar method to that found in [12]. The idea is to prepare identical thermal ³⁹K clouds with the field at a fixed initial value. At some point during time-of-flight the field should be switched to its measurement value. If the scattering length is tuned to a large positive value the increase in interaction energy causes a more rapid expansion of the cloud, while if it is tuned to large negative values three-body processes heat the cloud. Both of these processes produce an enhancement of the cloud size after further time-of-flight². This produces a sharper feature than our initial calibration as the density of the cloud after some TOF is greatly reduced, so an effect is only seen for very strong interactions.

Our first attempts at calibration in this manner failed due to the inability to ramp the Feshbach coils quickly and precisely, probably due to the large self-inductance of the coils. We therefore decided to use the vertical bias coils around the science cell to produce a small field offset which could be quickly switched off. A small current producing a field of about 3.5 G was turned on once a cold thermal gas had been produced in the CDT. The current in the Feshbach coils was then ramped to its measurement value, with the opposing bias field keeping the atoms away from resonance. After 4 ms TOF the bias field was switched off quickly, while the current in the Feshbach coils remained constant until being turned off just before imaging. Fig. 2.41 **b** shows the cloud-width measured with varying final field value. The sharp peak is close to the expected value. This measurement gives a resonance position of 402.3(3), in agreement with the pre-

¹We carried out several different calibrations with different initial fields and final hold times, and all produced similarly-shaped loss profiles.

²This effect is useful for calibrating the resonance position, but is detrimental when attempting to accurately determine temperature after TOF expansion, as discussed in §5.5.2.




Figure 2.41: Feshbach resonance calibration. **a** Surviving atom number after 100 ms hold time at the final in-trap field value. The centre of the broad asymmetric resonance feature is shifted to the attractive side of the expected resonance position of 402.50 G (dashed line) due to the enhancement of three-body processes at negative scattering lengths. **b** Cloud width at the end of 18 ms TOF after the field is switched to its final value 4 ms after release from the optical trap. This produces a sharp feature near the expected resonance position. The background rise in width with increasing field is due to the pre-TOF field increasing, being offset by 3.5 G from the measurement value, causing heating before release from the trap. Error bars are statistical.

cisely determined value of 402.50(3). Having satisfactorily reassured ourselves in this manner we chose to use the latter for calibration of the scattering length.

2.13.3 Zero crossing measurement

To independently calibrate for the width of the resonance we looked for the zero crossing position. For attractive interactions the mean-field interaction energy decreases with increasing density, so BECs are unstable against collapse [83, 84]. We therefore created cold atomic clouds of ^{39}K with a small condensed fraction at our standard evaporation scattering length of $135 a_0$, before tuning the field close to the zero-crossing position. Fig. 2.42 shows a typical series of five images where the final field strength is decremented through the crossing from 351.0 down to 349.3 G. The variation in scattering length across the images is less than a single a_0 (see Table 1.1), so the sharp drop in size observed from image **b** to image **c** is attributed to having just passed into the attractive regime. The location of the zero crossing is therefore somewhere in this region,

2.13 Field and resonance calibration




Figure 2.42: Sample images of Bose-Einstein condensates after the Feshbach field has been swept close to the zero-crossing point. Images **a** and **b** are representative of all images seen close to but on the repulsive side of the zero crossing. A sharp drop in size is seen between images **b** and **c**, suggesting the zero crossing lies between the field magnitudes for these two images. The collapse of the condensate is more obvious in images **d** and **e**. The colour bar shows the conversion between optical density and colour, though in these images the optical density saturates at the centre of the BEC due to high atomic column densities.




Figure 2.43: The complete setup on the vacuum chamber table used to reach quantum degeneracy of ^{39}K .

i.e. $B_0 = 350.4(2)\text{ G}$, consistent with $\Delta = -52\text{ G}$. Images **d** and **e** more clearly show the collapse of the condensate.

The system is shown in its completed state in Fig. 2.43.

Chapter 3

Experimental sequence

“ You know my methods, Jeeves. Apply them. ”

– *P. G. Wodehouse (Right Ho, Jeeves)*

Having outlined the techniques involved and some of our particular experimental parameters in the previous chapter, our process for creating Bose-Einstein condensates and the optimisation of experimental stages are discussed, as summarised in [85]. Much of the process is similar to that used to produce the first ^{39}K BEC in the Inguscio group at the University of Florence [7], although our original intention was to carry out sympathetic cooling in a quadrupole trap, as discussed later in the chapter. As the QUIC trap wasn't in the original plan it is not perhaps as well integrated as if we had intended to use this trap from the outset. However, the long trap lifetime it affords leads to very efficient sympathetic cooling, resulting in large ^{39}K condensates in the crossed dipole trap. At the end of the chapter a summary of the experimental stages is provided in the form of Fig. 3.13.

3.1 Laser cooling

The first stage of the experiment is to load both ^{39}K and ^{87}Rb into the MOT from the background vapour provided by the getters. Often the atom sources are only fired once or twice a day to keep the pressure at a usable level. When taking data though we switch to applying a low current 30s burst during every run to keep experimental conditions as stable as possible. The current is applied only while the experiment is being carried out in the science cell so as not to interfere with the loading of the MOT.

3.1.1 Monitoring MOT loading

In order to monitor the loading of the MOT during experimental sequences we use the standard fluorescence detection method. A lens is positioned close to a spare window of the MOT chamber to collect fluorescent light and direct it onto a photodiode, producing a current which is converted to a voltage for display on an oscilloscope. If the MOT

3.1 Laser cooling

parameters (i.e. magnetic field gradient, light frequency and intensity) remain constant the number is roughly proportional to the photodiode current. The absolute number can be gauged from the equation

$$N = \left(\frac{4\pi}{\Omega}\right) \left(\frac{I_{\text{pd}}/R_{\text{pd}}}{\bar{\Gamma}_{\text{scatt}}h\nu}\right) \left(\frac{1}{T}\right), \quad (3.1)$$

where Ω is the solid angle subtended by the aperture of the light-collecting lens, I_{pd} is the photodiode current, R_{pd} is the responsivity of the photodiode (i.e. the conversion factor of incident light power to photodiode current), $\bar{\Gamma}_{\text{scatt}}h\nu$ is the mean fluorescent power of an atom in the MOT, and T is the transmittance of the optics between the atomic cloud and the photodiode. This relationship holds well for low atomic densities, but breaks down when there is a significant amount of photon reabsorption within the cloud. At this point the fluorescence to number conversion becomes sub-linear. Early comparisons with absorption imaging in the MOT suggest that the fluorescence acts well as a rough guide in our system.

The conversion is fairly straightforward for ^{87}Rb as $\bar{\Gamma}_{\text{scatt}}$ is relatively simple to calculate, but the small level splitting in ^{39}K complicates matters [86]. The number can be calibrated by absorption imaging, but most of the time relative day-to-day comparisons are all that we are interested in. It should be noted that at high densities the linear conversion of fluorescence to number is lost due to photon reabsorption in the cloud, which also limits MOT densities and temperatures [87]. Additionally a *Sony* digital video camera (using ‘nightshot’ mode where the infra-red filter is removed) is used to provide constant live footage of the MOT on a monitor. This facilitates observation of the size and position of the trapped cloud for balancing the MOT beams and optimising the optical molasses cooling stage.

3.1.2 MOT optimisation

Number

The capacity of the MOT is essentially limited by the proportion of particles in the background vapour that are capturable to those that are not. The number saturates as the capture and loss rates become equal. Collisions with the background vapour are the main cause for trap losses at low densities, while light-assisted collisions become more important at high densities. Re-absorption of spontaneously emitted photons produces a repulsive force between atoms which limits the cloud density. Full descriptions of MOT dynamics are fairly complex [88], but generally speaking the capacity can be maximised by using high intensity cooling beams with large detunings such that the capture velocity is increased. For bosonic potassium this is slightly complicated by the small hyperfine splitting, and both laser frequencies (as both contribute to the trapping force) must be red-detuned from the entire upper-state manifold [13, 54].

To optimise the number we look at the fluorescence detected at the monitoring




Figure 3.1: Example of ^{39}K MOT number optimisation as discussed in the text, where the cooling light frequency detuning is being varied. These data were taken with the repumping light detuned by 31 MHz.

photodiode. Changing MOT parameters generally alters the fluorescence to number conversion factor, so is not a reliable direct gauge for optimising the number of trapped atoms. However, if an alteration is made quickly with respect to the MOT loading time, one may observe whether the MOT fluorescence consequently grows or shrinks, indicating whether the situation has improved or worsened. This is the standard method for rough optimisation of beam alignment, field gradient and so on. For a more precise comparison we optimised the frequency detunings by letting the MOT saturate for a given frequency detuning before quickly switching to resonance, recording the transient peak fluorescence. An example set of data for a given detuning of the repumping light is shown in Fig. 3.1. We found detunings of -7.6 and -4.6Γ ($\Gamma = \Gamma_{\text{K}} \approx \Gamma_{\text{Rb}} \approx 2\pi \times 6.0\text{ MHz}$) for the cooling and repumping light respectively¹ and $B' = 4\text{ G cm}^{-1}$ (the field gradient along the weak axis) maximised the number of trapped ^{39}K atoms. The total power in the six MOT beams is 310 mW where each beam has a diameter of around 3 cm, resulting in intensities of several times I_{sat} . When operated on its own the potassium MOT saturates at around 1×10^9 atoms.

Rubidium is more simple to work with having a large splitting between the $|F' = 2\rangle$ and $|F' = 3\rangle$ levels. The ^{87}Rb MOT is very robust in terms of number. We use a detuning of -3.1Γ for the cooling light, with the weak repumping light close to resonance. There is a total of 130 mW of cooling light in the cooling beams, and a total of 10 mW of repumping light. The number of atoms in the single-species rubidium MOT saturates at around 5×10^9 .

¹Detunings are measured from the specific transition frequency in question, e.g. for repumping light the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition.

3.1 Laser cooling




Figure 3.2: Photo of the fluorescence produced by the dual-species MOT. The horizontal compensation coils can also be seen, wound around the adjacent MOT chamber windows.

Temperature

The large detunings and low field-gradient used to maximise the amount of potassium in the MOT unfortunately result in high temperatures due to the small damping coefficient and low densities. This is the opposite of what is desired for evaporative/sympathetic cooling, so 40 ms from the end of MOT loading B' is suddenly increased to 6 G cm^{-1} , while the detunings of the cooling and repumping light are jumped to -1.5 and $-4.1 \text{ } \Gamma$ respectively. The duration of this stage is short enough to avoid any noticeable loss of atoms while the temperature is greatly reduced [54]. During this stage the intensity of the repumping light is also reduced slightly. The cooling frequency ends up within the excited state manifold in a more analogous situation to standard MOT setups. The effect is to increase the frictional Doppler force on the trapped atoms, cooling them down.

3.1.3 Dual-species MOT

When both species are simultaneously loaded the peak number of potassium atoms is reduced by a factor of ~ 3 by interspecies light-assisted collisions [89], while the amount of rubidium is only reduced by around 15%. This is not a great problem as we desire a large ratio of rubidium to potassium for sympathetic cooling, though it should be possible to reduce the problem (at the expense of ideal alignment with the magnetic field) by displacing the two MOTs from one another using a ‘push beam’ [90] if required. In order to control the amount of trapped potassium, we initially load rubidium on its own for around 15 s before turning on the potassium light for between 0.5 and 3 s. A picture of the loaded MOT is shown in Fig. 3.2.

3.1.4 Optical molasses

To cool the atoms further before loading the magnetic trap a brief stage of optical molasses is required. At this point, it is particularly important for the six MOT beams to be well balanced so that when the magnetic field is turned off (and the atoms are no longer trapped) the expansion of the cloud is slow and symmetric. The standard technique of matching the MOT position for strong and weak gradients is used to achieve this, while a more quantifiable signal is provided by observing the decay of the fluorescence on the photodiode when suddenly switching off the magnetic field.

During the molasses stage the rubidium cooling light detuning is increased to -4.9Γ and is reduced in intensity to allow low temperatures to be reached by sub-doppler cooling processes. The potassium level-splitting once again complicates matters though. It seems logical that when the level-splitting, Δ , between the upper-state cooling level and its adjacent state is much larger than the natural linewidth Γ that the adjacent level will not interfere with sub-doppler processes, which is indeed the case. More surprisingly, it has been shown that when Δ is smaller than Γ , sub-doppler processes continue to work well, as demonstrated in [91]. However, when $\Delta \sim \Gamma$, as is the case for ^{39}K , the proximity of the neighbouring levels can disrupt the sub-doppler processes, and indeed at the time our work was carried out sub-doppler laser cooling was yet to be observed for ^{39}K .

We optimised the final temperature of ^{39}K within our cooling scheme by dropping the relative repumping light intensity to roughly a tenth of the cooling light intensity for 6 ms of molasses, observing a minimum final temperature in the absence of ^{87}Rb of $\sim 200\ \mu\text{K}$, not far off the Doppler limit $T_{\text{D}} = 145\ \mu\text{K}$. When rubidium is also present the temperature is limited to around $600\ \mu\text{K}$, probably due to a higher starting temperature at the end of the MOT phase. Rubidium is cooled well below the Doppler limit to $\sim 50\ \mu\text{K}$ in spite of the presence of potassium.

Recent findings [92, 93] have shown that it is in fact possible to reach the sub-Doppler region for bosonic potassium isotopes with laser cooling. This is achieved by controlling the natural pumping of atoms into a dark state, minimising heating due to the reabsorption of spontaneously emitted photons [94]. Our current method is not very dissimilar to the successful sub-Doppler schemes, the main development of [93] being a clever ramping of frequency and intensity to address the majority of the atoms¹. This should be relatively easy to implement, and is a major step on the way to directly cooling ^{39}K . This progress may induce a boom in ^{39}K experiments due to its more favourable properties.

¹It is possible that sub-Doppler cooling is achieved in our system for some fraction of the atoms, producing a bimodal velocity distribution like that shown in Fig. 3 (b) of [93]. In this picture, low velocity atoms are cooled further ('the rich get richer') while energetic atoms are unaffected or possibly even heated.

3.2 Magnetic trapping and transport

At the end of the optical molasses stage the repumping light stays on for both species for a millisecond after the cooling light is switched off to ensure all the atoms are in the $|F = 2\rangle$ state. At this point the pumping guide field is switched to a few gauss, before a $200\ \mu\text{s}$ burst of intense pumping light is applied. The circularly polarised pumping light is resonant with the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition for rubidium, while it is tuned slightly above ($\sim 2\Gamma$) for potassium due to the close proximity of the $|F' = 1\rangle$ level. Pumping to the $|F, m_F\rangle = |2, 2\rangle$ state is about 80% efficient for rubidium, 65% for potassium.

To capture the atoms, the quadrupole field produced by the transport coils is switched on at $B' = 20\ \text{G cm}^{-1}$ (weak axis) and then ramped linearly in 50 ms to $80\ \text{G cm}^{-1}$ to adiabatically compress the cloud. During this ramp up the track is energised to prepare for the transport. The coils then travel the 74 cm to the science cell in under two seconds as described in §2.5.3.

During transport the atomic cloud has to pass through a minimum aperture 10 mm in diameter. Any atoms that happen to touch the walls as the coils pass along the system adhere to the vacuum chamber and are lost. As only the most energetic atoms can reach the walls, the effect is analagous to an inefficient stage of evaporative cooling. Atoms with an energy greater than $\sim 2.7\ \text{mK}$ are lost, so it is mainly potassium that is affected due to its higher temperature. Around 40% of the potassium atoms are removed via this process. Losses due to Majorana spin-flips are minimal during this stage as the temperature of the atoms is still relatively high.

Once the atoms have reached the science cell there is a one second wait to ensure the track settles to the correct location within its precision bounds before the track is de-energised. During this pause one of the microwave generators is turned on ready to start evaporating rubidium. Due to some thermalisation between the two species (and the loss of high energy potassium atoms during transport) the rubidium temperature has increased slightly and the potassium temperature decreased by this time.

3.3 Sympathetic cooling

3.3.1 Choice of trapping potential

Because of our earlier experiences of direct cooling in a quadrupole trap (see Appendix B) we initially planned to use this potential for sympathetic cooling. The fact that potassium has to end up in a purely optical potential meant that, instead of using an optical plug to suppress Majorana spin-flips, we intended to offset the CDT from the magnetic zero, as done in [55]. In fact this method allowed us to apply a 10 s evaporation sweep when cooling rubidium alone, before transfer to the CDT for further forced evaporation, resulting in BECs of 5×10^5 atoms.

Unfortunately, we were unsuccessful in using this method for sympathetic cooling. This was due in part to the relatively small K-Rb interspecies background scattering length¹, which precludes the use of such short evaporation sweeps. Longer sweeps are also prohibited by inelastic collisions with ^{87}Rb atoms in the $|2, 1\rangle$ state. These are created by Majorana spin-flips, in spite of the presence of the ODT. If we were able to use a stronger quadrupole field the evaporation rate could be increased, perhaps making this method more viable.

While attempting to find the best solution for sympathetic cooling we were also able to condense ^{87}Rb in an optically-plugged quadrupole trap [4], a time-averaged orbiting potential trap [59], and directly in the QUIC trap. With our system we found the QUIC trap to be by far the most effective for sympathetic cooling because of the elimination of $|2, 1\rangle$ ^{87}Rb atoms, as discussed below. Before turning on the QUIC trap a short microwave evaporation sweep is applied in the quadrupole trap to aid the loading efficiency of both species. We use a two-second linear sweep from 6909.68 to 6879.68 MHz, corresponding to removing rubidium atoms with energies above 2.4 to 1.4 mK.

3.3.2 Transfer to the QUIC trap

Transferring the atoms to the QUIC trap is not quite so simple as just ramping up the Ioffe and bias coils to their final values. Initially, an increase in just the Ioffe coil current draws the atoms towards the coil due to the translation of the zero-field point, which can bring the cloud into contact with the wall of the vacuum chamber. Once the magnetic zero disappears, any further increase in Ioffe current pushes the atoms back away from the coil, as the point at which B' is cancelled by the Ioffe field gradient gets further from the Ioffe coil. Additionally, it is important to ensure that the trap does not become too shallow at any stage.

Taking these factors into consideration resulted in the ramping scheme shown in Fig. 3.3. The bias coils are used to counteract the initial tendency of the trap to move towards the Ioffe coil (**i-iii**) before the current in the quadrupole coils is reduced slightly to increase the height of the trap near the Ioffe coil. Once the magnetic zero is lost (**iv**) the coil currents are slowly ramped to their final values (**v**). The minimum total trap depth during this process is $\simeq 1$ mK, occurring at stage **iii**. Applying this method results in 4×10^8 rubidium atoms and up to 4×10^7 potassium atoms in the QUIC trap.

3.3.3 Cooling in the QUIC trap

With sympathetic cooling there is a natural tradeoff between the size of the sympathetically cooled load and the lowest achievable temperature. Our particular aim is to

¹ $a_{\text{K-Rb}} = 36 a_0$ [65], to be compared with the rubidium intraspecies scattering length $a_{\text{Rb}} = 99 a_0$ [95].

3.3 Sympathetic cooling

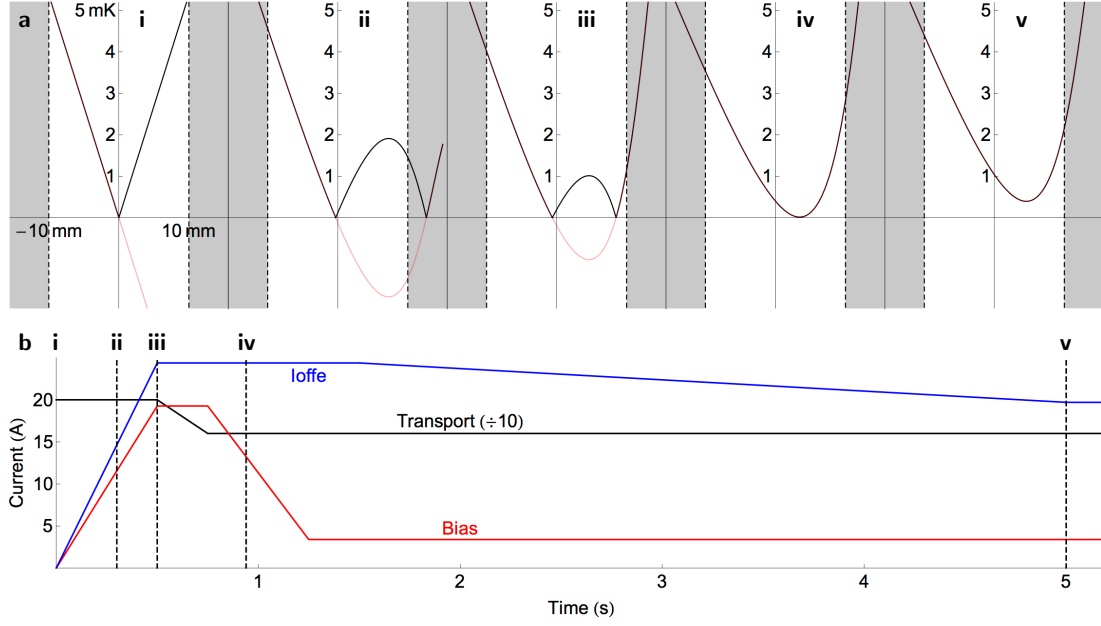


Figure 3.3: Transfer of atoms from the quadrupole trap to the QUIC trap. **a** Trapping potential along the Ioffe coil axis at five stages during transfer. In each plot the solid black line gives the potential $\propto |B|$ in terms of thermal energy, while the faint red line illustrates the variation of the magnetic field B . The dashed lines represent the sides of the science cell. **b** The current through the transport (black), Ioffe (blue) and bias (red) coils during transfer to the QUIC. **i** Initially current is running through just the transport coils forming a quadrupole trap. **ii** The currents in the Ioffe and bias coils are ramped up, reducing the total trap depth. **iii** The trap depth reaches its minimum value of ~ 1 mK. **iv** The point at which a QUIC trap is formed. **v** The currents reach their final values, ready for sympathetic cooling.

cool the largest possible number of ^{39}K atoms to a temperature of about $5\ \mu\text{K}$. This is mainly motivated by the $73\ \mu\text{m}$ difference in gravitational sags (discussed in §2.6) for the two species in our QUIC trap. This corresponds to the vertical thermal radii for the two species $R_T^{\text{Rb}} \approx R_T^{\text{K}}$ at a temperature of about $4\ \mu\text{K}$. Even for very small potassium loads sympathetic cooling will become inefficient at around this temperature due to the reduced spatial overlap of the clouds [96]. It is also motivated by the loading efficiency of the CDT from the QUIC, though this is a slightly cyclic argument as the CDT parameters themselves were also adjusted for this purpose.

The substantial field offset at the trap minimum of the QUIC trap eliminates the detrimental spin-flips of the quadrupole trap, but during evaporation a finite $|2, 1\rangle$ state population for rubidium is still observed. One possible mechanism that would cause this is illustrated in Fig. 3.4. Atoms that are successfully transferred to the $|1, 1\rangle$ state accelerate away from the trapped atoms, but the evaporating microwave frequency becomes resonant with the $|1, 1\rangle \rightarrow |2, 1\rangle$ transition at a particular field strength. There is some finite chance that atoms will be transferred to the $|2, 1\rangle$ state and once again become trapped. This mechanism is suggested in [97].

The field offset of the QUIC trap is once again beneficial however, as a second microwave frequency can be used to transfer any $|2, 1\rangle$ atoms to the $|1, 0\rangle$ state (illustrated




Figure 3.4: Possible mechanism by which rubidium atoms are transferred to the $|F, m_F\rangle = |2, 1\rangle$ state, and how they are removed. The variation of the energy levels with distance from the trap centre (dashed line) is shown (the gap between manifolds is not to scale). Microwave radiation used for evaporation to the $|1, 1\rangle$ state is also resonant with the $|1, 1\rangle \rightarrow |2, 1\rangle$ transition further from the trap centre (black arrows). The finite field offset at the trap centre allows a second microwave frequency to be applied to permanently remove atoms from the trap (red arrow).

by the red arrow in Fig. 3.4) which will then leave the trap permanently. We transmit this second microwave frequency with the same antenna, combining the signals before amplification. A saw-tooth sweep is used for the clean-up microwave signal, sweeping between 6838 and 6844 MHz with a 5 s period. This has no effect on the atoms in the $|2, 2\rangle$ state.

3.3.4 Sympathetic cooling results

The results of our sympathetic cooling experiments are summarised in Fig. 3.5. Firstly, as shown in Fig. 3.5 **a**, we characterised the rubidium cooling reservoir. Allowing for the fact that, for the same trap depth, the temperature of rubidium will inevitably be higher in the presence of potassium, we set the target temperature for rubidium to be $\sim 3 \mu\text{K}$. At the start of evaporation the remaining rubidium atoms have a temperature $\simeq 100 \mu\text{K}$ while the potassium atoms are at $\simeq 200 \mu\text{K}$. We therefore use a starting microwave frequency of 6873.68 MHz, corresponding to a thermal energy of $\sim 900 \mu\text{K}$, providing a high starting η . For efficient sympathetic cooling we evaporate the rubidium with a long exponential sweep lasting 56 s^{-1} with a decay constant of 14s, to a final

¹Shorter sweeps are found to be just as effective for cooling rubidium on its own due to its high rethermalisation rate, but are not sufficient when potassium is also present.

3.3 Sympathetic cooling

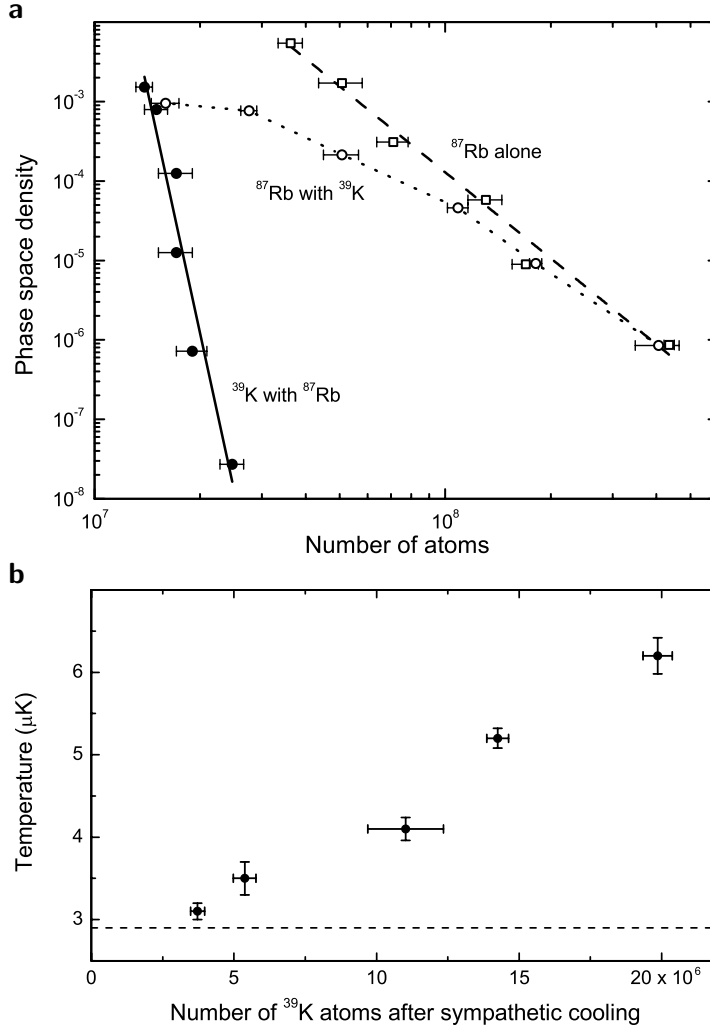


Figure 3.5: Evaporative and sympathetic cooling in the QUIC trap. **a** Cooling trajectories in single- (squares) and two-species (circles) experiments. Open and solid markers represent ^{87}Rb and ^{39}K respectively. The two fits show the difference between the steep trajectory for sympathetic cooling (solid line, $\gamma_{\text{sym}} = 20$) and that for efficient evaporative cooling (dashed line, $\gamma = 3.6$). Vertical error bars lie within the markers. **b** Final temperature reached by sympathetic cooling for different sizes of the ^{39}K load. The dashed line shows the temperature reached by ^{87}Rb for the same evaporative cooling sweep in the absence of ^{39}K .

frequency of 6847.24 MHz. This is followed by a 6 s period of plane evaporation where the frequency is held at its final value.

