

Bose-Einstein Condensation on an Atom Chip

M. P. A. Jones

Submitted for the degree of D. Phil.

University of Sussex

November, 2002

Declaration

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other university for a degree.

Acknowledgements

Firstly I would like to thank my supervisor Ed Hinds. His guidance and patience has been excellent at every step, and I have learnt an enormous amount from working with him.

A gigantic thank you goes to Chris Vale, who has worked as a postdoc on this project from the very beginning. His help and friendship made all the difference, especially when things weren't going so well. Appropriately he was at the helm when things finally started to work! It won't be the same without you Chris - good luck back down under.

Kentaro Furusawa at the ORC Southampton made a major contribution to the project, working long hours making beautiful atom chips that we subsequently destroyed. I am very grateful for his skill and patience.

Wenzel Salzmann worked on our experiment for a year, finally sorting out the wheat from the chaff in our CMOT data. His good humour, extremely steady hand and great beer were very much appreciated.

Thanks to Alan Butler and John Knight at the mechanical workshop for their great work and excellent blues tip-offs. I would also like to thank Ben Sauer, Mirek Walkiewicz, David Lau, Peter Smith and Daniel Sahagun Sanchez for their help with the experiment.

Although it sometimes seemed like it, these three (and a bit) years haven't all been about the lab. I thank my family and my friends at SCOAP and elsewhere for all of their support, without which I definitely wouldn't have got this far. Lastly, thanks to Jocelyn for being a fantastic housemate/labmate/officemate and thesis-mate (see [1]). You're a star.

Bose-Einstein Condensation on an Atom Chip

M. P. A. Jones

Submitted for the degree of D. Phil.

University of Sussex

November, 2002

Summary

This thesis describes the creation of one of the first Bose-Einstein condensates (BEC) to be made on an atom chip. The theory, construction and optimisation of the experiment are all detailed.

Like an electronic chip, an atom chip consists of an array of wires. The magnetic fields produced by the wires are used to trap and guide neutral atoms. Our atom chip was prepared in collaboration with the ORC at the University of Southampton. The atoms are trapped above a “guide wire” which sits in a channel formed on a glass substrate. Wires running underneath the surface of the chip close the ends of the guide to form a trap.

The atom chip was loaded with ^{87}Rb atoms cooled to $\sim 50 \mu\text{K}$ in a magneto-optical trap (MOT) formed 3.8 mm above the surface. After the loading process 1.5×10^7 atoms were trapped in a magnetic trap $220 \mu\text{m}$ above the surface. To reach Bose-Einstein condensation, the atoms were further cooled over 12.5 seconds using forced RF evaporative cooling. The condensate began to form at a temperature of 380 nK, and contained up to 1×10^4 atoms. The cloud was allowed to expand for 8 ms along a magnetic waveguide above the chip before it was imaged. By comparing these images to a theoretical model I show that the partially condensed clouds are well described by a classical thermal component and a Thomas-Fermi BEC component that do not interact, and obtain the static properties of the condensate in the trap. Good agreement between the theory and experiment shows that the BEC propagates in the waveguide as a coherent matter wave.

Contents

1	Introduction	1
1.1	Atom chips and Bose-Einstein condensation	1
1.2	Overview of the experiment	4
1.3	Overview of this thesis	6
1.4	The path to BEC in pictures	8
2	Background theory	10
2.1	Laser cooling and the MOT	11
2.1.1	Slowing atoms with light	11
2.1.2	Doppler cooling	12
2.1.3	Sub-Doppler cooling	13
2.1.4	The Magneto-Optical Trap	16
2.2	Magnetic guiding and trapping on an atom chip	19
2.2.1	The interaction of an atom with a static magnetic field	19
2.2.2	Magnetic potentials and the adiabatic condition	20
2.2.3	The single wire guide	22
2.2.4	The Ioffe-Pritchard microtrap	25
2.2.5	End wires	27
2.3	Evaporative Cooling	29
2.3.1	A simple model for evaporative cooling.	30
2.3.2	Collisional properties of rubidium	32
2.3.3	Optimising the RF trajectory	34
2.4	Thermal atom clouds	34
2.4.1	The thermodynamics of cold atoms	34
2.4.2	Thermal clouds in a harmonic trap	35
2.4.3	Expansion of thermal clouds	36
2.5	Bose-Einstein condensation in a harmonic trap	36
2.5.1	The formation of BEC	37
2.5.2	Quantum Mechanics and BEC - Mean Field Theory	38
2.5.3	The Thomas-Fermi approximation	40

2.5.4	Modelling axial expansion	41
3	Apparatus and methods	43
3.1	Overview	43
3.2	The vacuum system	43
3.2.1	The bottom flange	45
3.2.2	Assembly and bakeout	48
3.3	The Atom Chip	49
3.4	Magnetic fields and computer control	51
3.4.1	External coils	51
3.4.2	Current control	52
3.4.3	Computer control	54
3.5	The laser system	54
3.5.1	Overview	54
3.5.2	The lasers	58
3.5.3	Polarisation spectroscopy lock	62
3.5.4	Frequency offset lock	65
3.5.5	Servo loop electronics and frequency jumping	68
3.6	Imaging	71
3.6.1	The absorption imaging system	72
3.6.2	Analysing absorption images at low optical density	74
3.6.3	Absorption imaging at high optical density	76
3.6.4	Fluorescence imaging	77
4	Reaching BEC	79
4.1	The MOT	79
4.1.1	Making the MOT	79
4.1.2	Optimising the MOT	81
4.2	The Compressed MOT	82
4.2.1	The CMOT in theory	84
4.2.2	The CMOT in practice	86
4.3	Optical pumping	88
4.4	Loading the magnetic trap	90
4.4.1	Mode-matching in theory	91
4.4.2	Mode-matching in practice	91
4.5	The compressed trap	94
4.5.1	Properties of the compressed trap	95
4.5.2	The compressed cloud - initial conditions for evaporation.	96

4.6	Evaporative cooling	99
4.6.1	Looking for BEC	100
4.7	Evaluation	102
5	Properties of the BEC	103
5.1	A preliminary analysis	103
5.2	The density distribution	106
5.2.1	The cloud in the trap	106
5.2.2	Axial expansion	107
5.3	Imaging	107
5.3.1	The atom-light interaction	109
5.3.2	Recoil blurring	109
5.3.3	The optical resolution	110
5.4	Calculating the theoretical axial profiles	113
5.5	Comparison with the data	114
5.6	Summary	118
6	Conclusions	119
6.1	Summary	119
6.2	Future directions	119
	Bibliography	122
A	Polarisation spectroscopy lineshape	129

List of Figures

1.1	Our atom chip	6
1.2	Step 1 - The mirror MOT	8
1.3	Step 2 - The compressed MOT (CMOT)	8
1.4	Step 3 - The mode-matched trap	8
1.5	Step 4 Compressing the trap	9
1.6	Step 5 - Evaporative cooling	9
2.1	$Lin \perp lin$ polarisation gradient cooling.	15
2.2	The trapping force in the MOT	17
2.3	The geometry of a 6-beam MOT and a mirror MOT.	18
2.4	The Breit-Rabi diagram for the $5s^2S_{1/2}$ state of ^{87}Rb	20
2.5	Principle of the single-wire guide.	23
2.6	The single wire guide potential.	23
2.7	The Ioffe-Pritchard trap	26
2.8	Axial confinement with end wires underneath the surface.	28
2.9	The Z-trap	29
2.10	Threshold for runaway evaporation	33
3.1	A photograph of the experiment	44
3.2	Scale drawing of the vacuum system viewed from above.	46
3.3	The bottom flange assembly	47
3.4	The atom chip in position.	49
3.5	The structure of the atom chip	50
3.6	The current driver circuit.	53
3.7	The hyperfine structure of the $5s^2S_{1/2}$ and $5p^2P_{3/2}$ states of ^{87}Rb , showing the laser frequencies required for the experiment, and the intervals (in MHz) between the hyperfine levels.	55
3.8	Schematic diagram of the optics.	57
3.9	The diode laser	60
3.10	Block diagram of the diode laser system	61
3.11	Layout of the polarisation spectrometer	62
3.12	The reference laser polarisation spectrum	64

3.13	The repump laser polarisation spectrum	65
3.14	Schematic of the frequency offset lock electronics	66
3.15	The offset lock error signal	67
3.16	Power spectrum of the beat note	67
3.17	Piezo responses	68
3.18	Block diagram of the servo loop electronics	70
3.19	Jumping the Ti:Sapphire laser frequency	71
3.20	The response of the Ti:Sapphire laser to a 50 MHz frequency jump . .	72
3.21	Layout of the imaging system	73
4.1	MOT lifetime	80
4.2	MOT atom number.	81
4.3	MOT atom number versus intensity.	81
4.4	A fluorescence image of the MOT.	82
4.5	MOT atom number and lifetime versus height.	83
4.6	Rescattering force versus aspect ratio	85
4.7	Density versus detuning for two different aspect ratios	86
4.8	Phase space density versus detuning for two different aspect ratios .	87
4.9	The timing sequence for the CMOT.	88
4.10	The optical pumping (OP) sequence	90
4.11	Axial trap potential.	92
4.12	The radial (x and y) potential.	93
4.13	Timing diagram for the compression phase	95
4.14	Measurements of the axial and radial trap frequencies	96
4.15	The decay of the compressed cloud	99
4.16	Absorption cross-section for trapped atoms at $7\mu\text{K}$	101
4.17	The signature for BEC	101
5.1	Image of a partially condensed cloud.	104
5.2	Axial profiles after 8ms of expansion.	105
5.3	Scaling factors for axial expansion	108
5.4	Effect of recoil blurring on the optical density.	111
5.5	Effect of recoil blurring on the cloud profile.	111
5.6	Comparison of diffraction limited and measured resolution.	112
5.7	Image of the wire edge.	112
5.8	Comparison of model and experimental axial profiles.	115
5.9	Comparison of model and experimental axial profiles (2).	116

5.10	Comparison of the data and the model at 580 kHz for different temperatures.	116
5.11	Comparison of experimental and theoretical radial profiles at 590 kHz.	117
6.1	Fragmentation of a thermal cloud near the surface.	121

List of Tables

2.1	Weak field seeking states in ^{87}Rb	21
2.2	Properties of a single-wire guide	24
4.1	Properties of a typical mirror MOT	83
4.2	Properties of a typical compressed MOT.	88
4.3	Parameters of the “mode-matched” magnetic trap.	93
4.4	Typical parameters of a cloud in the “mode-matched” trap.	94
4.5	Parameters of the compressed magnetic trap	96
4.6	Typical parameters for a cloud in the compressed trap.	97
5.1	Scaling factors at $t = 8$ ms	107
5.2	Properties of the BEC in the trap.	118

Chapter 1

Introduction

1.1 Atom chips and Bose-Einstein condensation

The development of laser cooling techniques during the 1980s revolutionised atomic physics. Laser cooling is now used in almost every area of atomic physics, from the study of cold collisions to spectroscopy and ultra-precise atomic clocks. The importance of laser cooling was recognised with the award of the Nobel prize in physics in 1997 to Steven Chu, William Phillips and Claude Cohen-Tannoudji for their pioneering experimental and theoretical work.

The resonant interaction of an atom with laser light is very strong, and by harnessing this interaction, atoms can be cooled from room temperature to a few millionths of a degree above absolute zero in a fraction of a second. At these temperatures, atoms can be significantly deflected, or even trapped, by comparatively weak forces, such as gravity and the magnetic and electric dipole interactions. The control of the motion of cold atoms with these forces is known as atom optics. In analogy with light optics, a whole range of elements for controlling the trajectories of cold atoms has been developed.

Magnetic atom optics exploits the fact that atoms with an unpaired electron have a magnetic dipole moment $\boldsymbol{\mu}$. This gives rise to an interaction

$$U = -\boldsymbol{\mu} \cdot \boldsymbol{B} \tag{1.1}$$

with an external magnetic field \boldsymbol{B} . Atoms in a state where $\boldsymbol{\mu} \cdot \boldsymbol{B} < 0$ have a lower energy in regions of low-magnetic field, and are known as weak-field seeking. They can be reflected from regions of high magnetic field, and trapped and guided in local minima of the magnetic field. This has been used to make many different magnetic mirrors, guides and lenses for neutral atoms [2].

Research in this area has culminated in the development of “atom chips” [2, 3, 4]. On an atom chip, cold atoms are magnetically trapped and guided using the magnetic field produced by conductors on the surface. The power of these devices lies in the flexibility of the geometry of the magnetic field. Using lithographic techniques pioneered in the electronics industry, almost any planar arrangement of conductors can be created, allowing the creation of a huge range of devices. Guides [5, 6], traps [4, 7], beamsplitters [8] and even atomic “conveyor belts” [9] have already been demonstrated on atom chips.

Another important consequence of laser cooling is that the de Broglie wavelength of the atoms becomes significant. Louis de Broglie postulated that material particles as well as light can be described as a wave. The wavelike properties (wavelength λ_{dB}) and particle-like properties (momentum p) are related by his famous formula:

$$\lambda_{dB} = \frac{h}{p} \quad (1.2)$$

where h is Planck’s constant.

Perhaps the ultimate manifestation of the wave properties of cold atoms was the experimental observation in 1995 of Bose-Einstein condensation (BEC) in a dilute gas of cold atoms [10, 11]. In 1925, Bose and Einstein used statistical mechanics to predict that as a system of indistinguishable bosons is cooled below a critical temperature, the particles will condense into the ground state of the system. Picturing atoms as waves, rather than bosonic “particles” gives a different view. Here, Bose-Einstein condensation occurs when the de Broglie wavelength of the atoms becomes comparable with the interatomic spacing. As the atoms are bosons, they can all occupy the same quantum state, which means that the “atom waves” can add in phase. In this picture, a Bose-Einstein condensate is a giant, coherent matter wave. Shortly after the first observations of a Bose-Einstein condensate, this coherence was demonstrated in a beautiful experiment [12], by the observation of fringes in the density distribution of two overlapping condensates. Like laser cooling before it, Bose-Einstein condensation in dilute atomic gases has opened up a whole field of theoretical and experimental research, and it too was recognised by a Nobel prize in physics, which was awarded in 2001 to Eric Cornell, Carl Wieman and Wolfgang Ketterle.

Just as a coherent source of photons - the laser - revolutionised the field of optics, the realisation of coherent matter waves is transforming the field of atom optics. The first experiments in coherent atom optics were performed using the powerful technique of Bragg diffraction [13]. Here, the atom optical element is a phase grating formed by two overlapping, phase coherent far-detuned laser beams. Momentum is

coherently transferred to atoms in the condensate by Raman transitions. Using this technique, atoms can be coherently output coupled from the condensate to form an “atom laser” [14]. Bragg diffraction has also been used to demonstrate a number of atomic analogues of coherent optical phenomena, including coherent matter wave amplification [15, 16], four-wave mixing of matter waves [17] and coherent matter wave interferometers [18, 19]. Bose-Einstein condensates have also been combined with some of the optical dipole and magnetic atom optical elements described above. The coherence of the propagation of a BEC in a waveguide has been studied using a far detuned dipole beam guide [20], and experiments have shown that the coherence of a Bose-Einstein condensate can be preserved after reflection from optical dipole [21] and magnetic mirrors [22, 23].

An obvious next step is the combination of the power and flexibility of atom chips with Bose-Einstein condensates. Current research in this area is aimed at creating an atom interferometer on a chip. Our group at Sussex has proposed an atom chip interferometer based on the magnetic traps that can be formed above two parallel wires in a uniform externally applied magnetic field [24]. This interferometer operates in the time domain. The atomic wavefunction can be coherently split and recombined while the atoms remain trapped, simply by adjusting the strength of the externally applied field. Such an interferometer can be extremely sensitive, as very long interaction times can be used. A similar trapped atom interferometer was also proposed in [25]. Another design uses multi-mode interferometry, with atoms being split and recombined as they propagate above the chip [26].

The first steps towards coherent atom optics on an atom chip were realised in the summer of 2001, when two groups in Germany at Tübingen [27] and Munich [28] independently created Bose-Einstein condensates on an atom chip. Since then three more groups have successfully made BECs on an atom chip. The third group is at Heidelberg in Germany, and the fifth at JILA in Boulder, USA. Details of the latter two experiments can be found in ref. [29]. A different approach was used by [30], where they used a movable optical dipole trap to transport a Bose-Einstein condensate from a conventional magnetic trap onto an atom chip.

The creation of the fourth BEC on an atom chip is the subject of this thesis. We first observed Bose-Einstein condensation on the 19th of July, 2002. In this thesis, I will discuss the design and construction of the experiment, and the production and properties of our Bose-Einstein Condensate.

1.2 Overview of the experiment

With the very recent observation of BEC in caesium [29], Bose-Einstein condensation has now been achieved in all the alkali elements (except francium¹), in hydrogen, and in metastable helium. An excellent review of recent work and the condensation of various isotopes can be found in [32]. Two isotopes have become the workhorses for BEC experiments: ²³Na and ⁸⁷Rb. Their large repulsive scattering length leads to efficient evaporative cooling and the formation of large, stable, condensates. In terms of experimental simplicity ⁸⁷Rb has the edge, as solid state and diode lasers can be used, and atoms can be trapped from a room temperature vapour rather than a slowed atomic beam. All the BECs created on an atom chip have so far used ⁸⁷Rb, including this one. Although, the various isotopes and elements differ widely in the experimental complexity required, the path used to BEC in the alkali metals is broadly similar in all cases [33].

Atoms from an atomic beam or a thermal vapour are cooled and trapped by laser light in a magneto-optical trap (MOT). The light scattering process limits the phase space density in the MOT to about 10^{-5} , which is still five orders of magnitude away from the BEC transition. To make further progress, the atoms are usually transferred to a magnetic trap, and then compressed to increase the rate of elastic collisions between the trapped atoms². If the collision rate is high enough, the atoms can be further cooled using forced evaporative cooling. Radio-frequency transitions between magnetic sublevels are used to selectively output couple the hottest atoms from the trap. The remaining atoms collide with each other and rethermalise at a lower temperature. Formation of a Bose-Einstein condensate is usually observed by imaging the cloud with laser light. The condensate shows up as a sharp peak in the density distribution.

The experimental requirements for BEC follow from the above. A laser system is required that can produce light at the various frequencies needed for laser cooling, imaging and optical pumping. A stable magnetic trap is needed that can generate the tight confinement necessary for good evaporative cooling. Finally, an extremely good vacuum ($\leq 10^{-11}$ Torr) is also required, as it is important that the loss rate due to collisions with the background gas should be small compared to the rate of “good” collisions between the trapped atoms during evaporative cooling. To this end, many BEC experiments isolate the magnetic trapping region from the source

¹whose longest-lived isotope has a half-life of just 22 mins. Nonetheless, this hasn’t stopped it being laser cooled in a MOT! [31]

²An alternative all optical route to BEC has also been demonstrated, where the laser cooled atoms are loaded into an optical dipole trap and evaporatively cooled by reducing the laser intensity in the trap [34].

region using a low-conductance tube or a gate valve.

Making a Bose-Einstein condensate on an atom chip follows the same basic path, with the magnetic trap being formed by the atom chip rather than by the conventional large, external coils. The experimental requirements are also similar. Atom chips offer some advantages over the conventional route. The requirements on the loss rate are generally less stringent in atom chip experiments. This is because the proximity of the wires to the atom cloud means that extremely tight traps can be generated at modest currents, boosting the elastic collision rate. This also means that large water-cooled coils and their associated high-current power supplies and switches are no longer required, reducing the complexity of the experiment. However, atom chip BEC experiments must face two difficulties not present in a conventional BEC experiment. The first is a problem of mode-matching. The density limitation in the MOT means that the atom cloud is quite large (several mm), and this large cloud must subsequently be loaded into magnetic traps produced by wires whose scale is an order of magnitude smaller. The second is that an extremely good vacuum is harder to achieve. Rather than a completely empty chamber, the magnetic trapping region must contain the atom chip, its support structure, electrical feedthroughs etc. Great care must be taken during design and assembly to ensure that a good vacuum can be achieved, and that there is no significant outgassing when current is flowing.

Different approaches to these problems have been taken by the different atom chip BEC experiments. To illustrate this it is useful to briefly consider the approaches used by the first two atom chip BEC experiments. Both experiments used an atom chip fabricated with lithographically patterned wires. In the Tübingen experiment [27], almost all the atoms from the MOT are collected in a macroscopic magnetic trap formed by a pair of coils inside the vacuum chamber. The atoms are pre-cooled in this magnetic trap before being transferred to the atom chip, where the final cooling is performed. This intermediate step allows excellent mode-matching, and makes the overall process very efficient, and their condensates contain up to 4×10^5 atoms. In the Munich experiment [28], the atoms are transferred directly from the magneto-optical trap into the magnetic trap formed by the atom chip. Their magnetic traps are formed by $50 \mu\text{m}$ conductors, making mode-matching quite difficult. Although they load only 3×10^6 atoms into the magnetic trap, their extremely high trap frequencies (up to 6.2 kHz) allow them to evaporatively cool efficiently to BEC in just 900 ms. This method, while simpler, results in smaller condensates: their condensates typically contain 3×10^3 atoms.

The approach taken in our experiment is closer to the latter. We use a single




Figure 1.1: Our atom chip

vacuum chamber that contains the rubidium source and the atom chip. We achieve a base pressure of $\leq 10^{-11}$ Torr, and trap lifetimes of up to 30 s. The rubidium vapour is produced on demand by a resistively heated dispenser. Our atom chip is shown in figure 1.1. The conductors are macroscopic ($\geq 500 \mu\text{m}$) copper wires allowing us to use quite large currents (up to 15 A). The top surface of the chip is gold coated, which allows us to use the mirror MOT technique that is also used in [28] to collect cold atoms directly above the chip. Atoms are loaded directly on to the atom chip by compressing the MOT using the fields produced by the atom chip. A diode pumped Ti:Sapphire laser provides plenty of trapping light, allowing us to start with a large number of atoms (1×10^8) in the MOT. We load 1.5×10^7 into the magnetic trap, and our condensates contain 1×10^4 atoms. The highest current that we use during the experiment is 15 A, which means that simple, inexpensive laboratory power supplies can be used, and water cooling is not necessary. The relative simplicity of this single-chamber design should be contrasted with [30] where the atom chip is loaded with a BEC produced elsewhere. While the condensates are much larger (1×10^6 atoms), this experiment requires a Zeeman slower, a complex vacuum system, a high-current water-cooled magnetic trap and an extremely precise translation stage for the transfer process.

Our route to BEC is outlined in section 1.4

1.3 Overview of this thesis

The remainder of this thesis is structured as follows. Chapter 2 discusses the theory of making a BEC and the properties of a BEC in a harmonic trap. The order

follows that of the experimental sequence. Laser cooling and MOTs are discussed first, followed by magnetic trapping and atom chips. A theoretical treatment of evaporative cooling is presented, with the aim of highlighting the requirements for achieving BEC. The discussion of BEC theory finishes the chapter.

Chapter 3 presents the construction of the experiment, including the atom chip. Experimental techniques such as laser stabilisation and absorption imaging are also discussed.

Chapter 4 describes in detail the optimisation and operation of each step on the path to BEC outlined in the next section. The properties of the atom cloud after each step are also presented. The chapter concludes with an evaluation of the efficiency of the BEC creation process.

In Chapter 5, images of the atom cloud during condensate formation are analysed in detail. The images are compared to a model for the density distribution of a partially condensed cloud, and the static properties of the condensate in the trap are extracted.

The thesis ends with a summary in Chapter 6, along with a discussion of some future directions for the experiment.

1.4 The path to BEC in pictures




Figure 1.2: Step 1 - The mirror MOT

The process begins by collecting atoms in the mirror MOT. It forms ~ 4 mm above the chip, from a combination of laser light and a quadrupole magnetic field. The MOT begins to fill when the rubidium dispenser is heated. At the end of this stage the MOT contains 1×10^8 atoms at a temperature of $50 \mu\text{K}$.




Figure 1.3: Step 2 - The compressed MOT (CMOT)

To load the magnetic trap, the MOT is compressed by turning on the guide wire on the chip and the external bias field. This brings the MOT closer to the surface and compresses it. At this point, the cloud contains 1×10^8 atoms at a temperature of $85 \mu\text{K}$.




Figure 1.4: Step 3 - The mode-matched trap

At the end of the compressed MOT phase we optically pump the atoms into a weak-field seeking state. The “mode-matched” magnetic trap is turned on to collect as many of the atoms as possible. All the laser light is shuttered off. Here the cloud contains 1.5×10^7 atoms at $83 \mu\text{K}$.




Figure 1.5: Step 4 Compressing the trap

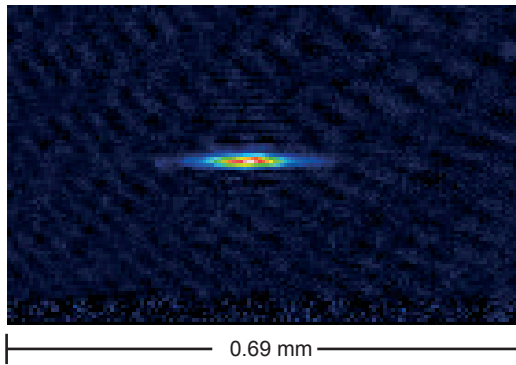


Figure 1.6: Step 5 - Evaporative cooling

Next we compress the magnetic trap to increase the elastic collision rate for evaporative cooling. The compression heats up the cloud, which extends beyond the field of view of our imaging system. This cloud also contains 1.5×10^7 atoms, but the temperature has increased to $690 \mu\text{K}$.

The final step is evaporative cooling to BEC. This reduces the temperature of the cloud by over three orders of magnitude in 12.5 seconds - note the change in scale. This cloud contains 2.6×10^4 atoms at just 270 nK , 48% of which are in the condensate.

Chapter 2

Background theory

This chapter covers the background theory underlying the creation and properties of a Bose-Einstein condensate on an atom chip. The treatment is illustrative rather than comprehensive; the theory of Bose-Einstein condensates and the methods used in their creation has been extensively reviewed elsewhere and references are provided in the appropriate section. The aim throughout is to cover the important results and show how they apply to the rest of the work presented in this thesis.

The ordering of the chapter follows the experimental BEC creation sequence. The first section looks at the theory of laser cooling and the MOT. Both Doppler and sub-Doppler cooling mechanisms are described, and a simple discussion of the origin of the confining force in the MOT is presented. Discussion of the limits on density and phase space density in the MOT is postponed to Chapter 4, where they are analysed in the context of our experimental data.

Once the atoms have been laser cooled, they can be magnetically trapped, which is the subject of the next section. This section also outlines the theory of atom chips.

The next stage in the experiment is to implement evaporative cooling to BEC. The discussion of evaporative cooling presented here concentrates on deriving the initial conditions required in the magnetic trap for successful evaporative cooling.

Before discussing the theory of Bose-Einstein condensates, it is useful to discuss the properties of thermal atom clouds. In section 2.4 I justify the use of Maxwell-Boltzmann statistics to describe laser-cooled and magnetically trapped atoms, and derive the corresponding density distribution in the magnetic trap.

The final section deals with the formation and properties of Bose-Einstein condensates. The formation of BEC in a harmonic trap is discussed using a statistical mechanical treatment. To examine the ground state properties of the condensate we

must turn to quantum mechanics. The mean field theory of BEC and the Thomas-Fermi approximation are briefly outlined. The remainder of this section deals with modelling what happens when the BEC is released from the trap and allowed to expand, in our case in one dimension. The results of this section are used to model the experimental data in Chapter 5.

2.1 Laser cooling and the MOT

The laser cooling stages of the path to BEC allow us to capture atoms from the vapour of rubidium from our dispenser (~ 500 K) and cool them to a temperature of ~ 50 μ K. The addition of a confining force in the MOT allows the atoms to be trapped at densities of over 10^{17} m^{-3} . This represents an astonishing gain in phase space density. The theory of laser cooling is extensively reviewed in [35] and [36].

2.1.1 Slowing atoms with light

The principle underlying the deflection and slowing of atoms with laser light is the conservation of momentum in the scattering process. Although the idea of radiation pressure emerges from Maxwell's equations, the most intuitive description for momentum transfer in light scattering comes from quantum mechanics. In 1917, Einstein proposed that light is quantised, with each quantum, or photon carrying energy $E = hf$ and momentum $\mathbf{p} = \hbar\mathbf{k}$, where f is the frequency and \mathbf{k} the wavevector of the light. When an atom absorbs a photon (from a laser beam, for example), the energy goes into the internal state of the atom as the atom is promoted to an excited state. The momentum is transferred to the centre of mass motion, as the conservation of momentum requires the atom to recoil in the opposite direction to the laser beam. Some time later, the excited atom will spontaneously decay to the ground state, emitting a photon and recoiling again. Spontaneous emission is a random process, with the probability of emission in any particular direction given by the symmetric dipole radiation pattern. Averaged over many absorption and emission cycles, the spontaneous emission results therefore in zero net change in the momentum of the atom. The result of shining a laser beam on the atom is therefore to change the momentum of the atom in the direction of the laser beam. The momentum carried by each photon is extremely small, but because on resonance an atom can scatter many photons the forces can be huge. With a laser beam it is easy to drive a strong atomic transition into saturation, so that the atom spends half its time in the excited state. In this case, each absorption-emission cycle takes twice the excited-state lifetime τ and transfers $\hbar\mathbf{k}$ of momentum. The maximum

force on the atom is then given by

$$F = \frac{\Delta p}{\Delta t} = \frac{\hbar k}{2\tau}. \quad (2.1)$$

For the Rb D_2 line $\tau = 27$ ns, and this force corresponds to a maximum acceleration of $1.1 \times 10^5 \text{ ms}^{-2}$, or $> 10^4 g$!

2.1.2 Doppler cooling

To move from slowing to cooling, it is necessary to find a way of narrowing the velocity distribution of the atoms. This requires a dissipative, velocity-dependent force. Doppler cooling was first suggested in 1975 by Hänsch and Schawlow [37]. Three counter-propagating pairs of beams, one along each of the Cartesian axes, are used to illuminate the atoms from all directions. The light is red-detuned slightly from resonance. As an atom moves through the light, the Doppler shift means that it sees the light it is moving towards as closer to resonance, and the light it moves away from as further from resonance. Therefore the atom scatters more photons from the light that propagates in the opposite direction to its motion, and is slowed down - regardless of which direction it is moving in. The Doppler effect for a moving atom makes the scattering force velocity dependent, and spontaneous emission provides the dissipation.

To get a more quantitative idea of how this works as a cooling mechanism, it is useful to consider a one-dimensional model with a single pair of counter-propagating laser beams and a two-level atom. If each beam has an intensity I , the total force on the atom is

$$F = \hbar k \frac{\Gamma}{2} \left[\frac{s}{1 + 2s + 4(\Delta - kv)^2/\Gamma^2} - \frac{s}{1 + 2s + 4(\Delta + kv)^2/\Gamma^2} \right] \quad (2.2)$$

where $s = I/I_{sat}$, Γ is the natural linewidth, and Δ is the detuning from resonance. For red detuning and close to $v = 0$ the force varies linearly with velocity, and we have a damping force

$$F = -\alpha v \quad (2.3)$$

with a damping coefficient given by

$$\alpha = -4\hbar k^2 s \frac{2\Delta/\Gamma}{[4\Delta^2/\Gamma^2 + 2s + 1]^2}. \quad (2.4)$$

Equation 2.3 describes motion in a viscous medium, and this cooling scheme is known as “optical molasses”.

The solution of the equation of motion 2.3 is exponential damping towards zero velocity. However, the damping force discussed above is only the average force.

The absorption and spontaneous emission process is a stochastic one, and so the force has fluctuations which produce heating. This situation, with a force term and a diffusion term can be analysed using the Fokker-Planck equation to yield the velocity distribution of the atoms. From this it can be shown that the lowest temperature is obtained when $\Delta = \frac{\Gamma}{2}$ and is given by

$$T_D = \frac{\hbar\Gamma}{2k_B} . \quad (2.5)$$

This is known as the Doppler temperature, and for rubidium $T_D = 144\mu K$.

2.1.3 Sub-Doppler cooling

Historically, the observation of sub-Doppler temperatures in optical molasses came as a complete surprise. The theoretical explanation was subsequently developed independently by two groups [38, 39], who showed that sub-Doppler cooling depends upon two features that were left out of the Doppler theory of optical molasses described above. The first is the multi-level nature of the atom. In practice, we do not have a two-level system, but the ground and first excited states are split into a number of sublevels. The second is the polarisation of the light field.

Consider a multi-level atom at rest, illuminated with light of a particular polarisation. Optical pumping will occur between the ground state sublevels. This results in an equilibrium distribution of the ground state population between the sublevels that reflects the local light field. Classically, this corresponds to orientation of the atomic dipole relative to the polarisation of the light field. When the atom moves through a spatially varying polarisation pattern, optical pumping will redistribute the population so that the dipole follows the field. However, optical pumping between the sublevels takes a finite time, and so the orientation of the moving atom always lags behind the polarisation state of the light. It is this non-adiabatic following that gives rise to the sub-Doppler cooling effects. This is known as polarisation gradient cooling.

Polarisation gradient cooling can be understood using simple one-dimensional models first presented in [40], and well reviewed in [35] and [36]. This treatment applies to a $\Delta J = +1$ transition (such as the $F = 2 \rightarrow F' = 3$ “trapping” transition in ^{87}Rb), with the light red-detuned from resonance as is the case in optical molasses. In one dimension, there are two basic types of polarisation gradient. In the *lin* \perp *lin* configuration the two counter-propagating beams have orthogonal linear polarisations. In this case the polarisation of the light field varies from linear to circular to orthogonal linear and back again over a distance of $\lambda/2$, as shown in

figure 2.1. In the $\sigma^+\sigma^-$ configuration the laser beams have opposite circular polarisations. In this case, the polarisation is linear everywhere, but the polarisation vector rotates around the beam axis with a pitch of $\lambda/2$. In the complex light field of three-dimensional molasses, both are present. To show how the non-adiabatic following gives rise to cooling we consider the two cases separately.

lin \perp lin cooling

The simplest atom susceptible to *lin \perp lin* polarisation gradient cooling has a $J = 1/2$ ground state and a $J = 3/2$ excited state. As the atom moves through the light field, the AC Stark shift of each of the two ground state levels varies, as shown figure 2.1. Consider an atom in the $m_J = +1/2$ sublevel at the origin. As it moves from left to right, kinetic energy is converted to potential energy as it climbs a potential hill. At the top of the hill, the light is σ^- and there is a strong probability that the atom will be optically pumped into the $m_J = -1/2$ state, and find itself at the bottom of a hill again. As the atom climbs these hills, kinetic energy is converted to potential energy. The potential energy is radiated away during the optical pumping process, as the spontaneous emission is at a higher frequency than the absorption, the difference being the difference in light-shift between the two states (figure 2.1). As with Doppler cooling, it is spontaneous emission that provides the dissipation. The cooling is most effective for atoms moving with a velocity such that they travel a distance $\lambda/4$ in the optical pumping time, as shown in figure 2.1.

$\sigma^+\sigma^-$ polarisation gradient cooling

In the case of $\sigma^+\sigma^-$ counter-propagating beams, the polarisation is linear everywhere, and the light-shift of each sublevel does not vary with position. The origin of the cooling in this case is more subtle. The simplest case here is a $J = 1 \rightarrow J' = 2$ transition. For such an atom at rest in a linearly polarised light field, the populations in the $m_J = \pm 1$ sublevels are equal. As the atom moves through the polarisation pattern, we can consider it as a dipole trying to follow the polarisation vector. As discussed above, due to the finite optical pumping time, the dipole does not follow the polarisation exactly, but lags slightly behind. This lag corresponds to an unequal population of the $m_J = \pm 1$ sublevels. The cooling arises because this unequal population causes the atoms to scatter more photons from one of the beams than the other. A more detailed consideration of the theory [40] shows that for red-detuned light, the atom scatters more photons from the beam it is moving toward, and hence is slowed down.




Figure 2.1: *Lin ⊥ lin* polarisation gradient cooling. An atom starting in the $m_J = +1/2$ state at the origin climbs a potential hill as it moves from left to right. As it reaches the top of the hill at $\lambda/4$, it is optically pumped into the other state, and starts climbing again. The kinetic energy of the atom is transferred to potential energy and radiated away during the optical pumping cycle.

The limits of polarisation gradient cooling

The temperature range over which these cooling mechanisms are active is roughly the same in both cases. Both schemes require the atom to be moving slowly enough that it travels roughly one wavelength in the optical pumping time. The optical pumping time τ_p can be defined in terms of the rate of scattering of incident photons R as

$$\tau_p = \frac{1}{R}. \quad (2.6)$$

For the limiting case where $R = \Gamma/2$, this gives an upper limit on the velocity of ~ 30 m/s for rubidium. This means that these cooling mechanisms are only active for atoms that have already been cooled by Doppler cooling. The fundamental limit on the lowest temperature that can be achieved is set by the single photon recoil velocity. As part of each optical pumping cycle, each atom spontaneously emits a photon, and recoils. The lower limit on the momentum of the atom is therefore set by the recoil momentum. By equating the kinetic energy corresponding to the recoil velocity to $1/2k_B T$ this limit can be define as a recoil temperature T_R :

$$T_R = \frac{(\hbar k)^2}{2mk_B}. \quad (2.7)$$

For rubidium, $T_R = 360$ nK. In practice, it can be shown that for both mechanisms optimal polarisation gradient cooling can achieve a lower limit of $\sim 10T_R$ [40].

2.1.4 The Magneto-Optical Trap

The preceding two sections have shown how atoms can be slowed and cooled until their temperature is only a few times the recoil limit. However, the forces involved in both Doppler and sub-Doppler cooling are dissipative, not confining. Eventually, the atoms will diffuse out of the laser beams, and be lost. What is needed is a method to trap the atoms while they are being cooled. This can be achieved by adding the magnetic field produced by a pair of anti-Helmholtz coils to the three intersecting pairs of laser beams used for optical molasses. If the laser beams have the correct polarisations, this adds a confining force to the dissipative cooling forces described above. The result is known as the magneto-optical trap [41].

To see how the confining force arises it is useful to consider a simple one-dimensional model. For simplicity we consider an atom that has a $J = 0$ ground state, and a $J' = 1$ excited state. The atom is considered to be at rest, so the Doppler shift can be ignored. The magnetic field is zero at the origin, and increases linearly with z . Its direction is such that it always points away from the origin (this is the one-dimensional analog of the three-dimensional quadrupole field produced




Figure 2.2: Schematic of the MOT in 1D. The laser beams are red-detuned by Δ_0 at the centre of the trap. The polarisation state is defined relative to the magnetic field direction. Due to the Zeeman shift an atom at z_0 scatters more photons from the σ^- beam, and hence is pushed towards the origin.

by anti-Helmholtz coils). The atom is illuminated from either side with circularly polarised light that is red-detuned slightly from resonance. To begin with we will consider just the $+z$ side of the origin as shown in figure 2.2. It is important to note that in this figure (2.2), the polarisation states of the two beams are defined relative to the magnetic field direction. The diagram shows that as we move away from the origin, the Zeeman shift brings the $m_J = -1$ sublevel closer to resonance, and the $m_J = +1$ level further away. This means that $\Delta m = -1$ transitions to the $m_J = -1$ sublevel are favoured. Hence an atom at z_0 will scatter more photons from the σ^- polarised beam coming from the right than it will from the σ^+ beam coming from the left. This results in a force pushing it to the left, towards the origin.

Now consider the situation on the $-z$ side of the origin. Here, the magnetic field points in the opposite direction. The light from the left is now σ^- polarised with respect to the magnetic field, and that from the right is σ^+ polarised. Thus we have the mirror image of figure 2.2, and away from the origin the atom feels a force that pushes it to the right, which is again towards the origin. Hence the result is a position-dependent restoring force that can be used to trap the atom. This treatment can be extended to any $J \rightarrow J' = J + 1$ transition, and applies along the beam axes shown in figure 2.3 with the appropriate choice of beam polarisations.

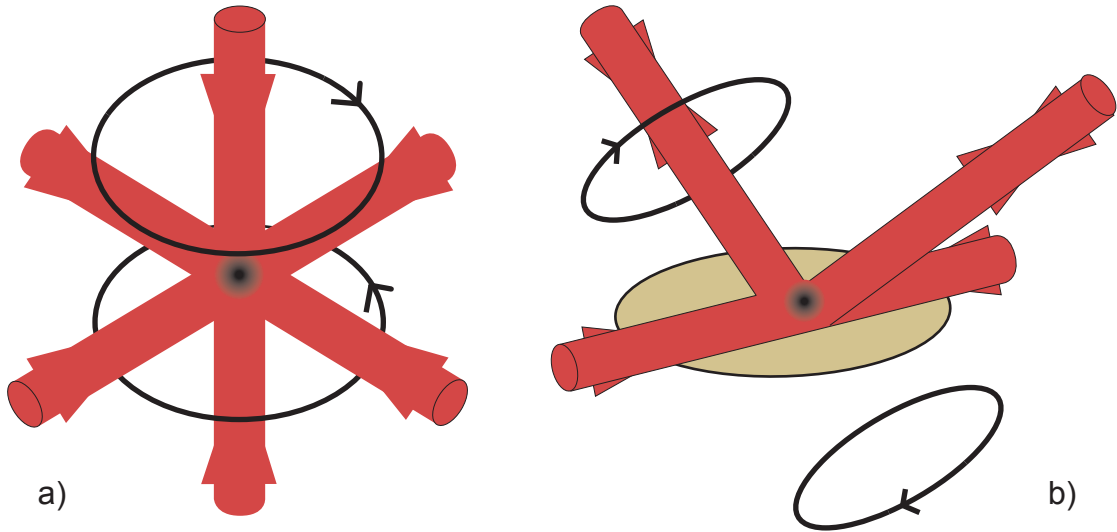


Figure 2.3: The laser beam and coil geometry for (a) the usual 6-beam MOT and (b) the mirror MOT. The beams are circularly polarised such that with respect to the local magnetic field the polarisation varies from σ^- to σ^+ on crossing the centre of the trap.

The magneto-optical trap was first realised in the three-dimensional geometry shown in figure 2.3. It is obvious that the red-detuned laser beams also from a three-dimensional Doppler molasses, and so the atoms that collect in the trap should be cooled to the Doppler temperature. It is less obvious whether the sub-Doppler cooling processes described above can still operate. The polarisation state of the light field is complex, and the Zeeman shift due to the magnetic field will perturb the pattern of light-shifts that these cooling mechanisms depend upon. In practice, for the heavier alkalis, sub-Doppler temperatures are routinely achieved in the MOT. There is always a small region in the centre of the trap where the Zeeman shift is low enough for the light shift to dominate, allowing polarisation gradient cooling to take place. This means that large clouds tend to be hotter than smaller ones. The limits of density and temperature that can be achieved in a MOT are discussed in more detail in section 4.2.1.

Trapping above a surface - the mirror MOT

In a mirror MOT, two of the 6 beams of a normal MOT are replaced by reflected light from a mirror, as shown in figure 2.3. As with a normal MOT, the polarisation of each laser beam is such that it varies from σ^- to σ^+ with respect to the local field on crossing the centre of the trap. The mirror MOT geometry allows the formation of a MOT very close to a surface [7]. This is extremely useful for loading an atom chip. Provided that the surface of the chip is smooth and reflective, atoms can be

collected directly above the chip, and transferred straight into the magnetic fields produced by the chip.

2.2 Magnetic guiding and trapping on an atom chip

2.2.1 The interaction of an atom with a static magnetic field

We start by considering an atom in a static magnetic field \mathbf{B} . We choose the basis where the total electronic angular momentum \mathbf{J} and the total nuclear angular momentum \mathbf{I} are uncoupled, $|J, m_J; I, m_I\rangle$. The hyperfine interaction Hamiltonian for an atom in the ground state is then

$$\hat{H} = A\mathbf{I} \cdot \mathbf{J} - \boldsymbol{\mu} \cdot \mathbf{B} . \quad (2.8)$$

The first term is the magnetic dipole contribution to the hyperfine interaction¹, and the second term the interaction with the applied magnetic field. The magnetic moment $\boldsymbol{\mu}$ is given by

$$\boldsymbol{\mu} = -(\mu_B g_J \mathbf{J} + \mu_N g_I \mathbf{I}) \quad (2.9)$$

where g_J and g_I are the electronic and nuclear Landé g-factors. This Hamiltonian (2.8) can be treated by perturbation theory in the case of weak and strong magnetic fields, but for intermediate fields it must be diagonalised exactly. The alkali metals turn out to be a rather fortunate case. As the ground state has $L = 0$ and $S = \frac{1}{2}$, the total electronic angular momentum $J = \frac{1}{2}$. The $J = \frac{1}{2}$ case can be diagonalised analytically for all I , giving the Breit-Rabi formula (see for example [42]). The result for the ground state of ^{87}Rb , which has $I = \frac{3}{2}$ is plotted in figure 2.4.

This shows that at low magnetic field (< 0.03 T), the total angular momentum F is a good quantum number, and we can use the basis $|F, m_F\rangle$. The Zeeman shift in this region is to a very good approximation linear with the applied magnetic field (i.e. the magnetic moment μ is independent of B), and we can write the Zeeman shift as a function of magnetic field as

$$\Delta E = \mu_B g_F m_F B \quad (2.10)$$

where g_F is an effective Landé g-factor given by

$$g_F = g_J \frac{F(F+1) + J(J+1) - I(I+1)}{2F(F+1)} - \frac{\mu_N}{\mu_B} g_I \frac{F(F+1) - J(J+1) + I(I+1)}{2F(F+1)} . \quad (2.11)$$

¹The electric quadrupole contribution to the hyperfine structure is ignored as it vanishes for our case of $J = \frac{1}{2}$




Figure 2.4: The Breit-Rabi diagram for the $5s^2S_{1/2}$ state of ^{87}Rb

In most cases, the second term can be neglected as $\frac{\mu_N}{\mu_B} \approx 10^{-3}$. For ^{87}Rb $I = \frac{3}{2}$ and so for the two hyperfine levels we obtain

$$\begin{aligned} F = 2 & \quad ; \quad g_F = +\frac{1}{2} \\ F = 1 & \quad ; \quad g_F = -\frac{1}{2} . \end{aligned}$$

It is useful to separate out two groups of states². States with $g_F m_F > 0$ have lower energy in low magnetic fields and so these states are called weak-field seeking. The opposite is true of states with $g_F m_F < 0$; these are called strong-field seeking. Atoms in weak-field seeking states can be trapped in a local minimum of the magnetic field. For ^{87}Rb these states are shown in table 2.1. The stretched state $F = 2, m_F = 2$ has the strongest interaction with the magnetic field, and will therefore be the most strongly trapped.

2.2.2 Magnetic potentials and the adiabatic condition

A weak field seeking atom trapped in a local minimum of the magnetic field oscillates in the trap. It therefore explores a region where the magnetic field varies in both magnitude and direction. In order to ensure that the atom remains trapped, it

²In the low-field approximation of a linear Zeeman shift, states with $m_F = 0$ form a third group that have no interaction with the applied magnetic field.

$ F, m_F\rangle$	$g_F m_F$
$ 2, 2\rangle$	1
$ 2, 1\rangle$	$\frac{1}{2}$
$ 1, -1\rangle$	$\frac{1}{2}$

Table 2.1: Weak field seeking states in ^{87}Rb

must stay in the same magnetic sublevel as it moves through the field. In other words, the projection of the magnetic moment on the local magnetic field must remain constant. We can get an idea of what this means by considering the classical picture of a magnetic moment precessing around the magnetic field at the Larmor frequency $\omega_L = g_F m_F \mu_B B / \hbar$. The precession angle will remain constant so long as any changes we make to the magnetic field are slow compared to the Larmor frequency.