Such a long evaporation sweep results in a very high cooling efficiency with $\gamma = 3.6$, where γ was defined in equation (2.33). Also shown in Fig. 3.5 **a** is the phase-space density evolution of a typical ^{39}K cloud used in our experiments. The near-vertical evaporation trajectory is the hallmark of effective sympathetic cooling. In this instance the cloud reaches a final temperature of about $5 \mu\text{K}$ while the atom number is reduced by less than a factor of two, from approximately 25×10^6 to 14×10^6 . This gives an effective cooling efficiency $\gamma_{\text{sym}} \approx 20$. Around a 20% reduction in N_{K} is expected over the evaporation period due to the trap lifetime, so the observed atom loss shows evidence for the presence of other inelastic processes, but the effect is not too large.




Figure 3.6: Sympathetic cooling in the QUIC trap. The field of view is 3.2×2.4 mm, the OD scale is given on the right. Images of ^{87}Rb (top) and ^{39}K (bottom) are taken at three points through the sympathetic cooling process after a fixed TOF. The potassium atom number does not decrease much, but the cloud is clearly cooled. The saturation of OD in the rubidium images is caused by imaging such large atomic samples on resonance.

Also shown in the same panel is the effect that the potassium load has on the cooling trajectory of ^{87}Rb . The path starts to deviate at the point at which the load becomes a significant fraction of the total atom number ($N_{\text{Rb}}/N_{\text{K}} \lesssim 10$), bending to almost coincide with the potassium cooling path towards the end of evaporation. This convergence also suggests that with this particular potassium load we are reaching the limits of the effectiveness of sympathetic cooling. Further evaporation past the point at which the heat capacities of the two species becomes comparable will have little benefit for the further cooling of potassium.

Fig. 3.5 b shows the temperature achieved at the end of sympathetic cooling for various final amounts of potassium in the QUIC trap. For low N_{K} the temperature tends to the $2.9 \mu\text{K}$ attained by rubidium when cooled alone. The quantity of potassium that we choose to load into the MOT is determined by the amount that can be loaded into the CDT. This in turn depends on the temperature of the atoms being transferred from the QUIC trap, as is discussed below. A sequence of images taken at different points through the sympathetic cooling process is shown in Fig. 3.6, highlighting the increase in phase-space density of ^{39}K with minimal atom loss.

3.4 Quantum degeneracy in the CDT

3.4.1 Cooling methods in the dipole trap

Once the atoms have been transferred to an optical trap they are no longer in a state-dependent potential. This gives us the opportunity to bring Feshbach resonances into play by transferring the atoms to a suitable state and applying a uniform magnetic field. We are thus provided with two clear options: continue with sympathetic cooling, or remove the remaining rubidium and directly evaporate potassium. Sympathetic cooling

3.4 Quantum degeneracy in the CDT




Figure 3.7: Illustration of equipotentials for a cut through the bisection point of a crossed optical dipole trap. The beams are horizontal and are perpendicular to one another, with the cut made vertically such that one beam is entering the page, the other is parallel to it. The lighter the colour the higher the total potential. The field of view is $2w_0$ square. **a** High-optical-intensity regime, where energetic atoms are lost along the beams. **b** Decreasing the intensity leads to a scenario where atoms may be lost along the beams or from the bottom of the trap. **c** Low-optical-intensity regime, where energetic atoms are lost from the bottom of the trap. **d** The trap eventually vanishes with further reduction of the intensity.

has the added enticement here of an interspecies Feshbach resonance that may be used to increase the cross-thermalisation rate [98]. Up until this point our experimental process has, in essence, closely followed that of the Inguscio group [7], where the former method is successfully applied. Given theirs is the only system previous to ours to have successfully condensed ^{39}K , it might seem sensible to bow to experience and continue to follow their course. This might also seem logical given the high efficiency observed in the sympathetic cooling stage in the QUIC trap. However, as might be deduced from the fact that we aimed to use up all of the rubidium in the previous step, we choose to directly cool potassium. The considerations that informed this decision are discussed below.

The main factor is the competition between the optical and gravitational potentials, shown in Fig. 3.7. In the high-optical-intensity regime the effect of gravity is negligible and energetic atoms are lost along the beams¹, while in the low-optical-intensity regime gravity distorts the Gaussian optical potential such that atoms fall from the bottom of the trap. The red detuning of the beam produces a stronger optical potential for rubidium than for potassium. The effect of gravity is also stronger for the more massive rubidium atoms. This means that above a certain laser power the trap depth will be greater for rubidium, while below this crossover point it will be greater for potassium. Sympathetic cooling is only possible in the latter regime, i.e. where mainly rubidium atoms are evaporated.

In the limit that the distance from the centre of the trap to the wall of the science cell is small compared to the Rayleigh length of our focussed trapping beams we may approximate the combined optical and gravitational potential for our trap geometry as

$$U(x, y, z) = mgw_0 \left[\frac{z}{w_0} - \frac{\alpha}{2} e^{-2(x^2+z^2)/w_0^2} - \frac{\alpha}{2} e^{-2(y^2+z^2)/w_0^2} \right], \quad (3.2)$$

¹The Rayleigh length, $z_R = \pi w_0^2/\lambda$, for typical focal beam waists, w_0 , and laser wavelengths, λ , is on the order of centimetres, so trapping along the beam axis is negligible.

where α is the ratio A/mgw_0 , A being the peak optical potential shift defined in §2.8.3. Here α has simply been substituted into equation (2.44). By switching to dimensionless variables $z' = z/w_0$ and $U' = U/mgw_0$ this reduces to

$$U'(0, 0, z') = z' - \alpha e^{-2z'^2} \quad (3.3)$$

and

$$U'(0, \infty, z') = z' - \frac{\alpha}{2} e^{-2z'^2} \quad (3.4)$$

for the relevant cases we will consider here. We can now find some useful system parameters. For example, by differentiating (3.3) twice with respect to z' we find that the minimum gradient of the total potential is at $z' = -1/2$. By setting the gradient at this point to be zero we obtain the value of α at which the atoms are no longer trapped, $\alpha_0 = \sqrt{e}/2 \approx 0.82$.

To find the crossover point where the trap depth is equal for our two atomic species we first note that $\alpha_{\text{Rb}} = (A_{\text{Rb}}m_{\text{K}}/A_{\text{K}}m_{\text{Rb}})\alpha_{\text{K}}$. The trap depth at this point for potassium is the difference in potential between the local minima z'_1 and z'_2 of (3.4) and (3.3) respectively, while for rubidium it is the difference between the local maximum z'_3 and local minimum z'_4 of (3.3). Casting the equations in terms of α_{K} , the condition that we are at extrema gives us

$$\frac{1}{2}z'_1 e^{-2z_1'^2} = z'_2 e^{-2z_2'^2} = \frac{A_{\text{Rb}}}{A_{\text{K}}} \frac{m_{\text{K}}}{m_{\text{Rb}}} z'_3 e^{-2z_3'^2} = \frac{A_{\text{Rb}}}{A_{\text{K}}} \frac{m_{\text{K}}}{m_{\text{Rb}}} z'_4 e^{-2z_4'^2} = -\frac{1}{4\alpha_{\text{K}}}, \quad (3.5)$$

while equating the trap depths at this point means we must also satisfy

$$\begin{aligned} & \left(z'_1 - \frac{\alpha_{\text{K}}}{2} e^{-2z_1'^2} \right) - \left(z'_2 - \alpha_{\text{K}} e^{-2z_2'^2} \right) \\ &= \frac{m_{\text{Rb}}}{m_{\text{K}}} \left[\left(z'_3 - \frac{A_{\text{Rb}}}{A_{\text{K}}} \frac{m_{\text{K}}}{m_{\text{Rb}}} \alpha_{\text{K}} e^{-2z_3'^2} \right) - \left(z'_4 - \frac{A_{\text{Rb}}}{A_{\text{K}}} \frac{m_{\text{K}}}{m_{\text{Rb}}} \alpha_{\text{K}} e^{-2z_4'^2} \right) \right]. \end{aligned} \quad (3.6)$$

From (2.44) we may use

$$\frac{A_{\text{Rb}}}{A_{\text{K}}} = \frac{\Gamma_{\text{Rb}}}{\Gamma_{\text{K}}} \frac{\lambda_{\text{Rb}}^4}{\lambda_{\text{K}}^4} \left(\frac{\lambda_{\text{K}}^2 - \lambda^2}{\lambda_{\text{Rb}}^2 - \lambda^2} \right) \approx 1.172 \quad (3.7)$$

for $\lambda = 1070$ nm. Solving (3.5) and (3.6) simultaneously, we find the value of α_{K} at which the trap depth is equal for the two species to be ≈ 4.01 (with the corresponding z'_i values).

Combining this result with the values found for z'_1 and z'_2 we find that the trap depth at the crossover point is given by

$$U_c \approx 1.97 m_{\text{K}} g w_0. \quad (3.8)$$

Considering this in terms of thermal energy, we find $U_c/w_0 \approx 0.091 \mu\text{K} \mu\text{m}^{-1}$. To have a reasonable starting η ($\gtrsim 5$) in the trap with an initial temperature of $6 \mu\text{K}$ we need

3.4 Quantum degeneracy in the CDT




Figure 3.8: Displacement of the trap minima for ^{39}K (solid blue line) and ^{87}Rb (solid red line) with respect to the centre of the optical potential for varying α_K . Displacement is given in terms of dimensionless distance $z' = z/w_0$. The atoms become untrapped at $z' = -0.5$, i.e. when the minimum has been displaced by half the beam waist. The dashed lines represent a single vertical thermal radius either side of each trap minimum for $\eta = 8$, showing that the clouds no longer significantly overlap for $\alpha_K \lesssim 1.86$.

a trap depth of at least $30 \mu\text{K}$. To also have a trap in which sympathetic cooling may be used straight from loading we must therefore use a beam waist of at least $30/0.091 = 330 \mu\text{m}$. However, that would require 40 W trapping beams to produce the required depth, far greater than the power available to us with our laser.

Also worth noting is the fact that, as in the QUIC trap, the difference in the gravitational sag between the two species prevents sympathetic cooling past a certain point. Fig. 3.8 shows how the trap minima for the two species vary with α_K , with the corresponding extent of each cloud to a thermal radius R_T , assuming $\eta = 8$. Sympathetic cooling ceases to work when the thermal radii no longer overlap, below $\alpha_K \approx 1.86$. Thus sympathetic cooling is only possible in the range $1.86 \lesssim \alpha_K \lesssim 4.01$. Fig. 3.9 shows a summary of these results for our chosen waist of $140 \mu\text{m}$ and laser wavelength of 1070 nm . The $140 \mu\text{m}$ waist is the largest we could use to have a $30 \mu\text{K}$ trap depth when using the full available optical power.

It is relatively simple to extend the above to calculate equivalent parameters for other species in crossed dipole traps, and it is hoped that the above discussion is useful for anyone considering sympathetic cooling between any two species in such a configuration. Note that the calculations are valid for horizontal beams only, but that the bisection angle of the two beams does not matter. The upshot is that we must abandon any hopes of sympathetically cooling ^{39}K in the CDT, and instead use the Feshbach resonance in the $|1, 1\rangle$ state to allow efficient forced evaporative cooling.




Figure 3.9: CDT depth versus laser power for both ^{39}K (solid line) and ^{87}Rb (dashed line) for a beam waist of $140\ \mu\text{m}$ and wavelength of $1070\ \text{nm}$. The ‘kinks’ in each plot (at **d** and **f**) occur at the boundary between two regimes where different mechanisms limit the trap depth. At high power the atoms are least confined along the beams, whereas at low power they are least confined vertically due to the distortion of the trap due to gravity (see Fig. 3.7). The table gives the values of α_K and corresponding power per beam P at six points of interest. **a** and **b** are the points at which ^{39}K and ^{87}Rb fall from the trap respectively, **d** and **f** the respective kinks already mentioned, while **c** is the lower limit of sympathetic cooling of ^{39}K (see Fig. 3.8) and **e** is the upper limit, i.e. the crossover point discussed in the text.

3.4.2 Transfer to the CDT

The centre of the QUIC trap during standard operation is offset from the centre of the science cell towards the Ioffe coil by several millimetres. Unfortunately, the coil mounts around the science cell limit the CDT position to a small region near the centre of the science cell, so the QUIC trap must be moved before transfer to the CDT can proceed. The current is reduced in the quadrupole coils and increased in the Ioffe and bias coils over the course of one second, reducing the tightness of the trap slightly, but repositioning it at the science cell centre. Sympathetic cooling is not carried out in this position as the shallowness of the trap would require the use of a longer evaporation sweep and the Ioffe coil would soon overheat with the increased steady current.

Once the QUIC trap is in position the CDT power is increased linearly to its maximum value over the course of 1 s. The currents in the three QUIC coils are ramped down by a factor of two in the subsequent second before being turned off abruptly, leaving the atoms trapped in the CDT, which at this point has a depth of $\sim 30\ \mu\text{K}$ for potassium and $\sim 35\ \mu\text{K}$ for rubidium (see Fig. 3.9).

The number of ^{39}K atoms that are transferred to the CDT as a function of the number in the QUIC is shown in Fig. 3.10. A constant transfer efficiency of $\sim 60\%$ is observed for samples with a temperature of up to $6\ \mu\text{K}$ (equivalent to $\eta \sim 5$), above which the efficiency rapidly drops off. This constant factor is due to the imperfect

3.4 Quantum degeneracy in the CDT




Figure 3.10: Transfer efficiency from QUIC to CDT. The numbers at each point are the temperatures in the CDT (QUIC).

overlap between the CDT and QUIC potentials. Given the cigar-like (prolate spheroid) shape of the QUIC potential and the squashed-sphere (oblate spheroid) shape of the CDT potential this mismatch is difficult to overcome, although using a large beam waist helps. The beam waist should be compared with twice the thermal radius of the cloud at the end of sympathetic cooling in the QUIC trap, roughly 160 (500) μm in the radial (axial) directions. For comparison, the loading efficiency for a 100 μm beam waist was found to be limited to $\sim 20\%$, saturating at around a third of the available CDT power.

In practice we load 8×10^6 potassium atoms into the CDT from the 13×10^6 that are left in the QUIC after sympathetic cooling. Any remaining rubidium atoms are removed from the trap to prevent the occurrence of spin-exchange collisions during the state transfer of the ^{39}K atoms from the $|2, 2\rangle$ to the $|1, 1\rangle$ state. This is accomplished with a strong 2 ms pulse of resonant imaging light, which does not noticeably affect the trapped potassium atoms.

3.4.3 Evaporation in the CDT

All of the atoms are transferred to the $|1, 1\rangle$ state by a 35 ms Landau-Zener sweep of the bias magnetic field while applying RF radiation at 469.3 MHz, as described in §2.9. To increase the thermalisation rate of the potassium we ramp up the current in the Feshbach coils in 5 ms to produce a uniform magnetic field of 393.3 G, resulting in a scattering length of $135 a_0$. The ramping time is kept short to avoid lingering near the narrow ($\Delta = -0.47$ G) Feshbach resonance at 25.9 G [81].

Forced evaporation is then carried out with a six-second exponential sweep of the CDT beam power to its final value with a decay constant of 1.4 s, typically followed by another second of plane evaporation. The beam power is only required to span a

single order of magnitude (see Fig. 3.9), so is easily stabilised with our PID feed-back loop. The condensation point is reached for a beam power of ~ 1.5 W, while further reduction increases the condensed fraction to create quasi-pure BECs.

3.4.4 Number calibration

In order to accurately determine the number of atoms in the condensate we need to calibrate the imaging absorption cross section as mentioned in §2.10.6. This requires a theoretical knowledge of the critical number N_c , which for an ideal gas in a harmonic trap is given by (1.17). We can tune the scattering length to a small value so that the shift of N_c due to interactions (discussed in Chapter 5) is well quantified within mean-field theory. The finite-size shift [99] (of around 2%) can also be taken into account. However, the combined optical and gravitational potential of the CDT is by no means perfectly parabolic. Firstly, the shape of the trap is different (or ‘anharmonic’), and secondly the trap depth is not infinite. To account for this we must use both the true expression for the trapping potential in our calculations and a truncated momentum distribution.

Going back to the semi-classical approximation of §1.2.1 we again set $\mu = 0$ in equation (1.6), but now use

$$\epsilon(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + V(\mathbf{r}) \quad (3.9)$$

where $V(\mathbf{r})$ is given by adding an offset to the potential given in (2.44) such that at the trap minimum (situated at $x = y = 0, z = z_0$) the potential is zero. This results in

$$V(\mathbf{r}) = mg(z - z_0) - \frac{A}{2} \left[e^{-2(x^2+z^2)/w_0^2} + e^{-2(y^2+z^2)/w_0^2} - 2e^{-2z_0^2/w_0^2} \right]. \quad (3.10)$$

Making the substitutions

$$q = \frac{p^2}{2mk_B T} \quad \text{and} \quad \gamma = \frac{V}{k_B T} \quad (3.11)$$

in (1.6) and integrating over both q and space gives the critical number as

$$N_c = \frac{1}{\lambda_T^{3/2}} \frac{2}{\sqrt{\pi}} \iiint_0^{q_{\max}} \frac{q^{1/2}}{e^q e^{\gamma(\mathbf{r})} - 1} dq d\mathbf{r} \quad (3.12)$$

within the bounding surface $\gamma(\mathbf{r}) = V_{\max}/k_B T = \eta$. V_{\max} is the trap depth, such that η is defined as for evaporative cooling. Assuming the cloud is collisionally transparent, the maximum momentum for a trapped atom at point \mathbf{r} is given by

$$\frac{p_{\max}^2}{2m} = V_{\max} - V(\mathbf{r}), \quad (3.13)$$

or equivalently

$$q_{\max} = \eta - \gamma(\mathbf{r}). \quad (3.14)$$

3.4 Quantum degeneracy in the CDT




Figure 3.11: The anharmonic shift in critical number for a set of typical experimental parameters. Black points are calculated with the momentum cutoff, red points represent the upper bound where $q_{\max} \rightarrow \infty$. The true potential has a larger volume than the extrapolated harmonic potential so this tends to increase the critical number. The momentum cutoff becomes important below $\eta \approx 10$ where the two lines diverge. As T tends to zero the critical number tends to that of the harmonic case as expected.

Setting this as the upper q limit we can now numerically integrate (3.12) and compare the critical number with that given by the ideal harmonic trap in (1.17).

In this comparison, the harmonic trap extrapolated from the quadratic terms of the Taylor expansion around the trap minimum is considered. The effect of the altered shape of the trap is to increase the trap volume and therefore increase the critical number, while the momentum truncation reduces the available states and therefore also N_c . The truncation only becomes important when the thermal energy per atom becomes a significant fraction of the trap depth, at an η of $\lesssim 10$. In reality, the expression for q_{\max} given above is not exact as collisions do occur within the cloud. The calculated value of the ‘anharmonic shift’ is therefore a lower bound, and an upper bound can be found by setting $q_{\max} = \infty$. However, under typical experimental conditions the values are similar and we are closer to the collisionless regime.

The resultant anharmonic shift of N_c for typical experimental conditions is positive and in the range 5-10%. The value depends strongly on the temperature of the cloud, while the shape of the trap itself also varies depending on the power in the dipole beams. The shift is therefore calculated for specific sets of experimental parameters as required. It is shown in Fig. 3.11 for a typical trap geometry over a range of temperatures. As expected, as $T \rightarrow 0$ the critical number tends to the value for the ideal harmonic trap.

The freshly determined N_c can be compared with the total optical density of an absorption image of a weakly interacting gas at the critical point, allowing calibration of the imaging absorption cross-section. The total OD can be found either by summing the OD values of the individual pixels or by fitting the image with the g_2 function (2.59). The latter method also allows the temperature of the cloud to be determined,




Figure 3.12: Condensation of ^{39}K atoms. From left to right the number of atoms in the BEC varies from 1.5×10^5 to 4.2×10^5 , while the condensate fraction grows from 20 % to 80 %. The bottom panels show the line densities of the clouds, obtained by integrating the images along the vertical direction, emphasising the signature bimodal distribution of a condensate and thermal gas. Images are taken after 20 ms time-of-flight for a fixed exposure duration. The changing aspect ratio of the condensate reflects the increasing importance of the gravitational potential at low CDT powers.

which is required to calculate N_c in the first place. Our initial investigations showed that the cross-section was reduced by factors of 1.9 ± 0.3 and 1.5 ± 0.3 for ^{39}K and ^{87}Rb respectively. These values are liable to change over time, with alterations in the imaging polarisation for example. A more careful calibration, carried out when investigating very weakly interacting ^{39}K clouds during the saturation investigations of Chapter 4, provided an updated value of 1.6 which has remained stable ever since.

Condensate and thermal atom numbers are extracted from absorption images where both are present using a bimodal fit, combining the g_2 function with the two-dimensional Thomas-Fermi distribution function (2.64). Typically all fit parameters (e.g. widths, amplitudes, centres) are allowed to vary within user-defined upper and lower bounds. These fits are used to assess the results of the final stage of evaporative cooling. Samples are imaged as described in §2.10, usually after around 20 ms TOF. This is near the limit of the TOF available to us for horizontal imaging as the imaging light passes through the centre of the Ioffe coil, such that further TOF would cause the atoms to fall out of the path of the imaging beam. Once the raw images have been acquired they are immediately processed into *OD* images, ready for fitting.

3.4.5 Condensation

The condensation point is typically reached with $\sim 10^6$ atoms at around 400 nK, at which point $\bar{\omega} \sim 2\pi \times 85$ Hz. Further reduction of the beam power results in quasi-pure ^{39}K condensates of over 4×10^5 atoms. Sample shots for various condensed fraction are shown in Fig. 3.12. These are the largest ^{39}K condensates to have been reported

3.4 Quantum degeneracy in the CDT

to date, representing a significant improvement on past work; the first ^{39}K condensates reported contained 3×10^4 atoms [7], later being increased to around 10^5 in the same system [12, 100]. Though condensate size isn't the be all and end all for a useful cold atoms system it is useful to be able to experiment on large samples, a basic example being the small fractional error on N_c which is useful for the experiments discussed in Chapters 4 and 5. The final atom numbers owe much to the high efficiency of sympathetic cooling in the QUIC trap.

We have created condensates by evaporation with scattering lengths of up to $\sim 350 a_0$, but the largest and most stable condensates are produced for scattering lengths of around $100 a_0$, similar to that of ^{87}Rb . In the same system we can directly cool rubidium to produce condensates of up to 8×10^5 atoms when the alignment of the CDT is optimised for the transfer of rubidium from the QUIC trap.

3.4.6 Summary

The experimental sequence for producing ^{39}K condensates is summarised in Fig. 3.13. Typical atom numbers and temperatures are specified at several stages of the process.


3.4 Quantum degeneracy in the CDT

a Magneto-optical trap loaded from vapour of K and Rb. Followed by 40 ms near-resonant MOT to reduce K temperature.

$$N_K \sim 2 \times 10^8, T_K \sim 1-2 \text{ mK}$$

$$N_{Rb} \sim 4 \times 10^9, T_{Rb} \sim 300 \mu\text{K}$$

$t \approx 15 \text{ s}$




b Magnetic field removed for short ‘optical molasses’ stage to increase phase-space density.

$$T_K \sim 600 \mu\text{K}$$

$$T_{Rb} \sim 50 \mu\text{K}$$

$t \approx 6 \text{ ms}$




c Optically pump atoms into $|F, m_F\rangle = |2, 2\rangle$ hyperfine ground state.

$$\text{K pumping efficiency} \sim 65\%$$


$$\text{Rb pumping efficiency} \sim 80\%$$

$t \approx 200 \mu\text{s}$



d Ramp up quadrupole coil for purely magnetic trapping, then transport atoms to the science cell using the translation stage.

$t \approx 2 \text{ s}$




e Load atoms into QUIC magnetic trap by ramping up Ioffe and anti-bias coils.

$$N_K \sim 2.5 \times 10^7, T_K \sim 200 \mu\text{K}$$

$$N_{Rb} \sim 4 \times 10^8, T_{Rb} \sim 100 \mu\text{K}$$

$t \approx 5 \text{ s}$



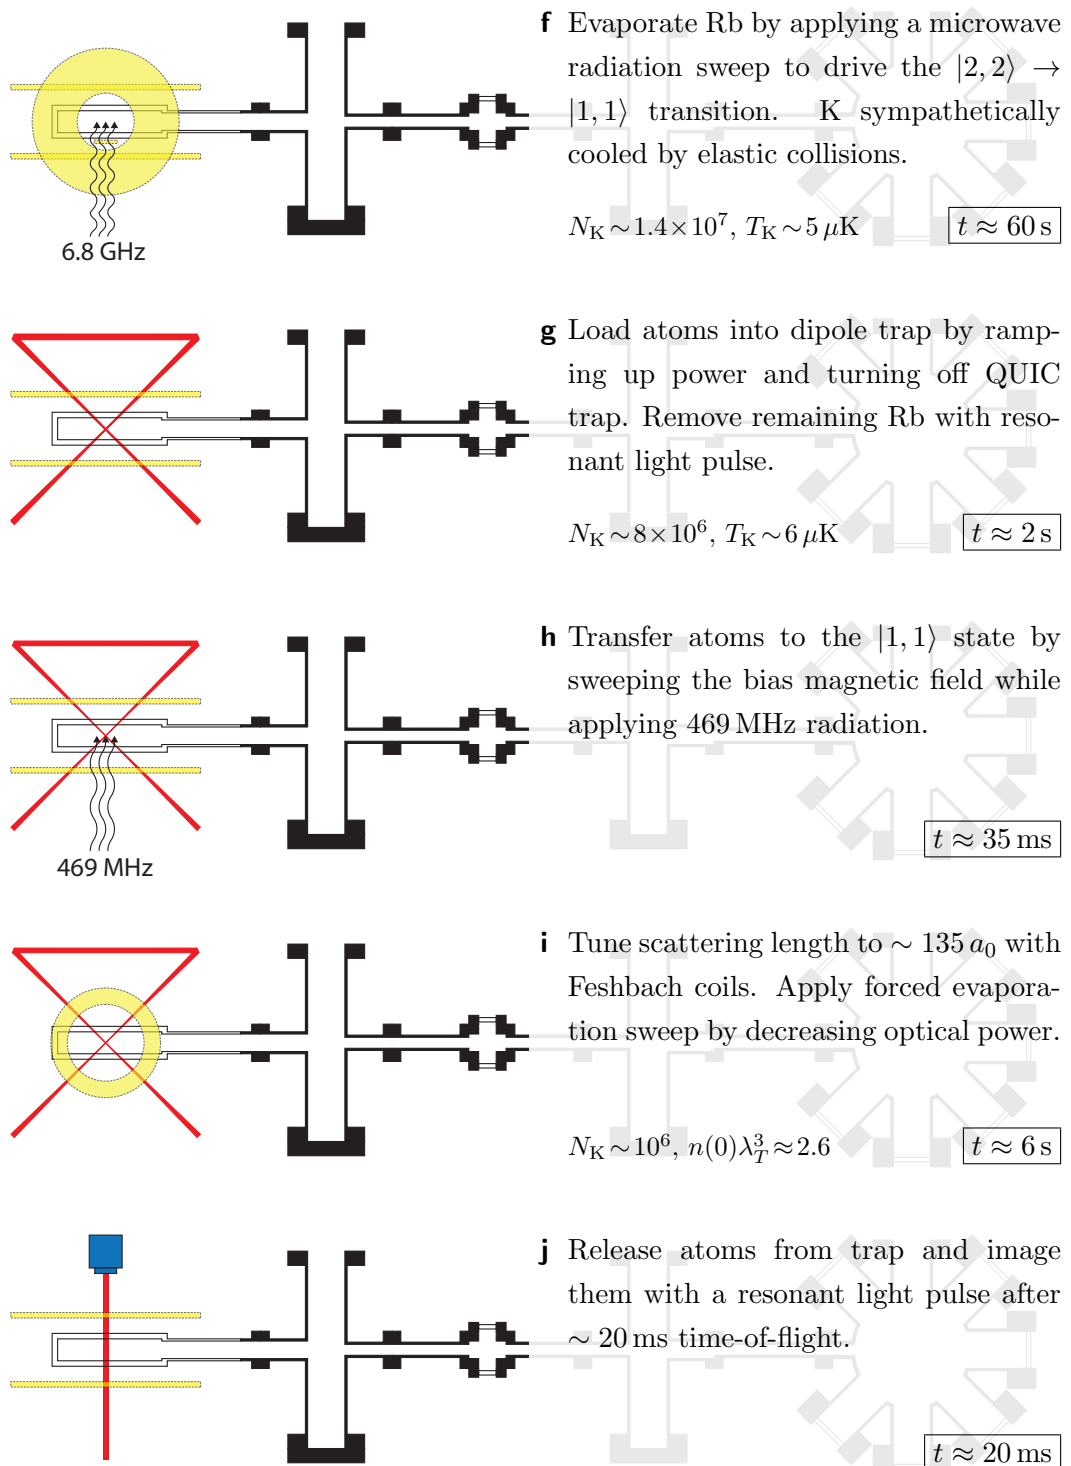


Figure 3.13: Summary of the main experimental stages to produce quantum degenerate ^{39}K . A brief description of each stage accompanies each illustration, along with typical parameters at the end of the stage. The approximate step duration t is framed in each caption.

Chapter 4

Lack of saturation in a Bose gas

“ The Right Hon. was a tubby little chap who looked as if he had been poured into his clothes and had forgotten to say “when”. ”

– P. G. Wodehouse (*Very Good, Jeeves*)

4.1 Saturation: another approach

The phenomenon of Bose-Einstein condensation is inextricably linked to the concept of saturation; that is, saturation of the available excited states that the atoms may occupy. If the number of identical bosons in the system exceeds the number of available excited states, the only place for surplus bosons to go is the ground state. At this critical point the excited states are said to be saturated, heralding the onset of macroscopic ground state occupation - the BEC.

How saturation comes about for an ideal gas in a three-dimensional harmonic trap was shown in Chapter 1 by considering the semi-classical distribution. This description has many uses as has already been made apparent (hence its early introduction). However, a simpler semi-classical concept may be used to clarify under what conditions a BEC may occur. We concern ourselves with the single-particle energy states of the system, but instead of considering the individual quantised states we ‘smooth out’ the fine details by considering a continuous density of states, allowing us to perform an integral rather than the more difficult sum. Of course, this approximation does not take account of the ground state, which is again treated as a special case.

To illustrate the concept we may consider once more the three-dimensional harmonic oscillator, precisely as treated in §2.1.1 of [29]. In this case the potential is given by

$$V(\mathbf{r}) = \frac{1}{2}m (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2) , \quad (4.1)$$

having the familiar quantised energy levels

$$\epsilon(n_x, n_y, n_z) = (n_x + \frac{1}{2})\hbar\omega_x + (n_y + \frac{1}{2})\hbar\omega_y + (n_z + \frac{1}{2})\hbar\omega_z . \quad (4.2)$$

4.1 Saturation: another approach

Here the n_k are positive integers. At this point we make the semi-classical approximation (recalling as previously stated that it is only good for $k_B T \gg \hbar\omega$) by introducing the continuous variables $\epsilon_k = \hbar\omega_k n_k$ and neglecting the zero-point motion (the inaccuracy introduced by this simplification is addressed in §5.2.1). The total number of excited states $G(\epsilon)$ up to an energy ϵ is then given by

$$G(\epsilon) = \frac{1}{\hbar^3 \omega_x \omega_y \omega_z} \int_0^\epsilon d\epsilon_x \int_0^{\epsilon - \epsilon_x} d\epsilon_y \int_0^{\epsilon - \epsilon_x - \epsilon_y} d\epsilon_z = \frac{\epsilon^3}{6\hbar^3 \omega_x \omega_y \omega_z}, \quad (4.3)$$

from which the density of states $g(\epsilon)$ naturally follows as

$$g(\epsilon) = \frac{dG(\epsilon)}{d\epsilon} = \frac{\epsilon^2}{2\hbar^3 \omega_x \omega_y \omega_z}. \quad (4.4)$$

In order to calculate the total number of available excited states N_T , we must multiply the density of states by the Bose occupation factor (or Bose-Einstein probability distribution function) $f(\epsilon)$, reproduced here for convenience:

$$f(\epsilon) = \frac{1}{e^{(\epsilon - \mu)/k_B T} - 1}, \quad (4.5)$$

before integrating the product over all energy space to give

$$N_T = \int_0^\infty g(\epsilon) f(\epsilon) d\epsilon. \quad (4.6)$$

By inserting (4.4) and (4.5) into (4.6) and taking the limit $\mu \rightarrow 0$ we retrieve the critical atom number N_c^0 for a given temperature T as

$$N_c^0 = \zeta(3) \left(\frac{k_B T}{\hbar\bar{\omega}} \right)^3, \quad (4.7)$$

exactly replicating (1.17), the result obtained via the semi-classical distribution.

To extend this description to other trap geometries we consider a generalised version of (4.1) describing a d -dimensional power-law potential,

$$V(\mathbf{r}) = \sum_{k=1}^d \frac{\hbar\omega_k}{2} \left| \frac{r_k}{a_k} \right|^{n_k}, \quad \text{where } a_k = \sqrt{\frac{\hbar}{m\omega_k}} \quad (4.8)$$

and the potential exponent along each axis k is given by n_k . It is clear that (4.1) is recovered simply by setting $d = 3$ and each $n_k = 2$. It is shown in [101] and [102] that in general the density of states for such a system follows a simple power law,

$$g(\epsilon) = \frac{C_s \epsilon^{s-1}}{\Gamma(s)}, \quad (4.9)$$

where

$$s = \frac{d}{2} + \sum_{k=1}^d \frac{1}{n_k} \quad (4.10)$$

and

$$C_s = \frac{2^s}{\pi^{d/2}} \prod_{k=1}^d \frac{\Gamma(1 + 1/n_k)}{(\hbar\omega_k)^{1/2+1/n_k}} . \quad (4.11)$$

So, for example, the three-dimensional harmonic oscillator with $d = 3$ and $n_k = 2$ gives an s of 3 and hence the density of states is proportional to ϵ^2 , as confirmed by (4.4). The corresponding value of C_s may be read from (4.4) (recalling that $\Gamma(3) = 2$) or calculated from (4.11). So for the commonly investigated three-dimensional scenarios of a harmonic trap and a hard-walled box we find, respectively,

$$C_s = \begin{cases} \left(\frac{1}{\hbar\bar{\omega}}\right)^3 & \text{for } d = 3, n_k = 2, s = 3 \\ \left(\frac{m}{2\pi}\right)^{3/2} \frac{V}{\hbar^3} & \text{for } d = 3, n_k = \infty, s = 3/2 . \end{cases} \quad (4.12)$$

In the latter equation V is the volume of the box, simply given by

$$V = (2\bar{a})^3 \quad \text{where } \bar{a} = (a_1 a_2 a_3)^{1/3} . \quad (4.13)$$

So where does this leave us with the conditions for saturation? The simple conclusion is that if the total number of available excited states is a finite value (however large or small) then it must be at least theoretically possible for a BEC to exist. We must therefore evaluate (4.6) using the density of states for the general power-law potential (4.9). Employing the substitution $x = \epsilon/k_B T$ with $\mu \rightarrow 0$, this evaluates to

$$N_T = \frac{C_s (k_B T)^s}{\Gamma(s)} \int_0^\infty \frac{x^{s-1}}{e^x - 1} dx \quad (4.14)$$

$$= C_s (k_B T)^s \sum_{p=1}^\infty \frac{1}{p^s} , \quad (4.15)$$

where (1.9) has been used. The sum in (4.15) is the well-known hyperharmonic series (a generalisation of the divergent harmonic series) which converges to $\zeta(s)$ *only* for $s > 1$, and is divergent otherwise¹. This is illustrated by Fig. 4.1, where the integrand of equation (4.14) is shown for various values of s . Therefore, macroscopic occupation of the ground state only occurs for $s \leq 1$ at $T = 0$, while the critical number for $s > 1$ is given by

$$N_c^0 = C_s (k_B T)^s \zeta(s) . \quad (4.16)$$

This leads directly to a simple expression for condensed fraction,

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c^0}\right)^s , \quad (4.17)$$

a relationship that is often used for temperature calibration. The conclusions of the above discussion are simple but far-reaching; BECs can occur in any three-dimensional trap, but not in a uniform two-dimensional system, while in one dimension the trap

¹This is easily shown using the integral test for convergence.

4.2 Context




Figure 4.1: Integrand of equation (4.14) for $s = 0 \rightarrow 9/2$ in steps of $1/2$. Lines representing values of s under which the area is infinite ($s = 0, 1/2, 1$) are coloured red, the others black. The area under these curves is $\Gamma(s)\zeta(s)$. Though the value of the function tends to infinity for $s = 3/2$, the bounded area is finite.

must be stronger than a harmonic trap (i.e. $n_x < 2$).

This standard view of Bose-Einstein condensation, as presented in statistical physics textbooks (see e.g. [19, 29, 30, 103]) and indeed in §1.2 of this thesis, is naturally introduced along the lines of the original predictions of Einstein. In this picture the number of ideal bosons in excited states, N_T , reaches a limiting value, N_c , after which any additional particles must occupy the ground state, forming a Bose-Einstein condensate. This comes about because the available excited states become saturated. The number of particles occupying the ground state N_0 is microscopic until this saturation point is reached, resulting in the simple occupation dependence on total number N shown in Fig. 4.2.

4.2 Context

In reality we can never produce the exact conditions prescribed by Einstein. This ideal scenario requires non-interacting atoms in thermal equilibrium, but interactions are necessary for thermal equilibrium to exist. In fact, because of this the ideal statistical phase transition proposed by Einstein has never been experimentally confirmed. However, weakly interacting Bose gases in harmonic traps are generally believed to closely portray this picture. It is accepted that the presence of weak repulsive interactions produces slight modifications, such as the small shift in T_c discussed in Chapter 5, but that the essential description holds true.

As such, a saturated thermal component is often heralded as a signature of Bose-Einstein condensation, as it was for the groundbreaking first observations of condensa-

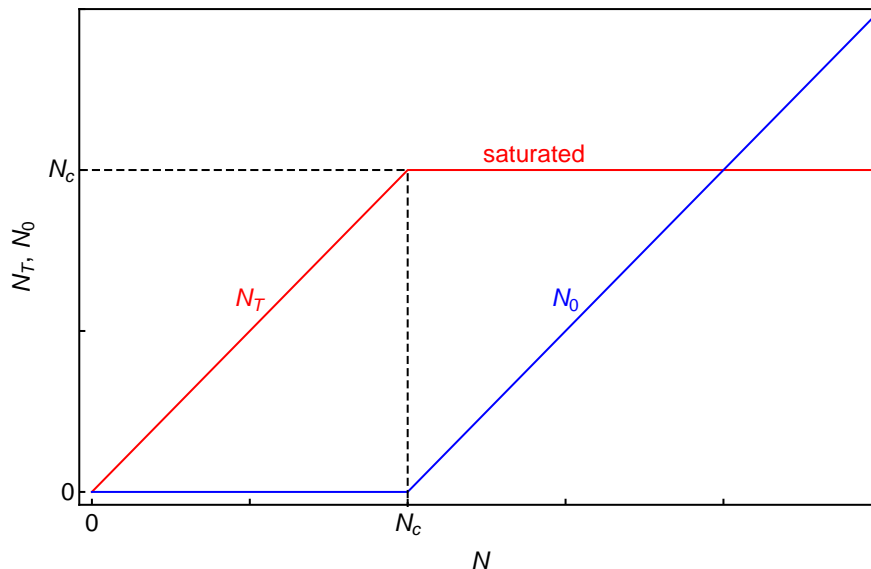


Figure 4.2: State populations of an ideal Bose gas. At fixed temperature, all the atoms accumulate in the thermal component until all available excited states are filled, after which all further atoms populate the ground state.

tion in a solid-state system [104] and of photons in a microcavity [105]¹. An attempt to closely follow Einstein’s idea of a phase transition brought about at constant temperature while increasing the atom number was also believed to directly show saturation behaviour [106]².

However, strong deviations from the ideal case have already been shown to exist. The T^3 dependence of N_c in a harmonic trap leads to the expression

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c^0} \right)^3 \quad (4.18)$$

for the condensed fraction in an ideal system - equation (4.17) with $s = 3$. Investigations carried out shortly after the first demonstrations of condensation in dilute atomic gases seemed to show consistency with (4.18) [107, 108]. More recently though it has been shown that negative shifts in the detected condensed fraction below T_c of $\sim 20\%$ (larger in relative terms) can be seen in a standard setup [109]. This is far from a small correction, and directly corresponds to non-saturation of the gas. Also, while studying the critical point of a harmonically trapped two-dimensional Bose gas at ENS [110], the authors compared their findings with a three-dimensional gas [111]. They were surprised to see a value of dN_0/dN substantially lower than unity past the condensation point.

¹However, Fig. 3 of [104] and Fig. 2 of [105] show growing thermal components as a condensate is formed.

²In this reference, a spinor gas was studied in which condensation was brought about in one spin state by transferring atoms from two others in which condensates were already present. The thermal atoms therefore experienced a potential already modified by the existing condensates, probably obscuring any saturation effects due to the emerging condensate.

4.3 Experimental procedure

The purpose of the investigation described in this chapter is twofold: firstly to show just how large an effect repulsive interactions have on the saturation picture of Bose-Einstein condensation, and secondly by studying a range of interaction strengths (at which thermal equilibrium is ensured) to extrapolate to the non-interacting limit to see if the saturation picture is recovered. This is made possible by our ability to decouple the scattering length from other parameters, a feature which is of course not available in many systems.

Much of the confusion surrounding saturation stems from a frequent lack of differentiation between saturation of atom number and saturation of available excited states. In an ideal gas these terms are equivalent, but when repulsive interactions are included they are not. Considering a mean-field description, the presence of the dense condensate leads to a flattening of the effective potential seen by the thermal atoms such that the trapping volume is increased, i.e. more excited states become available. In this picture a condensate does indeed form because of saturation of the excited states, but as atoms are added at constant temperature the thermal atom number must carry on increasing along with condensate number. We take ‘saturation’ to refer to atom number, and therefore ‘non-saturation’ to refer to deviations from the ideal theory.

4.3 Experimental procedure

The actual experimental concept is a simple one. Partially condensed clouds of ^{39}K are created in the CDT as described in Chapter 3, using the Feshbach resonance to tune the scattering length to $135 a_0$ for the final stage of evaporative cooling. At this point the optical power is ramped to a final fixed value, after which the scattering length is tuned to its measurement value (Fig. 4.3). The sample is then held in the trap for a variable amount of time, up to several tens of seconds. During this period atoms are gradually lost from the trap due to collisions with the background vapour, photon scattering from the optical trap and three-body recombination. The trap depth for each series is chosen such that the heating associated with these processes is cancelled by residual cooling due to plane evaporation, producing a near-constant temperature (Fig. 4.4). Consequently the trapping frequencies vary between datasets, lying in the range $\bar{\omega} = 2\pi \times 60\text{-}80$ Hz. Elastic scattering allows the remaining atoms to rethermalise at a far greater rate than particles are lost, and an equilibrium distribution between the thermal and condensed components is maintained. The cloud is then released from the trap and imaged after ~ 20 ms time-of-flight to allow number and temperature data to be extracted.

4.3.1 Image fitting for saturation data

Condensate population N_0 and thermal population N_T are extracted using a bimodal fit to the OD images, while the total number N is found separately by a direct summation




Figure 4.3: Scheme of the relevant final stages of the experimental sequence. The solid lines show the time dependence of the CDT beam power and Feshbach field for a sample sequence, while the dashed lines show the range of values covered by the complete set of experimental series. At the end of the exponential evaporative cooling ramp in the CDT the power is held for one second before it is ramped to its final value in half a second. The Feshbach field is then ramped to its measurement value, after which the atom number is allowed to decay over a varying period before time-of-flight imaging.




Figure 4.4: Variation of system parameters with trap hold time for the sample data series shown in Fig. 4.5 and considered throughout this chapter. **a** Total atom number in the trap N versus hold time. The data is fitted with a double exponential to guide the eye. **b** The corresponding temperature T . The temperature shows no drift for $N > N_c$ and has a standard deviation of only 3 nK.

4.3 Experimental procedure

over the density distribution. The value of $N_0 + N_T$ was found to be within 0.5% of N for every usable image. A separate g_2 fit is used to find the temperature, as interactions affect the ideal thermal distribution function. A region around the centre of the distribution is excluded from the fit as interaction effects are strongest in the high density regions and around the condensate [78, 108]. We choose to exclude the atoms within one thermal radius of the centre having found this to give stable temperature measurements after similar studies to those found in §II.B of [73]. A variety of data was fitted using a range of exclusion radii, and a plateau region was observed where temperature measurements remained constant. Excluding too large a region produces noisy data due to the loss of information, so the chosen radius was the minimum that safely lay within the plateau. The BEC lies within the exclusion region for all our images. The bimodal fit nevertheless gives accurate number information because of the quantity of free parameters.

4.3.2 Data processing

A total of 18 data series were taken with ^{39}K at temperatures in the range 115-284 nK and scattering lengths in the range 40-356 a_0 . Each series consists of up to ~ 100 data points, so to speed up processing images were automatically removed according to a set of criteria. First any blank images and those where the fit is clearly poorly matched to the density distribution were removed manually. The remaining images were then masked if the total peak optical density exceeded 3, the detected N_0 was lower than a nominal value, and if the condensate or thermal cloud widths lay further than several standard deviations from their averages over the entire series.

Though the temperature is fairly stable over each series as previously discussed, we apply small corrections to account for the remaining drifts. Instead of assigning each image its individually measured temperature for this correction, we use a two-stage fitting process based on image index and N_0 . The latter method is used because of the likelihood that any residual temperature fluctuations recorded are due to the fitting process rather than real effects. The image index designates the ordinal position of each image in the series in terms of the order in which they were acquired. This is needed as the temperatures of clouds produced by the system tend to decay slightly over time¹. N_0 and image index do not scale together as the hold times for each series were taken in a random order in an attempt to reduce systematic errors.

The temperature assignment works as follows. The dependence on index tends to be strongest, so firstly the measured temperature T_{meas} is plotted against index i . An exponential decay is used to fit the data, resulting in $T_i(i)$. We then plot $T_{\text{meas}}/T_i(i)$ against N_0 to reveal to what degree the temperature shifts from the index correction with condensate size. A linear fit is used to give the factor $f(N_0)$. Each image is then assigned a temperature $T'(N_0, i) = f(N_0) \times T_i(i)$. The component atom numbers are

¹This is most likely due to the gradual heating of optical components in the system.




Figure 4.5: Lack of saturation of a Bose gas for a sample series taken with $a = 135 a_0$ and $T = 177$ nK. N_T (red points) and N_0 (blue points) against total number N . The solid lines show the predictions for a saturated gas (as displayed in Fig. 4.2). The critical number N_c is marked by the vertical dashed line.

scaled with T' to the series temperature T , where T is calculated by plotting the raw temperature against N_0 and using a linear fit to give the $N_0 \rightarrow 0$ limit. The form of the scaling is based on a theoretical model as discussed later.

4.3.3 Results

A sample dataset taken with $a = 135 a_0$ and $T = 177$ nK is compared with the results for an ideal gas in Fig. 4.5. The difference between the two is conspicuous; as the total atom number is increased from the critical number of $\sim 200,000$ to 450,000 only about half of the additional atoms form a condensate.

The full set of 18 data series is shown in Fig. 4.6, where N_0 is plotted against the number of excess atoms above the critical point, $N - N_c$. The solid line shows the situation for a saturated gas, $N_0 = N - N_c$. Deviation from the ideal case is clear for every set of data, with the effect becoming more pronounced as both T and a are increased. The dependence of N_0 with N is also clearly not linear. To start interpreting these findings more clearly we look to mean-field theory.

4.4 Hartree-Fock Approximation

The simplest way to include interaction effects within a theoretical construct is by use of the Hartree-Fock approximation [112]. Thermal atoms are treated as a non-interacting gas with density $n_T(\mathbf{r})$ confined by the combination of the trapping potential $V(\mathbf{r})$ and the self-consistent interaction potential $2g(n_T(\mathbf{r}) + n_0(\mathbf{r}))$. The situation may be simplified by taking advantage of the small density of the thermal component. If

4.4 Hartree-Fock Approximation




Figure 4.6: Deviation from the saturation picture for a variety of interaction strengths and temperatures. We plot N_0 versus $N - N_c$ where N_c is experimentally determined for each of the 18 series. The temperature and scattering length for each series is encoded in the colour of the markers and are given in the accompanying legend. The solid line marks the prediction for a saturated gas, $N_0 = N - N_c$.

the effect of thermal atoms on the condensate is neglected then $n_0(\mathbf{r})$ is given by the Thomas-Fermi distribution of equation (1.26). If the mean-field energy of the thermal atoms $2gn_T(\mathbf{r})$ is also neglected¹ the resulting effective potential is simply given by

$$V_{\text{eff}}(\mathbf{r}) = V(\mathbf{r}) + 2gn_0(\mathbf{r}) = |V(\mathbf{r}) - \mu_0| + \mu_0 , \quad (4.19)$$

where the mean-field chemical potential μ_0 is given by (1.29) and the $T \rightarrow 0$ limit is indicated. μ_0 is thus the relevant energy scale parameterising the effects of interactions. This potential has a three-dimensional ‘Mexican-hat’ shape as shown in Fig. 4.7. The number of thermal atoms in the presence of N_0 condensed atoms can then be estimated using the semi-classical approximation. Equation (1.6) is reproduced here for convenience:

$$N_T = \frac{1}{(2\pi\hbar)^3} \iint \frac{1}{e^{\epsilon(\mathbf{r},\mathbf{p})-\mu}/k_B T} - 1} d\mathbf{p} d\mathbf{r} . \quad (4.20)$$

By setting

$$\epsilon(\mathbf{p}, \mathbf{r}) = \frac{p^2}{2m} + V_{\text{eff}}(\mathbf{r}) , \quad (4.21)$$

and noting that the chemical potential is μ_0 we can perform the momentum integration of (4.20) to obtain

$$n_T(\mathbf{r}) = \frac{1}{\lambda_T^3} g_{3/2} \left[e^{(\mu_0 - V_{\text{eff}}(\mathbf{r}))/k_B T} \right] , \quad (4.22)$$

just as (1.15) was found in §1.2.2.

¹Making this approximation means we lose the small mean-field shift in T_c and we find $N_c = N_c^0$.




Figure 4.7: Effective mean-field potential V_{eff} for $\alpha_0 = \mu_0/k_B T = 0.1$ (solid line) and $\alpha_0 = 0$ (dashed line), equivalent to the bare parabolic trapping potential. In practice most of the thermal atoms lie outside the condensate, where $V > \mu_0$ and $V_{\text{eff}} = V$.

When the temperature of the system is relatively high (compared with μ_0/k_B) the majority of thermal atoms lie outside the condensate in the region where $V_{\text{eff}}(\mathbf{r}) > \mu_0$ and $V_{\text{eff}}(\mathbf{r}) = V(\mathbf{r})$. It is therefore reasonable to approximate the full effective potential as the bare trapping potential and consider only the region outside the condensate. This does not mean that the effect of interactions may be neglected, as the chemical potential has a value that differs substantially from the ideal value. Substituting the harmonic form of $V(\mathbf{r})$ into (4.22) and switching to dimensionless variables

$$\alpha_0 = \frac{\mu_0}{k_B T} \quad \text{and} \quad R = \frac{r}{R_T}, \quad \text{where} \quad R_T = \sqrt{\frac{k_B T}{m\omega^2}} \quad (4.23)$$

we can write

$$N_T \approx 4\pi \left(\frac{R_T}{\lambda_T} \right)^3 \int_{\sqrt{2\alpha_0}}^{\infty} R^2 g_{3/2} \left[e^{-(R^2/2 - \alpha_0)} \right] dR. \quad (4.24)$$

Here it is sensible to introduce the variable Q , where

$$\frac{Q^2}{2} = \frac{R^2}{2} - \alpha_0 \quad (4.25)$$

to rewrite (4.24) as

$$N_T \approx 4\pi \left(\frac{R_T}{\lambda_T} \right)^3 \int_0^{\infty} Q^2 \left(1 + \frac{2\alpha_0}{Q^2} \right)^{1/2} g_{3/2} \left[e^{-Q^2/2} \right] dQ \quad (4.26)$$

$$\approx 4\pi \left(\frac{R_T}{\lambda_T} \right)^3 \int_0^{\infty} (Q^2 + \alpha_0) g_{3/2} \left[e^{-Q^2/2} \right] dQ, \quad (4.27)$$

where the binomial expansion has been evaluated to first order in α_0 . The summation form of the polylogarithmic function given by (1.9) can then be used to evaluate the

4.5 Interpretation of results

integral of (4.27):

$$N_T \approx 4\pi \left(\frac{R_T}{\lambda_T} \right)^3 \sum_{k=1}^{\infty} \frac{1}{k^{3/2}} \int_0^{\infty} (Q^2 + \alpha_0) e^{-kQ^2/2} dQ \quad (4.28)$$

$$= (2\pi)^{3/2} \left(\frac{R_T}{\lambda_T} \right)^3 \left(\sum_{k=1}^{\infty} \frac{1}{k^3} + \alpha_0 \sum_{k=1}^{\infty} \frac{1}{k^2} \right) \quad (4.29)$$

$$= (2\pi)^{3/2} \left(\frac{R_T}{\lambda_T} \right)^3 [\zeta(3) + \alpha_0 \zeta(2)] . \quad (4.30)$$

The first term on the right-hand side is simply the non-interacting critical number N_c^0 . Dividing through by N_c^0 gives the important result

$$\frac{N_T}{N_c^0} \approx 1 + \frac{\zeta(2)}{\zeta(3)} \alpha_0 \approx 1 + 1.37 \alpha_0 . \quad (4.31)$$

The approximations used here are the same as those used to derive equations (120) and (122) in [112], from which (4.31) can directly be deduced. In this reference the process is geared towards a revised expression for the condensed fraction below T_c , a concept that is far more prevalent in the field, though strongly linked with non-saturation.

4.5 Interpretation of results

The basic mean-field result (4.31) gives a guide to how parameters in the system should scale, as well as providing a quantitative result to which we can directly compare our measurements. Firstly, within each series there is a predicted scaling of N_T above N_c with $\mu_0 \propto N_0^{2/5}$, indicating that plots of the relationship between N_T and $N_0^{2/5}$ should be more informative than the linear-linear format used thus far. Secondly, it gives the temperature correction alluded to at the end of §4.3.2 from the T dependence of N_c^0 and α_0 . This amounts to plotting $N_T(T/T')^3$ versus $N_0^{2/5}(T/T')$. Finally, to directly compare our results with the prediction (4.31) we find the expected slope for such a plot within this Hartree-Fock approximation:

$$S_{\text{HF}} = \frac{dN_T}{d(N_0^{2/5})} \approx 1.37 X , \quad (4.32)$$

where

$$X = \xi T^2 a^{2/5} , \quad \xi = \frac{\zeta(3)}{2} \left(\frac{k_B}{\hbar\omega} \right)^2 \left(\frac{15}{a_{\text{ho}}} \right)^{2/5} . \quad (4.33)$$

Fig. 4.8 shows the result of plotting the data of Fig. 4.5 with this new approach. The blue line shows the slope $S_{\text{HF}} = 699$ with a y -intercept given by the experimental value of N_c . The initial linear slope of N_T with $N_0^{2/5}$ is in impressive agreement with the findings of the Hartree-Fock description. As the condensate grows past $\sim 10,000$ atoms, corresponding to $\mu_0/k_B T \sim 0.1$, this agreement starts to break down, and the lack of saturation becomes even more extreme.