Quantum mechanically, the atom must remain in the same eigenstate of the Hamiltonian (m_F sublevel) as the magnetic field varies. In other words, in order to avoid exciting transitions to other states, the Hamiltonian must change slowly enough that we can solve the stationary Schrödinger equation

$$\hat{H}u_n = E_n u_n \quad (2.12)$$

at each instant of time. Then we expect that an atom in the eigenstate $u_n(0)$ with energy $E_n(0)$ at time $t = 0$ will be in the eigenstate $u_n(t)$ with energy $E_n(t)$ at time t . This is an example of the adiabatic approximation [43], which requires that the rate of change of the Hamiltonian must be small compared to the transition frequency ω_{if} between the eigenstates:

$$\frac{1}{\hbar\omega_{if}} \langle i | \frac{\partial \hat{H}}{\partial t} | f \rangle \ll \omega_{if}. \quad (2.13)$$

In our case, ω_{if} is the Larmor frequency ω_L , and we obtain³

$$\frac{1}{|B|} \frac{\partial \mathbf{B}}{\partial t} \ll \omega_L \quad (2.14)$$

which means that the magnetic field must change slowly compared to the Larmor frequency, as we inferred from the classical argument. It can be seen from equation 2.14 that the adiabatic condition is most likely to be violated in regions of very low magnetic field. In particular, in regions of zero magnetic field, the adiabatic

³In this treatment, the atom is stationary in a field whose direction and magnitude are varying in time. This is equivalent to being in the frame of the atom for a moving atom.

criterion (2.14) cannot be satisfied, as the projection of the magnetic moment is no longer well defined. If the adiabatic condition is violated, the spin state of the atom can flip. These spin flip transitions are known as Majorana transitions [44]. If the atom flips into an untrapped spin state, it will be lost, and so Majorana transitions lead to a loss rate from magnetic traps and guides.

If the adiabatic criterion holds, then the projection of the magnetic moment onto the local field is always well-defined. The interaction energy $-\boldsymbol{\mu}\cdot\mathbf{B}$ can be replaced by $-\mu|B|$, and we can consider the atom to be moving in a scalar potential proportional to the magnitude of the magnetic field $U = -\mu|B|$. The Stern-Gerlach force comes naturally out of this picture as the gradient of the potential:

$$\mathbf{F} = -\nabla U = \mu\nabla|B|. \quad (2.15)$$

It should be noted that this adiabatic elimination of the spin is not rigorous, even when the rate of change of the magnetic field is small enough to satisfy 2.14. There is a geometric phase associated with the motion of an atom around a magnetic field minimum - the spin and the centre of mass motion are not separable. The quantum mechanics of an atom propagating in a magnetic quadrupole guide has been examined in detail in [45]. For the purposes of this thesis the adiabatic approximation is sufficient.

2.2.3 The single wire guide

A magnetic waveguide for weak field seeking atoms is realised by creating a minimum of magnetic field along a line or a plane. Linear magnetic waveguides have been demonstrated using arrays of permanent magnets [46], interleaved solenoids [47] and current carrying wires [48, 49, 5, 6]. Waveguides using current carrying wires are particularly attractive, as they are simple, switchable and allow great flexibility in the geometry of the guiding potential. A number of different designs have been used. The simplest geometry (based on wires alone) is the a two-wire guide [5, 6] which consists of two parallel wires carrying current in the same direction. This results in a line of zero magnetic field between the two wires, which can be used to guide atoms. Previous experiments at Sussex on magnetic waveguides [49, 50] used a four-wire guide. Here, four parallel wires are arranged on the corner of a square. Adjacent wires carry current in opposite directions. This results in a two-dimensional quadrupole field, with a zero along the axis of the guide. This kind of guide offers very high transverse confinement. These guides both suffer from the disadvantage that the atoms are guided between the wires, which makes optical access for imaging very difficult. Also, as these structures are miniaturised in the




Figure 2.5: Principle of the single-wire guide.




Figure 2.6: Sections through the potential: (a) (x, y_0) and (b) $(0, y)$. $I = 8.07$ A and $B_{bias} = 1$ mT

quest for tighter guiding potentials, it becomes increasingly difficult to load atoms into the space between the wires.

With the addition of a uniform magnetic “bias field”, it is possible to make a waveguide using only a single wire [48, 51]. The bias field exactly cancels the field of the wire along a line parallel to the wire. Since the atoms are guided above the wire, the wire can be located on a surface, and atoms can be guided parallel to the surface. This configuration is the basic building block for all atom chip experiments. It also has the advantage that the position and shape of the guide are easily adjusted using the external magnetic field and the current of the wire.

Consider the geometry shown in figure 2.5. An infinitely long wire carries current I in the $-z$ direction (into the page), and a uniform bias field B_{bias} is applied in the $-x$ direction. If we take the centre of the wire as the origin, the combined magnetic

Height	$y_0 = \frac{kI}{B_{bias}}$	1.61 mm
Depth	B_{bias}	1 mT
Gradient	$B' = \frac{2\pi B_{bias}^2}{kI}$	0.62 T/m

Table 2.2: Summary of the properties of a guide formed by a single wire and a bias field. The right-hand column shows typical values for a current of 8.07 A and a bias field of 1 mT, which correspond closely with the parameters of our “mode-matched” magnetic trap

field of the wire and the bias field is:

$$\mathbf{B}(x, y) = kI \begin{pmatrix} \frac{y}{(x^2+y^2)} \\ \frac{-x}{(x^2+y^2)} \end{pmatrix} + \begin{pmatrix} -B_{bias} \\ 0 \end{pmatrix} \quad (2.16)$$

where $k = \mu_0/2\pi$. The field lines are shown in figure 2.5. The magnetic field zero forms on the y-axis ($x = 0$) at a height $y_0 = kI/B_{bias}$ above the centre of the wire. From the field lines 2.5, we can see that the magnetic field around the zero has quadrupole symmetry.

The potential is proportional to the modulus of the magnetic field:

$$B(x, y) = \sqrt{\left(kI \frac{y}{(x^2+y^2)} - B_{bias}\right)^2 + \left(kI \frac{-x}{(x^2+y^2)}\right)^2}. \quad (2.17)$$

Sections through the potential along $x = 0$ and $y = y_0$ are shown in figure 2.6. Close to the centre of the guide, the magnitude of the magnetic field increases linearly as we move radially outwards. This is characteristic of a quadrupole potential. By performing a Taylor expansion of the potential around $(0, y_0)$ we find that the gradients of the potential in the x and y directions are equal, and proportional to the square of the bias field. At large distances the magnetic field becomes equal to the bias field, which means that it is the bias field that sets the depth of the potential well. The properties of the single-wire guide are summarised in table 2.2.

A waveguide of this type has a region of zero magnetic field along its centre. This means that atoms can be lost from the guide due to Majorana transitions. The zero can be eliminated by adding a magnetic field B_z along the z direction. The potential becomes

$$B(x, y, z) = \sqrt{\left(kI \frac{y}{(x^2+y^2)} - B_{bias}\right)^2 + \left(kI \frac{-x}{(x^2+y^2)}\right)^2 + B_z^2} \quad (2.18)$$

Looking again at the Taylor expansion around the centre of the guide we find that $B'_x = B'_y = 0$. The radial potential has changed from linear to harmonic close to

the centre of the guide. The strength of the confinement can now be described by the curvature of the potential:

$$B''_{x,y} = \frac{B'^2}{B_z} \quad (2.19)$$

where B' is the radial field gradient of the original quadrupole guide. The depth of the potential well becomes $B_{bias} - B_z$.

The quadrupole symmetry of the field lines makes the wire and bias field configuration useful for more than just magnetic guiding and trapping, as in two dimensions it is the correct magnetic field for making a MOT. For example, consider a wire mounted on or just below the surface of the mirror used to make the mirror MOT. With correct choice of current direction, the field lines of the guiding potential can be made to have the same direction as the field lines of the MOT coils. Combining the two magnetic fields gives a single, three dimensional quadrupole minimum, whose radial gradient is dominated by the field of the guide, and whose axial field gradient is provided by the MOT coils. Changing the bias field moves this minimum up and down. We can use the same laser beams that are used to make the mirror MOT to make a MOT in this combined field. This is how the compressed MOT stage of the experiment is operated.

2.2.4 The Ioffe-Pritchard microtrap

In order to turn the magnetic guide described above into a trap, the atoms must be confined axially. One way to achieve this is to superimpose a three-dimensional quadrupole field as described above, to give a linear confining potential in the axial direction as well. While it is possible to achieve tight confinement and a large trap depth with this configuration, it has one major disadvantage. The minimum in a trap of this type is always a zero of the magnetic field, and so there is a region at the centre of the trap where Majorana losses will occur. As the atom cloud is cooled towards the BEC temperature, a greater fraction of the cloud will be in this region, and the loss becomes increasingly significant.

Rather than adding a quadrupole field to confine the atoms axially, we can add an axial magnetic field that has a non-zero minimum. This kind of trap is called a Ioffe-Pritchard trap [52]. In its original configuration, shown in figure 2.7, the radial two-dimensional quadrupole field is generated by four current carrying rods. The axial confinement is provided by a pair of coils carrying current in the same sense and separated by a distance greater than their radius. In the same way, small coils can be used to turn miniature quadrupole waveguides into miniature Ioffe-Pritchard traps. For example, in the four-wire waveguide developed at Sussex, the four wires are exactly equivalent to the current carrying rods described above. By adding a




Figure 2.7: The Ioffe-Pritchard trap in its original configuration. The radial quadrupole field is generated by four rods each carrying a current I . Adjacent rods carry current in opposite directions, as indicated above. The axial confinement is provided by a pair of coils which carry current in the same sense.

pair of coils, the guide was “pinched off”, forming a miniature Ioffe-Pritchard trap which was used to trap cold atoms for several hundred milliseconds [49, 50]. As the single wire guide also produces a radial quadrupole field, it too can be turned into a miniature Ioffe-Pritchard trap by adding a pair of coils [51].

In our case, the axial confinement is produced by wires rather than coils. Nonetheless, close to the centre of the trap, the form of the magnetic potential is very similar. The original Ioffe-Pritchard configuration (figure 2.7) is considered here as its cylindrical symmetry allows simple approximate results for the trapping potential to be derived. Axial confinement with wires is the subject of the next section (2.2.5).

The trap potential is obtained by a Taylor expansion about the centre of the trap of the magnetic field produced by the rods and the coils. The four rods carrying current in the directions shown in figure 2.7 produce a radial quadrupole field given by

$$\mathbf{B}_{quad}(x, y, z) = B' \begin{pmatrix} y \\ x \\ 0 \end{pmatrix} \quad (2.20)$$

to leading order. This Taylor expansion also applies to the quadrupole produced in the single wire guide shown in figure 2.5. The Taylor expansion of the magnetic

field produced by the coils gives

$$\mathbf{B}_{coils}(x, y, z) = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2 + y^2) \end{pmatrix}. \quad (2.21)$$

Combining these two results 2.20 and 2.21, results in the following expression for the magnetic potential close to the origin⁴

$$B_{IP}(x, y, z) = B_0 + \frac{B''}{2}(x^2 + y^2) + \frac{B''}{2}z^2 \quad (2.22)$$

where the radial curvature B''_r is

$$B''_r = \frac{B''}{B_0} - \frac{B''}{2}. \quad (2.23)$$

The potential 2.22 is an anisotropic harmonic oscillator potential, with different curvatures in the radial and axial directions. Harmonic oscillator potentials can be characterised by their oscillator frequency ω , which is related to the curvature of the potential by

$$\omega = \sqrt{\frac{g_F m_F \mu_B B''}{m}}. \quad (2.24)$$

For good evaporative cooling, the trap needs to be as tight as possible. From equation 2.23 it follows that to get a high radial curvature the magnetic field at the bottom of the trap B_0 should be as small as possible. Therefore most of the field at the bottom of the trap is cancelled by adding a uniform magnetic field in the opposite direction. The increase in radial curvature due to this bias field cancellation is critical in reaching BEC. The aim is to make the axial bias field as small as possible, while avoiding Majorana losses. The net B_0 at the bottom of the trap is usually $< 10^{-4}$ T.

2.2.5 End wires

While coils have been successfully used to create miniature Ioffe-Pritchard traps, they are not very compatible with the idea of an atom chip, as they must protrude above the surface that the wires are patterned on. Also, it is difficult to make coils on the scale of a few tens of microns. To make a fully microscopic atom trap based on a chip, another method of confining the atoms axially is required. One way of doing this is using current carrying wires that run perpendicular to the guide wire.

⁴This Taylor expansion only applies close to the origin. Far from the origin, the radial potential is dominated by the gradient B' . As this gradient can be very high in a microtrap, the harmonic approximation to the potential often only applies in a very small region around the centre of the trap.




Figure 2.8: Axial confinement with end wires underneath the surface. The graph shows $B_z(z)$ 2 mm above the centres of the end wires for $I = 15$ A and $d = 2.65$ mm, with no cancellation of B_0 .

On our atom chip, the axial confinement is provided by two “end wires” that run underneath the surface of the chip, and carry current in the same direction (figure 2.8). Figure 2.8 also shows the z component of the magnetic field produced by these wires a distance h above their centres. Confining the atoms axially this way leads to a Ioffe-Pritchard trap that is harmonic axially and radially as before. However, as the wires are long in the x -direction, there is no curvature of the field in the x -direction, unlike for the coil pair.

Both the axial trap depth and trap frequency depend on the wire separation. The trap depth is the difference between the field above each wire and the field in the centre of the trap. Therefore the trap gets deeper as the wires are moved further apart. Moving the end wires too far apart gives a very “flat” minimum and a low trap frequency. The trap depth and frequency also depend on the height of the trap. The Taylor expansion for the magnetic field produced by the end wires at the centre of the trap is complicated, and so as the field due to a long straight wire is easily calculable analytically in all space, a computer program such as Mathematica

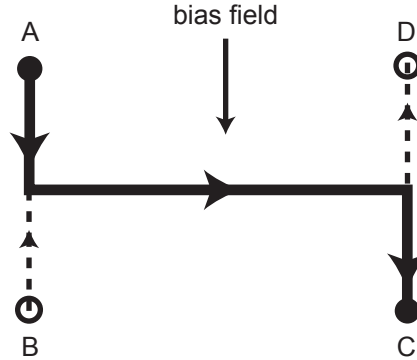


Figure 2.9: The Z-trap configuration. Axial confinement is produced by bending the ends of the wires into a “Z”, forming a Ioffe-Pritchard type trap. The bias field is in the plane of the wires. A three-dimensional quadrupole field can be produced instead by making the Z into a “U”, by using terminal B instead of A or C instead of D as shown.

is used evaluate the true potential shape and trapping frequencies. The shape of the actual magnetic trap is discussed in more detail in section 4.4.2.

Running wires underneath the guide wire in this way eliminates the problem of coils protruding above the surface, but ideally we would like all our wires to be patterned in a single plane, so that the whole chip can be fabricated using microelectronic techniques. An excellent solution to this problem is the “Z”-trap [7]. Here, the guiding wire is bent through 90 degrees at each end to form a Z, as shown in figure 2.9. With the addition of a homogenous bias field parallel to the to the ends of the Z, this forms a Ioffe-Pritchard trap. Apart from its simplicity and planar geometry, the Z configuration has several other advantages for producing a BEC. The axial depth can be made much larger than with the wires underneath. This is because at the ends of the trap B_{bias} adds to the field of the wire, increasing the depth. Having only one current flowing to produce the trap gives some common-mode type rejection of the magnetic field noise. By adding an extra connection, as shown in figure 2.9, the Z can become a “U”, which produces a three-dimensional quadrupole field suitable for a MOT. Its disadvantage is that it is less flexible, as the axial confinement cannot be varied independently of the radial confinement.

2.3 Evaporative Cooling

Magneto-optical traps work extremely well. However, as discussed in section 2.1.4, the temperature that can be achieved is limited by the recoil momentum. The density is also limited - by radiation pressure inside the cloud. Taken together, these limits combine to set a limit of $\sim 10^{-5}$ on the phase space density that can be

achieved - five orders of magnitude away from the BEC transition.

The only technique that has been developed to cross the gap is evaporative cooling. Evaporative cooling is a well-known and widely observed phenomenon ⁵. For a system of particles with a finite binding energy, such as molecules at the surface of a liquid, or in our case, atoms in a magnetic trap that has a finite depth, the hottest particles are the most likely to escape. When they leave, they take away more than their share of the energy. After the remaining energy has been redistributed, the average energy of the remaining particles decreases, and the system cools.

In these experiments, evaporative cooling is implemented in the magnetic trap, using forced RF evaporation. The hottest atoms are coupled out of the trap in an energy-selective way by using RF radiation to drive transitions between the magnetic sublevels. The transition rate is only significant when the atoms are in resonance, which occurs when the Zeeman splitting of the magnetic sublevels is equal to the RF frequency:

$$g_F \mu_B B = \hbar \omega_{RF} . \quad (2.25)$$

The atoms are driven into untrapped spin states and lost from the trap. Only atoms with sufficiently high energy can climb the trapping potential to a point where they are Zeeman-shifted into resonance. In three dimensions, this resonance condition forms a three-dimensional “evaporation surface” around the atoms.

For this to be a cooling process rather than simply a loss rate, the energy must be redistributed amongst the atoms that remain. In the magnetic trap, the energy is redistributed by elastic collisions between the atoms. To continue the cooling, the trap depth is reduced continuously by ramping down the frequency of the RF radiation. As the atoms cool down, they occupy a smaller and smaller volume at the bottom of the trap. Hence, if the cooling is working well, the density increases, even though atoms are being lost. As the atom cloud is getting both colder and denser, its phase space density increases.

2.3.1 A simple model for evaporative cooling.

The theory of evaporative cooling is an interesting application of classical kinetic theory. A detailed treatment can be found in [53]. I will use a simpler model that treats the evaporation process as a series of steps [54, 55]. Each step starts with N atoms in the trap at a temperature T . The RF cut then removes all the atoms with an energy greater than $\eta k_B T$ from the trap, where η is known as the truncation

⁵It’s the reason why Ph.D. students sweat when they cycle into the lab in the morning, and it’s what cools their cup of coffee while they check their email.

parameter. The atoms are then allowed to rethermalise, and at the end of each step, new values for N and T are calculated.

The equilibrium energy distribution of the atoms in the trap is given by the Boltzmann distribution (see section 2.4), which can be written as

$$P(\epsilon) = A\rho(\epsilon)e^{-\epsilon} \quad (2.26)$$

where $\epsilon = E/k_B T$, $\rho(\epsilon)$ is the density of states, and A is a normalisation constant. Here I will treat the case of a 3D harmonic trap where $\rho(\epsilon) \propto \epsilon^2$; the theory can also be applied to other power-law potentials.

After truncation, the most energetic atoms are lost and the energy distribution is well approximated by a truncated Boltzmann distribution [53], i.e.

$$P(\epsilon) = \begin{cases} A\rho(\epsilon)e^{-\epsilon} & \epsilon \leq \eta \\ 0 & \epsilon > \eta \end{cases} \quad (2.27)$$

After the rethermalisation time τ_{th} , the truncated distribution rethermalises. The fraction of atoms lost by the truncation at is

$$\frac{\Delta N}{N} = \frac{\int_{\eta}^{\infty} P(\epsilon)d\epsilon}{\int_0^{\infty} P(\epsilon)d\epsilon} = \frac{\Gamma(3, \eta)}{2!} \quad (2.28)$$

where $\Gamma(n, x) = \int_x^{\infty} t^{n-1}e^{-t}dt$ is the incomplete gamma function, and $\Gamma(n) \equiv \Gamma(n, 0)$. This allows us to write down a rate equation for the number of atoms:

$$\frac{\dot{N}(t)}{N} = -\frac{\Gamma(3, \eta)}{2!} \frac{1}{\tau_{th}} - \frac{1}{\tau_{loss}}. \quad (2.29)$$

The lifetime of the magnetic trap τ_{loss} includes the effect of undesirable losses from the trap, that are not due to the truncation process. It includes the effect of collisions with background gas, as well as losses due to inelastic collisions in the trap. Here, this loss rate is assumed to be constant; this assumption is justified in section 2.3.2.

Equation 2.29 shows the essential feature of evaporative cooling; it is an exponential process. By considering the energy lost from the cloud, a similar rate equation for the energy of the trapped cloud can be derived:

$$\frac{\dot{E}(t)}{E} = -\frac{\Gamma(4, \eta)}{3!} \frac{1}{\tau_{th}} - \frac{1}{\tau_{loss}}. \quad (2.30)$$

To make progress, it is necessary to look at what determines the characteristic time for rethermalisation τ_{th} . Energy is redistributed within the cloud by elastic collisions. The rethermalisation rate is therefore the rate of elastic collisions divided

by the mean number of collisions per atom needed for rethermalisation to take place, which I will take to be four, following the theoretical estimate in [56]. The elastic collision rate γ_{el} is given by

$$\gamma_{el} = n_0 \sigma_{el} \bar{v} \sqrt{2} \quad (2.31)$$

where n_0 is the peak density, σ_{el} is the cross-section for elastic collisions (see section 2.3.2) and $\bar{v}\sqrt{2}$ is the average relative velocity between two atoms. Using the results in section 2.4 for a thermal cloud at equilibrium in a harmonic potential, the elastic collision rate can be written in terms of the geometric mean trap frequency $\bar{\omega}$, the number of atoms N and the temperature T :

$$\gamma_{el}(t) = \sigma_{el} \bar{\omega}^3 \frac{\sqrt{6}m}{(2\pi)^{3/2}k_B} \frac{N(t)}{T(t)}. \quad (2.32)$$

The rate of change of the elastic collision rate $\dot{\gamma}_{el}(t)$ during evaporation is of critical importance. By combining equations 2.29 and 2.30 and writing $\tau_{th} = 4/\gamma_{el}$, a rate equation for the collision rate can be derived:

$$\frac{\dot{\gamma}_{el}(t)}{\gamma_{el}(t)} = \frac{-\gamma_{el}(t)}{4} \left(\Gamma(3, \eta) - \frac{\Gamma(4, \eta)}{3!} + \frac{4}{\gamma_{el}(t)\tau_{loss}} \right). \quad (2.33)$$

If $\dot{\gamma}_{el}(t) < 0$, the collision rate decays rapidly with time, the rethermalisation time $\tau_{th} = 4/\gamma_{el}(t) \rightarrow \infty$ and the evaporation process grinds to a halt. In contrast, exponential growth of the collision rate (i.e. $\dot{\gamma}_{el} > 0$) allows a faster than exponential growth in the phase space density. This is known as ‘‘runaway evaporation’’.

Equation 2.33 sets the requirements on the initial conditions for the evaporation. In order to have runaway evaporation and a sustainable increase in the phase space density, we must have $\dot{\gamma}_{el}(0) > 0$. The threshold for runaway evaporation as a function of η and $\gamma_{el}\tau_{loss}$ is shown in figure 2.10. This shows that to cool effectively and reach BEC we must have

$$\boxed{\gamma_{el}\tau_{loss} \gtrsim 150 \text{ at } t = 0.}$$

2.3.2 Collisional properties of rubidium

To see exactly what this condition means in terms of the experimental parameters, the collisional properties of rubidium must be considered. The collisional properties of rubidium are examined in detail in [57]. They show that in the temperature range 0 to 1 mK, the elastic cross-section for two body collisions in the $|2, 2\rangle$ state of ^{87}Rb is roughly constant at $\sim 10^{-11} \text{ cm}^2$. At low temperatures ($< 30 \mu\text{K}$) the scattering is dominated by the s-wave contribution and the elastic cross-section can be written in terms of the s-wave scattering length a :

$$\sigma_{el} = 8\pi^2 a^2. \quad (2.34)$$

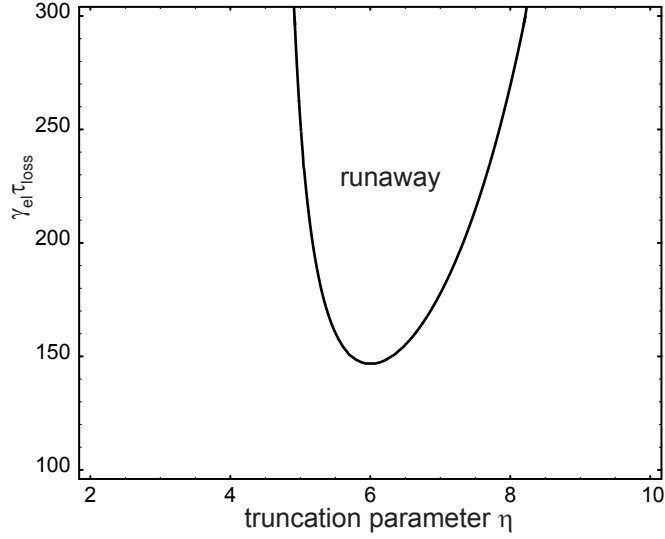


Figure 2.10: Threshold for runaway evaporation. The solid line shows the threshold $\dot{\gamma}_{el}(t) = 0$ as a function of the truncation parameter η and the product $\gamma_{el}\tau_{loss}$. Runaway can only occur if $\dot{\gamma}_{el}(t) > 0$, i.e. in the region shown

The s-wave scattering length is determined in [57] to be $(+104.5 \pm 2.5)a_0$ where a_0 is the Bohr radius. The positive value for the scattering length means that the interactions are repulsive.