Figure 4.8: Quantifying the lack of saturation. The data of Fig. 4.5 is now displayed in the form N_T versus $N_0^{2/5}$ as guided by Hartree-Fock theory. The dotted line shows the experimentally determined N_c . The blue line shows the predicted slope $S_{\text{HF}} = 1.37 \xi T^2 a^{2/5} = 699$. This is in close agreement with the data up to $N_0 \sim 10,000$. The red line shows a linear fit to the data in the range $0.1 < \mu_0/k_B T < 0.3$ between the dashed lines. This gives a non-saturation slope $S = 1283 \pm 84$. The black line provides a guide to the eye based on a second order polynomial fit, the slope of which is indistinguishable from the Hartree-Fock slope at low N_0 .

4.5.1 Popov approximation

The first port of call to account for the divergence from first-order mean-field theory for large N_0 is to use a more sophisticated theoretical description. Higher order terms can be retained and the effects of the finite thermal density included in the self-consistent Hartree-Fock approach, but to go beyond this it is necessary to include the effects of low-energy excitations in the system. This is done by using the so-called Popov approximation [113], in which the Bogoliubov terms for creation and destruction of particle pairs are included (see e.g. [19]).

The first thing to check with this model is the low- N_0 behaviour. A shift in N_c from the ideal value is of course now detected, but aside from this the initial slope of $\approx 1.37 X$ is indeed recovered. The results of using the Popov approximation for the most strongly and weakly interacting conditions used in our experiments are shown on top of the data for our sample series and the first-order Hartree-Fock prediction in Fig. 4.9. The computational method used to implement the theory is described in Appendix C. As condensate size increases it is clear that the Popov theory predicts a departure from the basic mean-field picture, but the effect is not great enough to fully explain the behaviour seen in our experiments.

To deal with these larger (and experimentally more typical) condensates we therefore take a heuristic approach, quantifying the degree of non-saturation by defining the

4.5 Interpretation of results




Figure 4.9: The predictions of Popov theory compared with first-order Hartree-Fock theory (black line) and the sample data series (grey points) of Fig. 4.5. Points are now normalised to the relevant N_c values and plotted against $\alpha_0 = \mu_0/k_B T$. The blue points show the prediction for the most weakly interacting series, $T = 122$ nK $a = 40 a_0$, the red points the prediction for the most strongly interacting series, $T = 284$ nK $a = 356 a_0$. The Popov predictions were calculated as described in Appendix C.

slope $S = \Delta(N_T)/\Delta(N_0^{2/5})$ for the range $0.1 \leq \alpha_0 \leq 0.3$. This slope is shown for the data series in Fig. 4.8 by the solid red line. Though the data isn't perfectly linear, it is well described by S over this broad experimental range. We therefore have to investigate the non-saturation behaviour in both this regime and that of low N_0 .

4.5.2 Small condensate limit

In order to test the findings of (4.31) for small condensates, it was necessary to take further series of data concentrating on the low N_0 limit, using a slightly altered experimental method. Interactions are 'turned off' during time-of-flight, limiting the expansion of the condensate to make it clearly stand out against the background of thermal atoms in absorption imaging. This technique is described in more detail in Chapter 5. Unfortunately, this method can not be applied when detecting larger condensates, as the peak optical density becomes extremely large. It is for these same reasons that we use the slope S to quantify non-saturation for larger condensates. The alternative method of a second-order polynomial fit to the data requires information at both very-low and high N_0 .

The low N_0 data series were taken for $a = 56$ - $274 a_0$ and $T = 177$ - 317 nK, combining to give X values in the range ~ 500 - 1300 . The experimental non-saturation slope $S_0 = dN_T/d(N_0^{2/5})$ was found in the $N_0 \rightarrow 0$ limit for each set of data for direct comparison with S_{HF} . The results are shown in Fig. 4.10 **a**, where the blue line shows $S_{\text{HF}} = 1.37 X$. The agreement between the experiment and theory to within a few




Figure 4.10: Recovery of saturation in the non-interacting limit. The non-saturation slopes S_0 and S are plotted against the dimensionless parameter $X = \xi T^2 a^{2/5}$ described in the text. The S_0 data are directly compared with Hartree-Fock theory, $S_{\text{HF}} = 1.37 X$ (blue line), with no free parameters. A linear fit to the S data gives $dS/dX = 2.6 \pm 0.3$ and an intercept $S(0) = -20 \pm 100$, consistent with complete saturation in the ideal gas limit. The S values are extracted from the 18 ^{39}K data series shown in Fig. 4.6 and use the same colour code. Two ^{87}Rb data points are also shown as black squares. All vertical error bars are statistical. The systematic uncertainty in N_0 and N_T of $< 10\%$ corresponds to $< 6\%$ uncertainty in S_0 and S values. The horizontal error bars include the 3 Hz uncertainty in the trapping frequencies $\bar{\omega}/2\pi$ and (for potassium) an allowance of 0.1 G uncertainty in the position of the Feshbach resonance.

percent shows that, in the non-interacting limit (where $\mu_0 = 0$ for all N_0), S_0 would indeed vanish. This satisfactory agreement with theory doesn't complete the attempted proof of saturation in the non-interacting limit, as we must now look at the large N_0 regime.

4.5.3 Large condensate limit

The values of S extracted from the 18 data series of Fig. 4.6 are shown in Fig. 4.10 b. The data falls in a straight line as plotted against X , supporting the assumption that $\mu_0/k_B T$ may still be used as the relevant interaction parameter. The red line shows an unconstrained linear fit to this data. In addition, data series were also taken with ^{87}Rb at two temperatures: 175 and 203 nK. The scattering length for this data is fixed to $99 a_0$ by the background triplet scattering length of ^{87}Rb . The good agreement of these two points with the ^{39}K data serves to validate our experimental approach.

The extrapolated y -intercept of the red line in Fig. 4.10 gives the prediction that in the $X \propto T^2 a^{2/5} \rightarrow 0$ limit, $S = -20 \pm 100$, consistent with zero. The bounds on the value are given by the confidence margin of the linear fit. This limit can be reached in

4.6 Conclusions

two very different ways. In the rather mundane $T \rightarrow 0$ limit, N_T trivially vanishes for any value of N_0 , automatically giving $S = 0$. However, the dependence of S only on X means that in the more pertinent $a \rightarrow 0$ limit (at finite T) we expect the slope to tend to zero, with a finite thermal fraction.

4.6 Conclusions

Combining the result of the heuristic approach for larger condensates with the success of Hartree-Fock theory for small condensates confirms the picture of a non-interacting saturated gas for a broad range of experimentally relevant parameters. We expect this behaviour to be universal for (relatively) weakly-interacting Bose gases in three-dimensional harmonic potentials, characterised by the short-range s -wave scattering length a . This strongly suggests that condensation is indeed the result of a purely quantum-statistical phase transition.

The Hartree-Fock prediction varies significantly with the form of the trapping potential. We can repeat the derivation of §4.4 for power law traps (in a three-dimensional system) of the form $V(r) \propto |r^n|$ for general n with interesting results. Once we make the approximation that the thermal atoms reside only outside the condensate, the remaining process is exact for a linear trap ($n = 1$), and we find

$$\frac{N_T}{N_c^0} \approx 1 + \frac{\zeta(7/2)}{\zeta(9/2)}\alpha_0 + \frac{\zeta(5/2)}{2\zeta(9/2)}\alpha_0^2, \text{ where } \alpha_0 \propto N_0^{1/4} \quad (4.34)$$

and N_c^0 is defined for a linear trap. This demonstrates a stronger dependence on N_0 than the findings for a harmonic trap. For a cubic trap ($n = 3$) the assumption that the thermal atoms reside mainly outside the condensate is no longer a good one¹, and the central region must be included explicitly. This gives the prediction

$$\frac{N_T}{N_c^0} \approx 1 + \frac{\zeta(3/2)}{\zeta(5/2)}\alpha_0 - \frac{4\sqrt{\pi}}{3\zeta(5/2)}\alpha_0^{3/2}, \text{ where } \alpha_0 \propto N_0^{1/2}, \quad (4.35)$$

where the dependence is now weaker on N_0 . The $\alpha_0^{3/2}$ term is included due to its relatively large effect. From purely geometric considerations we would also expect stronger non-saturation effects for lower dimensional systems. It would therefore be interesting to study the low N_0 limit for different trap dimensions and potentials, and in particular disordered potentials [114] where the effective dimensionality can take on non-integer values. It is, therefore, possible that non-saturation scaling could act as a method for classifying the geometries and interactions in many-body systems.

The high N_0 behaviour we see is not, however, due to altered mean-field effects because of the slight anharmonicity of our system. Numerical simulations are observed to be consistent with those for a harmonic trap. We hesitate to draw any further

¹In fact using this assumption would lead to the conclusion that the gas is actually saturated.

conclusions from the perceived enhanced non-saturation, aside from noting that beyond-mean-field effects due to critical correlations might be expected to scale with the size of the critical region. For a harmonically trapped gas in the Thomas-Fermi limit, this region is the surface of an ellipsoid with semi-axes $\propto \sqrt{\alpha_0}$, such that the size of the critical region $\propto \alpha_0$. The magnitude of these effects might therefore be expected to also scale with this parameter.

The hope is that this work has dispelled any expectation that the atom number should saturate during condensation at constant temperature, and has provided a solid yard-stick for the scale of non-saturation in a harmonic trap. Indeed, in recent work the lack of saturation of the thermal component has been acknowledged and attributed to repulsive interactions [115]. Our results also highlight the need to be careful when using ideal laws for calibration of experimental systems, for example using the condensed fraction to measure the temperature of a trapped gas. Perhaps most significantly though, we have leant the first direct experimental weight to Einstein's textbook picture of a saturated gas. The results of this chapter are summarised in [116].

Chapter 5

Effect of interactions on the critical temperature

“ I laughed heartily to myself, but I was the only one. At the point where the aunts should have rolled out of their seats as one aunt there occurred merely a rather ghastly silence as of mourners at a death-bed, which was broken by Aunt Charlotte asking what I had said. ”

– P. G. Wodehouse (*The Mating Season*)

5.1 History

Having investigated the effects of repulsive interactions in a Bose gas beyond the point of condensation, we now turn to the behaviour of the critical point itself. The question of how interactions affect this standard textbook result is, in contrast to the saturation question, one that has been hotly debated for more than half a century, since the pioneering work of Lee and Yang [117, 118]. This debate was reinvigorated with the advent of the first experimental realisation of a BEC, creating a surge in theoretical output on the topic since 1995. Though these gases are harmonically confined in the lab, most theoretical work has historically focussed on uniform systems, where the critical temperature for a non-interacting system T_c^0 is given by (1.12). The simplicity of the derivation given in §1.2.1 belies the complexity of the lowest-order correction due to interactions. For several decades there was no general agreement on the form, magnitude or even the sign of the shift, as illustrated by the summary of past work in Table 5.1. The T_c shift is defined as $\Delta T_c = T_c - T_c^0$.

The difficulty stems from the fact that a purely mean-field picture is not sufficient when considering the critical point. Mean-field theories have numerous useful applications for the description of interacting Bose systems, such as the successful prediction of the broadening of the condensate profile, the dispersion law for elementary excitations and other implementations reviewed in [112], as well of course as the initial lack of saturation of a harmonically trapped gas as discussed in Chapter 4. However, near the

5.1 History

Year	Prediction		Ref.
1957	$\Delta T_c/T_c^0 = C(an^{1/3})^{1/2}$	$C = 2.49$	Lee and Yang [117]
1958	$\Delta T_c/T_c^0 = can^{1/3}$		Lee and Yang [118]
1960	$\Delta T_c/T_c^0 = C(an^{1/3})^{1/2}$	$C > 0$	Glassgold <i>et al.</i> [119]
1964	$\Delta T_c/T_c^0 = D(an^{1/3})^{3/2}$	$D > 0$	Huang [120]
1971		$\Delta T_c < 0$	Fetter and Walecka [121]
1982	$\Delta T_c/T_c^0 = C(an^{1/3})^{1/2}$	$C = -3.5$	Toyoda [122]
1992	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 4.66$	Stoof [123]
1996		$\Delta T_c > 0$	Bijlsma and Stoof [124]
1997	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 0.34 \pm 0.06$	Grüter <i>et al.</i> [125]
1999	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 0.7$	Holzmann <i>et al.</i> [126]
1999	$\Delta T_c/T_c^0 = C(an^{1/3})^{1/2}$	$C = 3.5$	Huang [127]
1999	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 2.3 \pm 0.25$	Holzmann and Krauth [128]
1999	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 2.9$	Baym [129]
2000	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 2.33$	Baym <i>et al.</i> [130]
2000	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 5.1 \pm 0.9$	Reppy <i>et al.</i> [131]
2000	$\Delta T_c/T_c^0 = can^{1/3}$	$c = -0.93$	Wilkins <i>et al.</i> [132]
2000	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.71$	Arnold and Tomášik [133]
2001	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.32 \pm 0.02$	Arnold and Moore [134]
2001	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.29 \pm 0.05$	Kashurnikov <i>et al.</i> [135]
2003	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.14 \pm 0.11$	Kleinert [136]
2003	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.3 \pm 0.4$	Davis and Morgan [137]
2004	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.27 \pm 0.11$	Kastening [138]
2004	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.23 \pm 0.12$	Ledowski <i>et al.</i> [139]
2004	$\Delta T_c/T_c^0 = can^{1/3}$	$c = 1.32 \pm 0.14$	Nho and Landau [140]
2010	$\Delta T_c/T_c^0 = can^{1/3}$	$c = -2.33$	Betz and Ueltschi [141]

Table 5.1: Various predictions for the T_c shift since midway through the twentieth century. Results are given in terms of the diluteness parameter $an^{1/3}$ as used in the majority of the references. All results are theoretical, apart from the experimental work of Reppy *et al.* It was later shown that the estimation of the scattering length was incorrect in this work [135]. For the most part this data was collated by Seiringer and Ueltschi [142] and in the review by Andersen [143].

critical point (in what is known as the ‘critical region’) long-range correlations develop which must be taken into consideration. Correlations play an essential role in phenomena such as superfluidity [144] and associated vortex properties, as well as explaining shifts in the frequencies of collective oscillations [145] and quantum depletion of the condensate [146], and are as such a topic of great theoretical interest.

5.1.1 Uniform system

In a uniform system there is no shift within mean-field theory, and the lowest-order shift in T_c is solely due to correlations. The lack of a mean-field correction in a uniform system can be understood by observing that, in this description, interactions lead to a universal shift by $\epsilon_0 = 2gn$ of the effective potential V_{eff} . Note that ϵ_0 is now also the lowest energy state of the system. In this chapter we shall always consider the case where there is no condensate, such that $n \equiv n_T$. By setting $\epsilon(\mathbf{r}, \mathbf{p}) = p^2/2m + \epsilon_0$ in (1.6) we obtain a modified version of (1.7),

$$n = \frac{1}{\lambda_T^3} g_{3/2} \left[e^{(\mu - \epsilon_0)/k_B T} \right]. \quad (5.1)$$

At $\mu = 0$ there is clearly a shift in n from the ideal critical value, but the upper bound on the chemical potential, according to (1.1), is the lowest single-particle energy state ϵ_0 . Macroscopic occupation of the ground state now occurs when μ approaches this value. The fact that interactions shift both the critical chemical potential μ_c and the effective potential by the same amount results in the ideal n_c (and equivalently T_c) being recovered.

Calculation of the shift induced by correlations requires non-perturbative methods which are difficult to compute [143, 147], hence the long and varied theoretical background. Effects are considered in terms of the small diluteness parameter, defined as the ratio between the scattering length and the typical inter-particle separation

$$an^{1/3} \equiv \zeta(3/2)^{1/3} \frac{a}{\lambda_0}, \quad (5.2)$$

where in the equivalent form λ_0 is the thermal wavelength at T_c^0 ,

$$\lambda_0 = \left(\frac{2\pi\hbar^2}{mk_B T_c^0} \right)^{1/2}. \quad (5.3)$$

A general consensus for ΔT_c has only been reached in the last ten years [143, 147–149], thanks to the Monte Carlo simulations carried out by Arnold and Moore [134] and Kashurnikov *et al.* [135] in 2001. It turns out that long range correlations act to aid the establishment of coherence between particles, so that condensation sets in at a higher temperature. The result is

$$\frac{\Delta T_c}{T_c^0} \approx 1.3 an^{1/3} \approx 1.8 \frac{a}{\lambda_0}, \quad (5.4)$$

5.1 History

where the functional dependence is that proposed by Lee and Yang in 1958 [118]. Though the matter is generally considered settled, it is still not universally accepted, as may be seen from the latest entry in Table 5.1. It may require direct experimental evidence to finally put the affair to bed.

5.1.2 Harmonic confinement

In the harmonically trapped case, interactions affect μ_c and V_{eff} differently due to the inhomogeneity of the system, and the critical temperature *is* shifted within mean-field theory. Note though that $\mu_c = V_{\text{eff}}(0)$, and the critical phase-space density of $\zeta(3/2)$ is preserved at the centre of the trap. The mean-field shift for a harmonic potential was first calculated by Giorgini *et al.* [150] and is discussed in §5.2. The first-order result is given by

$$\frac{\Delta T_c}{T_c^0} \approx -3.426 \frac{a}{\lambda_0}, \quad (5.5)$$

predicting a negative shift in this case. For a harmonically trapped gas, the mean-field result dominates over the effect due to correlations in a manner that goes beyond the difference in pre-factors between (5.4) and (5.5). This is because for a confined gas at the condensation point, the critical region at the centre of the trap is vanishingly small, so correlation effects are expected to effect T_c at higher order in a/λ_0 . The mean-field result of (5.5) is therefore expected to be exact to first order.

It is pointed out in [150] that critical correlations are expected to affect the result at the order of a^2 and above. This is as a result of the Ginzburg-like criterion [103, 151], predicting that correlations become significant when the correlation length of the Bose gas becomes comparable to the length-scale of nonperturbative physics in the system. This condition is calculated in §II.C of [134] to be

$$\delta\mu = |\mu - \mu_c| \lesssim \frac{\hbar^2 a^2}{m\lambda_0^4}, \quad (5.6)$$

where μ_c is the value of μ at the critical point. This condition is equivalent to the one that appears in the original mean-field T_c shift paper [150]. The conclusion that critical correlations come in only at second order and above is most simply attained by dimensional analysis, as is carried out in [134], along with other more thorough derivations. The investigation further shows that a logarithmic term is also expected to appear, resulting in a shift of the form

$$\frac{\Delta T_c}{T_c^0} = c_1 \frac{a}{\lambda_0} + \left(c'_2 \ln \frac{a}{\lambda_0} + c''_2 \right) \left(\frac{a}{\lambda_0} \right)^2 + \dots \quad (5.7)$$

where $c_1 \approx -3.426$ is the mean-field result. Attempts have been made to theoretically quantify the beyond-MF corrections to the MF result [134, 149, 152–154], but no consensus has been reached beyond the expectation that the additional beyond-MF shift should be positive.




Figure 5.1: Summary of the past experimental determinations of the T_c shift by Ensher *et al.* [108], Gerbier *et al.* [77] and Meppelink *et al.* [155]. The dashed lines, representing the data of Gerbier *et al.*, show the single standard deviation bounds as found in Fig. 3 of [154]. The data covers the range $a/\lambda_0 = 0.007$ - 0.024 and is consistent with the MF prediction (blue line).

5.1.3 Past experiments

Since the MF prediction of [150] in 1996, experiments in harmonic traps have ruled out the ideal prediction T_c^0 with ever increasing certainty, as summarised in Fig. 5.1. The first measurement performed by Ensher *et al.* [108] found $\Delta T_c/T_c^0 = -0.06 \pm 0.05$ for fixed $a/\lambda_0 \approx 0.015$, just excluding the ideal case. Gerbier *et al.* [77] managed to probe a range of a/λ_0 from 0.016 to 0.024 by varying the density of the cloud, excluding T_c^0 by more than two standard deviations. Most recently in 2010 Meppelink *et al.* [155] measured the T_c shift with deft precision at $a/\lambda_0 \approx 0.007$, finding $\Delta T_c/T_c^0 = -0.026 \pm 0.001$. All of these investigations found results consistent with the mean-field theory, but none could discern any corrections beyond this description.

5.1.4 Aims of the investigation

Though in the experiments carried out by Gerbier *et al.* they were able to tune the interaction parameter a/λ_0 by varying the sample density (and hence critical temperature), it is difficult to span a wide range with this method due to the weak dependence $a/\lambda_0 \propto n^{1/3}$. These experiments were carried out with ^{87}Rb (and in a magnetic trap), entailing a fixed scattering length. Their concluding remarks from this work point the way to future progress: “We note to conclude that measuring corrections to T_c beyond the mean field for our typical experimental parameters would require thermometry with an accuracy of 1% or better. A more direct route to investigate such effects might be to measure the critical density near the center of the trap, directly sensitive to the presence of critical fluctuations. Alternatively, these many-body effects could be enhanced in the vicinity of a Feshbach resonance.”

Using the above as a manifesto, we use an improved thermometry technique to

5.2 Mean-field shift in a harmonic trap

increase measurement accuracy, while using the available Feshbach resonance in ^{39}K to investigate from the very-weakly interacting limit up to a regime where we hope to see beyond-MF effects. With the versatility of the Feshbach resonance we can also probe more extreme regimes, where thermodynamic equilibrium is lost. Direct probing of the in situ density is beyond our current means, but remains an exciting area of investigation. This technique has recently been used to show strong deviations from MF theory in excellent agreement with quantum Monte Carlo simulations of the effects of correlations [156]. We proceed with a discussion of the mean-field effects that usually dominate the T_c shift in typical BEC setups, before moving on to the experimental methods geared towards revealing the beyond-MF behaviour.

5.2 Mean-field shift in a harmonic trap

The negative mean-field shift of equation (5.5) can be qualitatively understood by considering the effect of repulsive interactions near the trap centre. As the mean-field potential is proportional to the density of the cloud the effect is strongest here. The net effect then is that interactions spread the atoms out, reducing the central density in comparison to an ideal Bose gas. In the non-interacting case the critical point is reached when N_c^0 atoms are in the trap and, as found in §1.2.2, a critical density equal to that of a uniform system is reached at the trap centre. For the interacting Bose gas this critical density is no longer reached with N_c^0 atoms due to the flattening of the distribution at the trap centre, and further atoms must be added before the critical phase-space density is attained. There is therefore a positive shift in the critical number at constant temperature ($\Delta N_c > 0$), or correspondingly a negative shift in the critical temperature for constant atom number ($\Delta T_c < 0$). The first-order shift is derived analytically below, before a self-consistent numerical method is presented to help both illustrate the effect and for use later in the chapter.

5.2.1 Analytic solution

To calculate the first order effect of interactions on the critical temperature in a harmonic trap we look again at the Hartree-Fock description in the semi-classical limit. In this case the single-particle energy levels are given by

$$\epsilon(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + V(\mathbf{r}) + \epsilon_{\text{int}} , \quad (5.8)$$

where $\epsilon_{\text{int}} = 2gn(\mathbf{r})$, and $V(\mathbf{r})$ takes the parabolic form given by (1.14). We want to find the shift in critical temperature for a particular value of N , so we consider the

differential dN , given by

$$dN = \frac{\partial N}{\partial T} dT + \frac{\partial N}{\partial \mu} d\mu + \int \frac{\partial n(\mathbf{r})}{\partial \epsilon_{\text{int}}} d\epsilon_{\text{int}} \quad (5.9)$$

$$= \frac{\partial N}{\partial T} dT + \frac{\partial N}{\partial \mu} d\mu + \int \frac{\partial n(\mathbf{r})}{\partial \epsilon_{\text{int}}} 2gn(\mathbf{r}) d\mathbf{r} . \quad (5.10)$$

The last term must be kept in this form because of the spatial dependence of ϵ_{int} . We are interested in the small variations from the ideal case where $T = T_c^0$, $\mu = 0$ and $\epsilon_{\text{int}} = 0$, so the partial derivatives are evaluated at this point. Considering the first order shift from the ideal case in terms of the small quantities $\Delta T_c = T_c - T_c^0$, ϵ_0 and $2gn(\mathbf{r})$ allows us to write

$$0 = \frac{\partial N}{\partial T} \Delta T_c + \frac{\partial N}{\partial \mu} \epsilon_0 + 2g \int \frac{\partial n(\mathbf{r})}{\partial \epsilon_{\text{int}}} n(\mathbf{r}) d\mathbf{r} . \quad (5.11)$$

The quantity ϵ_0 is the true zero-point energy of the system, which is shifted by an amount $\approx 2gn(0)$ by interactions¹. Up until now we have also neglected the zero-point motion of atoms in the harmonic oscillator ground state. Its inclusion results in a shift of the critical temperature even for a non-interacting gas, though it is independent of N (as we shall see) and so in the limit $N \rightarrow \infty$ the relative shift tends to zero. Condensation actually sets in when the chemical potential μ is equal to

$$\epsilon_0 = \frac{3}{2} \hbar \omega_m + 2gn(0) , \quad (5.12)$$

where the first term is the zero-point energy of the harmonic oscillator, and $\omega_m = (\omega_x + \omega_y + \omega_z)/3$ is the algebraic mean of the trap frequencies.

The first term on the right-hand side of (5.11) can be found from the ideal equation for N at T_c^0 given by (1.18). We immediately find

$$\frac{\partial N}{\partial T} = 3 \frac{N}{T_c^0} . \quad (5.13)$$

To calculate the second term we need to use the familiar thermal distribution function

$$n(\mathbf{r}) = \frac{1}{\lambda_T^3} g_{3/2} \left[e^{-(V(\mathbf{r}) + \epsilon_{\text{int}} - \mu)/k_B T} \right] . \quad (5.14)$$

It is helpful to use the sum form of the $g_{3/2}$ function, so that by taking the partial derivative with respect to μ and integrating over space we obtain

$$\frac{\partial N}{\partial \mu} = \frac{1}{\lambda_T^3} \sum_{k=1}^{\infty} \frac{1}{k^{3/2}} \int \frac{\partial}{\partial \mu} \left[e^{-k(V(\mathbf{r}) + \epsilon_{\text{int}} - \mu)/k_B T} \right] d\mathbf{r} . \quad (5.15)$$

¹This approximation is good as long as $a_{ho} \ll R_T$, i.e. the extent of the lowest single particle wavefunction is much smaller than the size of the thermal cloud.

5.2 Mean-field shift in a harmonic trap

At the ideal critical point this can be written as

$$\frac{\partial N}{\partial \mu} = \frac{1}{k_B T_c^0} \frac{1}{\lambda_0^3} \sum_{k=1}^{\infty} \frac{1}{k^{1/2}} \int e^{-kV(\mathbf{r})/k_B T_c^0} d\mathbf{r} , \quad (5.16)$$

where

$$\lambda_0 = \left(\frac{2\pi\hbar^2}{mk_B T_c^0} \right)^{1/2} \quad (5.17)$$

is the thermal wavelength at T_c^0 . Evaluating the above expression and substituting in N from (1.18) gives the result

$$\frac{\partial N}{\partial \mu} = \frac{\zeta(2)}{\zeta(3)} \frac{N}{k_B T_c^0} , \quad (5.18)$$

giving the first-order saturation result of (4.31). The final term of (5.11) can be found by following a similar process. Using the above expression for $n(\mathbf{r})$ and using the sum form of the $g_{3/2}$ function allows us to write

$$\begin{aligned} \int \frac{\partial n(\mathbf{r})}{\partial \epsilon_{\text{int}}} n(\mathbf{r}) d\mathbf{r} = \\ \frac{1}{\lambda_T^6} \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \frac{1}{j^{3/2}} \frac{1}{k^{3/2}} \int \frac{\partial}{\partial \epsilon_{\text{int}}} \left[e^{-j(V(\mathbf{r})+\epsilon_{\text{int}}-\mu)/k_B T} \right] e^{-k(V(\mathbf{r})+\epsilon_{\text{int}}-\mu)/k_B T} d\mathbf{r} , \end{aligned} \quad (5.19)$$

which, at the critical point, can be evaluated as

$$\int \frac{\partial n(\mathbf{r})}{\partial \epsilon_{\text{int}}} n(\mathbf{r}) d\mathbf{r} = -\frac{1}{k_B T_c^0} \frac{1}{\lambda_0^6} \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \frac{1}{j^{1/2}} \frac{1}{k^{3/2}} \int e^{-(j+k)V(\mathbf{r})/k_B T_c^0} d\mathbf{r} \quad (5.20)$$

$$= -\frac{\mathcal{S}}{\zeta(3)} \frac{N}{\lambda_0^3 k_B T} , \quad (5.21)$$

where

$$\mathcal{S} = \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \frac{1}{j^{3/2} k^{1/2} (j+k)^{3/2}} \approx 1.208 . \quad (5.22)$$

Inserting the expressions (5.13), (5.18) and (5.21) into (5.11) and collating similar terms gives the first order result

$$\frac{\Delta T_c}{T_c^0} = -\frac{\zeta(2)}{3\zeta(3)} \left(\frac{\hbar\omega_m}{k_B T_c^0} \right) - \frac{4}{3\zeta(3)} [\zeta(2)\zeta(3/2) - \mathcal{S}] \frac{a}{\lambda_0} \quad (5.23)$$

$$\approx -0.684 \left(\frac{\hbar\omega_m}{k_B T_c^0} \right) - 3.426 \frac{a}{\lambda_0} . \quad (5.24)$$

The first term on the right-hand side results solely from the zero-point motion of the harmonic oscillator, and is known as the finite-size shift. From this it is clear that the absolute shift is indeed independent of the atom number. With the number of atoms in our system the relative shift of T_c is typically $\sim 1\%$. The second term is the shift due

to interactions of (5.5), characterised by the interaction parameter a/λ_0 . In typical cold atoms systems this takes a value on the order of 0.01, shifting T_c by about 3%. Equation (5.24) is often written in the equivalent form

$$\frac{\Delta T_c}{T_c^0} \approx -0.728 \frac{\omega_m}{\bar{\omega}} N^{-1/3} - 1.326 \frac{a}{\bar{a}_{\text{ho}}} N^{1/6}, \quad (5.25)$$

where the conversion factor $N^{1/6}/a_{\text{ho}} = \sqrt{2\pi}\zeta(3)^{1/6}/\lambda_0$ has been employed for the interaction term. We choose to work in terms of the original form of (5.24). We are only interested in the shift due to interactions, so must cater for the finite-size shift in some manner.

It is worth noting that the value $2gn(0)$ used for the zero-point interaction energy is the mean-field value. The effect of correlations is to suppress density fluctuations near the critical point, decreasing $\epsilon_{\text{int}}(0)$. Allowing this to propagate through the above calculation results in the expected positive shift with respect to the mean-field result (5.24).

5.2.2 Numerical solution

To visualise the effect on the atomic density distribution we return to the semi-classical Hartree-Fock approximation and equation (1.6). The interaction term $2gn(\mathbf{r})$ must be added to the expression for the single-particle energy spectrum,

$$\epsilon(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + V(\mathbf{r}) + 2gn(\mathbf{r}), \quad (5.26)$$

where for simplicity we consider a spherical harmonic potential such that

$$V(\mathbf{r}) = \frac{1}{2}m\omega^2 r^2. \quad (5.27)$$

By conversion to the dimensionless variables

$$A = \frac{4a}{\lambda_T}, \quad D = n\lambda_T^3, \quad R = \frac{r}{R_T} \text{ and } \alpha = \frac{\mu}{k_B T} \quad (5.28)$$

and integrating (1.6) over momentum space we obtain the implicit equation

$$D(R) = g_{3/2} \left(e^{\alpha - R^2/2 - AD(R)} \right). \quad (5.29)$$

In mean-field theory we know that at the critical point $D(0) = g_{3/2}(1) = \zeta(3/2)$. This allows us to write down the value of α at the critical point,

$$\alpha_c = A\zeta(3/2). \quad (5.30)$$

5.2 Mean-field shift in a harmonic trap




Figure 5.2: Mean-field theory predicted critical density distribution for a harmonically trapped Bose gas with repulsive interactions. The black line shows the non-interacting case (identical to Fig. 1.2). The blue lines show the distribution for increasing interaction strength, in successive steps of $0.01 a/\lambda_T$ up to $0.05 a/\lambda_T$. Here it is clear that N_c increases with increasing a/λ_0 . Note in each case the central phase-space density is $\zeta(3/2)$.

Combining this with (5.29) we obtain an expression for the phase-space density distribution at the critical point

$$D_c(R) = g_{3/2} \left(e^{A\zeta(3/2) - R^2/2 - AD_c(R)} \right). \quad (5.31)$$

Equation (5.31) must be solved self-consistently for a given value of A . This may be done numerically, and the results are shown for various values of $a/\lambda_0 = A/4$ in Fig. 5.2. This clearly shows how the central density distribution is increasingly spread out for higher a/λ_0 . The mean-field N_c shift can be determined by integrating (5.31) over R -space. In Fig. 5.3 this result is compared with the first-order result found from (5.5):

$$\frac{\Delta N_c}{N_c^0} \approx -3 \frac{\Delta T_c}{T_c^0} \approx 10.278 \frac{a}{\lambda_0}. \quad (5.32)$$

5.2.3 Introducing critical effects

When critical correlations are also taken into account there are now two competing effects; mean-field repulsions act to hinder condensation, while correlations act to aid it. Qualitatively the result can be visualised as in Fig. 5.4. Thinking for now in terms of the atom number at a fixed temperature, the MF effect reduces the central density, such that more atoms are required to reach a particular peak density than for an ideal gas. However, correlations act to reduce the critical density, so fewer atoms are needed to reach the onset of condensation. We do not simply expect the critical density to be described by the uniform result (5.4), as long-range correlations are not a local phenomenon. We therefore proceed to determine the effect experimentally, as related




Figure 5.3: Comparison between the first order analytic result $\Delta N_c/N_c^0 = 10.278a/\lambda_0$ (blue line) and the numerical method (blue points) used to find the mean-field N_c shift described in the text.

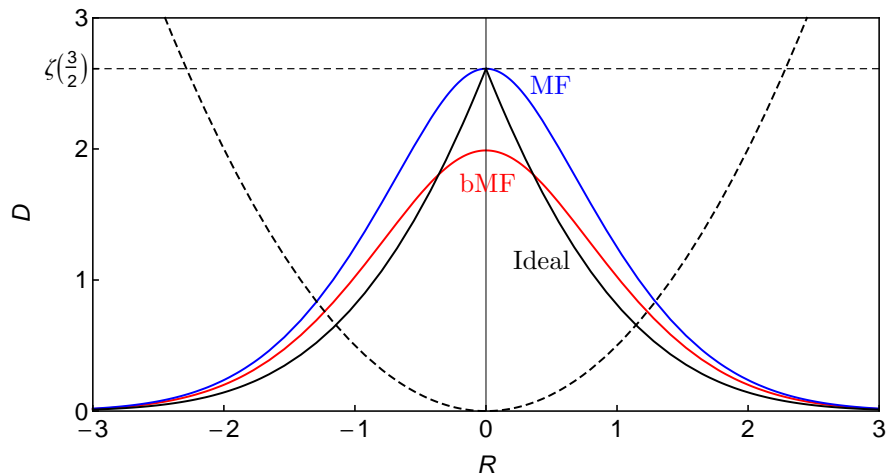


Figure 5.4: Illustration of the competing effects induced by repulsive interactions on the critical point of a harmonically trapped Bose gas. Compared with an ideal gas (black line) with the same T_c , repulsive interactions reduce the critical density, but also broaden the density distribution (red line). Mean-field theory (blue line) captures only the latter effect, thus predicting an increase in the critical atom number N_c at fixed T , corresponding to a reduction in T_c at fixed N .

below.

5.3 Experimental procedure

To measure the critical point we create clouds with various condensed fractions in order to extract the value of N_T as N_0 tends towards zero. This involves a basic method closely following the one used in Chapter 4, with each data series made up of ~ 100 images (see Fig. 5.5). Condensates are prepared once more using $a = 135 a_0$ before the trap depth and Feshbach field are ramped to their final values, the fixed trap depth ensuring good temperature stability over each experimental series. In this

5.3 Experimental procedure




Figure 5.5: Sample measurement series taken for a scattering length of $274 a_0$. The cloud is held in the trap for between 4 and 22 s at the measurement scattering length. Hold time increases from left to right and from top to bottom. The data is shown with our standard *OD* colour map. The red box indicates the image shown in Fig. 5.6, where the range of the colour map has had to be reduced to be able to distinguish the tiny condensed fraction from the thermal cloud. All clouds have a condensed fraction of under 2%.




Figure 5.6: **a** A sample optical density image of a cloud with 450,000 atoms and a condensed fraction of 0.14 %, taken after 19 ms TOF with interactions turned off. The range of the colour map has been restricted to allow the condensate to be discerned. **b** Azimuthally averaged column density of the same data, displayed for illustrative purposes. The full 2D optical-density profile is used for fitting.

set of experiments we use trapping frequencies in the range $\bar{\omega} = 2\pi \times 75\text{-}85\text{ Hz}$. The clouds are held while the atom number decays and elastic collisions redistribute atoms between the two components. This is interrupted at a -dependent intervals (loss rates being much greater for high a) by switching off the trap for imaging after 19 ms time-of-flight. In contrast to the saturation measurements we are interested only in the very-low N_0 regime, studying condensed fractions no higher than $\sim 2\%$.

Other groups, as discussed in §5.1.3, have measured the T_c shift for values of a/λ_0 between 0.007 and 0.024. The ability to tune the scattering length with a Feshbach resonance allows us to expand this range to explore the region $0.001 < a/\lambda_0 < 0.06$. In addition to this, the resonance provides two major experimental advantages that are essential for the precision and accuracy of our measurements. Firstly, for each measurement series, taken at a particular a and λ_0 , we concurrently take a series of reference data with a different scattering length a_{ref} , using the same $\bar{\omega}$ and very similar N and hence λ_0 . We use $a_{\text{ref}}/\lambda_0 \approx 0.005$, which is small enough that beyond-MF effects are expected to be negligible, but also large enough that the system remains close to thermal equilibrium. This allows us to directly measure the small T_c shift due to the difference in a/λ_0 between the two series, essentially eliminating all a -independent systematic errors that would otherwise affect the absolute measurement of $T_c(N, \bar{\omega}, a)$. These include the finite-size shift of equation (5.24), absolute number calibration and the correction due to the trap anharmonicity (see §3.4.4). This method is described in more detail later.

The second experimental advantage is the ability to turn interactions off during time-of-flight. This minimises the expansion of the condensate, allowing it to stand out against the thermal background. A sample image is shown in Fig. 5.6. We achieve this by ramping the Feshbach field to 350 G (the zero crossing value) in under 2 ms once the atoms have been released from the optical potential. Along with our image fitting technique, this allows the reliable detection of condensed fractions as small as $\sim 10^{-3}$. This allows us to home in on the critical point with high precision.

5.3 Experimental procedure

5.3.1 Image fitting routine

The details of the Feshbach-field ramp during TOF are discussed in the context of temperature errors in §5.5.2. The narrow condensates this method produces show up much more clearly than without the ramp, but the images must still be carefully processed to accurately detect the condensate numbers. This is important to allow accurate extrapolation to the critical point. The smallest condensates can easily be missed with standard fitting techniques, so a method tailored to the purpose is used. This must be combined with the careful measurements of the cloud temperature and thermal atom number used in our saturation measurements. The result (summarised in Fig. 5.7) is a seven-step process:

1. Data are excluded within a $60\ \mu\text{m}$ radius of the cloud centre (an estimate is entered manually) to be sure to remove the BEC from the fit. A simple g_2 fit is then applied to the thermal cloud.
2. The g_2 profile from the first step is fixed, and a Thomas-Fermi profile matched to the condensate.
3. The condensate width found from the previous step is used to provide a tighter exclusion region to mask the condensate. The exclusion radius is set to $1.1 R_k$ (where R_k is the previously fitted Thomas-Fermi radius), with the cloud centre now defined by the value found in step 1. A new g_2 fit is used to extract the thermal atom number N_T .
4. Step 2 is repeated using the new g_2 profile.
5. The central region out to one thermal radius R_T is now excluded and the thermal atoms fitted with a g_2 to find T . This is the most reliable method to find the temperature as previously discussed.
6. Now just the region between a radius of $35\ \mu\text{m}$ and one R_T is used for a g_2 fit. This provides the best fit to the thermal atoms near the cloud centre.
7. N_0 is found by summing the difference between the data and the thermal profile of step 6 within $0.5 R_T$ of the centre.

The value of N_0 found in the last step is consistent with the value obtained from step 4 for $N_0 \gtrsim 2,000$, but for smaller condensates the direct summation is more reliable. The shape of the condensate is not preserved when $a = 0$ during TOF, so a Thomas-Fermi fit is no longer applicable in any case.

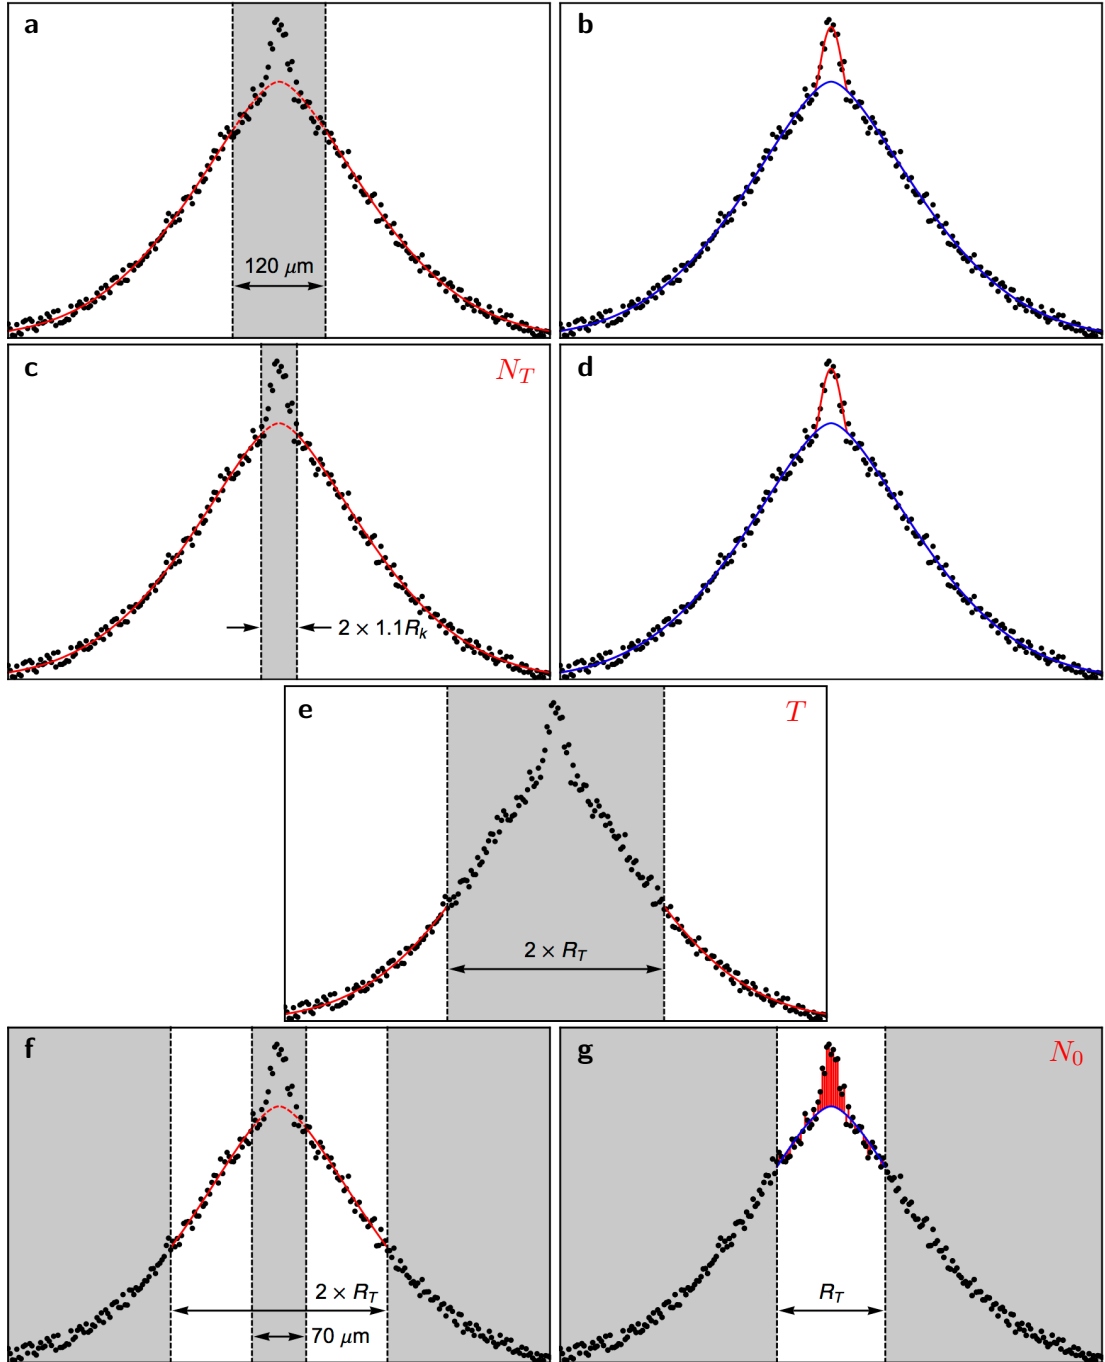


Figure 5.7: Illustration of our fitting routine. All fitting is carried out with full 2D column densities. Red curves indicate fits to the data, blue curves are fixed profiles. Greyed regions are excluded from the fitting step. **a** A g_2 fit to the thermal atoms is carried out, excluding data inside a $60 \mu\text{m}$ radius (with a manually provided centre) to be sure to exclude the BEC. **b** A Thomas-Fermi profile is used to fit the BEC (with fixed thermal profile). **c** The Thomas-Fermi radius R_k provided by the previous step allows us to more stringently exclude data within $1.1R_k$ of the centre, with a new g_2 fit giving N_T . **d** As in **b**, but with the new fixed g_2 profile. **e** A g_2 fit to the wings of the cloud (outside one thermal radius R_T) gives us T . **f** We focus on the critical region, fitting a g_2 profile to the thermal atoms around the BEC. **g** N_0 is found by summing the difference between the data and the thermal profile around the critical region.

5.4 Critical temperature shift results

5.3.2 Differential measurement process

The first stage of data processing is to compensate for any residual temperature drifts through each series, using the Hartree-Fock scaling of Chapter 4. This time the process of assigning a temperature to each image is simplified, as images were taken sequentially with decreasing atom number. Temperature is assigned by a fit to the variation in measured image temperature with N_0 . This method is used because we are only interested in the critical point, which can be done by finding both T and N_T in the $N_0 \rightarrow 0$ limit.

The next stage draws upon the N_T scaling with N_0 found in the Hartree-Fock regime from the saturation investigation. In order to accurately determine N_c we plot N_T versus $N_0^{2/5}$ and fit the data with

$$N_T = N_c + S_{\text{HF}} N_0^{2/5}, \quad (5.33)$$

using the definition of $S_{\text{HF}}(T, \bar{\omega}, a)$ given by (4.32). This method is illustrated in Fig. 5.8. Here the importance of the Hartree-Fock scaling and data at low N_0 is clear; linear extrapolation in Fig. 5.8 **a** would produce vastly different results.

To correct for the small (few per cent) difference in N_c between the measurement and reference series, we apply the ideal gas scaling $T_c \propto N_c$ to the reference series. The second order correction in ΔT_c due to the small ($< 2\%$) T_c shift at $a/\lambda_0 \approx 0.005$ is much smaller than our statistical error bars. The result is a measurement for T_c at two values of a/λ_0 and equal N . The measurement at a_{ref}/λ_0 is expected to be well described by mean-field theory. This provides an accurate measure of $\Delta T_c/T_c^0$, even if the absolute system parameters are not as accurately known.

5.4 Critical temperature shift results

The 21 series for which thermal equilibrium was deemed to hold are shown in Fig. 5.9, where data were taken for $N \approx (2-8) \times 10^5$ (corresponding to $T_c^0 \approx 180-330$ nK) in order to verify that the results depend only on a/λ_0 . The results show very good agreement with the mean-field result of equation (5.5) for $a/\lambda_0 \lesssim 0.01$. For large a/λ_0 we see a clear positive deviation from this prediction. All data points are fitted well with a second order polynomial,

$$\frac{\Delta T_c}{T_c^0} = c_1 \frac{a}{\lambda_0} + c_2 \left(\frac{a}{\lambda_0} \right)^2, \quad (5.34)$$

with $c_1 = -3.5 \pm 0.3$ (consistent with MF theory) and $c_2 = 46 \pm 5$. Note that the polynomial fit was used to calibrate the y -axis of Fig. 5.9 by setting $\Delta T_c/T_c^0 = 0$ where the fit cuts the axis. However, our results are unchanged within errors if we assume the MF prediction for our reference data at $a/\lambda_0 \approx 0.005$ is exact. By fixing the slope c_1 to the mean-field result we can improve the estimate for the quadratic result to




Figure 5.8: **a** N_0 versus N_T for two concurrently taken series. The blue points were taken at $a_{\text{ref}} = 56 a_0$, the red points at $a = 274 a_0$ (where $\lambda_0 \approx 10^4 a_0$ for both series). The data has been scaled to the same temperature ($T = 240$ nK) for illustrative purposes, directly showing the shift of the critical number $\Delta N_c(T)$. All points shown correspond to a condensed fraction of under 2%. The solid lines show the extrapolation to $N_0 = 0$ used to extract N_c . **b** The same data is shown plotted in the form N_T versus $N_0^{2/5}$, showing the extrapolation more clearly. In both plots the filled red square is the data point calculated from the image shown in Fig. 5.6.




Figure 5.9: Shift of T_c with repulsive interactions. Data points were taken with $N \approx 2 \times 10^5$ (blue circles), 4×10^5 (black squares) and 8×10^5 (red triangles) atoms. The dashed line is the mean-field result $\Delta T_c / T_c^0 = -3.426a / \lambda_0$. The solid line shows the result of a second-order polynomial fit to the data, discussed in the text. Vertical error bars are statistical, while horizontal error bars reflect the uncertainty in the position of the Feshbach resonance.

$$c_2 = 42 \pm 2.$$

The value of c_2 strongly excludes zero, and is positive as expected for beyond-MF correlations. As previously mentioned, this correction is not expected to be purely quadratic, with the pre-factor expected to take the form $c_2 = [c'_2 \ln(a/\lambda_0) + c''_2]$ according to (5.7). However, our experimental precision is not high enough to discern the exact form of the correction. Fitting with the logarithmic term included would only produce a better fit to the data due to the increased number of fitting parameters.

These results rely heavily on our differential measurement approach to remove many sources of error. However, any errors that scale with a (or a/λ_0) will systematically




Figure 5.10: Modification of the effective trapping potential caused by repulsive interactions. The black dashed line shows the harmonic trapping potential, while the red line shows the effective potential produced for N_c atoms at $a/\lambda_0 = 0.05$. The solid lines show the phase-space density distribution of the ideal (black) and interacting (red) clouds.

distort our findings. The following section describes the sources of such errors and how they have been quantified. The results have been incorporated in the error bars of Fig. 5.9.

5.5 Assessment of errors

As we are looking for variations on the per cent scale, it is crucial to ensure systematic errors are handled correctly. The results depend heavily on the precise relative measurement of temperature, as most other sources of error are eliminated by using the differential measurement method, while the dependence of our result on the recorded atom number is less strong and less susceptible to systematic variation. Two a -dependent sources of temperature error have been identified, and require careful attention. These are discussed below. Two other effects (due to the finite imaging time and residual Feshbach field curvature) that provide only absolute errors are still of general interest, and are assessed in Appendix D.

5.5.1 In-trap momentum distribution

The mean-field modification of the ideal in-trap density distribution caused by interactions has been discussed already, and this has a knock-on effect for the momentum distribution. The effective trapping potential for the atoms is a flattened version of the harmonic potential due to the addition of the interaction potential. This is illustrated in Fig. 5.10. The flatter trap means more low momentum states are occupied, so fitting the resultant time-of-flight spatial distribution with the ideal g_2 function will naturally produce errors. Significantly, the error here is dependent on a/λ_0 , so the magnitude of the effect will differ between our reference data and our measurement data.

The modified momentum distribution can be found by using the self-consistent $D_c(R)$ of equation (5.31) and plugging this back into (1.6), integrating only over space to give $n(\mathbf{p})$. Using dimensionless coordinates once more this may be written

$$n(P) \propto \int_0^\infty \frac{R^2 dR}{\exp\left[\frac{P^2}{2} + \frac{R^2}{2} + AD_c(R) - A\zeta(3/2)\right] - 1}, \quad (5.35)$$

where the T -dependent prefactors have been dropped for clarity as we are only interested in the shape of the distribution. Here we have introduced another dimensionless variable P , given by

$$P = \frac{p}{P_T}, \quad (5.36)$$

where $P_T = \sqrt{mk_B T}$. To see what effect the modified distribution has on our image fitting we need to integrate along one axis (cf. the imaging axis) to give

$$n_{2D}(P) \propto \int_0^\infty \int_0^\infty \frac{R^2 dR dP_y}{\exp\left[\frac{P^2}{2} + \frac{P_y^2}{2} + \frac{R^2}{2} + AD_c(R) - A\zeta(3/2)\right] - 1}, \quad (5.37)$$

such that P is now in the x - z plane. Taking $A \rightarrow 0$ it is possible to recover the ideal g_2 distribution. n_{2D} is fitted with a g_2 function, and the extracted radius in terms of P_T gives an estimate of the expected temperature measurement error. P_T is used to avoid confusion between the thermal radius in-trap and in time-of-flight. For large TOF the thermal radius $R_T(t) \rightarrow P_T \times (t/m)$. As shown in Fig. 5.11, it is clearly important to exclude data around the cloud centre from the fit, as it is here where the distribution is most altered. The temperature error expected for the experimental range of a/λ_0 is shown in Fig. 5.12, where the fits to the high-energy wings of the distributions exclude data within P_T of the centre. If the fits are carried out on the entire distribution the errors are about three times as large. We can see that the error is roughly linear in a/λ_0 , and at $a/\lambda_0 = 0.04$ is $\approx -2.5\%$. This is too large to ignore, and must be factored in to our error calculations for the data series.

5.5.2 Interaction energy

When the trapping potential is suddenly turned off to allow the expansion of the cloud, the stored interaction energy can be converted into kinetic energy. The cloud effectively forces itself apart such that on average the atoms travel further than their initial kinetic energies would have taken them, giving an apparent temperature higher than the true value. This produces another a -dependent error, which will not cancel out in our differential measurement.