The cross-section for inelastic processes is important in determining τ_{loss} . From the same paper, the inelastic rate constant for collisions between ^{87}Rb atoms in the $|2, 2\rangle$ state is $2 \times 10^{-16} \text{ cm}^3\text{s}^{-1}$. Typically, the decay rate due to collisions with background gas is $\sim 0.1 \text{ s}^{-1}$, and so inelastic collisions are not important until the density reaches $\sim 5 \times 10^{-14} \text{ cm}^{-3}$. In our trap, densities close to this size are only reached in the condensate itself. Background gas collisions dominate during the evaporation, and the assumption made above of a constant τ_{loss} during the evaporation is justified.

In contrast, in ^{85}Rb , the elastic cross-section is at least an order of magnitude lower over the above temperature range. This because of a zero in the s-wave contribution to the cross-section. The inelastic cross-section shows no such suppression, and so it is extremely difficult to produce a ^{85}Rb condensate in a standard magnetic trap by simple evaporative cooling. In order to achieve a high enough collisional cross-section to reach BEC, the scattering length must be tuned by using a Feshbach resonance in high magnetic field [58]. This is the reason why ^{87}Rb was chosen for these experiments.

2.3.3 Optimising the RF trajectory

In the model above, the evaporation has been considered as proceeding at a constant truncation parameter η . In this case the trajectory for the RF frequency is obtained from the magnetic field that corresponds to the truncation energy:

$$E = \eta k_B T = \mu_B g_F m_F (B - B_0) \quad (2.35)$$

where B_0 is the magnetic field at the bottom of the trap. Using the resonance condition equation 2.25 this yields the following equation for the RF trajectory:

$$f(t) = f_0 + \frac{\eta k_B T(t)}{h m_F} \quad (2.36)$$

with $f_0 = \mu_B g_f / h B_0$. The RF trajectory follows the evolution of the temperature of the atoms.

The goal of optimising the evaporation trajectory is to cool to BEC while losing as few atoms as possible. This leads to a natural definition of the efficiency K as the gain in phase space density per atom lost:

$$K = - \frac{d(PSD)/PSD}{dN/N} . \quad (2.37)$$

It has been shown that the condition for optimisation of the evaporation trajectory is equivalent to maximising the efficiency at every point [59]. The efficiency of our evaporation ramp is discussed in section 4.6.1.

2.4 Thermal atom clouds

This section briefly considers the thermodynamics of a gas of cold atoms far from quantum degeneracy. We begin by justifying the use of Maxwell-Boltzmann statistics to describe laser-cooled and magnetically trapped atoms. This distribution is then used to derive the density distribution of atoms in the magnetic trap. The expansion of a thermal cloud after release from the magnetic trap is also discussed.

2.4.1 The thermodynamics of cold atoms

In the above discussion, both laser and evaporative cooling were discussed in terms of a temperature. The underlying assumption in both cases is that the velocity distribution of the atoms is described by a Maxwell-Boltzmann distribution. Each component of the velocity v_j follows a Gaussian distribution:

$$P(v_j) = \frac{1}{v_{j,rms} \sqrt{2\pi}} \exp\left(-\frac{v_j^2}{2v_{j,rms}^2}\right). \quad (2.38)$$

where $v_{j,rms} \equiv \sqrt{\langle v_j^2 \rangle}$. The temperature T is defined by its relation to the mean square velocity:

$$T \equiv \frac{m \langle v_j^2 \rangle}{k_B} \quad (2.39)$$

In this thesis, a cloud whose temperature can be defined according to equation 2.39 will be referred to as a thermal atom cloud.

The Maxwell-Boltzmann distribution is usually derived by considering a classical gas of non-interacting particles in thermal equilibrium with a heat bath at temperature T . This is quite different from the case of laser cooled or magnetically trapped atoms, where the atoms are isolated from their room temperature surroundings.

In the magneto-optical trap, the atoms interact strongly with the light field, which certainly cannot be described as a heat bath with a well-defined temperature. The justification for the use of the Maxwell-Boltzmann distribution in this case follows from a consideration of the forces acting on the atoms [36]. The force can be described in terms of two components: an average force arising from the radiation pressure of the light beams, and a fluctuating force due to spontaneous emission. The statistical mechanics of this situation are described by the Fokker-Planck equation. If the average force is proportional to the velocity (i.e. a damping force), and the fluctuations are independent of velocity (which is the case for spontaneous emission), then the steady state solution of the Fokker-Planck equation is the Maxwell-Boltzmann distribution. It is therefore reasonable to define a temperature according to equation 2.39, even though the concept of thermal equilibrium does not apply.

In the magnetic trap, the atoms are completely isolated from the environment. In this case, the concept of a temperature only applies in the presence of collisions, which allow the system to reach an equilibrium state. The temperature of the ensemble is again defined using equation 2.39. As the collision rate in the magnetic trap can be very low, the system can take several seconds to reach equilibrium, and non-equilibrium states can easily be produced. This is the case during the compression of the magnetic trap, as discussed in section 4.6.

2.4.2 Thermal clouds in a harmonic trap

For a particle in a one-dimensional harmonic potential, the time averaged kinetic and potential energies are equal, i.e.

$$\frac{1}{2} m \overline{v_j(t)^2} = \frac{1}{2} m \omega^2 \overline{x_j(t)^2}. \quad (2.40)$$

This means that the Gaussian Maxwell-Boltzmann velocity distribution of the atoms in a thermal cloud gives rise to a Gaussian spatial distribution,

$$n_{th}(\mathbf{r}) = n_0 \exp\left(-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2} + \frac{z^2}{2\sigma_z^2}\right)\right) \quad (2.41)$$

where the peak density n_0 is

$$n_0 = \frac{N}{(2\pi)^{(3/2)}\sigma_x\sigma_y\sigma_z} . \quad (2.42)$$

The rms width σ_j in each direction is given by

$$\sigma_j = \sqrt{\frac{k_B T}{m\omega_j^2}} . \quad (2.43)$$

2.4.3 Expansion of thermal clouds

The rms velocity of the atoms in each direction can be expressed in terms of the cloud size and the trap frequencies by combining equations 2.39 and 2.43:

$$v_{j,rms} = \sqrt{\frac{k_B T}{m}} = \sigma_j \omega_j . \quad (2.44)$$

If the thermal cloud is suddenly released from the trap, its initial velocity distribution is “frozen in”. This means that after an expansion time t its size in each direction becomes

$$\sigma_j(t) = \sqrt{\sigma_j^2(0) + v_{j,rms}^2 t^2} = \sigma_j(0) \sqrt{1 + \omega_j^2 t^2} . \quad (2.45)$$

Thus the evolution of the width of the expanding cloud can be described by a scaling parameter $\lambda_{j,th}$

$$\lambda_{j,th} = \sqrt{1 + \omega_j^2 t^2} . \quad (2.46)$$

It is also useful to write the above results in a different way. Equation 2.46 can also be written as

$$\sigma_j^2(t) = \sigma_j^2(0) + v_{j,rms}^2 t^2 = \sigma_j^2(0) + \frac{k_B T}{m} t^2 . \quad (2.47)$$

Thus the temperature of a thermal cloud can be measured by allowing it to expand and measuring its width as a function of time. This applies to the cloud released from the MOT as well as from the magnetic trap,

2.5 Bose-Einstein condensation in a harmonic trap

The intention of this section is to establish a theoretical framework for describing the properties of our BEC, and to relate these properties to parameters that can

be measured experimentally. Rather than complete, rigorous derivations, I will concentrate on presenting the main results and an argument as to where they come from. Many reviews of BEC theory have been written; the treatment in this section mainly follows refs. [33, 60].

2.5.1 The formation of BEC

Bose-Einstein condensation is a macroscopic occupation of the ground state of the system as the temperature is lowered below a critical temperature T_c . The statistical mechanics for the case of a uniform gas of non-interacting bosons is most easily treated in the grand-canonical ensemble. This is a textbook problem [61], and yields simple results for thermodynamic quantities such as the critical temperature and the fraction of atoms in the ground state. In this section I will discuss the extension of this treatment to the case of non-interacting bosons in a harmonic external potential [60, 62]:

$$U(\mathbf{r}) = \frac{1}{2} \sum_{j=x,y,z} m\omega_j^2 r_j^2 . \quad (2.48)$$

In the grand canonical ensemble, the mean number of bosons in the state i of energy ε_i is given by the Bose-Einstein distribution:

$$\bar{N}_i = \frac{1}{e^{\beta(\varepsilon_i - \mu)} - 1} \quad (2.49)$$

where μ is the chemical potential and $\beta = 1/(k_B T)$. In the case of our harmonic oscillator potential, the energy of the i th state is given in terms of the quantum numbers of the energy levels n_x, n_y, n_z by

$$\varepsilon_i = \varepsilon_{n_x, n_y, n_z} = (n_x + \frac{1}{2})\hbar\omega_x + (n_y + \frac{1}{2})\hbar\omega_y + (n_z + \frac{1}{2})\hbar\omega_z . \quad (2.50)$$

We are interested in the occupation number of the ground state N_0 . From equation 2.49 it follows that this will become large when the chemical potential approaches the energy of the lowest state:

$$\mu \rightarrow \frac{1}{2}(\omega_x + \omega_y + \omega_z) . \quad (2.51)$$

The total number of atoms N can be written as the sum of the number of atoms in the ground state N_0 and the number of atoms in all the excited states, $\sum_{i \neq 0} N_i$. Taking the limit (2.51), this gives

$$N = N_0 + \sum_{n_x, n_y, n_z \neq 0} \frac{1}{\exp[\beta\hbar(\omega_x n_x + \omega_y n_y + \omega_z n_z)] - 1} . \quad (2.52)$$

To obtain N_0 in terms of N , we must evaluate the sum. By making a “semi-classical” approximation, this sum can be replaced by an integral. This requires two assumptions. The first is that the available energy is large compared to the spacing of the harmonic oscillator energy levels. This means that instead of summing over the discrete energy levels, we can integrate over a smooth density of states $\rho(\varepsilon)$. The second assumption is that we can take the thermodynamic limit $N \rightarrow \infty$, which sets the upper bound on the integral. Equation 2.52 becomes

$$N = N_0 + \int_0^\infty \frac{\rho(\varepsilon)d\varepsilon}{\exp(\beta\varepsilon) - 1} . \quad (2.53)$$

The density of states for the harmonic oscillator potential (2.48) is

$$\rho(\varepsilon) = \frac{1}{2} \frac{\varepsilon^2}{(\hbar\bar{\omega})^3} \quad (2.54)$$

where $\bar{\omega}$ is the geometric mean trapping frequency $\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}$. On integrating, we obtain

$$\frac{N_0}{N} = 1 - \frac{\zeta(3)}{N} \left(\frac{k_B T}{\hbar\bar{\omega}} \right)^3 \quad (2.55)$$

where $\zeta(n)$ is the Riemann Zeta-function. This can be written as

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_C} \right)^3 \quad (2.56)$$

where T_C the critical temperature for Bose-Einstein condensation is defined by

$$T_C = \frac{\hbar\bar{\omega}}{k_B} \left(\frac{N}{\zeta(3)} \right)^{1/3} = 0.94 \frac{\hbar\bar{\omega}}{k_B} N^{1/3} . \quad (2.57)$$

The results obtained above are a very good approximation, both to more detailed theoretical results and to experimental data [63]. The main weakness of this treatment is the semiclassical approximation; and the taking of the limit $N \rightarrow \infty$ in particular. For 10^4 atoms, finite size corrections to the thermodynamic relations 2.56 and 2.57 are of the order of 5%. It is also approximate in a different sense, as we have completely neglected interatomic interactions. The rigorous treatment of the statistical mechanics of a finite number of trapped, interacting bosons is complicated theoretical problem (see [60] and references therein).

2.5.2 Quantum Mechanics and BEC - Mean Field Theory

To describe the statics and dynamics of the BEC, we have to treat the quantum mechanics of a system consisting of N interacting bosons confined in an external

potential $U(\mathbf{r}, t)$. A strong analogy can be drawn here with quantum optics. Photons are also bosons, and the electromagnetic field can be quantised in terms of harmonic oscillator modes (see for example [64]). To generalise this treatment to BEC, we replace the photon creation and annihilation operators with the boson field operators $\hat{\psi}(\mathbf{r})$ and $\hat{\psi}(\mathbf{r})^\dagger$, which respectively annihilate and create a particle at position \mathbf{r} . One advantage of applying this second-quantised approach is that the creation and annihilation operators provide a straightforward way to describe the interactions. This allows us to write down the complete Hamiltonian for the system

$$\hat{H} = \int d\mathbf{r} \hat{\psi}^\dagger(\mathbf{r}) \left[\frac{\hbar^2}{2m} \nabla^2 + U(\mathbf{r}) \right] \hat{\psi}(\mathbf{r}) + \frac{1}{2} \int d\mathbf{r} d\mathbf{r}' \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}^\dagger(\mathbf{r}') V(\mathbf{r} - \mathbf{r}') \hat{\psi}(\mathbf{r}') \hat{\psi}(\mathbf{r}) \quad (2.58)$$

The first term describes the kinetic energy and the external potential energy. The second term describes interactions between the atoms. As the BEC is a dilute system, only two body interactions are included⁶ via an interaction potential $V(\mathbf{r} - \mathbf{r}')$. For cold, dilute alkali atoms, the interactions are well described by a “hard-sphere” interaction characterised by a single parameter - the s-wave scattering length a . This means that we can write an effective interaction potential

$$V(\mathbf{r} - \mathbf{r}') = g\delta(\mathbf{r} - \mathbf{r}') \quad (2.59)$$

where g is related to the scattering length a by $g = 4\pi\hbar^2 a/m$. While it is possible to calculate parameters of the condensate by attacking this Hamiltonian 2.58 directly with Monte-Carlo methods, this strategy becomes harder and harder as the number of atoms increases. A common approach to interacting many-body problems is to use a mean-field description. For a large number of particles we can approximate the (time-dependent) field operators as

$$\hat{\Psi}(\mathbf{r}, t) = \Phi(\mathbf{r}, t) + \hat{\Psi}'(\mathbf{r}, t) . \quad (2.60)$$

The function $\Phi(\mathbf{r}, t)$ represents a mean field. It is defined as the expectation value of the field operator: $\Phi(\mathbf{r}, t) = \langle \hat{\Psi}(\mathbf{r}, t) \rangle$. The density of the condensate is given by its modulus $n(\mathbf{r}, t) = |\Phi(\mathbf{r}, t)|^2$, and it has a well-defined phase which describes the first-order coherence. As a result, $\Phi(\mathbf{r}, t)$ is often referred to as the “wave function” of the condensate. The operator $\hat{\Psi}'(\mathbf{r}, t)$ describes the fluctuations around this mean field. An equation for $\Phi(\mathbf{r}, t)$ can be derived [60] by writing the time evolution of the field operator $\hat{\Psi}(\mathbf{r}, t)$ using Heisenberg’s equation of motion with the Hamiltonian

⁶Higher order processes do occur, but processes such as three body recombination cause the particles to be lost from the condensate, and so do not affect the ground state properties to first order. They do mean however that the condensate has a finite lifetime.

(2.58). The result is known as the time-dependent Gross-Pitaevskii (G-P) equation

$$i\hbar\frac{\partial\Phi(\mathbf{r},t)}{\partial t} = \left[-\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r},t) + g|\Phi(\mathbf{r},t)|^2 \right] \Phi(\mathbf{r},t) \quad (2.61)$$

This has the form of a non-linear Schrödinger equation. The time-independent version of this equation is obtained by writing the order parameter as $\Phi(\mathbf{r},t) = \phi(\mathbf{r})e^{(-i\mu t/\hbar)}$, where Φ is real. The result is the time-independent Gross-Pitaevskii equation:

$$\left(-\frac{\hbar^2\nabla^2}{2m} + U(\mathbf{r}) + g\phi^2(\mathbf{r}) \right) \phi(\mathbf{r}) = \mu\phi(\mathbf{r}) \quad (2.62)$$

2.5.3 The Thomas-Fermi approximation

To find the ground state of the condensate, we need to solve this time-independent G-P equation 2.62. As the atom number becomes large, the kinetic energy term becomes less and less important with respect to the interaction energy. For large enough N , we can neglect the kinetic energy term completely⁷. This is called the Thomas-Fermi approximation and it yields a simple formula for the density profile of the ground state in a harmonic trap:

$$n_{TF}(\mathbf{r}) = |\Phi^2| = \max \left(\frac{\mu - \frac{1}{2} \sum_{j=1,2,3} m\omega_j^2 r_j^2}{g}, 0 \right). \quad (2.63)$$

The density profile is an inverted parabola, mirroring the shape of the potential, which goes to zero in each direction at the Thomas-Fermi radius given by

$$R_{j,TF} = \frac{1}{\omega_j} \sqrt{\frac{2\mu}{m}}. \quad (2.64)$$

From the normalisation of the density profile (2.63), we can derive a useful expression for the chemical potential in terms of the number of atoms in the condensate N_0 , the scattering length a and the geometric mean trapping frequency $\bar{\omega} = [\omega_1\omega_2\omega_3]^{1/3}$:

$$\mu = \frac{1}{2}\hbar\bar{\omega} \left(15N_0a\sqrt{\frac{m\bar{\omega}}{\hbar}} \right)^{2/5}. \quad (2.65)$$

⁷How large is large enough? The relevant quantity [65] is the dimensionless parameter $\zeta = (8\pi Na/a_r)^{1/5}$ where $a_r = \sqrt{\hbar/m\omega_r}$ is the radial harmonic oscillator length. The ratio of the kinetic energy to the interaction energy is of the order $\zeta^{-4} \sim N^{-4/5}$. This means that the kinetic energy is only 0.4% of the interaction energy for just 1000 ⁸⁷Rb atoms.

2.5.4 Modelling axial expansion

The formation of a Bose-Einstein condensate is usually observed by releasing the atoms from the trap and allowing the cloud to expand. The condensate shows up as a sharp peak, as it expands much more slowly than the thermal cloud. In our case, rather than using a three-dimensional expansion, the cloud is just released axially, and allowed to expand along the waveguide. In both cases, relaxing the trap amounts to a perturbation of the condensate. One approach to calculating its subsequent evolution is to directly solve the time-dependent G-P equation (2.61). Although this equation can be solved numerically in many cases, it isn't a trivial problem. Castin and Dum [66] introduced an alternative method for calculating the evolution of a BEC in the Thomas-Fermi approximation in a time-dependent harmonic potential:

$$U(\mathbf{r}, t) = \frac{1}{2} \sum_{j=1,2,3} m\omega_j^2(t)r_j^2 \quad (2.66)$$

They show that the evolution can be described by a scaling of the widths of the condensate in each direction r_j by a time-dependent scaling parameter $\lambda_j(t)$. They derive the equations describing the evolution of the scaling parameters using a classical argument, and then show that the result is also a solution of the quantum mechanical equation 2.61. The classical argument is as follows. Consider a classical gas of density n_{cl} in which each particle experiences a force

$$\mathbf{F}(\mathbf{r}, t) = -\nabla (U(\mathbf{r}, t) + gn_{cl}(\mathbf{r}, t)) \quad (2.67)$$

At $t = 0$, the gas is in equilibrium and so $\mathbf{F} = 0$. This equilibrium condition is satisfied if the classical density is replaced by the density in the Thomas-Fermi approximation $n_{TF}(\mathbf{r}, t)$. Thus the classical and quantum steady-state solutions coincide in the Thomas-Fermi approximation. For $t > 0$, the solution we seek has the form of a dilatation of the gas. This means that each particle should move along a trajectory given by

$$r_j(t) = \lambda_j(t)r_j(0) \quad (j = 1, 2, 3) \quad (2.68)$$

We obtain the equations of motion for the particles by substituting the force (equation 2.67) into Newton's second law, giving

$$m\ddot{r}_j(t) = -\frac{\partial U(\mathbf{r}, t)}{\partial r_j} - \frac{\partial(gn_{cl}(\mathbf{r}, t))}{\partial r_j} \quad (2.69)$$

From the equation for the trajectory (2.68) it follows that the density will evolve as

$$n_{cl}(\mathbf{r}, t) = \frac{1}{\lambda_1(t)\lambda_2(t)\lambda_3(t)} n_{cl}(\mathbf{r}, 0) \quad (2.70)$$

Substituting this into (2.69) and using the equilibrium condition at $t = 0$ to write $\nabla(gn_{cl}(\mathbf{r}, 0)) = -\nabla U(\mathbf{r}, 0)$ gives

$$m\ddot{\lambda}_j(t)r_j(0) = -\frac{\partial U(\mathbf{r}, t)}{\partial r_j} + \frac{1}{\lambda_j(t)\lambda_1(t)\lambda_2(t)\lambda_3(t)} \frac{\partial U(\mathbf{r}(0), 0)}{\partial r_j(0)} \quad (2.71)$$

By combining this with the equation for the trapping potential (2.66) we obtain the result that the trajectory (2.68) is a solution of the equation of motion provided the scaling parameters satisfy

$$\ddot{\lambda}_j(t) = \frac{w_j^2(0)}{\lambda_j(t)\lambda_1(t)\lambda_2(t)\lambda_3(t)} - \omega_j^2(t)\lambda_j(t) \quad (2.72)$$

The evolution of our BEC can now be obtained by solving these coupled differential equations, which is a far easier task than solving the time-dependent G-P equation. The initial conditions are that $\lambda_j(0) = 1$ and that $\dot{\lambda}_j(0) = 0$, as the gas is initially at rest. To describe the axial expansion, we are interested in the case where the trap frequencies are suddenly altered ($\omega \rightarrow \beta\omega$) at $t = 0$. Assuming that the trapping potential is radially symmetric, then the equations for the scaling parameters in the radial and axial directions are

$$\ddot{\lambda}_r(t) = \beta_r\omega_r^2\lambda_r(t) - \frac{\omega_r}{\lambda_r(t)^3\lambda_z(t)} \quad ; \quad \ddot{\lambda}_z(t) = \beta_z\omega_z^2\lambda_z(t) - \frac{\omega_z}{\lambda_r(t)^2\lambda_z(t)^2} \quad (2.73)$$

Chapter 3

Apparatus and methods

3.1 Overview

The experimental requirements for BEC were discussed in section 1.2. To recap, we need a laser system that produces light for cooling, imaging and optical pumping, a very good vacuum and a tight, stable magnetic trap. This chapter will describe in detail the apparatus that was constructed to meet these requirements. Details of general experimental techniques such as imaging and laser locking are also presented.

The business end of the experiment is shown in figure 3.1. The vacuum chamber containing the atom chip is at the centre of the picture. The vacuum system is described in section 3.2 and the atom chip in section 3.3. In the photograph they are obscured by coils that generate various magnetic fields used in the experiment. For example, the magnetic trap is formed from the fields produced by the wires on the atom chip, the bias coils, and the z -cancelling coils. Production and control of the magnetic fields is the subject of section 3.4, which also describes the computer control of the timing.

Also shown in figure 3.1 are the paths of some of the laser beams as they enter the vacuum chamber. A new laser system was designed and constructed as part of the work for this thesis, and so it is described in some detail in section 3.5. This is followed by section 3.6 which describes the imaging system and the simple analysis of absorption images.