To get an idea of the magnitude of the effect, the interaction energy per particle at the critical point ε_{int} can be calculated from the average self-consistent mean-field

5.5 Assessment of errors




Figure 5.11: Two-dimensional momentum distribution for an ideal gas and an interacting gas in a harmonic trap. The black line shows the ideal g_2 function, while the red line shows the enhancement of low-momentum states caused by interactions ($a/\lambda_0 = 0.05$). The distributions are normalised to the same atom number. This illustrates the need to exclude data from the central region when fitting the images. The dashed line shows the experimental exclusion region for temperature fitting, equivalent to a single thermal radius in time-of-flight.




Figure 5.12: ΔT error associated with the modified momentum distribution. The error is roughly linear in a/λ_0 . The shift is negative and ~ 1 -2% for most of our experimental data.




Figure 5.13: Variation of ϵ_{int} per $k_B T$ with interaction strength. The black points show the results of equation 5.40, where the modification of the density distribution is taken into account. The red dashed line shows the result assuming a fixed ideal gas distribution.

value of $2gn_c(\mathbf{r})$. This is given by

$$\epsilon_{\text{int}} = \frac{E_{\text{int}}}{N} = 2g \int_0^\infty 4\pi r^2 n_c(r)^2 dr \Big/ \int_0^\infty 4\pi r^2 n_c(r) dr \quad (5.38)$$

$$= \frac{2g}{\lambda_0^3} \int_0^\infty R^2 D_c(R)^2 dR \Big/ \int_0^\infty R^2 D_c(R) dR \quad (5.39)$$

$$= Ak_B T \int_0^\infty R^2 D_c(R)^2 dR \Big/ \int_0^\infty R^2 D_c(R) dR . \quad (5.40)$$

As the integrals of equation (5.40) depend only on A , the quantity $\epsilon_{\text{int}}/k_B T$ also depends only on A , so we expect the error to scale with a/λ_0 once more. The dependence is superlinear as the average density increases with a/λ_0 due to the central spreading of the distribution. The dependence is shown in Fig. 5.13, along with the linear scaling that would be obtained if the distribution remained ideal. At $a/\lambda_0 = 0.04$ the interaction energy per particle is over 10% of $k_B T$, a considerable fraction. The effect is limited however by using the zero-crossing of the Feshbach resonance to turn off interactions during time-of-flight. Unfortunately we cannot completely eliminate the effect because of the finite ramping time of the Feshbach field. We also want to avoid changing the interaction strength before the trap is turned off which could invalidate our results.

To estimate the residual effect another mean-field model was used. Randomly distributed starting positions and momenta for 100,000 atoms were generated according to the Bose distribution (for our particular trapping potential). The final positions were then numerically calculated given the time-varying interaction potential experienced during time-of-flight. These positions were then projected along one-axis to produce a fake image for fitting. This process was repeated to produce many such images. To keep it reasonably simple the interaction potential during time-of-flight was modelled on the




Figure 5.14: Example of our Feshbach field ramp during TOF, used to minimise interaction broadening of the cloud. Black points show the field strength and red points the corresponding ^{39}K scattering length. In this example the field is ramped down from close to resonance towards the zero crossing value of 350 G immediately after the trap is turned off (dashed line). The apparent quantisation in the scattering length is due to the resolution of the current meter used to monitor the ramp, and the strong scaling of scattering length with field near the resonance.

density distribution of an ideal gas released from a harmonic trap, given in equation (2.56). The interaction strength could also be varied in time to compare what happens when interactions persist during the whole TOF, when they are suddenly switched off at some point, or when they are gradually ramped off.

Under typical experimental conditions the model predicts (when fitted using an exclusion radius of R_T) that if interactions are sharply switched off after 2 ms TOF there is no discernible temperature shift for $a = 56 a_0$ (i.e. our reference data), while for $a = 357 a_0$ there is a residual shift of $\sim 2\%$. This latter value corresponds to $a/\lambda_T \approx 0.034$ for the temperature used in the model of 250 nK. This 2% shift is a rough upper limit, as can be seen by looking at the actual field variation during time-of-flight given in Fig. 5.14.

It was also possible to investigate this effect experimentally by comparing the fitted temperatures for interactions on during TOF with those attained when ramping off interactions for various a . The results, shown in Fig. 5.15, confirm the general scaling of the effect with a/λ_T , and also helped to validate the above model. Note that here $\Delta T = T_{\text{on}} - T_{\text{off}}$, the difference between the observed temperature with interactions left on during time-of-flight and when they are ramped off, so the residual systematic error (which we are trying to quantify) is not included.

Fortuitously, this positive shift tends to counteract the negative shift caused by the change to the momentum distribution discussed in the previous section, resulting in a net error of at most $\pm 1\%$ at $a/\lambda_0 = 0.04$. This is large enough to take account of in our error assessments, but not so large as to greatly modify the results themselves.




Figure 5.15: Experimental observation of the shift in measured temperature due to interactions during time-of-flight. ΔT is the difference between the temperature measured when interactions are left on until just before imaging and the temperature measured when interactions are tuned to zero directly after the trap is turned off. For the right-most data point this turn-off follows the ramp of Fig. 5.14. All of the data was taken at a temperature of ~ 260 nK, so multiplying the x -scale by 10^{-4} gives a rough conversion to a/λ_T . The error bars show the standard error, while the red line is an unweighted linear fit to the data points.

5.6 Non-equilibrium effects

As previously stated, data was taken in the range $0.001 < a/\lambda_0 < 0.06$, but only that for which $0.005 < a/\lambda_0 < 0.04$ was shown for the T_c shift results. This is because outside this region the atom-loss rate becomes comparable to the re-thermalisation rate, and the condition $\alpha \gg 1$ discussed in §2.7.2 is violated. As thermal equilibrium is never fully established in a system with continuous dissipation (we can only be ‘close’ to equilibrium) we need a better way to quantify this condition, while also confirming that global (rather than just local) equilibrium can be achieved.

The relevant quantities are the elastic collision rate Γ_{el} , the hold time¹ in the trap t_{hold} , the trapping frequency $\bar{\omega}$ and a characteristic atom-number decay time τ . The definition of τ is guided by our measurement precision. We measure N_c within $\sim 1\%$, so we require that the gas continuously re-equilibrates on a time-scale τ corresponding to 1% atom loss. As long as this occurs the distribution of atoms between the thermal and condensed components should reflect the conditions of equilibrium to within the measurement precision.

This condition is therefore parameterised by the quantity $\Gamma_{\text{el}}\tau$. Equilibrium is usually deemed to require about 3 collisions per particle [67]. For all the data used for

¹The hold time t_{hold} is taken to be the time from which the scattering length a reaches its measurement value to the moment at which the atoms are released from the trap.

5.6 Non-equilibrium effects

the T_c shift investigation we made sure that $\Gamma_{\text{el}}\tau > 5$, where Γ_{el} is given by

$$\Gamma_{\text{el}} = n(0)v_{\text{th}}\sigma_0/2, \quad (5.41)$$

where

$$n(0) = N \left(\frac{m\bar{\omega}^2}{2\pi k_{\text{B}}T} \right)^{3/2}, \quad v_{\text{th}} = \sqrt{\frac{8k_{\text{B}}T}{\pi m}} \text{ and } \sigma_0 = 8\pi a^2, \quad (5.42)$$

and τ was found from our experiments near N_c . The expression for the elastic scattering rate is the classical result in a harmonic trap used in [157].

For global equilibrium to hold, the condition $t_{\text{hold}} > \tau > 1/\bar{\omega}$ must be satisfied. This ensures there is enough time for the measurement conditions to be reflected across the full extent of the cloud. Global equilibrium will also only hold outside the hydrodynamic regime (in which the cloud becomes collisionally opaque), governed by the condition $\Gamma_{\text{el}}/\bar{\omega} < 1$.

The results of our T_c measurement for $N \approx 4 \times 10^5$ ($\lambda_0 \approx 10^4 a_0$) over the full range of a/λ_0 values studied is displayed in Fig. 5.16 **a**. All data shown still satisfy the conditions $t_{\text{hold}} > \tau > 1/\bar{\omega}$ and $\Gamma_{\text{el}}t_{\text{hold}} > 5$, sufficient for energy to be redistributed through the system once the measurement conditions have been established. However, if $\Gamma_{\text{el}}\tau$ is too small the elastic collisions are unable to ‘keep up’ with the continuous dissipation. This cannot be rectified by simply increasing t_{hold} , and the system is inherently out of equilibrium.

In Fig. 5.16 **b** we plot $\Gamma_{\text{el}}\tau$ for each T_c measurement, suggesting that quasi-equilibrium appears to be lost for a universal value of $\Gamma_{\text{el}}\tau$, despite $\Gamma_{\text{el}} \approx 0.7\text{-}1000 \text{ s}^{-1}$ and $\tau \approx 2\text{-}1000 \text{ ms}$ individually varying by around three orders of magnitude. Also shown is the hydrodynamic parameter $\Gamma_{\text{el}}/\bar{\omega}$.

The recorded values of $\Delta T_c/T_c^0$ where $\Gamma_{\text{el}}\tau > 5$ is violated have no quantitative meaning, as they depend heavily on our initial conditions, as is intrinsically true to all non-equilibrium phenomena. We can qualitatively interpret the results based on our experimental procedure. In the small- a limit there is a smooth apparent rise in T_c above T_c^0 for $a \rightarrow 0$. The loss mechanism in this regime is predominantly single-body, affecting N_0 and N_T equally. Elastic collisions act to redistribute the atoms between the two components to maintain an equilibrium population distribution. In this case this corresponds to a net transfer of atoms from the condensate to the thermal cloud. However, as a decreases so does $\Gamma_{\text{el}}\tau$, and a condensate is observed for total N lower than the expected equilibrium N_c (corresponding to a positive measurement of $\Delta T_c/T_c^0$). Strictly speaking of course T is not defined out of equilibrium, but the observed effect is sufficiently small that the fitted equilibrium distribution still provides a good measure of the energy content of the cloud.

In the large- a limit, ‘super-heated’ condensates surviving at a temperature above the equilibrium T_c are also observed, but the explanation here is less trivial. Three-body losses affect the condensed and thermal fractions differently, the thermal component is




Figure 5.16: **a** $\Delta T_c/T_c^0$ for $N \approx 4 \times 10^5$ atoms over the full experimental range, where all data is taken using the procedure which assumes thermal equilibrium. The vertical dashed lines show the bounds of thermal equilibrium. At both low and high a there is sharp deviation from the smooth equilibrium curve (red line). **b** Equilibrium criteria, showing the values of $\Gamma_{el}\tau$ (solid squares) and $\Gamma_{el}/\bar{\omega}$ (open circles) that correspond to each data point in **a**. The horizontal dashed line shows $\Gamma_{el}\tau = 5$. The onset of the hydrodynamic regime occurs at $\Gamma_{el}/\bar{\omega} = 1$.

far from saturation and the cloud enters the hydrodynamic regime. Without further investigation it is hard to give even a proper qualitative explanation for the outcome observed in our system.

5.7 Conclusions

Using the beneficial properties of a Feshbach resonance in ^{39}K and a novel differential thermometry technique, we have taken high precision measurements of the T_c shift over a far broader range of a/λ_0 than has previously been achieved. In doing so we have shown good agreement with mean-field theory below $a/\lambda_0 \approx 0.01$, and above this value observed positive deviation from mean-field theory, showing clear evidence for the condensation-enhancing effects of critical correlations for the first time. In addition we have probed beyond the regime in which equilibrium measurements can be taken. In both very-weakly and very-strongly interacting limits non-equilibrium effects have been observed. The results of this chapter are summarised in [158].

It is hoped that the new information provided by our investigation will act to guide and encourage further theoretical studies of the difficult problem of beyond-mean-field effects in a harmonically trapped gas. It would once again be of interest to

5.7 Conclusions

study the behaviour in a variety of trapping geometries, tending towards a hard box potential. Combined with direct investigation of local behaviour of a trapped gas using in situ measurements, this could provide a clearer pathway to investigate the build-up of coherence across a Bose gas.

It should be possible to extend the range of a/λ_0 over which equilibrium measurements can be taken by altering other system parameters in combination with a (with which the three-body collision rate scales strongly). However, in our picture more precise measurements necessarily restrict the range that equilibrium measurements may be taken over, so it is not clear whether it will be possible to reveal logarithmic terms in the form of the T_c shift with this method.

The non-equilibrium behaviour seen will also vastly benefit from dedicated investigation in order to disentangle the competing effects in the strongly-interacting regime. This should prove fruitful for comprehending condensation in intrinsically out-of-equilibrium systems, such as polariton gases [104, 159].

Chapter 6

Conclusion

“ Well, it was worth it, Jeeves. ”

– *P. G. Wodehouse (The Code of the Woosters)*

In this concluding chapter a brief summary of the accomplishments from each aspect of the work covered in this thesis is given, before the direction of more recent and future work is discussed.

6.1 Summary

A timeline of key events in the history of the group since its foundation is shown in Fig. 6.1. From this it is clear that the majority of experimental time over the course of this PhD was spent creating the two working sets of apparatus to produce Bose-Einstein condensates. Once a stable system for producing ^{39}K BECs was in place, the scientific investigations progressed rapidly and promise to continue to do so. The primary achievements of the work described in this thesis are as follows.

- A single-chamber system has been created, allowing the rapid production (~ 15 s cycle) of ^{87}Rb BECs of over 10^5 atoms in an optically plugged quadrupole trap. This system is briefly described in Appendix B.



Figure 6.1: Timeline of key events over the course of this thesis. The suffixes I and II refer to the single-chamber and two-chamber experimental systems respectively. Dates are marked for the first occurrence of magneto-optical trapping and Bose-Einstein condensation in each system and for each atomic species. The publication dates of the papers [85, 116, 158] are also marked.

6.2 Outlook

- A two-chamber system has been created, and is only the second system worldwide to have produced ^{39}K BECs. These are held in a crossed-dipole trap, allowing tuning of the strength of interactions in the system via a Feshbach resonance. The maximum condensate number of over 4×10^5 is the highest ever reported for this atomic species. In the same system it is possible to produce ^{87}Rb BECs of over 8×10^5 atoms. The setup and operating procedure are described in Chapters 2 and 3 respectively.
- The non-saturation of an interacting Bose gas in a harmonic trap has been quantified, both near the critical point and for larger condensed fractions. Extreme deviations from the behaviour of an ideal gas are seen, even for what are generally considered weak interactions. This goes against preconceptions that are commonly held in the field. Excellent quantitative agreement with mean-field theory is seen in the low- N_0 limit, while extrapolation to the non-interacting limit provides the most compelling experimental evidence to date of ideal saturation as predicted by Einstein. This work is described in Chapter 4.
- The interaction induced T_c shift for the BEC transition has been experimentally recorded at high precision over a greater range of interaction strengths than ever before. Good agreement with mean-field theory is observed in the weakly-interacting limit, while for stronger interactions beyond-mean-field effects have been observed for the first time. The positive direction of this shift is in agreement with widely accepted theoretical prediction. In addition to this, in both the very-weakly and very-strongly interacting limits non-equilibrium behaviour has been detected. This work is described in Chapter 5.

6.2 Outlook

It is hoped that the work carried out in the group thus far is just the tip of the experimental iceberg. Now that the essential machinery is in place there is a wealth of possible topics to explore, both as direct extensions to this work and other more diverse projects. Several such areas for exploration at various stages of conception are discussed below.

6.2.1 Condensed fraction induced by critical correlations

An investigation directly related to the work on the T_c shift has already taken place [160]. The downside to the earlier study is that no straightforward link can be made to the uniform system due to the intrinsic dependence of T_c on the global properties of a non-uniform system. In the more recent work this problem is negotiated by observing the condensed fraction $f_0 = N_0/N$ in a harmonic trap at the mean-field critical number $N = N_c^{\text{MF}} > N_c$. The condensed fraction observed at this point is then solely

due to the effects of critical-correlations. Moreover, the effect is limited to the quasi-uniform central region of the harmonic trap with volume $\propto (a/\lambda_0)^3$. Combined with the expected scaling $n_0/n \propto \Delta n_c \propto a/\lambda_0$ in a uniform system leads to a predicted scaling $f_0 \propto (a/\lambda_0)^4$. This scaling has been experimentally observed. Furthermore, the measured effect has been quantitatively compared (via the local density approximation) with the Monte Carlo simulated prediction for a uniform system, and shows excellent agreement.

The ability to tune the strength of interactions with the ^{39}K Feshbach resonance was crucial to this investigation. Not only was it necessary to cover a wide range of a/λ_0 values and to use the differential measurement scheme, but also the largest value of f_0 recorded at $N = N_c^{\text{mf}}$ was only $\sim 1\%$. It was therefore essential to also use the methods for precise measurement of small condensed fractions developed during the saturation and T_c shift investigations.

6.2.2 Out-of-equilibrium studies

The non-equilibrium behaviour observed during the T_c shift measurement has sparked further interest in this area. Studies are well underway into the dynamics of a Bose gas equilibrating from an out-of-equilibrium starting point. The method involves the preparation of an ultra-cold cloud at a relatively high scattering length (where there is only a small condensed fraction) and then instantaneously reducing the scattering length. The system that was in thermodynamic equilibrium suddenly is no longer, as the value of N_c suddenly drops. At this point elastic collisions act to return the system to equilibrium by the redistribution of energy between the atoms, resulting in a rapid rise in the condensed fraction. The energy released by condensation results in the heating of the remaining thermal atoms, which increases N_c , so that we do not simply acquire the condensed fraction expected for the new scattering length at the original temperature. The details of this process have recently been submitted for review [161]. It is also planned to start the process from a position where there is no initial BEC, to observe the phase transition as driven by the change of interaction strength.

6.2.3 Two-dimensional Bose gases

The planned extension of the system to incorporate two-dimensional trapping has been prepared, and will hopefully be implemented in the near future. This will allow direct extension of the saturation experiment to reduced dimensionality, and allow detailed studies of the Berezinskii-Kosterlitz-Thouless (BKT) transition [162, 163] with the ability to tune interactions. Two-dimensional Bose gases come with a rich variety of new phenomena, where in a uniform system superfluidity exists in the absence of a BEC. The addition of 2D trapping to the system will therefore open up another branch of future research topics [164].

6.2 Outlook

6.2.4 Annular Bose-Einstein condensates

In a brand new direction, work has been carried out on the two-chamber system to allow transfer of a BEC from the CDT to a ring-shaped trap, created by the intersection of a red-detuned hollow Laguerre-Gauss beam with a thin light sheet. A two-photon Raman process [165] is used to impart angular momentum to the ^{87}Rb condensate. In a toroidal geometry the BEC total angular momentum is quantised. The superfluid flow of the BEC has been shown to perpetuate for long periods, showing the high metastability of the state. There is currently much interest in the decay mechanism of the multiply-charged quantised angular momentum, with possible parallels in systems ranging in size from superconducting nano-wires to rotating (pulsar) neutron stars. The results of these investigations have recently been submitted for review [166].

6.2.5 Power-law traps

The different scaling of physical parameters (such as non-saturation) with the trapping potential make the study of gases in controllable power-law traps a topic of interest [167]. This should be possible using overlapped blue-detuned beams with various Laguerre-Gauss modes, enabling trapping from a typical harmonic potential through to a near-homogeneous hard-walled potential [168]. This would increasingly enhance the effects of critical correlations in a three-dimensional Bose gas, and may even lead to direct global measurements of a close-to-uniform system. The tools necessary to achieve such a scenario are well under way to being developed, with Laguerre-Gauss beams already in use as described above. The direct measurements will involve new interesting challenges, such as accurate temperature determination from the greatly modified time-of-flight distributions expected from these traps. New techniques developed in this direction should prove interesting to the general community as well of course as the results of the investigations themselves.

6.2.6 Superconducting magnetic trap

The single-chamber system is being redeveloped with the planned addition of a dipole trap. It is also being used to test magnetic coils wound from high-temperature superconducting tape, which is now commercially available through *AMSC*. When cooled with liquid nitrogen, this tape can support current densities of close to $10,000 \text{ A cm}^{-2}$. The coils in their present prototype form do not have any advantages over normal copper-wound coils, but it is hoped that the proof-of-principal will allow advances in the future to the advantage of cold atoms physicists. Possible benefits might include the ability to position the coils further from the system than equivalent copper coils, allowing greater optical access to experiments, or producing large magnitude magnetic fields to, for example, access high-field Feshbach resonances. Another intriguing possibility would be to use persistent flow in the coils to produce highly-stable magnetic

fields for experiments.

6.2.7 Final remarks

The outlook is promising for the group and the field as a whole, with many diverting questions still open. While being proud of the scientific achievements of the last four years, the real lasting legacy will hopefully be the two functional BEC systems created during this period. The use of ^{39}K in our system has been key to most of our experiments thus far, and fully vindicates our choice to use this difficult species. The experimental advantages of an accessible Feshbach resonance open up myriad possibilities for the future of many-body research with quantum degenerate gases.

Appendix A

Useful Physical Quantities

In this appendix some useful physical quantities are tabulated that are relevant to the rest of the thesis. The physical constants in Table A.1 are the recommended CODATA values [169]. The majority of the properties of ^{39}K and ^{87}Rb were collated in [23, 170] (see references therein), while the scattering length data is taken from [65, 95, 171].

Quantity	Symbol	Value
Speed of light in free space	c	$2.99792458 \times 10^8 \text{ m s}^{-1}$
Permeability of free space	μ_0	$4\pi \times 10^{-7} \text{ H m}^{-1}$
Permittivity of free space	$\epsilon_0 = 1/(\mu_0 c^2)$	$8.854187817\dots \times 10^{-12} \text{ F m}^{-1}$
Planck's constant	$h = 2\pi\hbar$	$6.62606957(29) \times 10^{-34} \text{ J s}$
Electronic charge	e	$1.602176565(35) \times 10^{-19} \text{ C}$
Electron rest mass	m_e	$9.10938291(40) \times 10^{-31} \text{ kg}$
Atomic mass unit	u	$1.660538921(73) \times 10^{-27} \text{ kg}$
Bohr radius	$a_0 = 4\pi\epsilon_0\hbar^2/(m_e e^2)$	$5.2917720859(36) \times 10^{-10} \text{ m}$
Boltzmann's constant	k_B	$1.3806488(13) \times 10^{-23} \text{ J K}^{-1}$
Bohr magneton	$\mu_B = e\hbar/(2m_e)$	$9.27400968(20) \times 10^{-24} \text{ J T}^{-1}$

Table A.1: Values for some fundamental physical constants, as recommended by the National Institute of Standards and Technology [169].

μ_B/k_B	$\approx 2/3 \text{ K T}^{-1}$	$= 200/3 \mu\text{K G}^{-1}$
μ_B/h	$\approx 1.4 \times 10^{10} \text{ s}^{-1} \text{ T}^{-1}$	$= 1.4 \text{ MHz G}^{-1}$
h/k_B	$\approx 4.8 \times 10^{-11} \text{ K s}$	$= 48 \mu\text{K MHz}^{-1}$

Table A.2: Useful quantities when considering RF evaporation in a magnetic trap.

Property	Symbol	³⁹ K	⁸⁷ Rb
Atomic number	Z	19	37
Mass number	A	39	87
Relative natural abundance	η	93.2581(44) %	27.83(2) %
Atomic mass	m	38.96370668(20) u	86.909180520(15) u
Vapour pressure at 25 °C	P_v	2.42×10^{-8} mbar	5.23×10^{-7} mbar
Nuclear spin	I	3/2	3/2
Electronic ground state		$4^2S_{1/2}$	$5^2S_{1/2}$

Table A.3: General properties of ³⁹K and ⁸⁷Rb.

Property	Symbol	³⁹ K	⁸⁷ Rb
Frequency	$\omega_0/2\pi$	389.286058716(62) THz	377.107463380(11) THz
Wavelength	λ	770.10838505(12) nm	794.978851156(23) nm
Lifetime	τ	26.37(5) ns	27.679(27) ns
Natural linewidth	$\Gamma/2\pi$	6.03(1) MHz	5.7500(56) MHz

Table A.4: Optical properties of the ³⁹K and ⁸⁷Rb D₁ lines.

Property	Symbol	³⁹ K	⁸⁷ Rb
Frequency	$\omega_0/2\pi$	391.01617003(12) THz	384.2304844685(62) THz
Wavelength	λ	766.700921822(24) nm	780.241209686(13) nm
Lifetime	τ	26.37(5) ns	26.2348(77) ns
Natural linewidth	$\Gamma/2\pi$	6.035(11) MHz	6.0666(18) MHz
Recoil velocity	v_{rec}	13.357 mm s ⁻¹	5.8845 mm s ⁻¹
Recoil temperature	T_{rec}	418.06 nK	361.96 nK
Doppler temperature	T_D	144.82 μ K	145.57 μ K
2, 2⟩ → 3, 3⟩ cycling transition, σ^+ polarisation:			
Saturation intensity	I_{sat}	1.750 mW cm ⁻²	1.669 mW cm ⁻²
Cross-section	σ_0	2.807×10^{-9} cm ⁻²	2.907×10^{-9} cm ⁻²

Table A.5: Optical properties of the ³⁹K and ⁸⁷Rb D₂ lines.

Property	Symbol	^{39}K	^{87}Rb
Electron spin g -factor	g_S		2.00231930436153(53)
Electron orbital g -factor	g_L		0.99999369
Landé g -factor	g_J	2.00229421(24)	2.00233113(20)
Nuclear g -factor	g_I	-0.00014193489(12)	-0.000 9951414(10)
Hyperfine splitting	$\Delta E_{\text{hf}}/h$	461.7197202(6) MHz	6834.68261090429(9) MHz

Table A.6: Ground state properties for ^{39}K ($4^2\text{S}_{1/2}$) and ^{87}Rb ($5^2\text{S}_{1/2}$), which may be used in the Breit-Rabi equation (1.33).

Property	Symbol	^{39}K	^{87}Rb
Singlet scattering length	a_s	138.49(12) a_0	90.0(2) a_0
Triplet scattering length	a_t	-33.48(18) a_0	98.99(2) a_0
Interspecies singlet scattering length	a_s^{KRb}		$824_{-70}^{+90} a_0$
Interspecies triplet scattering length	a_t^{KRb}		$35.9(7) a_0$

Table A.7: Singlet and triplet scattering lengths for ^{39}K and ^{87}Rb .

Appendix B

Small system

Much of the first year after our group was founded was spent getting to grips with the experimental techniques required for atom trapping and cooling, taking the form of the construction of a small single-chambered system. The original intention was to create a magneto-optical trap for ^{87}Rb , which would entail setting up a laser system, a vacuum chamber, magnetic coils and computer control. The lessons learned while developing this small system would then aid in the construction of the large system. It soon became apparent though that we might push the limits of this small system, and it was decided that we should attempt to use it to create a BEC. This was achieved about a year after experimental work commenced in the lab.

The vacuum system is shown in Fig. B.1. The quartz optical chamber was made by *Triad Technology*, and features seven radial 1" windows and two vertical 2" windows for optical access, with a ~ 1 cm diameter connection to the rest of the vacuum system. The atom sources are of the same style as the large system, and vacuum is maintained with a 40 l s^{-1} ion pump. The original laser system consisted of a stripped-down version of that shown in Fig. 2.23. When the large system was first built, light was shared between the two systems. The small system has since been moved to a different lab with a new dedicated laser setup¹.

The scheme for magneto-optical trapping is shown in Fig. B.2. The splitting of the cooling and repumping light into six beams was initially achieved using the *Schäfter+Kirchhoff* fibre-port cluster used in the main setup. This has now been replaced with a home-built version. The field for the MOT is provided by a pair of water-cooled magnetic coils situated just next to the 2" windows of the chamber. These are capable of producing a quadrupole field gradient of up to 200 G cm^{-1} along the weak axis, and are also used for capturing and holding the atoms magnetically after an optical molasses stage.

The principal drawback to the single-chamber design is the short trapping lifetime resulting from the necessity to have a reasonable ^{87}Rb pressure for MOT loading. Once the magnetic trap has been loaded the cooling process must therefore proceed rapidly.

¹The compact nature of the small system meant that we could fit it on a small 'bread-board', so it is easily transportable while still under vacuum.




Figure B.1: Small system vacuum setup. **a** Turbomolecular pump. **b** Ion pump. **c** All-metal right-angle valve. **d** Pressure gauge. **e** Optical chamber. **f** Atom sources. There is a window on the opposite side of the vacuum system to the optical chamber to allow additional optical access.

We thus choose to evaporatively cool the atoms using RF radiation directly in the quadrupole trap, which provides tight confinement, aiding the thermalisation rate of the cloud. This requires some means of suppressing Majorana spin flips at the magnetic zero of the quadrupole trap. This we achieve with an ‘optical plug’, this being a tightly focussed ($\sim 30 \mu\text{m}$ waist) blue-detuned gaussian beam directed at the trap centre. This is the method that was used by the Ketterle group to create the first ever sodium BEC [4]. We use plug light with wavelength in the range 700-740 nm produced by a *Coherent* MR110 titanium sapphire laser pumped with a Verdi 18 W laser at 532 nm. Despite the relatively close detuning of the plug beam there is very little heating and atom loss through light scattering, as the tight focus of the laser produces a very steep barrier near the trap centre.

The potential produced by the combination of the quadrupole field, optical plug and evaporating RF field is shown in Fig. B.3. The potential is displayed as a cross-section perpendicular to the optical axis of the plug and around the centre of the quadrupole field. The plug beam is fired horizontally at the magnetic trap such that the asymmetrical magnetic gradients in the vertical and radial directions result in two distinct trap minima (the red regions in Fig. B.3). If the azimuthally symmetric gaussian beam were instead directed along the vertical axis a ring-shaped minimum would be formed. Small imperfections in the laser profile would then lead to one or

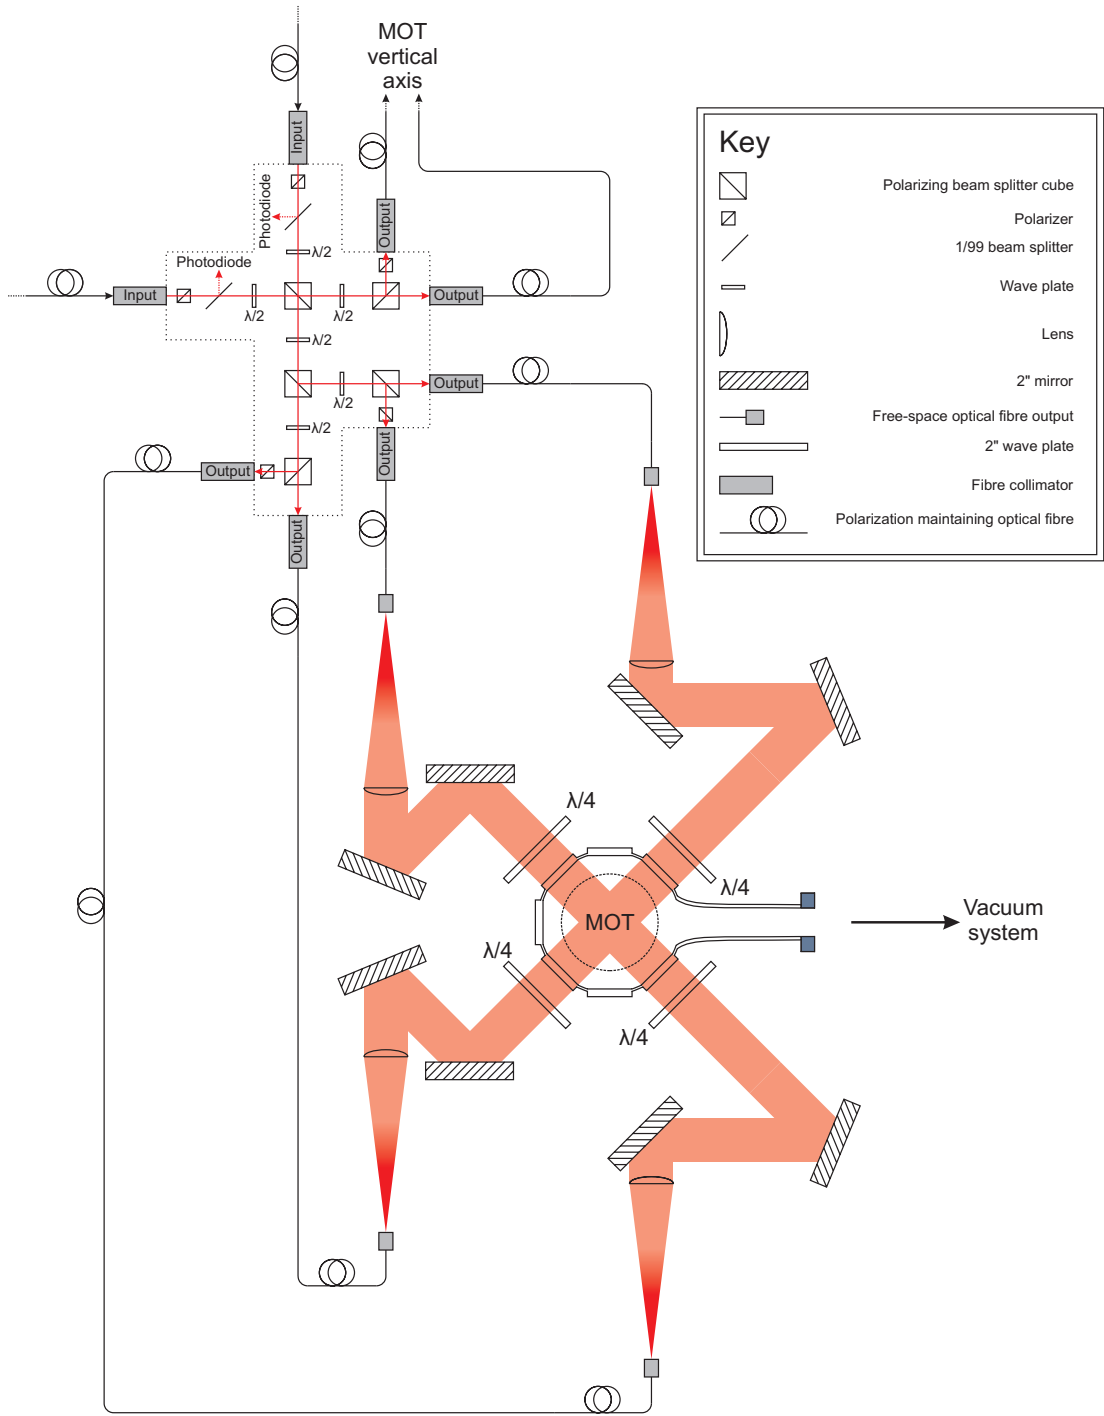


Figure B.2: Schematic of the splitting of the MOT beams and their delivery to the optical chamber.




Figure B.3: Potential produced by the combination of a quadrupole field, optical plug and RF radiation. The plug beam is directed along one of the weak quadrupole axes. This produces two distinct trap minima (shown by the red regions). The RF radiation used for evaporation causes the quadrupole potential to ‘turn over’ as atoms are transferred into untrapped states.

more minima at random locations [172], so it is preferable to have the control provided by our trap geometry. The resultant anisotropic harmonic trap frequencies for this type of potential are given in [4].

The standard experimental sequence is as follows. We load rubidium atoms into the MOT for around 5 s before turning off the magnetic field for 2.5 ms of optical molasses. A $20 \mu\text{s}$ burst of pumping light drives the atoms into the $|2, 2\rangle$ state for magnetic trapping. The atoms are captured by suddenly turning the quadrupole field on with a weak-axis field $B' = 32 \text{ G cm}^{-1}$ before the cloud is compressed by increasing B' to 100 G cm^{-1} in 0.5 s. An RF field is then turned on and held at 10 MHz for 0.3 s of plain evaporation. At this point there are typically 10^8 atoms at $100 \mu\text{K}$. The RF frequency is swept linearly to 2 MHz in 3 s, reducing both the atom number and temperature by a factor of 10. To slow the loss rate due to three-body collisions we then decompress the trap to $B' = 40 \text{ G cm}^{-1}$ in 200 ms. Finally, we apply a final exponential evaporation sweep from 1.4 MHz to a chosen final value in 0.5 s. This results in quasi-pure BECs of over 10^5 atoms¹. A sample sequence with increasing time-of-flight for a condensed fraction of $\sim 30\%$ at $\sim 500 \text{ nK}$ is shown in Fig. B.4. The bimodality of the cloud is revealed as TOF increases and the thermal atoms expand more rapidly than the BEC. The fact that we see only one BEC is due to the fact that the plug beam is aligned

¹When BECs were being created in this manifestation of the small system we did not calibrate the imaging fudge-factor described in §2.10.6. This atom number is therefore likely to be an underestimation of the true value.




Figure B.4: Optical density images of BECs with $\sim 30\%$ condensed fraction, released from the optically-plugged quadrupole trap and imaged after 3-12 ms time-of-flight. As time-of-flight increases, the bimodal distribution becomes apparent.




Figure B.5: Photograph of the trial system as it is now. The optical chamber can be seen amongst the magnetic coils to the right-hand side of the picture. The original plug beam was aligned with the axis of the connection between the optical chamber and the rest of the vacuum system.

slightly off centre such that one trap minima is higher than the other. During the final evaporation sweep all of the atoms in the higher minima are removed.

After a period of neglect the small system is now under new management and is nearing a usable state once more. The disadvantage of the short trap lifetime is partially compensated for by the extremely fast condensation process and overall simplicity of the system. It is planned to add an optical dipole trap to the system for a comparison with the plug, which has proved problematic at times due to the high alignment accuracy required. The current incarnation of the system is shown in Fig. B.5.

Appendix C

Popov Approximation

The Popov approximation extends Hartree-Fock theory to include small deviations from standard Hartree-Fock theory in the form of Bogoliubov-like excitations. The background of this theory may be found in several text-books, but the notation used below follows that found in [29] most closely. Here it is used to predict deviations from the Hartree-Fock description of non-saturation in a harmonically-trapped interacting Bose gas. The semi-classical limit is taken once again, so the excited state density in this picture (following equation (1.6)) is given as the familiar starting point:

$$n_T(\mathbf{r}) = \frac{1}{(2\pi\hbar)^3} \int \frac{1}{e^{\epsilon(\mathbf{r}, \mathbf{p})/k_B T} - 1} d\mathbf{p} . \quad (\text{C.1})$$

The chemical potential has been absorbed into $\epsilon(\mathbf{r}, \mathbf{p})$ for the purposes of describing the excitations. The process below was laid out by Jean Dalibard.

In the absence of a BEC the excitation spectrum takes the standard form

$$\epsilon(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + 2gn_T(\mathbf{r}) + V(\mathbf{r}) - \mu . \quad (\text{C.2})$$

When a BEC is present we must first generalise the Thomas-Fermi relationship of (1.26) to take account of the effects of the thermal atoms on the condensate distribution. This gives the chemical potential as also required for the self-consistent Hartree-Fock approach,

$$\mu = V(\mathbf{r}) + gn_0(\mathbf{r}) + 2gn_T(\mathbf{r}) . \quad (\text{C.3})$$

The factor of 2 before the excited state density appears due to the inclusion of the exchange term for atoms in different states, which is thus absent for the condensate. The form of $\epsilon(\mathbf{r}, \mathbf{p})$ now reflects the revised excitation spectrum, found by diagonalising a modified form of the Hartree-Fock Hamiltonian¹. This is given by (from equation (8.119) of [29]):

$$\epsilon(\mathbf{r}, \mathbf{p}) = \left[\epsilon_{\text{HF}}^2(\mathbf{r}, \mathbf{p}) - (gn_0(\mathbf{r}))^2 \right]^{1/2} = \left[\left(\frac{p^2}{2m} + 2gn_0(\mathbf{r}) \right) \frac{p^2}{2m} \right]^{1/2} , \quad (\text{C.4})$$

¹The added term allows for creation and annihilation of momentum-conserving pairs of particles.

where

$$\epsilon_{\text{HF}} = \frac{p^2}{2m} + 2g(n_0(\mathbf{r}) + n_T(\mathbf{r})) + V(\mathbf{r}) - \mu = \frac{p^2}{2m} + 2gn_0(\mathbf{r}) . \quad (\text{C.5})$$

In this case the relation (C.1) is modified by a factor to reflect the number of non-condensed particles associated with an excitation:

$$n_T(\mathbf{r}) = \frac{1}{(2\pi\hbar)^3} \int \frac{\epsilon_{\text{HF}}(\mathbf{r}, \mathbf{p})}{\epsilon(\mathbf{r}, \mathbf{p})} \frac{1}{e^{\epsilon(\mathbf{r}, \mathbf{p})/k_{\text{B}}T} - 1} d\mathbf{p} . \quad (\text{C.6})$$

The above equations are local, so we can simplify the situation by omitting the spatial dependence for now and switching to dimensionless variables

$$P = \frac{p}{\sqrt{mk_{\text{B}}T}}, \quad A = \frac{4a}{\lambda_T}, \quad D = n\lambda_T^3, \quad \text{and} \quad \kappa = \frac{\mu - V}{k_{\text{B}}T} . \quad (\text{C.7})$$

Working in terms of the phase-space density D , we find in the purely thermal case (cf. §5.2)

$$D_T = g_{3/2}(e^{\kappa - AD_T}) . \quad (\text{C.8})$$

When a condensate is present this is replaced by

$$D_T = \sqrt{\frac{2}{\pi}} \int_0^\infty \frac{P(P^2 + AD_0)}{\sqrt{P^2 + 2AD_0}} \frac{1}{e^{P\sqrt{P^2 + 2AD_0}/2} - 1} dP , \quad (\text{C.9})$$

where

$$AD_0 = 2(\kappa - AD_T) \quad (\text{C.10})$$

via equation (C.3). For computational purposes it is advantageous to use the substitution $x^2 = P^2(P^2 + 2AD_0)/4$ to rewrite (C.9) as

$$D_T = I(\kappa - AD_T), \quad \text{where} \quad I(z) = \frac{2}{\sqrt{\pi}} \int_0^\infty \frac{[(z^2 + x^2)^{1/2} - z]^{1/2}}{e^x - 1} dx , \quad (\text{C.11})$$

and to use the integral form of $g_{3/2}(e^z)$ given in (1.9) for when there is no condensate. Using $z = \kappa - AD_T$, we must find the value $z = z_s$ that solves the equation

$$\frac{\kappa - z}{A} = F(z) , \quad (\text{C.12})$$

where $F(z)$ is defined for all z by

$$F(z) = \begin{cases} g_{3/2}(e^z) & \text{for } z \leq 0 , \\ I(z) & \text{otherwise.} \end{cases} \quad (\text{C.13})$$

If z_s is negative there is no BEC, and $D_T = (\kappa - z_s)/A$. If z_s is positive, $D_T = (\kappa - z_s)/A$ and $D_0 = 2z_s/A$. Fig. C.1 gives a graphical representation of equation (C.12).

When $\kappa \leq 0$ the solution is straightforward and $z_s \leq 0$, corresponding to $D_0 = 0$. When $\kappa > 0$ there are three possible cases, dependent on A , labelled **i-iii**:

- **i**: If A is larger than some non-trivial value, there is a single solution where no




Figure C.1: Graphical solutions for the self-consistent Popov approximation. Note that the function $F(z)$ is continuous at $z = 0$. When $\kappa \leq 0$ there is only one solution (as we consider only positive A). When $\kappa > 0$ there are three possible cases, discussed in the text. For case **ii**, where there are multiple solutions, we choose the negative z result.

condensate is present.

- **ii:** For some intermediate value of A smaller than this threshold but larger than $\kappa/\zeta(3/2)$ there are three possible solutions, two with a condensate and one without.
- **iii:** For $A < \kappa/\zeta(3/2)$ there again exists only one solution, this time giving a finite condensate density.

We choose to omit the solutions of case **ii** for which a condensate is present, based only on the assumption that the local phase-space density reaches the ideal $\zeta(3/2)$ before a condensate appears. We therefore obtain solutions where $D_0 = 0$ for $\kappa \leq \zeta(3/2)A$, and $D_0 > 0$ otherwise.

To find bulk solutions for a harmonically-trapped gas we consider the simple isotropic case where the confining potential is given by $V(\mathbf{r}) = m\omega^2 r^2/2$. To carry on working with dimensionless variables we set

$$R = \frac{r}{R_T}, \quad \text{where } R_T = \sqrt{\frac{k_B T}{m\omega^2}}, \quad (\text{C.14})$$

and reinstate the position dependence of the other system parameters. κ is therefore given by

$$\kappa(R) = \alpha - \frac{R^2}{2}, \quad \text{where } \alpha = \frac{\mu}{k_B T}, \quad (\text{C.15})$$

and the number of atoms for each component is related to the phase-space density by

$$N_{0,T} = \frac{1}{(2\pi)^{3/2}} \left(\frac{k_B T}{\hbar\omega} \right)^3 \int_0^\infty 4\pi R^2 D_{0,T}[\kappa(R), A] dR . \quad (\text{C.16})$$

A final simplification is provided by setting $S = R^2/2$ and giving the results in terms of the ideal critical number N_c^0 , resulting in

$$\frac{N_{0,T}}{N_c^0} = \frac{2}{\zeta(3)\sqrt{\pi}} \int_0^\infty \sqrt{S} D_{0,T}[\alpha - S, A] dS . \quad (\text{C.17})$$

To calculate the data of Fig. 4.9 this procedure was carried out for the two most extreme values of A used in experimental series. The data was then normalised for comparison with both Hartree-Fock theory and the data by dividing by N_c , the value of N_T when N_0 becomes finite in the Popov approximation, incorporating the inherent shift of N_c for finite A .

Appendix D

Assessment of absolute errors

D.1 Finite imaging time

When taking a photo of a moving object with a stationary camera (such as a tennis ball in flight) the finite shutter time on the camera introduces motion blur, and the object appears to be spread along its axis of motion. The same is true when imaging the cloud of potassium atoms after $t_{\text{TOF}} = 20$ ms time-of-flight, where it has a finite velocity gt_{TOF} , while the image is taken with a $\tau_{\text{image}} = 80 \mu\text{s}$ pulse of light (Fig. D.1). The effect can be characterised by the ratio of the distance moved in this time, $d \approx \tau_{\text{image}} \times gt_{\text{TOF}}$, to the thermal width of the cloud R_T ,

$$d' = \frac{d}{R_T} . \quad (\text{D.1})$$

In our experiment $d \approx 16 \mu\text{m}$, while $R_T \approx 130 \mu\text{m}$ for the lowest temperature data series, giving $d' = 0.12$. d here is a reasonable fraction of the thermal radius, and instinctively one might think this would produce a large distortion in the measured temperature.

To estimate the magnitude of the effect, the imaged cloud was modelled as an ideal g_2 exponential function. This was convolved with a unit box function of width d in the vertical z direction, and the resulting distribution was fitted with a g_2 function once more where the x and z widths were permitted to vary independently. The data within a single thermal radius R_T was excluded from the fit, which results in a perceived temperature closer to the true value. The temperature shift is extracted by squaring the widths given by the fit. The results of this investigation are shown in Fig. D.2. Temperature shifts ΔT_x and ΔT_z are found for the horizontal and vertical directions, and the average ΔT_{av} is also shown. Of course the effect is predominantly along the direction of motion z , but there is also a small shift in the horizontal direction.

Surprisingly the effect is negligible for even our coldest clouds, only exceeding the per cent level for $d' \gtrsim 0.5$. The scaling of d' shows that the effect is enhanced for increasing τ_{image} or for particularly cold clouds and massive atoms. It is interesting to note that for long time-of-flight (when the cloud width is much greater than its in-trap width) $R_T \approx \sqrt{k_{\text{B}}T/m} \times t_{\text{TOF}}$, meaning d' is essentially independent of time-of-flight.




Figure D.1: Illustration of the effect on the perceived cloud width for an atom cloud that is falling during imaging. The effect is analogous to motion blur in photography. **a** A tennis ball falling at $\sim 2.5 \text{ ms}^{-1}$ captured with a shutter time of $1/800 \text{ s}$ is sharp, and the size is easily measured. **b** When the shutter time is increased by a factor of 8 the image of the ball is noticeably blurred vertically. **c** A cloud imaged with a short imaging pulse (i.e. $d' \approx 0$) shows the characteristic circular symmetry of free thermal expansion. **d** If the imaging time is increased the image is blurred along the z direction, leading to inaccuracies in the temperature measurement. Here $d' = 0.75$ to emphasize the effect.




Figure D.2: Temperature measurement error ΔT caused by motion blur with finite imaging time. The error is dominated by the spreading of the distribution along the z -axis. The greyed region represents the range of d' covered by our measurements. The errors of $\ll 1\%$ may be safely ignored.

The key point for us though is that the shift is $\ll 1\%$ for all our measurements, so it can safely be ignored.

D.2 Feshbach field curvature

The Feshbach coils were designed to be as close to Helmholtz configuration as possible, but they nonetheless produce a finite radial field curvature, calculated to be $\sim 0.00325 \text{ G cm}^{-2} \text{ A}^{-1}$. The zero crossing of the Feshbach resonance is reached with a current of 129 A, so during time-of-flight the atoms experience a field curvature $-B''/2 \approx 0.42 \text{ G cm}^{-2}$. This causes radial deceleration of atoms in the expanding cloud (acceleration for negative curvature) as they fall along the z -axis. Assuming this simple quadratic form for the field magnitude holds over the region in which the atoms move we may again use a simple model to calculate the magnitude of the effect of this curvature. For positive curvature we simply have a very weak harmonic trap with trapping frequency ω_{FB} and it is straightforward to calculate the radial motion of the atoms. Assuming all the particles start at the centre of this weak potential the final position is given (in a one-dimensional description) by

$$x = \frac{v_i}{\omega_{\text{FB}}} \sin(\omega_{\text{FB}} t_{\text{TOF}}), \quad (\text{D.2})$$

where v_i is the initial velocity of the atom. The position the particle would have reached without the field curvature is $x_0 = v_i t_{\text{TOF}}$. The perceived temperature T divided by the real temperature T_0 is then simply given by

$$\frac{T}{T_0} = \left(\frac{x}{x_0} \right)^2 = \frac{\sin^2(\omega_{\text{FB}} t_{\text{TOF}})}{(\omega_{\text{FB}} t_{\text{TOF}})^2}. \quad (\text{D.3})$$

For negative field curvature a similar expression involving \sinh rather than \sin is found. The value for ω_{FB} with the above field curvature for our potassium atoms in the $|1, 1\rangle$ state is $\sim 2\pi \times 1.2 \text{ Hz}$, so for our 20 ms TOF the field curvature should cause a perceived shift in T_x of $\sim -0.8\%$. The cloud temperature is calculated from T_z as well, so this shift can be halved. This can't accurately be accounted for as the field won't be purely quadratic and all of the atoms aren't released from the centre of the Feshbach field. However, the rough magnitude of the effect is still small enough to be ignored, though it is large enough to be wary of in other experimental setups using large Feshbach bias fields.

Bibliography

- [1] A. Einstein. Quantentheorie des einatomigen idealen gases. *Sitzungsber. Preuss. Akad. Wiss.* **1925**, 3 (1925).
- [2] S. N. Bose. Plancks gesetz und lichtquantenhypothese. *Z. Phys.* **26**, 178 (1924).
- [3] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman & E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science* **269**, 198 (1995).
- [4] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn & W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.* **75**, 3969 (1995).
- [5] I. Bloch, J. Dalibard & W. Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80**, 885 (2008).
- [6] L. De Sarlo, P. Maioli, G. Barontini, J. Catani, F. Minardi & M. Inguscio. Collisional properties of sympathetically cooled ^{39}K . *Phys. Rev. A* **75**, 022715 (2007).
- [7] G. Roati, M. Zaccanti, C. D'Errico, J. Catani, M. Modugno, A. Simoni, M. Inguscio & G. Modugno. ^{39}K Bose-Einstein condensate with tunable interactions. *Phys. Rev. Lett.* **99**, 010403 (2007).
- [8] G. Modugno, G. Ferrari, G. Roati, R. J. Brecha, A. Simoni & M. Inguscio. Bose-Einstein condensation of potassium atoms by sympathetic cooling. *Science* **294**, 1320 (2001).
- [9] G. Roati, F. Riboli, G. Modugno & M. Inguscio. Fermi-Bose quantum degenerate ^{40}K - ^{87}Rb mixture with attractive interaction. *Phys. Rev. Lett.* **89**, 150403 (2002).
- [10] S. E. Pollack, D. Dries, M. Junker, Y. P. Chen, T. A. Corcovilos & R. G. Hulet. Extreme tunability of interactions in a ^7Li Bose-Einstein condensate. *Phys. Rev. Lett.* **102**, 090402 (2009).
- [11] J. Stenger, S. Inouye, M. R. Andrews, H.-J. Miesner, D. M. Stamper-Kurn & W. Ketterle. Strongly enhanced inelastic collisions in a Bose-Einstein condensate near Feshbach resonances. *Phys. Rev. Lett.* **82**, 2422 (1999).
- [12] M. Zaccanti, B. Deissler, C. D'Errico, M. Fattori, M. Jona-Lasinio, S. Muller, G. Roati, M. Inguscio & G. Modugno. Observation of an Efimov spectrum in an atomic system. *Nature Physics* **5**, 586 (2009).
- [13] T. Kishimoto, J. Kobayashi, K. Noda, K. Aikawa, M. Ueda & S. Inouye. Direct evaporative cooling of ^{41}K into a Bose-Einstein condensate. *Phys. Rev. A* **79**, 031602 (2009).
- [14] N. R. Claussen, S. J. J. M. F. Kokkelmans, S. T. Thompson, E. A. Donley, E. Hodby & C. E. Wieman. Very-high-precision bound-state spectroscopy near a ^{85}Rb Feshbach resonance. *Phys. Rev. A* **67**, 060701 (2003).

- [15] S. Dürr, T. Volz & G. Rempe. Dissociation of ultracold molecules with Feshbach resonances. *Phys. Rev. A* **70**, 031601 (2004).
- [16] A. D. Lange, K. Pilch, A. Prantner, F. Ferlaino, B. Engeser, H.-C. Nägerl, R. Grimm & C. Chin. Determination of atomic scattering lengths from measurements of molecular binding energies near Feshbach resonances. *Phys. Rev. A* **79**, 013622 (2009).
- [17] C. Chin, R. Grimm, P. Julienne & E. Tiesinga. Feshbach resonances in ultracold gases. *Rev. Mod. Phys.* **82**, 1225 (2010).
- [18] N. N. Bogoliubov. On the theory of superfluidity. *J. Phys. (USSR)* **11**, 23 (1947).
- [19] L. Pitaevskii & S. Stringari. *Bose-Einstein Condensation* (Oxford University Press, 2003).
- [20] C. J. Foot. *Atomic Physics* (Oxford University Press, 2005).
- [21] E. Arimondo, M. Inguscio & P. Violino. Experimental determinations of the hyperfine structure in the alkali atoms. *Rev. Mod. Phys.* **49**, 31 (1977).
- [22] N. Bendali, H. T. Duong & J. L. Vialle. High-resolution laser spectroscopy on the D₁ and D₂ lines of ^{39,40,41}K using RF modulated laser light. *J. Phys. B: At. Mol. Phys.* **14**, 4231 (1981).
- [23] D. A. Steck. Rubidium 87 D line data. <http://steck.us/alkalidata/> (23rd December 2010).
- [24] G. Breit & I. I. Rabi. Measurement of nuclear spin. *Phys. Rev.* **38**, 2082 (1931).
- [25] R. Grimm, Weidemüller & Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. *Adv. At. Mol. Opt. Phys.* **42**, 95 (2000).
- [26] L. D. Landau & E. M. Lifshitz. *Quantum Mechanics* (Elsevier, 1981), third edn.
- [27] B. H. Bransden & C. J. Joachain. *Physics of Atoms and Molecules* (Prentice Hall, 2003), second edn.
- [28] J. Dalibard. Collisional dynamics of ultra-cold atomic gases. In M. Inguscio, S. Stringari & C. Wieman (eds.) *Proceedings of the International School of Physics Enrico Fermi, Course CXL: Bose-Einstein condensation in gases, Varena 1998* (Addison-Wesley, 1998).
- [29] C. J. Pethick & H. Smith. *Bose-Einstein Condensation in Dilute Gases* (Cambridge University Press, 2008), second edn.
- [30] K. Huang. *Statistical Mechanics* (Wiley, 1987), second edn.
- [31] E. Braaten & H. W. Hammer. Universality in few-body systems with large scattering length. *Phys. Rep.* **428**, 259 (2006).
- [32] L. S. Butcher, D. Stacey, C. J. Foot & K. Burnett. Ultracold collisions for Bose-Einstein condensation. *Phil. Trans. R. Soc. Lond., Ser. A* **357**, 1421 (1999).
- [33] H. R. Thorsheim, J. Weiner & P. S. Julienne. Laser-induced photoassociation of ultracold sodium atoms. *Phys. Rev. Lett.* **58**, 2420 (1987).