3.2 The vacuum system

Ultra High Vacuum (UHV) is essential for all atom optics experiments. The need for a long magnetic trap lifetime makes the requirement even more stringent for BEC




Figure 3.1: A photograph of the heart of the experiment. The experimental axes used throughout the rest of this thesis are also shown. The vacuum chamber containing the atom chip is at the centre of the picture under all the coils. The chip surface is in the horizontal $x - z$ plane, and the guide wire runs in the z direction. Orange lines indicate how some of the laser beams enter the chamber. The 45° beams (see section 2.1.4) for the mirror MOT (top left and top right) enter the vacuum chamber through a large viewport on the top. The horizontal MOT beams (one of which is shown bottom left) and the imaging beam (bottom right) enter through smaller viewports on the side of the chamber.

experiments. A new vacuum system was designed and constructed to accommodate the mirror MOT design, and to allow us to reach pressures of the order of $\lesssim 1 \times 10^{-11}$ Torr.

The vacuum system and its internal components are illustrated in figure 3.2. The chamber is an 8" "spherical octagon" supplied by Kimball Physics Inc, and is made from 304 grade stainless steel. The chamber has 8 inch (DN160 CF) Conflat sealing surfaces on the top and bottom and a $2\frac{3}{4}$ inch (DN40 CF) Conflat sealing surface on each side. All of the internal components including the electrical feedthroughs are mounted on the bottom flange. This can be removed as a single unit and worked on separately. The top port is sealed with a viewport, which allows access for the mirror MOT 45° beams as shown in figure 3.1. The remaining ports are used for pumps, gauges and viewports as shown in figure 3.2. All viewports are standard uncoated glass.

The system is pumped by three pumps. The primary pump is a 20 l/s ion pump (Varian VacIon Plus 20 model 919-0235). This is mounted on the end of a 150 mm nipple to reduce the stray field at the centre of the chamber due to the ion pump magnets. The chamber is also pumped by a non-evaporable getter pump (SAES Sorb-AC GP50). When the experiment is operating, the system is also pumped by a liquid nitrogen cold finger. This consists of a dewar connected by a pipe to a piece of OFHC copper rod which enters the chamber just behind the dispenser. The aim of the cold finger was to rapidly pump untrapped rubidium atoms after the dispenser pulse. Using the nude ion gauge (Varian UHV 24p) we have measured the pressure in the chamber as a function of time following the dispenser pulse. The cold finger does not change the decay time constant, but it does reduce the peak pressure that is reached after a pulse. We also find that it makes a big difference to the base pressure in the chamber, typically gaining us a factor of 3, and it also helps maintain this pressure during a day of operating the experiment.

The gauge has an X-Ray limit of 5×10^{-12} Torr. In practice, the gauge controller stops displaying a reading below 1×10^{-11} Torr. When the vacuum is at its best, we obtain this condition without adding the liquid nitrogen to the cold finger. Under these conditions we measure an initial magnetic trap lifetime of 30s.

3.2.1 The bottom flange

The bottom port of the vacuum chamber is sealed by a 304 stainless steel blank flange, which carries the chip, the rubidium dispenser, the electrical feedthroughs and the MOT coils. The complete assembly is shown in figure 3.3. Two mini-Conflat flanges milled directly on to the flange take 10-pin electrical feedthroughs. Three




Figure 3.2: Scale drawing of the vacuum system viewed from above.




Figure 3.3: The bottom flange assembly

blind tapped holes in the bottom flange are used to mount a stainless steel baseplate, which carries the other components.

The rubidium dispenser consists of a small stainless steel pocket containing rubidium chromate and a reducing agent. The dispenser is heated by passing current through it. Above a threshold temperature of $\sim 500^\circ\text{C}$, the rubidium chromate is reduced, producing rubidium vapour which escapes through a slit. The threshold nature of the chemical process helps ensure that the rubidium flux drops rapidly when the dispenser is switched off. The dispenser is placed close to the trapping region to give a large rubidium flux in the capture region [67]. By applying a short, high current pulse, the dispenser is rapidly heated up, raising the local rubidium partial pressure and filling the MOT. Once the MOT has filled, the dispenser is switched off, allowing the vacuum to recover. Two bolts and large stainless steel washers secure each end of the dispenser to a MACOR pillar, and provide the electrical connections. A piece of copper wire positioned in front of the dispenser shields the MOT and magnetic trap from atoms streaming directly out of the dispenser.

The MOT coils have an internal diameter of 3 cm and are wound from Kapton insulated copper wire. Each coil has 81 turns. We find that with sufficient baking, the outgassing from the Kapton wire is negligible. Running current through the wire during the bakeout, helps drive off any trapped gas. The coils are wound on stainless steel formers, which position them at the correct angle. The remaining

part of the bottom flange assembly is the atom chip, which is described in section 3.3.

3.2.2 Assembly and bakeout

Before assembly, all of the vacuum components were thoroughly cleaned as follows:

- scrub thoroughly with detergent
- ultrasonic cleaning for 10 minutes in detergent
- rinse with tap water and then distilled water
- dry
- rinse with methanol
- ultrasonic cleaning for 10 minutes in methanol
- rinse with acetone
- ultrasonic cleaning for 10 minutes in acetone
- dry

During and after cleaning, all the components are only handled with gloves to avoid fingerprints.

During bakeout, the system is placed in an oven which can maintain a constant temperature of up to 270 °C . A bellows connects the UHV valve to a pumping station consisting of a turbo pump backed by a diaphragm pump. A residual gas analyser and an ion gauge are also mounted on the system, giving excellent diagnostics on the progress of the bakeout.

The maximum temperature that we can use during bakeout is limited by the glue used in the assembly of the chip to about 125 °C . We typically bake the system at this temperature for four days. Before the UHV valve is closed to isolate the system from the turbo, the dispenser and getter pump must be activated. The getter pump is activated first by passing current through its integral heater. During this process, large amounts hydrogen are liberated from the getter. The dispenser must then be activated. Following the procedure laid out in [67], the dispenser is degassed at 5A for 5 hours. Finally, a second short getter activation is performed. The ion gauge is also degassed at this stage. Once all the activations have been performed, the ion pump is started and the chamber is isolated from the turbo pump. The heaters are




Figure 3.4: The atom chip assembled and in position. The upper and lower MOT coils are visible at the top and bottom of the picture. Also visible are the ceramic blocks that hold the chip in place and the end wires covered with ceramic beads.

switched off when the pressure has dropped below 10^{-8} Torr. When the system is cold, the bellows is disconnected, and a blank flange attached to a roughing valve is used to blank off the UHV valve (see figure 3.2). By roughing the space behind this valve to 10^{-1} Torr the leak rate through the UHV valve is dramatically reduced.

3.3 The Atom Chip

Our atom chip is shown in figure 3.4. It was designed and constructed in collaboration with the Optoelectronics Research Centre at the University of Southampton, UK. All of the “glasswork”, including the preparation of the fibre and the difficult and time-consuming polishing was undertaken at Southampton by Kentaro Furusawa.

The atom chip was originally designed for a two-wire interferometry experiment. The two guide wires were held in a specially prepared silica fibre, which was then glued to a 3 mm thick glass substrate (figure 3.5a). A mirror-quality surface was then obtained by gluing two microscope cover slips to the substrate either side of the fibre and polishing the surface flat. The polishing process left only $15\ \mu\text{m}$ of glass above the wires, to enable the atoms to be trapped very close to the wires. To make the mirror for the mirror MOT, the surface was gold coated. Axial confinement was provided by four “end-wires”, two of which pass through holes in the substrate underneath the guide, and two of which are placed at the ends of the chip (figure 3.5b).

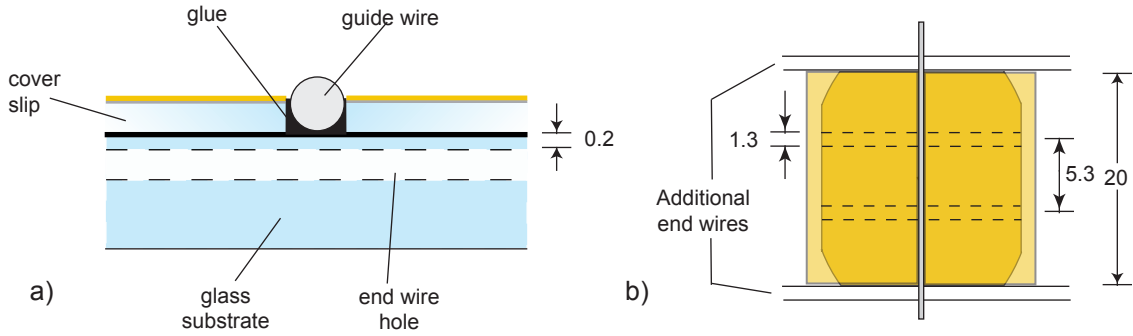


Figure 3.5: The structure of the atom chip, (a) cross-sectional view and (b) top view (not to scale). All dimensions are in millimetres.

Unfortunately, during early experiments with this atom chip, the fibre was destroyed due to the failure of the electrical connection to the guide wires. The two guide wires were $80\ \mu\text{m}$ diameter copper wires. Despite the relatively poor thermal conductivity of silica, the fibre provided good heatsinking of the wires, and currents of up to $4\ \text{A}$ per wire could be carried in vacuum with no significant heating. However, to avoid hotspots forming where the wires emerged from the fibre, it was vital to make an electrical connection to a much thicker wire right at the end of the fibre. This was done using silver DAG, a suspension of silver particles in an organic solvent, which was painted over the whole area of the connection. When the solvent is driven off by heating the connection, the deposit of silver that remains forms a low-resistance electrical connection. Its drawback is that the structure is highly porous, and it is easy to form air bubbles. One of these large bubbles under the surface of the silver left part of one of the guide wires exposed, forming a hotspot that rapidly lead to the failure of the one of the guide wires. During attempts to repair the chip, the fibre became detached from the surface and was destroyed. In the chip used for the experiments described in this thesis, the fibre was therefore replaced by a single wire.

The glue used to assemble the atom chip is a two-component epoxy adhesive (BYLAPOX 7285). Its low viscosity allows thin, even layers to be produced, and as our experiments show, it has excellent vacuum properties. The metal coating on the atom chip is made up of two layers, deposited in a single run. The first is a $100\ \text{nm}$ thick layer of chromium, deposited using RF sputtering with argon gas. This helps the gold adhere to the glass. The gold layer is $60\ \text{nm}$ thick and was deposited by thermal evaporation. Ceramic insulated copper wire was used for the guide wire, to ensure electrical isolation from the gold surface. The copper wire is $0.5\ \text{mm}$ in

diameter, and the insulation is $15\ \mu\text{m}$ thick. The end wires are 0.8 mm diameter OFHC copper wire.

The mount for the atom chip is made from stainless steel and MACOR ceramic. Ceramic end pieces hold the chip in place. The guide wire is located using V-grooves on the upper surface of these end pieces, and held in place by stainless steel clamps as shown in figure 3.4. These end pieces also hold in place the end wires at each end of the atom chip. Two more ceramic blocks either side of the atom chip help locate the end wires. The wires are connected directly to the pins of the vacuum feedthrough using gold plated UHV connectors. Ceramic beads are used to insulate the end wires. In figure 3.4, the chip is mounted in position on the bottom flange, and the upper MOT coil is visible at the top of the picture.

3.4 Magnetic fields and computer control

In this section I will discuss the how the fields that make up the magnetic trap are generated and controlled. The computer control of the timing of the experiment is also described.

3.4.1 External coils

The coils outside the vacuum chamber are shown in figure 3.1. Each coil that is labelled in the figure is one of a pair of coils that are arranged as close as possible to the Helmholtz configuration.

As well as the magnetic fields produced by the atom chip, two external magnetic fields are needed to produce the magnetic trap. The first is the bias field B_{bias} which in combination with the guide wire forms the single wire guide on which the trap is based. This is produced by a pair of rectangular Helmholtz coils mounted around the vacuum chamber, which produce a magnetic field in the $-x$ direction. The coils ($22\ \text{cm} \times 29\ \text{cm}$) are wound from 30 turns of enamelled copper wire wound on aluminium formers, which are discontinuous to eliminate eddy currents. Perspex spacers hold the coils 13 cm apart. The second is the z -cancelling field, which is applied in the opposite direction to the z component of the field produced by the end wires. This field reduces the axial bias field at the bottom of the trap B_0 , and increases the radial frequency. It is produced by two 9 cm diameter 100 turn coils wound on plastic formers that are mounted around the z viewports. The magnetic field produced by these sets of coils is calibrated as a function of current *in situ*. For the bias coils, this is done by measuring the height above the surface of the atom cloud in the waveguide. The field produced by the z -cancelling coils is measured by

looking at the variation of the field at the bottom of the trap as a function of the current in the coils, using RF evaporation.

Other magnetic fields are used during the experiment for imaging, controlling the height of the MOT etc. These are produced by three additional pairs of coils producing fields in the x , y and z directions that were originally designed for cancelling the earth's magnetic field for optical molasses. They are labelled as the x , y , z -earth coils in figure 3.1. These coils are made from 25-core computer ribbon cable. A strip of the cable can be formed into a coil by adding a D-connector at each end, offsetting the connectors by one strand. This has the advantage that the coils can easily be unplugged and removed to allow access to the vacuum system. The y -coils are wound directly around the top and bottom flanges, producing a near-perfect Helmholtz configuration. The x and z coils are rectangular and supported by a framework of aluminium posts that also support the bias coils.

3.4.2 Current control

To control the magnetic fields during the experiment, currents through the various coils and the wires of the atom chip have to be switched and/or ramped. The current control circuit that we use is shown in figure 3.6. Six of these circuits were built for the experiment. The switching element is a field effect transistor (FET). In the steady state, the circuit provides a current that is proportional to an analogue control voltage. The switching characteristics are determined by the characteristics of the load and the time constant of the integrator. The circuits are built on aluminium heat sinks to dissipate the heat produced by the FET and the sense resistor.

In its original form with purely integral gain, this circuit works well when driving resistive loads such as the guide and end wires. The current through these wires can be switched on and off in $\sim 100 \mu\text{s}$. The voltage of the power supply that drives the load must be carefully adjusted so that it is just sufficient to drive the required current. Any additional voltage is dropped across the FET, which may then have to dissipate excessive power. Operating for long in this regime is bad news for the FET, which meets a smoky end. In the case of coils and wires that are inside the vacuum, this is disastrous, as the FETs blow to a short circuit¹.

Driving inductive loads such as the bias and z -cancelling coils is more difficult. Here the load has an inherent time-constant which depends on the inductance. Trying to drive the load faster than this with the circuit as shown in figure 3.6 results in substantial overshoot. Increasing the voltage on the power supply helps

¹A relay-based protection circuit was developed that trips on sensing the over-current that occurs when the FET fails. However, this is not foolproof either, and so extreme vigilance is currently used to protect the wires in the vacuum.

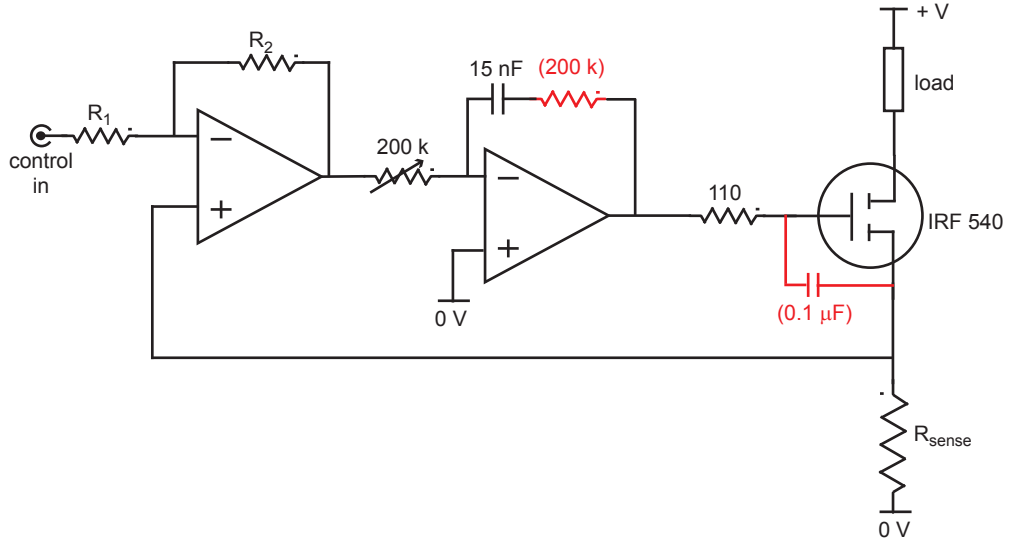


Figure 3.6: The current driver circuit. Components in red are additional components used when driving inductive loads. The $I(V_{in})$ transfer function is set by R_1 , R_2 and R_{sense} . The sense resistor is typically 0.24Ω .

by allowing a higher short-term current which speeds up the response of the load. However, this strategy is limited by the power that the FET can dissipate, as once the circuit had reached its steady state, the additional voltage is dropped across the FET. We get around this by adding some proportional gain to the circuit using a resistor as shown in red in figure 3.6. This allows much higher switching speeds at short times, while the integrator ensures the correct, stable value of the current at longer times. An inductive load also makes the circuit much more likely to oscillate, as stray capacitance can cause the circuit to behave as an LCR resonant circuit. For example, the bias coils have a resonance at 234 kHz. A capacitor between the gate and drain of the FET rolls off the gain at high frequencies to reduce this problem. With these modifications, the current in the coils can be switched in ~ 1 ms.

The analog control voltages come from two types of source. For simple switching, an analog multiplexer circuit gives out a preset voltage that depends on the state(s) of TTL logic input(s). During the experiment, the guide current, bias field and z -cancelling field must all be ramped slowly. The analog control voltages in this case are generated by programmable arbitrary waveform generators. On receiving a TTL trigger, these output a preprogrammed voltage sequence.

Low current noise is essential for reaching BEC, as noise in the magnetic field can lead to heating of the trapped cloud, especially if it is at a frequency close to the trap frequencies. When the FET circuit is operating correctly, the noise level is

limited by the noise on the control voltage. The RMS level of current noise is less than 0.04% over the bandwidth 20 Hz to 300 kHz in all cases.

3.4.3 Computer control

During an experimental run, shutters, laser detunings, AOMs and magnetic fields etc. must all be precisely switched. To ensure good repeatability, a high timing accuracy ($\sim 1 \mu\text{s}$) is required. This precludes the use of a Windows based program to control the timing, as Windows system interrupts can introduce significant timing error. Our experiment is controlled by a 286 PC running DOS. The computer controls an Amplicon PC214E digital I/O and counter-timer board, which can provide 48 TTL outputs, 24 of which are used in these experiments. Each digital output is buffered by a 50Ω line driver. The board is controlled using a Pascal program written by Ben Sauer, which allows programming of the timing sequence for the experiment. The system works well, providing a repeatability of better than 500 ns and a minimum pulse length of $\sim 25 \mu\text{s}$.

It has two main drawbacks. The first is that current control beyond simple switching must be delegated to external waveform generators as described above. These are expensive, must be programmed separately and are not necessarily well synchronised with the computer as they have their own internal clocks. The second is that DOS limits us to a fairly primitive user interface. The experiment is sufficiently complex that this is time-consuming to reprogram for different experiments, and mistakes are easily made. Improved computer control would be a useful future modification to the experiment.

3.5 The laser system

3.5.1 Overview

Laser light at several different frequencies is required during the BEC experiment. These are illustrated in figure 3.7 along with the hyperfine structure of the ground $5s^2S_{1/2}$ and excited $5p^2P_{3/2}$ states of the $D2$ transition in ^{87}Rb . The process starts with the magneto-optical trap. This requires light at two frequencies, for trapping and repumping. The trapping light must be red detuned by a few linewidths from resonance with the $F = 2 \rightarrow F' = 3$ transition. To collect a large number of atoms, we would like as much light as possible. With the high intensity and the mixed polarisation of the MOT beams there is a high probability of off-resonant excitation of the $F = 2 \rightarrow F' = 2$ transition. The $F' = 2$ state can decay by spontaneous emission into the lower ground state ($F = 1$), which is dark. A repump laser is




Figure 3.7: The hyperfine structure of the $5s^2S_{1/2}$ and $5p^2P_{3/2}$ states of ^{87}Rb , showing the laser frequencies required for the experiment, and the intervals (in MHz) between the hyperfine levels.

required to prevent all the atoms being rapidly pumped into this state and lost from the trap. This is tuned to the $F = 1 \rightarrow F' = 2$ transition as shown in figure 3.7.

During the CMOT phase, the detuning of the trapping light must be increased to allow the cloud to compress. After the CMOT, the atoms must be prepared in the $F = 2, m_F = 2$ state for magnetic trapping. This is done by optical pumping, using σ^+ light on the $F = 2 \rightarrow F' = 2$ transition. The $F = 2, m_F = 2$ final state that we desire is dark, which minimises the heating of the atom cloud during optical pumping.

Finally, to extract the properties of the atom cloud it must be imaged. This is done using two methods as described in section 3.6. Clouds in the magneto-optical trap are imaged using the fluorescence due to the MOT light itself. Clouds in free expansion and in the magnetic trap are imaged using the shadow that they cast in a probe beam. To get the highest signal, this probe beam is resonant with the $F = 2 \rightarrow F' = 3$ transition.

The light at these various frequencies is generated using three separate lasers, as shown in figure 3.8. The reference laser is an external cavity diode laser locked on the trapping transition ($F = 2 \rightarrow F' = 3$ in ^{87}Rb) using a polarisation spectroscopy locking technique. The MOT laser is a single mode Ti:Sapphire ring laser, pumped by a diode pumped solid state laser at 532 nm and locked to an internal reference cavity. We control its centre frequency by beating it against the reference laser and controlling the reference cavity to maintain the desired beat frequency. This is called frequency offset locking (section 3.5.4). The repump laser is an external cavity diode laser of the same design as the reference laser. It is locked close to the repump transition ($F = 1 \rightarrow F' = 2$ in ^{87}Rb) by polarisation spectroscopy.

The MOT light is obtained from the Ti:Sapphire laser using an AOM, which allows it to be switched on and off rapidly. The MOT laser is detuned to red of the reference laser by 106 MHz. Upshifting it by 88 MHz with the AOM, the light is finally detuned by -18 MHz as we require for the MOT. The detuning is varied simply by adjusting the beat frequency between the MOT and reference lasers. This makes the system very flexible. For example, a ^{85}Rb MOT can be made instead, simply by locking the repump and reference laser to the corresponding transitions. No re-alignment or adjustment of AOM frequency is necessary.

Repump light for the MOT simply comes straight from the repump laser as shown in figure 3.8. An optical isolator prevents unwanted feedback into the laser, and a glass plate splits off a few percent of the light for the polarisation spectroscopy lock. The MOT and repump beams are combined on a polarising beam cube. The 45 degree beams enter the chamber via the top viewport, and the horizontal beams

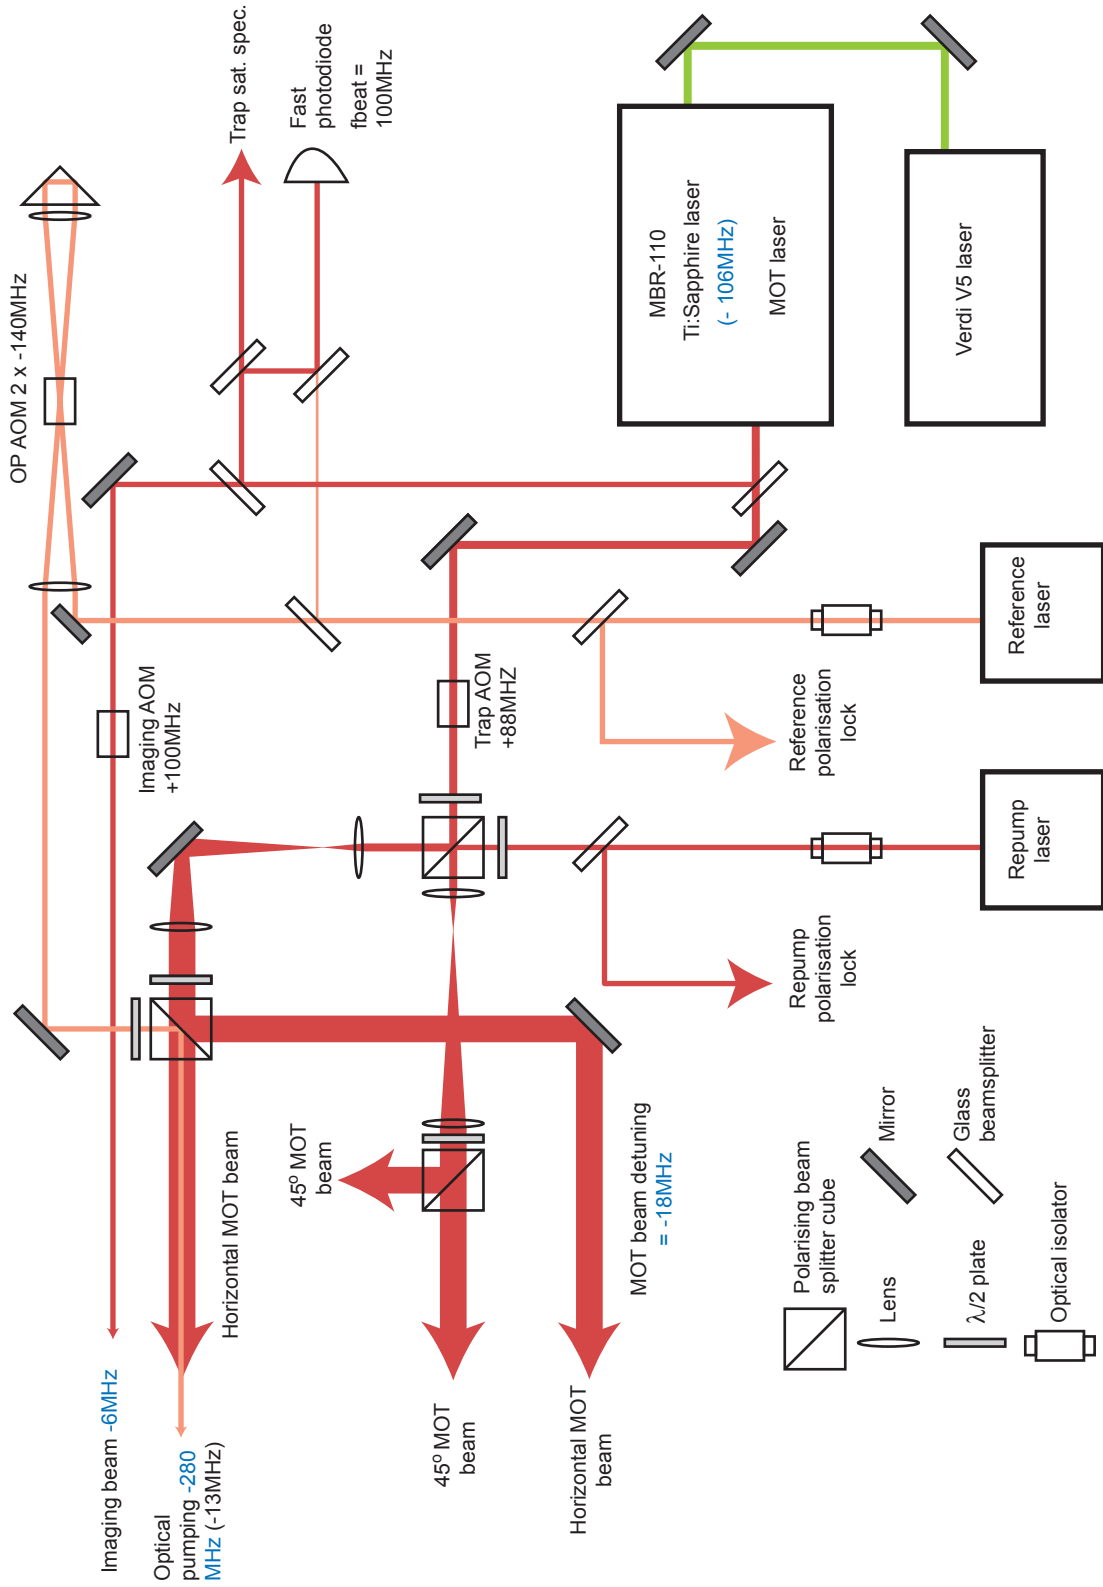


Figure 3.8: Schematic diagram of the optics. The frequencies in blue are the detunings from the $F = 2 \rightarrow F' = 3$ transition, as they are set for the MOT phase. For on-resonant absorption imaging, the detuning of the Ti:Sapph is adjusted to -100 MHz.

via the z viewports as shown in figure 3.1. Polarising beam splitter cubes and half wave plates allow the intensity in each beam to be varied independently. The light is circularly polarised by quarter-wave plates just before the beams enter the chamber

The optical pumping light is obtained from the reference laser by a double pass through a 140 MHz AOM as shown in figure 3.8. Again, an optical isolator just after the laser prevents unwanted optical feedback. The optical pumping beam enters the chamber along the same path as one of the horizontal MOT beams. This allows it to be collinear with the direction of the residual axial bias field of the trap B_0 , which makes optical pumping more straightforward. The optical pumping beam is discussed in detail in section 4.3.

The light used for imaging is derived from the MOT laser, upshifted by 100 MHz using a fixed-frequency AOM. For imaging, the offset frequency is set to -98.5 MHz, which gives a net detuning of +1.5 MHz after the AOM. This small detuning compensates for the effect of the imaging magnetic field (see section 3.6.1). The beam propagates perpendicular to the guide wire and enters the chamber via the x viewports as shown in figure 3.1

To achieve a long magnetic trap lifetime, we have to ensure that the atoms are not scattering any laser light. The attenuation provided by the AOMs used to switch the beams is at best a factor of 10^{-4} on a single pass. Therefore, mechanical shutters have to be used on all the beam paths that lead to the chamber. To eliminate less direct light, the laser system is completely covered by a black enclosure, with holes for the MOT, optical pumping, and imaging beams. Blackout curtains around the vacuum chamber further reduce the amount of stray light.

The rest of this section describes the lasers and the locking schemes in detail.

3.5.2 The lasers

The diode lasers

External cavity diode lasers are now commonly used in atomic physics experiments. Free-running laser diodes typically have a linewidth of 100 MHz, and so some form of frequency selective feedback must be used to narrow the linewidth before they are suitable for use in laser cooling experiments [68]. The most popular design uses a diffraction grating to provide the frequency selection. The laser runs in an “external cavity” formed by the back facet of the laser diode and the diffraction grating. The diffraction grating is mounted at the Littrow angle, so that light in the first diffracted order is fed back into the laser diode, and the zeroth order forms the output beam [69] (see figure 3.9). At Sussex University, this technique has evolved into a simple and reliable laser design based around a commercial mirror mount [70]. Part of the

work for this thesis consisted of developing a standardised design for the whole laser system. Several systems have now been produced to this design and are in use in laser cooling and trapping experiments in the Sussex group. This laser system is briefly discussed here.

The laser itself is constructed exactly as laid out in [70]. The only modification is the addition of a piezoelectric stack element (Thorlabs AE0203D04) underneath the precision screw that adjusts the horizontal angle of the grating, as shown in figure 3.9. This operates at a lower voltage (0 to 150V) than that typically used with the piezo discs, but provides a large extension ($4.6\ \mu\text{m}$), allowing the laser to be tuned over a very large range without opening the enclosure to adjust the precision screws. An integrated temperature sensor was also embedded into the laser mount, giving an accurate calibrated reading of the laser temperature.

All the lasers use Sanyo DL7140-201 diodes which have a nominal wavelength of 785 nm and a maximum CW output power of 80 mW. These diodes work very well. At a temperature of 16-18°C and an injection current of $\sim 60\ \text{mA}$, we get an output off the grating of 30 mW for 70 mW intracavity power at 780 nm. It is fairly easy to tune the laser so that it will scan single mode over 3 GHz at the D_2 transition wavelength (see for example figure 3.12), although the size of the obtainable tuning range does depend on the individual diode. These diodes work better on the trapping transitions than on the repump transitions. The laser will scan further between mode hops, and we have had one diode that would not lase single mode at full power on the repump transition.

A block diagram of the laser system is shown in figure 3.10. In order to provide a stable output and a narrow linewidth, the laser temperature, current and cavity length must all be precisely controlled. In addition it is useful to be able to monitor the laser temperature, current and intracavity power. The laser box is shown in figure 3.9. It holds the laser, its current driver and protection electronics, and some monitoring electronics. A small perspex cover encloses just the laser itself, reducing the effects of acoustic noise and temperature changes. A single shielded cable connects the laser box to a control/interface box. The laser current and temperature are set here, and a panel meter displays the laser current, temperature and intracavity power. This box also contains the commercial temperature control electronics. Connectors on the rear allow connection to the piezo stack, and permit modulation of the injection current. A single unit provides separate power supplies for the temperature control and for the laser diode and control electronics. The remaining component of the system is a driver for the piezo stack.

The laser system as a whole performs well, and is very easy to use. The temper-




Figure 3.9: A photograph of the diode laser box. The perspex box covering the laser itself has been removed. This photograph was taken by Dr. Peter N. Smith.

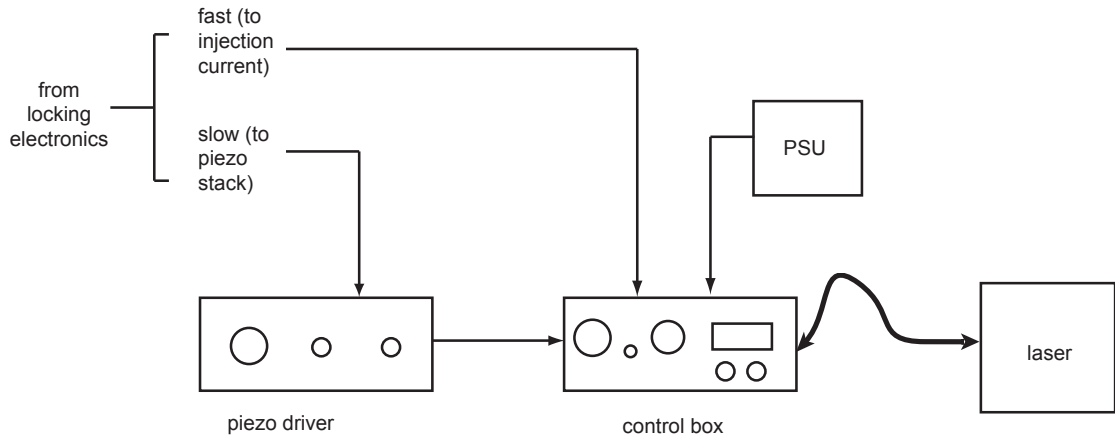


Figure 3.10: Block diagram of the diode laser system

ature stability provided by the two enclosures and the large tuning range provided by the piezo stack mean that we almost never have to open the laser box to adjust the precision screws. In an air-conditioned lab, the passive stability is good enough that the lasers will remain on the correct mode and transition for several days.

The Ti:Sapphire Laser

The MOT and imaging beams are provided by a Coherent MBR-110 Ti:Sapphire laser system, pumped by a 5 W Coherent Verdi V-5 diode-pumped solid state laser at 532 nm. This provides much more power than the diode lasers, which allows us to collect many more atoms in the MOT.

The Ti:Sapphire laser has a ring cavity. Coarse wavelength selection is provided by a birefringent filter, and fine tuning by a piezo and galvanometer controlled etalon. The laser also houses a temperature controlled reference cavity. The control electronics supplied with the laser include servo loops to control the etalon transmission, and to lock the laser to the reference cavity. With the laser locked to the reference cavity, the laser frequency can be scanned by up to 40 GHz by scanning the reference cavity length. The reference cavity servo lock gives the laser an RMS linewidth of 100 kHz measured relative to the reference cavity. With 5 W of pump power, we obtain a single mode output power of 400 mW with the laser running on the rubidium D2 transition. The output beam spatial mode is TEM_{00} , and the beam is horizontally polarised.

The passive stability of this laser is excellent. We find that for day to day operation, a wavemeter is not necessary. The birefringent filter almost never requires adjustment. We tune the laser using a saturated absorption spectrometer. Once the

output power has been peaked using the pump mirror alignment, the rubidium $D2$ lines are located by looking for the fluorescence in the rubidium cell using a CCD camera. The saturated absorption spectrum then allows us to pick out the ^{87}Rb trapping transition. The etalon and reference cavity servo loops are then closed, and fine tuning of the laser frequency is performed using the “scan offset” option on the control electronics.

3.5.3 Polarisation spectroscopy lock

The most common technique for sub-Doppler resolution spectroscopy is saturated absorption spectroscopy [71]. Here, the change in absorption of a weak probe beam is monitored as the laser frequency is scanned across the Doppler-broadened transition. A stronger counter-propagating pump beam depletes the population in the lower state, causing a dip in the absorption of the probe beam. The counter-propagating beam geometry ensures that only those atoms in a narrow velocity class near $v = 0$ interact resonantly with both beams at the same time, allowing sub-Doppler resolution. A dispersion signal can be obtained by modulating the laser frequency over a transition peak, and using lock-in detection.

Polarisation spectroscopy is a related technique that offers a higher inherent signal-to-noise ratio [72]. Instead of monitoring the change in absorption of the probe beam, we look for a change in its polarisation. It allows a dispersion curve to be obtained directly, without the need for modulation of the laser frequency [73]. This has been used elsewhere to stabilise the frequency of diode lasers in ^{87}Rb trapping experiments [74]. The polarisation spectrometer we use is shown in figure 3.11.




Figure 3.11: Layout of the polarisation spectrometer

The incoming laser light is split into a weak probe beam and a stronger pump beam. A quarter-wave plate circularly polarises the pump beam, while the probe beam is linearly polarised. The beams overlap in the rubidium cell. The probe polarisation is adjusted so that far from the atomic transition, the signal on each of the photodiodes is equal. When the laser frequency is tuned to an atomic transition, the circularly polarised pump beam causes optical pumping, leading to a non-uniform population of the magnetic sublevels. This anisotropy makes the sample birefringent and dichroic for the incoming probe light. After passing through the cell, the plane of polarisation of the probe light is rotated slightly, and its ellipticity is changed. These changes in the polarisation state of the probe light are detected as changes in the intensity on each photodiode. As shown in Appendix A, taking the difference of the two photodiode signals yields a dispersion curve as the laser is scanned across an atomic transition. No modulation of the laser frequency is required. The counter-propagating beam geometry allows sub-Doppler resolution, in the same way as for standard saturated absorption spectroscopy. In the case of the reference laser spectrometer, the rubidium cell is placed in a mu-metal shield, to reduce the effects of Faraday rotation on the observed signal.

Figure 3.12 shows the error signal from the reference laser polarisation spectrometer as the laser frequency is scanned across the transitions from the upper ground state hyperfine level for both isotopes. Large dispersion features are visible on the trapping transition for both isotopes, providing excellent signal to noise for locking to. There are weaker dispersive-type signals on some of the other transitions, as well as on some of the crossovers. Power broadening means that not all of the transitions are fully resolved. For the repump transition, the technique is less successful. The experimental spectrum for the ^{87}Rb repump transition is shown in figure 3.13. Despite the severe power broadening, it is clear that the largest signal is obtained on the cross-over resonances. If the laser power is reduced, we see that the dispersive signal on the $F = 1 \rightarrow F' = 2$ transition is extremely weak. The power broadened signal in figure 3.13 is preferred as it gives us a reasonable slope close to the transition, which we can lock to.

The performance of this lock is excellent; the reference laser can remain locked for several days. Despite the somewhat messy error signal, the repump laser will remain locked for a day. Even severe shocks to the optical table do not cause the laser to jump out of lock. The technique is simple to implement, and offers an error signal with excellent signal to noise without the need for modulation of the laser frequency. The only drawback is that the zero-crossing of the dispersion signal is not fixed - it depends on the orientation of the linear polariser. Rotating this

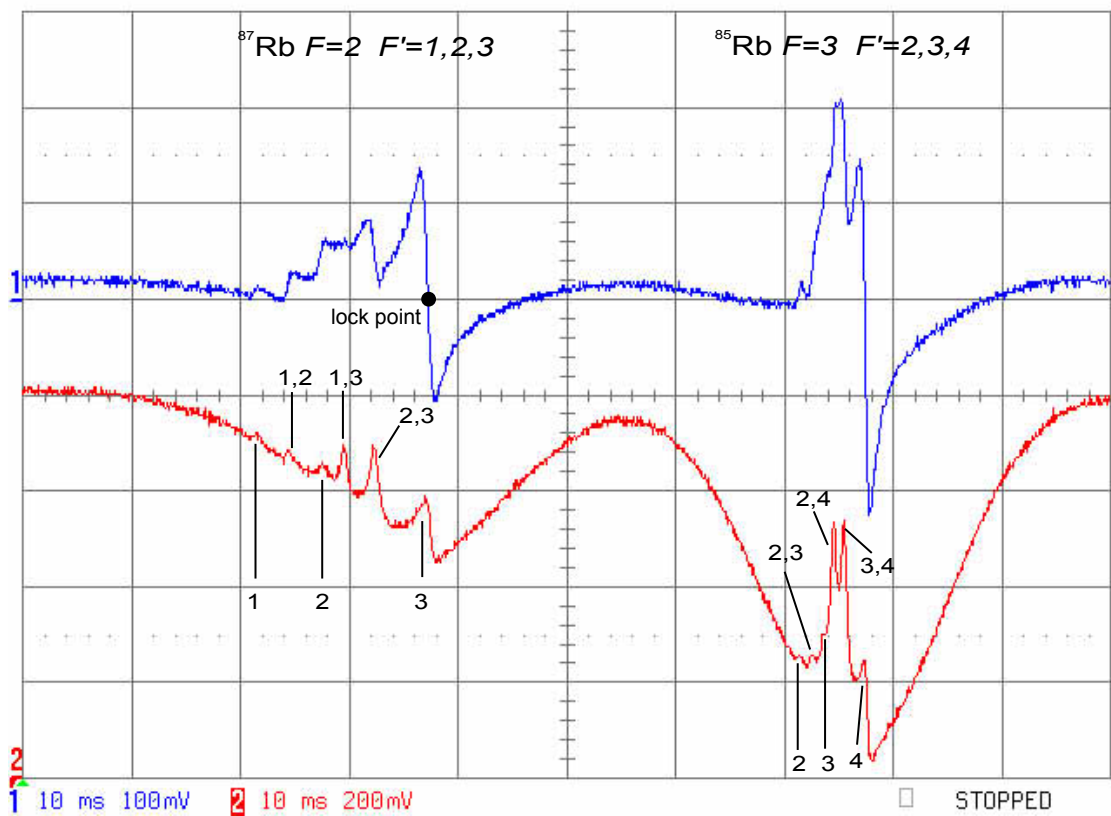


Figure 3.12: The output of the reference laser polarisation spectrometer, in a single scan across the “trapping” transitions for both isotopes. The red trace shows the signal from a single photodiode. The blue trace shows the error (difference) signal. The transitions and cross-over resonances are labelled by F' of the upper state(s). The horizontal scale is 10 ms/div.

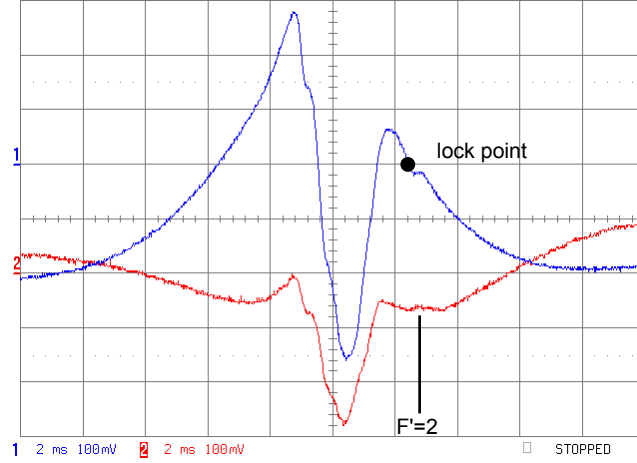


Figure 3.13: The output of the repump laser polarisation spectrometer, in a single scan across the ^{87}Rb $F = 1$ transition. The red trace shows the signal from a single photodiode. The blue trace shows the error (difference) signal. The horizontal scale is 2 ms/div.

polariser introduces an offset, which shifts the whole difference signal up or down. By comparing the signal with a standard saturated absorption signal from the trap spectrometer, the reference laser frequency can be set with an uncertainty of around 1 MHz. This is adequate for the experiments described in this thesis. However, for experiments that require the absolute frequency to be known more accurately, this offset would have to be addressed, or an alternative technique used.

3.5.4 Frequency offset lock

The frequency offset lock stabilises the beat frequency between the MOT laser and the reference laser. The scheme we use follows [75], with some modifications for operation at lower difference frequencies ($\simeq 100\text{MHz}$), and a steeper error signal. It works by comparing the frequency of the beat note between the two lasers to a local oscillator. An error signal is derived which is used to control the MOT laser reference cavity. The beat frequency, and hence the MOT laser, is tuned by adjusting the local oscillator frequency.

A schematic of the offset lock electronics is shown in figure 3.14. To measure the beat note, we overlap 2 mW of light from each laser. The light is focused onto an amplified fast photodiode (Newport 818-BB-21A), which gives a signal of +4 dBm at $f_{\text{beat}} \simeq 100\text{MHz}$. The beat signal is mixed with the output of a voltage controlled oscillator ($f_{\text{vco}} = 50 - 100\text{MHz}$), giving the difference frequency $\Delta f = |f_{\text{beat}} - f_{\text{vco}}|$. A 90 MHz low-pass filter blocks the output of the mixer at the sum frequency

$f_{beat} + f_{vco}$. The difference signal is split into two equal parts. One part is delayed by 5.5m of coaxial cable. This delay line introduces a phase shift between the two parts of the signal of $\phi = 2\pi\Delta f\tau$, where τ is the time delay introduced by the cable. The two parts are then recombined on a second mixer which acts as a phase detector. The output of the second mixer varies as $\cos\phi$. The output of the mixer at the sum frequency of the two arms is removed by a low-pass filter which has a cutoff frequency of 100 kHz. The resulting signal is buffered by an op-amp follower to give the error signal.

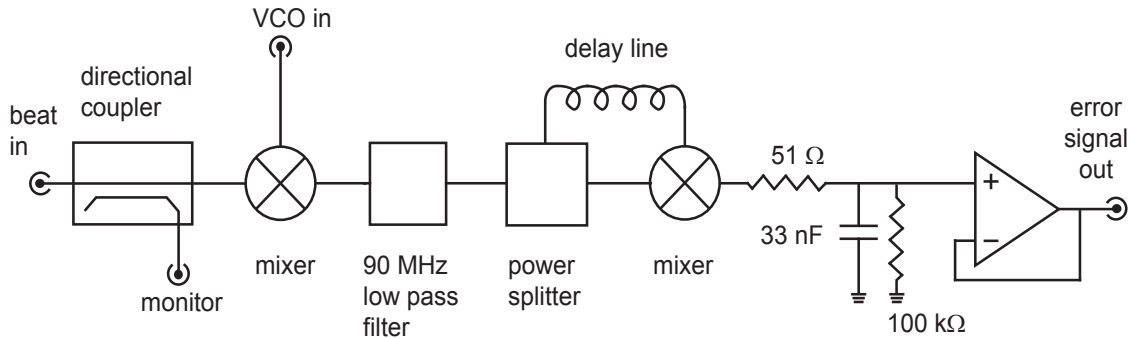


Figure 3.14: Schematic of the frequency offset lock electronics

The error signal as a function of the beat frequency is a series of cosine fringes, with a period that is set by the length of the delay line. For a cable of 1m length $\tau \simeq 5$ ns, and so in our case 5.5 m of cable gives a period $1/\tau$ of 36 MHz. This was chosen to give a similar capture range to the polarisation lock ($\simeq 18$ MHz). The blue curve in figure 3.15 shows the experimental error signal as the MOT laser is scanned, with the reference laser locked to the $F = 2 \rightarrow F' = 3$ transition. The envelope of the fringes is set by the 90 MHz low pass filter, which limits Δf to less than 90 MHz.