- [34] A. J. Moerdijk, B. J. Verhaar & A. Axelsson. Resonances in ultracold collisions of ${}^6\text{Li}$, ${}^7\text{Li}$, and ${}^{23}\text{Na}$. *Phys. Rev. A* **51**, 4852 (1995).
- [35] T. Köhler, K. Góral & P. S. Julienne. Production of cold molecules via magnetically tunable Feshbach resonances. *Rev. Mod. Phys.* **78**, 1311 (2006).
- [36] S. Inouye, M. R. Andrews, J. Stenger, H.-J. Miesner, D. M. Stamper-Kurn & W. Ketterle. Observation of Feshbach resonances in a Bose-Einstein condensate. *Nature* **392**, 151 (1998).
- [37] P. O. Fedichev, M. W. Reynolds & G. V. Shlyapnikov. Three-body recombination of ultracold atoms to a weakly bound s level. *Phys. Rev. Lett.* **77**, 2921 (1996).
- [38] H. M. J. M. Boesten, A. J. Moerdijk & B. J. Verhaar. Dipolar decay in two recent Bose-Einstein condensation experiments. *Phys. Rev. A* **54**, R29 (1996).
- [39] J. F. O'Hanlon. *A User's Guide to Vacuum Technology* (John Wiley and Sons Ltd, 2003).
- [40] M. Greiner, I. Bloch, T. W. Hänsch & T. Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Phys. Rev. A* **63**, 031401 (2001).
- [41] H. J. Lewandowski, D. M. Harber, D. L. Whitaker & E. A. Cornell. Simplified system for creating a Bose-Einstein condensate. *J. Low Temp. Phys.* **132**, 309 (2003).
- [42] A. N. Nesmeyanov. *Vapor Pressure of the Chemical Elements* (Elsevier, 1963).
- [43] K. M. Birnbaum. *Cavity QED with multilevel atoms*. Ph.D. thesis, California Institute of Technology (2005).
- [44] J. H. De Boer. *The Dynamic Character of Absorption* (Clarendon Press, 1953).
- [45] A. Chambers, R. K. Fitch & B. S. Halliday. *Basic Vacuum Technology, 2nd Edition* (Institute of Physics Publishing, 1998).
- [46] T. W. Hänsch & A. L. Schawlow. Cooling of gases by laser radiation. *Optics Communications* **13**, 68 (1975).
- [47] V. S. Letokhov, V. G. Minogin & B. D. Pavlik. Cooling and capture of atoms and molecules by a resonant light field. *Sov. Phys. JETP* **45**, 698 (1977).
- [48] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable & A. Ashkin. Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Phys. Rev. Lett.* **55**, 48 (1985).
- [49] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould & H. J. Metcalf. Observation of atoms laser cooled below the Doppler limit. *Phys. Rev. Lett.* **61**, 169 (1988).
- [50] J. Dalibard & C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B* **6**, 2023 (1989).
- [51] Y. Castin, J. Dalibard & C. Cohen-Tannoudji. The limits of Sisyphus cooling. In L. Moi, S. Gozzini, C. Gabbanini, E. Arimondo & F. Strumia (eds.) *Proceedings of the workshop "Light Induced Kinetic Effects on Atoms, Ions and Molecules"*, held in Elba Island, Italy (May 2-5 1990), 5 (ETS Editrice Pisa, 1991).

- [52] S. Chu. Nobel lecture: The manipulation of neutral particles. (1997).
- [53] E. L. Raab, M. Prentiss, A. Cable, S. Chu & D. E. Pritchard. Trapping of neutral sodium atoms with radiation pressure. *Phys. Rev. Lett.* **59**, 2631 (1987).
- [54] C. Fort, A. Bambini, L. Cacciapuoti, F. Cataliotti, M. Prevedellia, G. Tino & M. Inguscio. Cooling mechanisms in potassium magneto-optical traps. *Euro. Phys. J. D* **3**, 113 (1998).
- [55] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman & J. V. Porto. Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential. *Phys. Rev. A* **79**, 063631 (2009).
- [56] T. Esslinger, I. Bloch & T. W. Hänsch. Bose-Einstein condensation in a quadrupole-Ioffe-configuration trap. *Phys. Rev. A* **58**, R2664 (1998).
- [57] H. F. Hess. Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen. *Phys. Rev. B* **34**, 3476 (1986).
- [58] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews & W. Ketterle. Evaporative cooling of sodium atoms. *Phys. Rev. Lett.* **74**, 5202 (1995).
- [59] W. Petrich, M. H. Anderson, J. R. Ensher & E. A. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.* **74**, 3352 (1995).
- [60] W. Ketterle & N. J. van Druten. Evaporative cooling of trapped atoms. In B. Bederson & H. Walther (eds.) *Advances In Atomic, Molecular, and Optical Physics*, vol. 37, 181 (Academic Press, 1996).
- [61] O. J. Luiten, M. W. Reynolds & J. T. M. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. *Phys. Rev. A* **53**, 381 (1996).
- [62] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel & P. Bouyer. All-optical runaway evaporation to Bose-Einstein condensation. *Phys. Rev. A* **79**, 061406 (2009).
- [63] D. J. McCarron, H. W. Cho, D. L. Jenkin, M. P. Köppinger & S. L. Cornish. Dual-species Bose-Einstein condensate of ^{87}Rb and ^{133}Cs . *Phys. Rev. A* **84**, 011603 (2011).
- [64] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell & C. E. Wieman. Production of two overlapping Bose-Einstein condensates by sympathetic cooling. *Phys. Rev. Lett.* **78**, 586 (1997).
- [65] F. Ferlaino, C. D’Errico, G. Roati, M. Zaccanti, M. Inguscio, G. Modugno & A. Simoni. Feshbach spectroscopy of a K-Rb atomic mixture. *Phys. Rev. A* **73**, 040702 (2006).
- [66] M. Anderlini & D. Guéry-Odelin. Thermalization in mixtures of ultracold gases. *Phys. Rev. A* **73**, 032706 (2006).
- [67] C. R. Monroe, E. A. Cornell, C. A. Sackett, C. J. Myatt & C. E. Wieman. Measurement of Cs-Cs elastic scattering at $T = 30\ \mu\text{K}$. *Phys. Rev. Lett.* **70**, 414 (1993).
- [68] I. Gotlibovych. *Microwave manipulation of ultra-cold atoms*. Master’s thesis, University of Cambridge (2010).

- [69] E. Hecht. *Optics* (Addison Wesley, 2002), fourth edn.
- [70] P. Richman. *Development of laboratory equipment for the trapping and cooling of atomic gases*. CPGS report, University of Cambridge (2010).
- [71] T. Müller, B. Zimmermann, J. Meineke, J.-P. Brantut, T. Esslinger & H. Moritz. Local observation of antibunching in a trapped Fermi gas. *Phys. Rev. Lett.* **105**, 040401 (2010).
- [72] C. Sanner, E. J. Su, A. Keshet, R. Gommers, Y.-i. Shin, W. Huang & W. Ketterle. Suppression of density fluctuations in a quantum degenerate Fermi gas. *Phys. Rev. Lett.* **105**, 040402 (2010).
- [73] J. Szczepkowski, R. Gartman, M. Witkowski, L. Tracewski, M. Zawada & W. Gawlik. Analysis and calibration of absorptive images of Bose-Einstein condensate at nonzero temperatures. *Rev. Sci. Instrum.* **80**, 053103 (2009).
- [74] C. F. Ockeloen, A. F. Tauschinsky, R. J. C. Spreeuw & S. Whitlock. Detection of small atom numbers through image processing. *Phys. Rev. A* **82**, 061606 (2010).
- [75] W. Alt. An objective lens for efficient fluorescence detection of single atoms. *Optik* **113**, 142 (2002).
- [76] A. Dareau. *Optical devices for imaging, trapping and manipulating cold atoms*. Master's thesis, University of Cambridge (2010).
- [77] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer & A. Aspect. Critical temperature of a trapped, weakly interacting Bose gas. *Phys. Rev. Lett.* **92**, 030405 (2004).
- [78] W. Ketterle, D. Durfee & D. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In *Proceedings of the International School of Physics - Enrico Fermi* (IOS Press, 1999).
- [79] Y. Castin & R. Dum. Bose-Einstein condensates in time dependent traps. *Phys. Rev. Lett.* **77**, 5315 (1996).
- [80] A. Keshet. <http://akeshet.github.com/Cicero-Word-Generator/> (1st September 2011).
- [81] C. D'Errico, M. Zaccanti, M. Fattori, G. Roati, M. Inguscio, G. Modugno & A. Simoni. Feshbach resonances in ultracold ^{39}K . *New Journal of Physics* **9**, 223 (2007).
- [82] B. D. Esry, C. H. Greene & J. P. Burke. Recombination of three atoms in the ultracold limit. *Phys. Rev. Lett.* **83**, 1751 (1999).
- [83] C. C. Bradley, C. A. Sackett & R. G. Hulet. Bose-Einstein condensation of lithium: Observation of limited condensate number. *Phys. Rev. Lett.* **78**, 985 (1997).
- [84] J. M. Gerton, D. Strekalov, I. Prodan & R. G. Hulet. Direct observation of growth and collapse of a Bose-Einstein condensate with attractive interactions. *Nature* **408**, 692 (2000).

- [85] R. L. D. Campbell, R. P. Smith, N. Tammuz, S. Beattie, S. Moulder & Z. Hadzibabic. Efficient production of large ^{39}K Bose-Einstein condensates. *Phys. Rev. A* **82**, 063611 (2010).
- [86] R. S. Williamson III. *Magneto-optical trapping of potassium isotopes*. Ph.D. thesis, University of Wisconsin (1997).
- [87] C. G. Townsend, N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin & J. Dalibard. Phase-space density in the magneto-optical trap. *Phys. Rev. A* **52**, 1423 (1995).
- [88] J. Weiner, V. S. Bagnato, S. Zilio & P. S. Julienne. Experiments and theory in cold and ultracold collisions. *Rev. Mod. Phys.* **71**, 1 (1999).
- [89] W. Süptitz, G. Wokurka, F. Strauch, P. Kohns & W. Ertmer. Simultaneous cooling and trapping of ^{85}Rb and ^{87}Rb in a magneto-optical trap. *Opt. Lett.* **19**, 1571 (1994).
- [90] M. L. Harris, P. Tierney & S. L. Cornish. Magnetic trapping of a cold Rb-Cs atomic mixture. *J. Phys. B* **41**, 035303 (2008).
- [91] X. Xu, T. H. Loftus, J. W. Dunn, C. H. Greene, J. L. Hall, A. Gallagher & J. Ye. Single-stage sub-Doppler cooling of alkaline earth atoms. *Phys. Rev. Lett.* **90**, 193002 (2003).
- [92] V. Gokhroo, G. Rajalakshmi, R. K. Easwaran & C. S. Unnikrishnan. Sub-Doppler deep-cooled bosonic and fermionic isotopes of potassium in a compact $2d^+-3d$ MOT set-up. *Journal of Physics B: Atomic, Molecular and Optical Physics* **44**, 115307 (2011).
- [93] M. Landini, S. Roy, L. Carcagní, D. Trypogeorgos, M. Fattori, M. Inguscio & G. Modugno. Sub-Doppler laser cooling of potassium atoms. *Phys. Rev. A* **84**, 043432 (2011).
- [94] C. J. Cooper, G. Hillenbrand, J. Rink, C. G. Townsend, K. Zetie & C. J. Foot. The temperature of atoms in a magneto-optical trap. *Europhys. Lett.* **28**, 397 (1994).
- [95] E. G. M. van Kempen, S. J. J. M. F. Kokkelmans, D. J. Heinzen & B. J. Verhaar. Interisotope determination of ultracold rubidium interactions from three high-precision experiments. *Phys. Rev. Lett.* **88**, 093201 (2002).
- [96] G. Delannoy, S. G. Murdoch, V. Boyer, V. Josse, P. Bouyer & A. Aspect. Understanding the production of dual Bose-Einstein condensation with sympathetic cooling. *Phys. Rev. A* **63**, 051602 (2001).
- [97] D. Xiong, P. Wang, H. Chen & J. Zhang. Evaporative cooling rubidium atoms with microwave radiation. *Chin. Opt. Lett.* **8**, 351 (2010).
- [98] A. Simoni, M. Zaccanti, C. D'Errico, M. Fattori, G. Roati, M. Inguscio & G. Modugno. Near-threshold model for ultracold K-Rb dimers from interisotope Feshbach spectroscopy. *Phys. Rev. A* **77**, 052705 (2008).
- [99] W. Ketterle & N. J. van Druten. Bose-Einstein condensation of a finite number of particles trapped in one or three dimensions. *Phys. Rev. A* **54**, 656 (1996).

- [100] M. Zaccanti. *Tuning of the interactions in ultracold K-Rb quantum gases*. Ph.D. thesis, LENS and Università degli Studi di Firenze (2007).
- [101] Z. Yan, M. Li, L. Chen, C. Chen & J. Chen. Density of states and thermodynamic properties of an ideal system trapped in any dimension. *J Phys. A: Math. Gen.* **32**, 4069 (1999).
- [102] V. I. Yukalov. Modified semiclassical approximation for trapped Bose gases. *Phys. Rev. A* **72**, 033608 (2005).
- [103] L. D. Landau & E. M. Lifshitz. *Statistical Physics, Part 1* (Elsevier, 1980), third edn.
- [104] J. Kasprzak *et al.* Bose-Einstein condensation of exciton polaritons. *Nature* **443**, 409 (2006).
- [105] J. Klaers, J. Schmitt, F. Vewinger & M. Weitz. Bose-Einstein condensation of photons in an optical microcavity. *Nature* **468**, 545 (2010).
- [106] M. Erhard, H. Schmaljohann, J. Kronjäger, K. Bongs & K. Sengstock. Bose-Einstein condensation at constant temperature. *Phys. Rev. A* **70**, 031602 (2004).
- [107] M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. M. Kurn, D. S. Durfee & W. Ketterle. Bose-einstein condensation in a tightly confining dc magnetic trap. *Phys. Rev. Lett.* **77**, 416 (1996).
- [108] J. R. Ensher, D. S. Jin, M. R. Matthews, C. E. Wieman & E. A. Cornell. Bose-Einstein condensation in a dilute gas: Measurement of energy and ground-state occupation. *Phys. Rev. Lett.* **77**, 4984 (1996).
- [109] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer & A. Aspect. Experimental study of the thermodynamics of an interacting trapped Bose-Einstein condensed gas. *Phys. Rev. A* **70**, 013607 (2004).
- [110] P. Krüger, Z. Hadzibabic & J. Dalibard. Critical point of an interacting two-dimensional atomic Bose gas. *Phys. Rev. Lett.* **99**, 040402 (2007).
- [111] J. L. Miller. Ultracold Bose gases deviate from the textbook picture. *Physics Today* **64**, 16 (2011).
- [112] F. Dalfovo, S. Giorgini, L. P. Pitaevskii & S. Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.* **71**, 463 (1999).
- [113] V. N. Popov. *Functional Integrals and Collective Excitations* (Cambridge University Press, 1987).
- [114] L. Sabchez-Palencia & M. Lewenstein. Disordered quantum gases under control. *Nature Physics* **6**, 87 (2010).
- [115] J. Klaers, J. Schmitt, T. Damm, F. Vewinger & M. Weitz. Bose-Einstein condensation of paraxial light. *Appl. Phys. B* **105**, 17 (2011).
- [116] N. Tammuz, R. P. Smith, R. L. D. Campbell, S. Beattie, S. Moulder, J. Dalibard & Z. Hadzibabic. Can a Bose gas be saturated? *Phys. Rev. Lett.* **106**, 230401 (2011).

- [117] T. D. Lee & C. N. Yang. Many-body problem in quantum mechanics and quantum statistical mechanics. *Phys. Rev.* **105**, 1119 (1957).
- [118] T. D. Lee & C. N. Yang. Low-temperature behavior of a dilute Bose system of hard spheres. I. Equilibrium properties. *Phys. Rev.* **112**, 1419 (1958).
- [119] A. E. Glassgold, A. N. Kaufman & K. M. Watson. Statistical mechanics for the nonideal Bose gas. *Phys. Rev.* **120**, 660 (1960).
- [120] K. Huang. *Studies in Statistical Mechanics, Vol. II* (North-Holland, 1964).
- [121] A. Fetter & J. D. Walecka. *Quantum Theory of Many-Particle Systems* (McGraw-Hill, 1971).
- [122] T. Toyoda. A microscopic theory of the lambda transition. *Annals of Physics* **141**, 154 (1982).
- [123] H. T. C. Stoof. Nucleation of Bose-Einstein condensation. *Phys. Rev. A* **45**, 8398 (1992).
- [124] M. Bijlsma & H. T. C. Stoof. Renormalization group theory of the three-dimensional dilute Bose gas. *Phys. Rev. A* **54**, 5085 (1996).
- [125] P. Grüter, D. Ceperley & F. Laloë. Critical temperature of Bose-Einstein condensation of hard-sphere gases. *Phys. Rev. Lett.* **79**, 3549 (1997).
- [126] M. Holzmann, P. Grüter & F. Laloë. Bose-Einstein condensation in interacting gases. *Eur. Phys. J. B* **10**, 739 (1999).
- [127] K. Huang. Transition temperature of a uniform imperfect Bose gas. *Phys. Rev. Lett.* **83**, 3770 (1999).
- [128] M. Holzmann & W. Krauth. Transition temperature of the homogeneous, weakly interacting Bose gas. *Phys. Rev. Lett.* **83**, 2687 (1999).
- [129] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë & D. Vautherin. The transition temperature of the dilute interacting Bose gas. *Phys. Rev. Lett.* **83**, 1703 (1999).
- [130] G. Baym, J.-P. Blaizot & J. Zinn-Justin. The transition temperature of the dilute interacting Bose gas for N internal states. *Europhys. Lett.* **49**, 150 (2000).
- [131] J. D. Reppy, B. C. Crooker, B. Hebral, A. D. Corwin, J. He & G. M. Zassenhaus. Density dependence of the transition temperature in a homogeneous Bose-Einstein condensate. *Phys. Rev. Lett.* **84**, 2060 (2000).
- [132] M. Wilkens, F. Illuminati & M. Krämer. Transition temperature of the weakly interacting Bose gas: perturbative solution of the crossover equations in the canonical ensemble. *J. Phys. B* **33**, L779 (2000).
- [133] P. Arnold & B. Tomášik. T_c for dilute Bose gases: Beyond leading order in $1/N$. *Phys. Rev. A* **62**, 063604 (2000).
- [134] P. Arnold & G. Moore. BEC transition temperature of a dilute homogeneous imperfect Bose gas. *Phys. Rev. Lett.* **87**, 120401 (2001).

- [135] V. A. Kashurnikov, N. V. Prokof'ev & B. V. Svistunov. Critical temperature shift in weakly interacting Bose gas. *Phys. Rev. Lett.* **87**, 120402 (2001).
- [136] H. Kleinert. Five-loop critical temperature shift in weakly interacting homogeneous Bose-Einstein condensate. *Mod. Phys. Lett. B* **17**, 1011 (2003).
- [137] M. J. Davis & S. A. Morgan. Microcanonical temperature for a classical field: Application to Bose-Einstein condensation. *Phys. Rev. A* **68**, 053615 (2003).
- [138] B. Kastening. Bose-Einstein condensation temperature of a homogenous weakly interacting Bose gas in variational perturbation theory through seven loops. *Phys. Rev. A* **69**, 043613 (2004).
- [139] S. Ledowski, N. Hasselmann & P. Kopietz. Self-energy and critical temperature of weakly interacting bosons. *Phys. Rev. A* **69**, 061601 (2004).
- [140] K. Nho & D. P. Landau. Bose-Einstein condensation temperature of a homogeneous weakly interacting Bose gas: Path integral Monte Carlo study. *Phys. Rev. A* **70**, 053614 (2004).
- [141] V. Betz & D. Ueltschi. Critical temperature of dilute Bose gases. *Phys. Rev. A* **81**, 023611 (2010).
- [142] R. Seiringer & D. Ueltschi. Rigorous upper bound on the critical temperature of dilute Bose gases. *Phys. Rev. B* **80**, 014502 (2009).
- [143] J. O. Andersen. Theory of the weakly interacting Bose gas. *Rev. Mod. Phys.* **76**, 599 (2004).
- [144] R. Onofrio, C. Raman, J. M. Vogels, J. R. Abo-Shaeer, A. P. Chikkatur & W. Ketterle. Observation of superfluid flow in a Bose-Einstein condensed gas. *Phys. Rev. Lett.* **85**, 2228 (2000).
- [145] E. Braaten & J. Pearson. Semiclassical corrections to the oscillation frequencies of a trapped Bose-Einstein condensate. *Phys. Rev. Lett.* **82**, 255 (1999).
- [146] L. Pitaevskii & S. Stringari. Elementary excitations in trapped Bose-Einstein condensed gases beyond the mean-field approximation. *Phys. Rev. Lett.* **81**, 4541 (1998).
- [147] M. Holzmann, J.-N. Fuchs, G. A. Baym, J.-P. Blaizot & F. Laloë. Bose-Einstein transition temperature in a dilute repulsive gas. *C. R. Physique* **5**, 21 (2004).
- [148] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë & D. Vautherin. Bose-Einstein transition in a dilute interacting gas. *Eur. Phys. J. B* **24**, 107 (2001).
- [149] O. Zobay. Phase transition of trapped interacting Bose gases. *Laser Phys.* **19**, 700 (2009).
- [150] S. Giorgini, L. P. Pitaevskii & S. Stringari. Condensate fraction and critical temperature of a trapped interacting Bose gas. *Phys. Rev. A* **54**, R4633 (1996).
- [151] V. L. Ginzburg. Some remarks on phase transitions of the second kind and the microscopic theory of ferroelectric materials. *Sov. Phys. Solid State* **2**, 1824 (1961).

- [152] M. Houbiers, H. T. C. Stoof & E. A. Cornell. Critical temperature of a trapped Bose gas: Mean-field theory and fluctuations. *Phys. Rev. A* **56**, 2041 (1997).
- [153] M. Holzmann, W. Krauth & M. Naraschewski. Precision Monte Carlo test of the Hartree-Fock approximation for a trapped Bose gas. *Phys. Rev. A* **59**, 2956 (1999).
- [154] M. J. Davis & P. B. Blakie. Critical temperature of a trapped Bose gas: Comparison of theory and experiment. *Phys. Rev. Lett.* **96**, 060404 (2006).
- [155] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels & P. van der Straten. Thermodynamics of Bose-Einstein-condensed clouds using phase-contrast imaging. *Phys. Rev. A* **81**, 053632 (2010).
- [156] N. Navon, S. Piatecki, K. Günter, B. Rem, T. C. Nguyen, F. Chevy, W. Krauth & C. Salomon. Dynamics and thermodynamics of the low-temperature strongly interacting Bose gas. *Phys. Rev. Lett.* **107**, 135301 (2011).
- [157] H. Wu, E. Arimondo & C. J. Foot. Dynamics of evaporative cooling for Bose-Einstein condensation. *Phys. Rev. A* **56**, 560 (1997).
- [158] R. P. Smith, R. L. D. Campbell, N. Tammuz & Z. Hadzibabic. Effects of interactions on the critical temperature of a trapped Bose gas. *Phys. Rev. Lett.* **106**, 250403 (2011).
- [159] A. Amo *et al.* Collective fluid dynamics of a polariton condensate in a semiconductor microcavity. *Nature* **457**, 291 (2009).
- [160] R. P. Smith, N. Tammuz, R. L. D. Campbell, M. Holzmann & Z. Hadzibabic. Condensed fraction of an atomic Bose gas induced by critical correlations. *Phys. Rev. Lett.* **107**, 190403 (2011).
- [161] R. P. Smith, S. Beattie, S. Moulder, R. L. D. Campbell & Z. Hadzibabic. Condensate growth dynamics in a quantum-quenched super-saturated Bose gas (2011). [arXiv:1112.4457v1](https://arxiv.org/abs/1112.4457v1)[cond-mat.quant-gas].
- [162] V. L. Berezinskii. Violation of long range order in one-dimensional and two-dimensional systems with a continuous symmetry group. I. Classical systems. *Sov. Phys. JETP* **34**, 610 (1971).
- [163] J. M. Kosterlitz & D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C* **6**, 1181 (1973).
- [164] Z. Hadzibabic & J. Dalibard. Two-dimensional Bose fluids: An atomic physics perspective. In R. Kaiser, D. Wiersma & L. Fallani (eds.) *Proceedings of the International School of Physics Enrico Fermi, Course CLXXIII: Nano optics and atomics: transport of light and matter waves, Varena 2009* (IOS Press, Amsterdam and SIF, Bologna, 2011).
- [165] M. F. Andersen, C. Ryu, P. Cladé, V. Natarajan, A. Vaziri, K. Helmerson & W. D. Phillips. Quantized rotation of atoms from photons with orbital angular momentum. *Phys. Rev. Lett.* **97**, 170406 (2006).
- [166] S. Moulder, S. Beattie, R. P. Smith, N. Tammuz & Z. Hadzibabic. Quantised superflow glitches in an annular Bose-Einstein condensate (2011). [arXiv:1112.0334v1](https://arxiv.org/abs/1112.0334v1)[cond-mat.quant-gas].

- [167] V. Bagnato, D. E. Pritchard & D. Kleppner. Bose-Einstein condensation in an external potential. *Phys. Rev. A* **35**, 4354 (1987).
- [168] A. Jaouadi, N. Gaaloul, B. Viaris de Lesegno, M. Telmini, L. Pruvost & E. Charron. Bose-Einstein condensation in dark power-law laser traps. *Phys. Rev. A* **82**, 023613 (2010).
- [169] P. J. Mohr, B. N. Taylor & D. B. Newell. The 2010 CODATA internationally recommended values of the fundamental physical constants. <http://physics.nist.gov/cuu/Constants/> (31st December 2010).
- [170] T. G. Tiecke. Properties of potassium. <http://staff.science.uva.nl/~tgiecke/PotassiumProperties.pdf> (February 2010).
- [171] S. Falke, H. Knöckel, J. Friebe, M. Riedmann, E. Tiemann & C. Lisdat. Potassium ground-state scattering parameters and Born-Oppenheimer potentials from molecular spectroscopy. *Phys. Rev. A* **78**, 012503 (2008).
- [172] D. S. Naik & C. Raman. Optically plugged quadrupole trap for Bose-Einstein condensates. *Phys. Rev. A* **71**, 033617 (2005).