This locking scheme performs very well. Under normal laboratory conditions, the offset lock can be maintained until thermal drifts within the laser causes one of lasers to mode hop. This means the lasers will easily stay locked all day. The power spectrum of the beat note with the offset lock operating is shown in figure 3.16. This shows that over 1 second, the difference frequency is stable to better than 1 MHz. Indeed, the width is mostly accounted for by the linewidth of the reference laser. During normal operation, we monitor the beat note with a frequency counter. This shows that over long periods, the centre frequency is stable to better than 100 kHz.

Currently, the tunability is limited by the tuning range of the VCO to 50 MHz. A

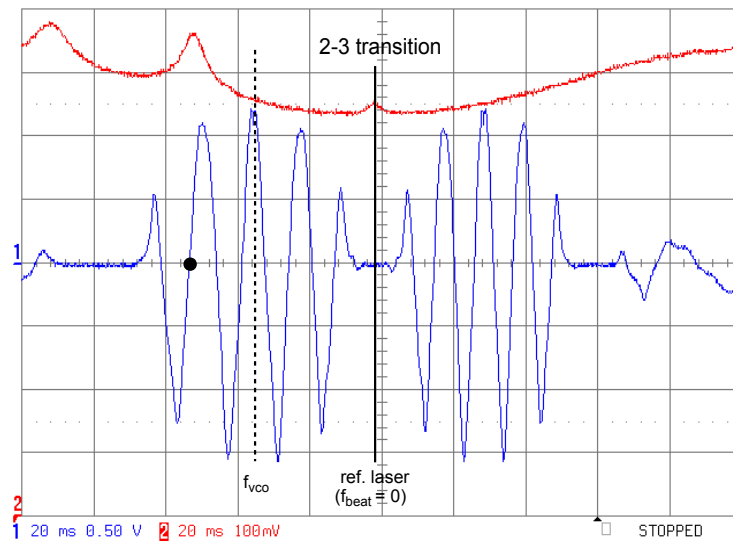


Figure 3.15: The offset lock error signal (blue) and saturated absorption spectrum (red) as the trap laser is scanned across the $F = 2 \rightarrow F'$ transition in ^{87}Rb . The reference laser is locked to the $F = 2 \rightarrow F' = 3$ transition. The dotted line indicates where $f_{\text{beat}} = f_{\text{VCO}} = 93.6\text{MHz}$. The dot marks the MOT laser lock point for a detuning of -137.2 MHz (-49.2 MHz after the trap AOM). The horizontal scale is 20 ms/div .




Figure 3.16: Power spectrum of the beat note (from the monitor output) with the MOT laser locked at the position marked in figure 3.15. Vertical scale is logarithmic, 10 dB per division. Horizontal scale is 2 MHz/div , and the sweep time is 1 s .

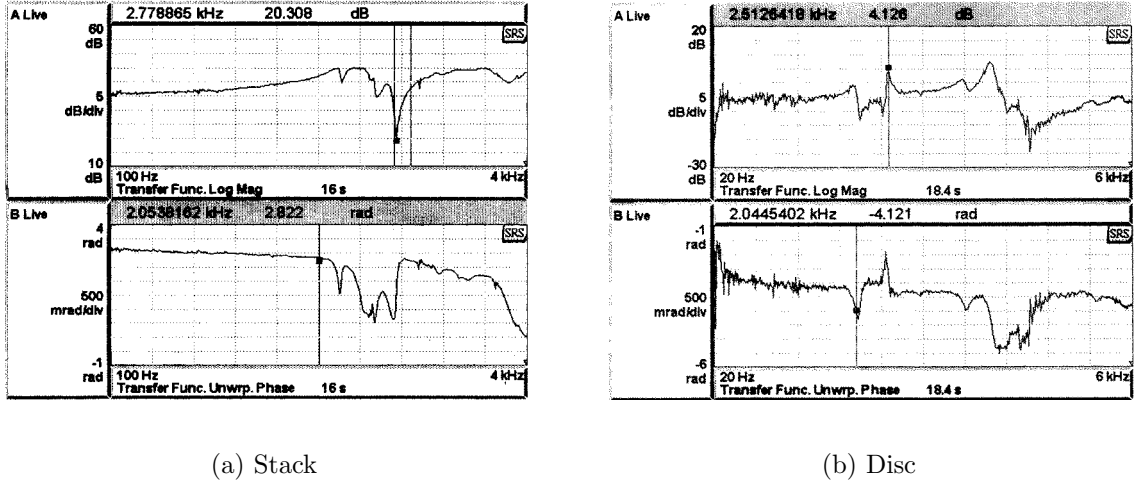


Figure 3.17: The transfer functions of the piezo stack and disc measured using a network signal analyser. A small amplitude swept sine wave is applied to the piezo, and the response of the laser frequency is monitored using the slope of the dispersion curve from the polarisation spectrometer

more fundamental limit is set by the low-pass filter, which must be able to separate $f_{beat} + f_{VCO}$ from $f_{beat} - f_{VCO}$. It follows therefore that the higher the operating frequency, the greater the tuning range. We are limited by the centre frequency of our trap AOM to an operating frequency of ~ 100 MHz.

3.5.5 Servo loop electronics and frequency jumping

The final link in the chain for frequency stabilisation is the servo loop electronics. Their job is to take each error signal as an input, and produce the necessary signals to correct the frequency of each laser.

The diode lasers have three possible feedback paths (see figure 3.9). The piezo stack under the precision screw and a piezo disc under the grating can both be used to control the laser frequency by changing the cavity length. The piezo stack also alters the angle of the grating, which enables coarse tuning of the frequency. The laser frequency can also be controlled by adjusting the injection current. The large tuning range of the piezo stack makes it ideal for dealing with DC drifts of the laser frequency. However, the relatively large mass of the grating and mount assembly means that its response rolls off at about 2 kHz (figure 3.17(b)). We have investigated using both the piezo disc and the injection current for high-frequency feedback.

The advantage of using the piezo disc is that no direct access to the injection

current is necessary, which makes it much harder to accidentally destroy the laser diode. However, the bandwidth over which feedback can usefully be applied is severely limited by the presence of mechanical resonances at frequencies as low as 2 kHz (figure 3.17(b)). We tried removing the gain at the lowest resonance frequency using a notch filter; however, the resonances are too closely spaced to make this really work - the laser merely oscillates on one of the higher frequency resonances. On the contrary, fast feedback to the injection current works extremely well, and is used on both the reference and repump lasers. Here there are no problems with mechanical resonances, and the bandwidth of the feedback is limited by the bandwidth of the electronics. In the case of the reference and repump lasers, it is the bandwidth of the photodiodes which is the limiting factor.

A block diagram of the electronics is shown in figure 3.18. For the polarisation lock, where we have two photodiodes, the input amplifier is a differential amplifier which produces the dispersion signal. In the case of the offset lock, this amplifier simply acts as a follower. The error signal is then sent into a “fast” proportional arm, and a “slow” integral arm, each of which has variable gain. In the fast arm, a high pass filter blocks the DC signal. The final stage is an amplifier whose bandwidth can be switched using DIP switches and capacitors to 16 values between 3.2 kHz and 159 kHz. When feeding back to the injection current, we set this limit to 159 kHz. In the slow arm, an integrator gives infinite gain at DC, ensuring that the laser is locked to the zero crossing. The frequency response of the slow arm is set by the time constant of the integrator. The final stage in this arm allows us to invert the sign of the slow feedback relative to the fast feedback. This is very useful, as the stack responds in antiphase to its input signal.

For the Ti:Sapphire laser, the requirements are different. Here, the reference cavity servo loop controls the linewidth, and we need only lock the reference cavity to the dispersion curve. However we also need to make fast (\sim ms) controlled changes to the laser frequency, such as detuning for the compressed MOT stage. Because of the design of the commercial control electronics, we only have one feedback path. The “external lock” input connects directly to the reference cavity. Although this input is quite sensitive (DC response 325 MHz/V), it is also very slow, with the response rolling off at 6 dB per octave above 12 Hz [76]. This response is fine for stabilisation at one constant frequency. The laser can be locked with the same electronics used for the diode lasers, using just the slow arm. Locking this way, the laser linewidth is less than 1 MHz (figure 3.16).

However, the slow response makes it hard to make rapid changes to the laser frequency. To get around this problem, we use a feed-forward. A specially shaped

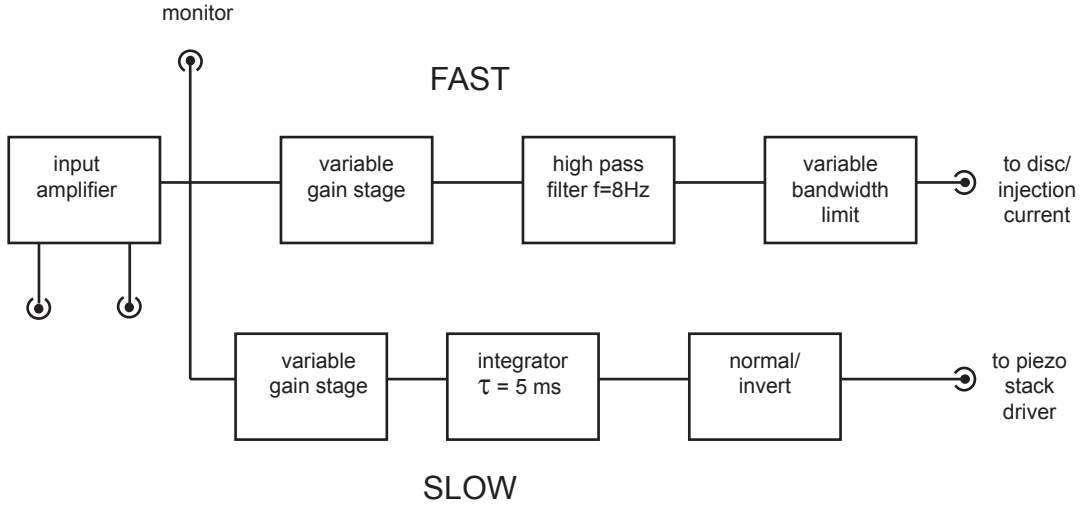


Figure 3.18: Block diagram of the servo loop electronics

pulse is applied to the external lock input to jump the laser frequency simultaneously with the change in frequency of the offset lock VCO. If the change in the laser frequency is matched to the change in the VCO frequency, the laser will remain on the correct fringe, and relock at the new frequency. Using the correct pulse shape, the jump in the laser frequency can be made very quickly (~ 5 ms). The correct pulse shape can be calculated by treating the external lock input as a simple low-pass RC filter, with a time constant $\tau_c = RC$. The resistor and capacitor form a potential divider for the input voltage, and so the output voltage is simply

$$V_{out} = V_{in} - IR = V_{in} - RC \frac{dV_{out}}{dt} \quad (3.1)$$

The output we require from this filter is a step response with a time constant $\tau \ll \tau_c$ that matches the time constant of the VCO frequency change:

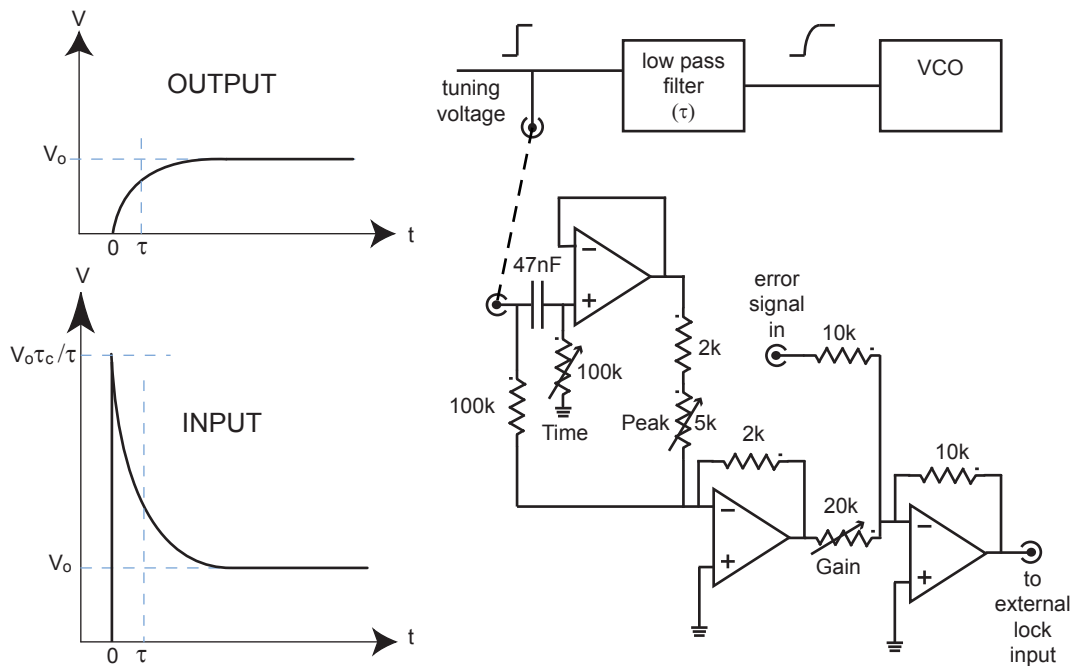
$$V_{out} = V_0(1 - e^{-t/\tau}) \quad (3.2)$$

Combining these two equations and rearranging gives the following expression for the corresponding input voltage:

$$V_{in} = V_0 \left(1 + e^{-t/\tau} \left(\frac{\tau_c}{\tau} - 1 \right) \right) \quad (3.3)$$

The input and output voltages are shown in figure 3.19(a).

The way this works in practice is shown in figure 3.19(b). A simple circuit generates the correct pulse shape. The time constant, overall pulse height (gain) and



(a) The input pulse shape that gives an output response time of τ .

(b) The electronics used to jump the laser frequency.

Figure 3.19: Jumping the Ti:Sapphire laser frequency

peak height are all adjustable to take into account the fact that the external lock input response is not quite that of a simple RC filter. The input is taken from the VCO tuning voltage, to ensure that the laser and VCO change frequency simultaneously. The output of this circuit is then summed with the output of the servo loop electronics and fed to the external lock input. This scheme works surprisingly well; as figure 3.20 shows, the laser frequency can be changed by 50MHz with a settling time of 6ms.

3.6 Imaging

All the information about the atom clouds prepared in our experiments is obtained by imaging them. We use two different techniques to probe the atoms, fluorescence imaging and absorption imaging.

In fluorescence imaging, the atoms are illuminated by laser light (usually from the MOT beams), and some of the light scattered by the atoms is collected to form an image. This technique is particularly useful for imaging atoms in the magneto-




Figure 3.20: The input signal is shown in blue, and the error signal in red. Dotted lines mark the maximum and minimum of the error signal fringes, showing that the laser remains on the same fringe. The horizontal scale is 2 ms/div.

optical trap, where the atoms are continuously scattering light. Imaging this scattered light provides a non-destructive measurement of the the MOT. By pulsing on some light, fluorescence imaging can also be used to image atoms in a magnetic trap or during ballistic expansion. In this case the imaging process is destructive, as the atoms are accelerated and heated by the probe light. A superior technique for these clouds is absorption imaging. Here, instead of imaging the light emitted by the atoms, the shadow cast in a probe beam due to the absorption by the atoms is imaged. This technique is also destructive for the same reasons. Its advantage over fluorescence imaging is a higher signal to noise ratio for short light pulses. The solid angle for light collection by the imaging system is small ($\ll 4\pi$), and so in fluorescence imaging, only a small fraction of the light scattered by the atoms is collected. This fact gives absorption imaging its strength; the whole of the absorption signal is collected, and as only a small fraction of the rescattered light is makes it into the imaging optics, the contrast of the absorption signal is very high.

3.6.1 The absorption imaging system

The imaging beam is shown at the top left-hand corner of figure 3.8. After the AOM, the polarisation of the beam is set with a polaroid linear polariser. The beam is then spatially filtered using a $40\ \mu\text{m}$ pinhole, before being recollimated with a spot size of $\sim 12\ \text{mm}$ FWHM. The large spot size ensures even illumination across the atom cloud. Gold-coated mirrors are used to direct the imaging beam into the chamber to preserve the polarisation. Just before the beam enters the chamber, it is circularly polarised by a quarter-wave plate, such that the light is σ^+ polarised




Figure 3.21: Layout of the imaging system

with respect to an imaging magnetic field of 1 G, which is applied in the direction of the light beam ($-x$) using the x earth cancelling coils. The total power is 2.5 mW in the beam entering the chamber.

The image is formed by a high-quality lens (Comar 03 TT 25), which consists of a pair of aplanatic doublets mounted back-to-back. This lens has a nominal focal length of 80.5 mm. We mount the lens approximately 150 mm from the cloud on a micrometer stage that is firmly bolted directly to the optical bench to reduce vibrations. The micrometer stage allows extremely fine adjustment of the focus, which is crucial in obtaining good images. The magnification of the imaging system deduced from the object and image distances is $M = 1.22(3)$.

The image is acquired by a Princeton Instruments PentaMax 5 CCD camera. The pixel size is $6.8 \mu\text{m}$ square, and the chip size is 1317×1035 pixels. The camera has a 12-bit DAC, and so each pixel gives a number of counts between 0 and 4095. The camera also has a mechanical shutter which opens each time an exposure is taken. The minimum exposure time was measured to be 10 ms. The CCD chip must be read out between frames, which takes 330 ms for a full-chip image. This sets a limit on how close together images can be acquired.

During our early attempts at absorption imaging the CCD chip was illuminated at normal incidence. The resulting images were covered with a pattern of closely-spaced fringes. These did not completely divide away during the image processing, probably due to small movements of the camera because of shutter vibrations. The cause of the fringes was eventually tracked down to multiple reflections between the two windows in front of the chip. As the imaging beam is highly coherent, these

windows act as a Fabry-Perot interferometer, producing the fringe pattern. In order to eliminate these fringes, the camera must be rotated away from the normal by an angle large enough to ensure that multiple reflections between the windows are not incident on the CCD array. The camera was tilted by $\theta = 27.4$ degrees to overcome this problem. Combined with the magnification, this means that each camera pixel corresponds to an area of $5.55 \mu\text{m}$ (vertical) by $4.93 \mu\text{m}$ (horizontal) in the object plane.

The CCD chip was carefully calibrated against a power meter to yield a sensitivity for each pixel of 0.0256 counts per photon at 780nm. This information can be used in conjunction with the imaging pulse duration and the pixel size to yield an accurate measurement of the intensity where the atoms are located. The imaging beam intensity and pulse duration are a trade-off between a number of factors. The percentage of light absorbed per atom is higher at lower intensities. Using a short imaging pulse reduces the effect of recoil blurring (see section 5.3.2). However, as the final image results from dividing one frame by another, we would like to use as much of the dynamic range of the camera as possible, to reduce the statistical noise and increase the “depth” resolution. A good compromise is an intensity of 2.0 mW cm^{-2} and a pulse duration of $30 \mu\text{s}$, yielding an average of 1700 counts per pixel.

3.6.2 Analysing absorption images at low optical density

If we neglect the finite resolution of the imaging system, and assume that the image perfectly reconstructs the intensity distribution in the object, then the signal $F(y, z)$ in each absorption image has the following form:

$$F(y, z) = I(y, z)T(y, z) + S(y, z) . \quad (3.4)$$

Here $I(y, z)$ is the intensity profile of the absorption imaging beam, $T(y, z)$ is the transmission of the atom cloud and $S(y, z)$ represents the background. The background will consist of dark counts, ambient light and leakage through the AOM during the exposure.

In the low intensity limit ($I < I_{sat}$), light absorption in the cloud is an exponential process, and hence we can write

$$T(y, z) = e^{-D(y, z)} \quad (3.5)$$

where $D(y, z)$ is the optical density of the cloud. The optical density is related to the column density of the cloud $n_{2D}(y, z)$ as follows.

$$D(y, z) = \sigma_L n_{2D} = \sigma_L \int n(x, y, z) dx \quad (3.6)$$

where σ_L is the cross-section for light absorption. When imaging untrapped clouds, a magnetic field is applied during imaging to give a well-defined quantisation axis. Since the imaging light is σ^+ polarised with respect to this quantisation axis, the atomic population is rapidly optically pumped into the $F = 2, m_F = +2$ state. For σ^+ polarised light this forms a closed two level system with the $F' = 3, m_F = 3$ sublevel. The cross-section can then be calculated using the two-level atom result:

$$\sigma_L = \hbar\omega \frac{\Gamma}{2} \frac{1/I_{sat}}{1 + I/I_{sat} + 4\Delta^2/\Gamma^2}. \quad (3.7)$$

The saturation intensity for this transition is $I_{sat} = 1.67 \text{ mWcm}^{-2}$. For our imaging beam parameters, $\sigma_l = 1.31 \times 10^{-13} \text{ m}^2$. The fact that the transition is closed means that repump light is not needed during imaging.

When we want to image clouds that are still confined in the magnetic trap, the situation is not so simple. Here, the quantisation axis is defined by the local magnetic field, and so the imaging beam no longer drives purely σ^+ transitions. In this case, the effective cross-section is hard to calculate. Instead, we determine it experimentally. Using clouds prepared under identical experimental conditions, we compare the amount of absorption obtained in the trap to that obtained after the cloud has been released. The ratio of the signals gives the ratio of the effective cross-section for images in the trap σ'_L to the cross-section σ_L calculated above. This cross-section can σ'_L then be applied to all the images of similar clouds in this particular magnetic trap. It is important to note that as the trapping field is highly inhomogeneous, the correction to the cross-section is different for a hot cloud compared to a cold cloud in the same trap.

In the low optical density limit, the transmission of the atom cloud $T(y, z)$ can be extracted from the absorption image $F(y, z)$ with the help of a background image taken under the same conditions but with no atoms present. The ratio of the two is

$$F'(y, z) = \frac{(I(y, z)T(y, z) + S(y, z))}{(I(y, z) + S(y, z))}. \quad (3.8)$$

In the limit of low optical density $I(y, z)T(y, z) \gg S(x, y)$ and

$$F'(y, z) \simeq T(y, z) = e^{-\sigma n_{2D}}. \quad (3.9)$$

On taking the log, the column density distribution is recovered:

$$n_{2D}(y, z) = -\frac{1}{\sigma_L} \ln(F'(y, z)). \quad (3.10)$$

In practice, shot-to-shot variations in the probe beam intensity mean that the images must be normalised before the log is taken. This is achieved by measuring the

average value per pixel in a region far from the atom cloud, and then dividing the entire image by this value. The number of atoms can then be obtained by integrating over the column density of the cloud:

$$N = \frac{A}{\sigma_L} \sum_{pixels} -\ln(F'(y, z)) \quad (3.11)$$

where A is the area of a pixel in the object plane. To extract other parameters such as density, temperature, phase space density etc., a model for the density distribution of the cloud must be assumed.

3.6.3 Absorption imaging at high optical density

The above analysis no longer applies when the $I(y, z)T(y, z) \lesssim S(y, z)$. In our case $\bar{S}/\bar{I} \simeq 0.04$ so this restriction is satisfied provided the optical density in equation 3.5 is less than three. Above this limit, the column density cannot be extracted directly from the images using equation 3.10. Instead, a more sophisticated model of the imaging process must be used.

In the experiments presented in this thesis, optically dense clouds are observed in two regimes. The first is for very large clouds containing a large number of atoms, such as a cloud that has been released from the MOT and allowed to expand ballistically. Here the cloud is physically “thick” along the line of sight of the imaging beam, and hence the column density is large. In this case the optical density can be reduced by detuning sufficiently from resonance. Quantitative information can still be obtained by accounting for the detuning in the formula for the cross-section (equation 3.7).

At the other end of the scale, clouds close to the BEC transition, where the density is extremely high, also have very high optical densities. Here, detuning to reduce the optical density does not work. This is because the dispersive part of the refractive index becomes significant away from resonance, and refraction of the imaging light leads to severe distortion of the cloud image. These distortions are much less important for the large clouds discussed above. In principle, if the imaging light is detuned far enough, the dispersive effects can be reduced to a satisfactory level. However, the cross-section then becomes very small making the signal very hard to detect. The analysis of absorption images of cold, dense trapped clouds is discussed in detail in Chapter 5.

It is worth noting that the limit of detectable optical density should be dramatically increased by taking an image without the light pulse, yielding $S(y, z)$. By subtracting this from both the data and normalisation images, the background can

be eliminated. In this case, the limit of the optical density that can be measured is set by the dynamic range of the camera. The smallest change in intensity we can measure is 1 count in 1700, which corresponds to an optical density of 7.4. However, we found that normalising the pictures in this way introduced unacceptable amounts of noise into the final image, due to small movements of the camera between frames.

3.6.4 Fluorescence imaging

The same lens and camera can be used for fluorescence imaging, by taking an image of the light scattered by the atoms. Light scattered by other sources (such as the wire) is eliminated by subtracting a background image without the atoms present. In our experiment, fluorescence imaging is only used to take images of the atoms when they are in the magneto-optical trap, and scattering light anyway. The imaging beam is not used for fluorescence imaging.

The number of atoms in a fluorescence image is given by

$$N = \frac{1}{d\Omega LPR\tau} \sum_{pixels} Count \quad (3.12)$$

where R is the scattering rate, L is the transmission fraction of the optics, P is the number of counts per photon per pixel, $d\Omega$ is the solid angle captured by the lens and τ is the exposure time. The uncertainty here is dominated by two factors. The first is the uncertainty in the solid angle. This can be difficult to calculate, as the solid angle can be limited by things like the edge of the chip whose exact location is hard to measure. The second factor is the uncertainty in the scattering rate. This arises because in the MOT, the saturation intensity is not well defined. The atoms are scattering on many different transitions, with different amplitudes. Some authors use a saturation intensity averaged over all the transitions, some the saturation intensity for the stretched state transition only. The saturation intensity in a MOT can be measured experimentally (see for example [55, 77]). These measurements agree closely with the stretched state value of 1.67 mW cm^{-2} , and so this value is used for fluorescence measurements in this thesis. It should be noted that both these factors will only affect the *absolute* atom number; relative atom numbers can be known much more accurately.

Fluorescence images were taken using two different lens systems. The first is the system used for absorption imaging described above. In the second system, a commercial zoom lens was attached directly to the camera using a ‘‘C’’ mount. This lens has a minimum F-number of 2.5, and a maximum magnification of 1.15. It is quite lossy at 780 nm - only 63% of the light is transmitted. As the scattered light is not coherent, the camera does not have to be mounted at an angle for fluorescence

imaging. This system was used in the study of density and phase space density in the compressed MOT.

Chapter 4

Reaching BEC

This chapter discusses in detail how each stage of the path to BEC outlined in the introduction was implemented. To reach BEC, the performance of each stage needed to be carefully optimised. Therefore the optimisation of the experimental parameters at each stage is described in some depth. The goal was to get the highest possible value of $\gamma_{el}\tau_{loss}$ in the compressed trap at the beginning of evaporative cooling. Once the trap frequencies were decided, this boiled down to reaching this point with as many atoms as possible. From there, we must then evaporatively cool as effectively as possible. Therefore the story of this chapter is really the quest for more atoms at every step of the way. The order of the sections reflects the order of the experimental sequence. The performance of each stage is summarised by a table of the important parameters of the cloud. At the end of the chapter, the path as a whole is evaluated.

4.1 The MOT

4.1.1 Making the MOT

To make the MOT, the trap and repump lasers are locked to the correct frequencies and the shutters are opened. The MOT coils are switched on at a current of 1.5 A, which gives a field gradient of 0.15 Tm^{-1} along the axis of the coils. The MOT is then filled by pulsing the dispenser. Between experimental runs, the current through the dispenser was set at 2 A. This is below the threshold for rubidium production, but keeps the dispenser warm so that it heats up more quickly when pulsed. During the pulse, the current through the dispenser is increased from 2 A to 7.5 A. The duration of the pulse varies slightly; a longer pulse gives a larger MOT, but it takes longer for the vacuum to recover. The dispenser pulse normally lasts 10-12 seconds. At the end of the pulse, the dispenser is switched off completely. The subsequent




Figure 4.1: The filling and decay of the MOT, plotted on a log scale. At $t = 0$ the dispenser is switched off. The subsequent evolution of the atom number is monitored by taking a sequence of fluorescence images at constant detuning. The fit shown gives a lifetime of 20s.

evolution of the MOT is shown in figure 4.1.

The number of atoms in the MOT depends on the balance between the loading rate R_f and the loss rate γ . If we assume that the loss rate is independent of N , then the number of atoms in the MOT is given by the solution of the following equation:

$$\frac{dN(t)}{dt} = R_f(t) - \gamma(t)N(t) \quad (4.1)$$

Two components to the loss rate are normally identified; collisions with atoms in the background gas, and collisions with untrapped rubidium atoms. In a normal vapour cell MOT, the rubidium partial pressure is constant, and the loading rate and decay rate are both independent of time. In this case, the solution of 4.1 is a simple exponential fill curve, with the steady state number of atoms set by the ratio of the fill and loss rates [35].

In the case of our pulsed MOT, the rubidium partial pressure in the trapping region is time-dependent and the situation is more complicated. Figure 4.1 shows that atoms are being collected for several seconds after the dispenser is switched off, which means that it takes several seconds for the rubidium pressure in the trapping region to decay to the background level. At long times, the decay of the MOT is exponential, and the loss rate γ is constant. This suggests that the untrapped rubidium background has decayed away, and the loss rate is not fixed by collisions with background rubidium. The fundamental limit on this loss rate is




Figure 4.2: MOT atom number versus detuning. In each case, the detuning was jumped to -40 MHz for imaging. Intensity $I = 114 \text{ mW cm}^{-2}$.




Figure 4.3: MOT atom number versus trap laser intensity. The detuning here was -18 MHz (3Γ).

set by collisions with non-rubidium background gas. However, as shown in the next section, the lifetime also depends on the parameters of the MOT, and we do not reach this limit.

4.1.2 Optimising the MOT

The number of atoms collected in a MOT depends on several parameters. Previous work on 6-beam MOTs has shown that the number of atoms increases with both increasing beam size and power [78], and that the largest number of atoms is collected at moderate field gradients ($\sim 1.5 \text{ Tm}^{-1}$) and detunings ($\Delta \sim -3\Gamma$). In our experiment, the maximum useful beam size is fixed by the area of the atom chip. The dependence of N on trap laser detuning and intensity is shown in figures 4.2 and 4.3 respectively. These are in good agreement with results for six beam MOTs already presented in the literature [78]. The peak atom number is obtained with a detuning of -3Γ (-18 MHz) and an intensity of 80 mW cm^{-2} (which corresponds to a total power of 170 mW), and so the MOT is normally operated with these parameters.

The number of atoms and the lifetime are quite sensitive to the alignment of the MOT beams. Empirically we find that perfect overlap and alignment does not always give us the optimum conditions for the MOT. This is partly due to interference fringes in the light field as shown in figure 4.4. However, we also observe a systematic increase in the number of atoms and the lifetime with height above the




Figure 4.4: A fluorescence image of a large MOT, containing 1.6×10^8 atoms. The fringes in the image of the cloud are due interference fringes in the MOT beams.

surface, as shown in figure 4.5.

This height dependence is probably due to the shadow that the guide wire casts in the MOT beams reflected from the surface. To lift the MOT as far above the surface as possible, the x “earth cancelling” coils are used to produce a magnetic field in the $+x$ direction during the MOT phase, which translates the centre of the quadrupole field in the $+y$ direction. These coils have a high resistance which limits the current to ~ 2 A. At this current, the MOT forms 3.8 mm above the surface of the wire. With the MOT alignment well optimised, the lifetime can be as long as 30 s at this height.

With the parameters described above, we can collect $> 1.6 \times 10^8$ atoms with a 14 s dispenser pulse (figure 4.4). This MOT is clearly visible with the naked eye. For the BEC experiments, we use a 12 second pulse and collect $\sim 1 \times 10^8$ atoms. The properties of such a MOT are summarised¹ The atoms were held in the MOT for 10 seconds after the end of the dispenser pulse so that the vacuum could recover before the magnetic trapping phase. At this point the MOT was compressed to load the magnetic trap.

4.2 The Compressed MOT

Once the MOT is full, as many of the atoms as possible must be transferred to the magnetic trap. The first part of the transfer process is the compressed MOT

¹Table 4.1 gives the properties of a single realisation of the MOT, under the conditions used to produce a BEC. Averaging over many experimental runs, the dominant uncertainty in the average properties is statistical, shot-to-shot variation in the cloud properties. The size of this variation depends a lot on how well the experiment is set up. Hence I have chosen to use data from a single image throughout this chapter. Systematic errors are discussed in section 3.6.

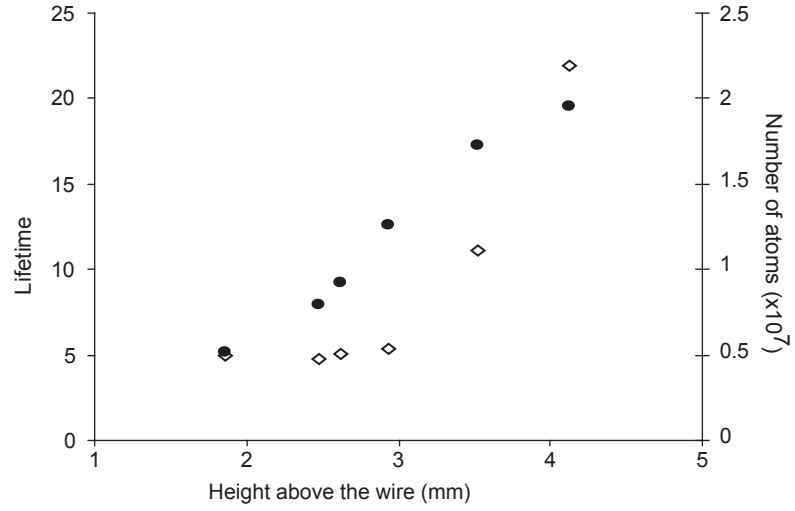


Figure 4.5: The dependence of MOT atom number (filled circles) and lifetime (open circles) on height above the surface of the wire.

Number of atoms	N	1.0×10^8
Radial size	σ_y (μm)	510
Axial size	σ_z (μm)	810
Peak density	n_0 (m^{-3})	3.0×10^{16}
Temperature	T (μK)	50
Phase space density	PSD	5.5×10^{-7}

Table 4.1: Properties of a typical mirror MOT. Here σ is the rms cloud radius.

(CMOT). The guide and bias field are switched on so that the quadrupole field they produce overlaps the MOT. The bias field is then ramped up, bringing the MOT closer to the surface and compressing it radially.

The compressed MOT stage fulfils several roles. The first is to bring the atoms closer to the surface. At the height of the MOT (3.8 mm), the fields produced by the wires on the chip are quite weak, and so only a shallow magnetic trap can be formed here. The strong confining force in the CMOT allows the atoms to be kept tightly trapped as they are lowered towards the surface. The compressed MOT geometry, where the radial field gradient is much higher than the axial field gradient, helps change the cloud shape from roughly spherical to an ellipsoid. This helps match the cloud shape to the shape of the magnetic trap. Finally, as dissipative forces are active in the compressed MOT, the cloud can be cooled as it is compressed, which means that the phase space density can be increased. This should be contrasted with compressing the magnetic trap, where the best one can hope for is to preserve the initial phase space density by compressing adiabatically.

4.2.1 The CMOT in theory

The effect of compressing the MOT by increasing the field gradient has been studied experimentally and theoretically for “spherical” 6-beam MOTs [77, 79]. The central observation is that merely increasing the magnetic field gradient does not lead to significant compression of the cloud. This is due to radiation pressure. Light scattered out of the trapping laser by one atom can be reabsorbed by another, pushing the atoms apart. This can be thought of as a repulsive force between the atoms. As the density increases, the probability of rescattering increases, and this repulsive force goes up. In the regime where each photon is rescattered only once before leaving the cloud, this force has an inverse square law dependence on the distance r between the atoms [80]:

$$\mathbf{F}_R = \frac{I\sigma_L}{c} \frac{\sigma_R}{4\pi r^2} \hat{\mathbf{r}} \quad (4.2)$$

where σ_L is the cross-section for scattering light from the trapping beams, σ_R is the cross-section for absorption of the rescattered light, I is the intensity of the trapping light and c is the speed of light. The density distribution of the cloud is the result of equilibrium between this rescattering force and the trapping force.

The force due to a single rescattering discussed above (equation 4.2) is important in determining the density in the MOT, but it does not give rise to the observed incompressibility. This arises because of multiple scattering. As the density increases, the mean free path of a scattered photon will drop. As soon as there is an appreciable probability for a photon to be rescattered more than once before leaving the cloud,




Figure 4.6: The ratio of the total rescattering force for cylindrical and spherical clouds of the same volume and atom number, as a function of the aspect ratio l/r of the cylinder. This is calculated in the regime where $F \propto 1/r^2$. This figure is reproduced with permission from [81]

the rescattering force increases dramatically. This sets an upper limit on the density. The rescattering force can be reduced by reducing the rate for scattering light out of the trapping beams, by reducing the detuning or the intensity. However, this also reduces the restoring force in the MOT, and eventually this strategy reaches a limit. As the temperature is also limited by the recoil temperature as discussed in section 2.1.2, this limit on the density places a limit on the maximum phase-space density that can be achieved. In practice, the density is limited to $\sim 5 \times 10^{17}$ atoms/m³, and the phase space density to 1.5×10^{-5} [77] for a “spherical” MOT [77, 79].

The rescattering force can also be altered by changing the geometry of the cloud. In a long, thin “cylindrical” geometry, photons do not have to travel so far to leave the cloud radially, reducing the probability of rescattering. For the $1/r^2$ case, the variation of rescattering force with aspect ratio is shown in figure 4.6. Thus by compressing the cloud in only the radial direction, it should be possible increase the limiting density and phase space density.

We performed a comprehensive study of the behaviour of these asymmetrically compressed MOTs, looking at the constraints on the achievable densities and phase space densities. A detailed analysis of the results and a theoretical study were performed by Wenzel Salzmann [81]. In particular, the dependence of density, and phase space density on the detuning and axial and radial field gradients was inves-

tigated. In these experiments, ^{85}Rb atoms were collected in a normal MOT, which was then compressed by ramping up the current in the guide wire and the bias field in 100 ms. The laser detuning was switched from the detuning for the MOT (-12 MHz in this case) to the new detuning for 20 ms before the cloud was imaged. Some of the results obtained are illustrated in figure 4.7 and 4.8. The highest values measured were a peak density of $6.7(9) \times 10^{17} \text{ m}^{-3}$ and a peak phase space density of $3.4(5) \times 10^{-5}$, which are a factor of two larger than the largest values reported for the spherical case. This data shows that increasing the aspect ratio can increase both the density and the phase space density. It also shows that in order to achieve significant compression of the MOT, the cross-section σ_L must be reduced, in this case by detuning further from resonance.

4.2.2 The CMOT in practice

The experiments described above were performed using a diode laser as the trap laser, and only a small number of atoms ($\sim 5 \times 10^6$) were collected. With the Ti:Sapphire laser, the intensity and the atom number are both much higher. To reduce the rescattering force and allow the MOT to compress, a large reduction in the scattering rate was necessary. The timing diagram for the CMOT phase used in the BEC experiments is shown in figure 4.2. The scattering rate is reduced by both increasing the detuning and lowering the intensity. Relaxing the axial confinement by lowering the MOT coil current reduces the temperature and increases the aspect ratio. The CMOT forms 1.5 mm above the centre of the wire and has a radial gradient of 0.29 Tm^{-1} and an initial axial gradient of 0.06 Tm^{-1} .

The properties of a typical CMOT are shown in table 4.2. Both the density and the phase space density are somewhat lower than the optimum values quoted above. The main reason for this is the sheer size of the cloud. For such a large cloud, only a small fraction of the atoms are in the region near the centre where polarisation gradient cooling is active [77]. This means that the cloud is hotter, and also less dense, as atoms outside this region feel a weaker confining force. The lifetime of the compressed MOT (~ 500 ms) is considerably shorter than that of the MOT. This is presumably because the compressed MOT is much closer to the region of the MOT beams that is shadowed by the wire (see figure 4.5).

Many BEC experiments use a different technique, known as a “dark MOT” [82] to increase the density and lifetime of the MOT. This technique works by excluding the repump light from part of the cloud. Atoms in this region are optically pumped into the lower hyperfine level, where they no longer scatter any light. This leads to a huge reduction in the rescattering force, and densities of up to 10^{18} m^{-3} can be




Figure 4.7: Density versus detuning for two different magnetic field aspect ratios. Open diamonds are with $B'_{radial} = 0.6 \text{ Tm}^{-1}$, $B'_{axial} = 0.06 \text{ Tm}^{-1}$, giving an aspect ratio of 10. Filled diamonds are with $B'_{radial} = 0.72 \text{ Tm}^{-1}$, $B'_{axial} = 0.015 \text{ Tm}^{-1}$, giving an aspect ratio of 48. Note that the spatial distribution of the atom cloud has a much lower aspect ratio (~ 5).




Figure 4.8: Phase space density for two different aspect ratios. Open diamonds are with $B'_{radial} = 0.6 \text{ Tm}^{-1}$, $B'_{axial} = 0.06 \text{ Tm}^{-1}$, giving an aspect ratio of 10. Filled diamonds are with $B'_{radial} = 0.72 \text{ Tm}^{-1}$, $B'_{axial} = 0.015 \text{ Tm}^{-1}$, giving an aspect ratio of 48.




Figure 4.9: The timing sequence for the CMOT.

Number of atoms	N	1.0×10^8
Radial size	σ_y (μm)	240
Axial size	σ_z (μm)	680
Peak density	n_0 (m^{-3})	1.7×10^{17}
Temperature	T (μK)	85
Phase space density	PSD	1.4×10^{-6}

Table 4.2: Properties of a typical compressed MOT.

achieved. The dark MOT can also be implemented temporally instead of spatially [83], by dropping the intensity of the repump light at the end of the MOT stage. In the future, this technique could be implemented on our experiment. A higher density would allow better mode-matching (see section 4.4.2). Also, by allowing the atoms to be pumped into the lower hyperfine level, it would enable us to magnetically trap atoms in the $F = 1, m_F = -1$ state. This might prove more favourable for preparing BEC because the inelastic loss rate is more favourable.

4.3 Optical pumping

At the end of the CMOT phase, all the magnetic fields and the trap AOM are switched off. The compressed MOT leaves the atoms in a distribution of m_F sub-levels. To magnetically trap as many as possible, optical pumping is used to transfer the atoms into the $F = 2, m_F = 2$ state.

In this experiment, the atoms are optically pumped into $F = 2, m_F = 2$ using the $F = 2 \rightarrow F' = 2$ transition with σ^+ light, for which this state is a dark state. This keeps the heating to a minimum during the optical pumping pulse. The optical pumping light is derived from the reference laser using a double pass through an AOM, as shown in figure 3.8. The separation of the $F' = 2$ and $F' = 3$ levels is 267 MHz. As the reference laser is locked on resonance with the $F = 2 \rightarrow F' = 3$ transition, a double pass through the AOM at -140 MHz leaves us 13 MHz red-detuned from the optical pumping transition. Red-detuned light was used as the cloud is optically thick for resonant light. Repump light is also needed during optical pumping; this was provided by the repump light in the MOT beams.

After the AOM, the optical pumping beam is spatially filtered and expanded to a 6 mm diameter (FWHM) spot, to ensure that the atom cloud is evenly illuminated. The beam enters the chamber along the same path as one of the horizontal MOT beams, as shown in figure 3.8. The beam paths are overlapped using a polarising beam splitter cube. A linear polariser before the beam splitter enhances the polarisation purity of the beam, and ensures that all the light leaves through the correct port of the beam splitter cube. The beam is circularly polarised by the MOT quarter-wave plate. The resulting polarisation purity is better than 99.5%. The optical pumping beam is σ^+ polarised with respect to a quantisation axis in the $+z$ direction, opposite to its propagation direction. This quantisation axis is provided by the combined magnetic field produced by the end wires and the z cancelling coils.

The optical pumping cycle lasts 2 ms, and begins as soon as the CMOT is switched off. Figure 4.10 illustrates the timing. At the end of the CMOT phase, all the magnetic fields are switched off. The trap light is switched off using the trap AOM, and the trap laser shutter closes. The optical pumping beam shutter opens. At the same time, the end wires and z -cancelling coils are triggered. The optical pumping light pulse, which lasts 250 μ s, is triggered 1.5 ms later. The beam is pulsed by switching the RF power to the AOM².

The efficiency of optical pumping was optimised in two ways. Firstly, the number

²If the RF is off except during the pulse, this AOM takes several seconds to turn on. This is a thermal effect - the AOM runs warm after a few seconds at full power. As a result, the optical pumping AOM was kept switched on until 20 ms before the pulse, with the beam blocked by a shutter, so that the AOM did not have time to cool down.




Figure 4.10: The optical pumping sequence. The horizontal scale is 0.5 ms/div. Red and blue traces show the current through the end wires and bias coils respectively, and the TTL signal for the OP AOM is shown in black. The z-cancelling field (not shown) is switched on simultaneously with the end wires. The magnetic trap turns on 2 ms after the CMOT has ended.

of atoms in the magnetic trap with and without the optical pumping pulse were compared. After optimisation, optical pumping gave a factor of ~ 2.5 improvement. The distribution of m_F sublevels after the CMOT is not known, but if we assume a completely random distribution, perfect optical pumping would indeed give an improvement of 2.5 in the population of trapped states ($m_f = +1, +2$). The second test of how well optical pumping was working was to reverse the helicity of the light polarisation. The atoms should then be pumped into untrapped m_F states, and no atoms should be trapped. With the helicity reversed, the number of atoms trapped was a factor ~ 5 lower than with no optical pumping. This suggests that optical pumping worked well. No heating of the atom cloud by the optical pumping pulse was observed; however if too much light is used the cloud could be “pushed” slightly in the direction of the beam. This suggests that the σ^+ polarisation of the light was not perfect and therefore the $F = 2$, $m_F = 2$ state was not entirely dark.

After the optical pumping pulse has finished, the optical pumping and repump shutters are closed, leaving the atoms in the dark. The guide wire and bias field switch on at the end of the cycle to complete the magnetic trap.

4.4 Loading the magnetic trap

In order to maximise the phase space density in the magnetic trap, the shape and size of the trap should be matched to the dimensions of the atom cloud after optical

pumping. This is called mode-matching, by analogy with the coupling of light into an optical fibre. If the trap is too tight or too weak, breathing oscillations of the cloud are excited. If the trap and the cloud are not properly overlapped, centre of mass oscillations will occur. Both excitations heat the cloud.

4.4.1 Mode-matching in theory

By measuring the size and temperature of the cloud after optical pumping, a mode-matched trap can be designed. To avoid heating, the trap frequency ω_j should be such that in the trap, a cloud of width σ_j and temperature T after optical pumping has a mean potential energy in the trap of $\frac{1}{2}k_B T$ in each direction, i.e.

$$\frac{1}{2}k_B T = \frac{1}{2}\kappa_j \sigma_j^2 = \frac{1}{2}m\omega_j^2 \sigma_j^2 \quad (4.3)$$

which gives the following conditions on the trap frequencies:

$$\omega_j = \frac{1}{\sigma_j} \sqrt{\frac{k_B T}{m}}. \quad (4.4)$$

On an atom chip however, there are other constraints on the trapping potential. The first of these is the trap position. If the trap centre does not coincide with the centre of the atom cloud, then the atom cloud will slosh around in the magnetic trap, causing heating. The second constraint is trap depth. To keep as many of the atoms as possible, the trap must be as deep as possible. In addition, the range of trap parameters that we can choose from is fundamentally constrained by the maximum currents that the wires can safely carry.

To see how these factors constrain how well the trap can be mode-matched, consider the radial case. The maximum safe guide current is 8.0 A and the height of the cloud above the wire centre after optical pumping is ~ 1.5 mm. These two constraints set the bias field B_{bias} to 1.0 mT. The maximum possible radial depth of the trap is B_{bias} , which means that for the trap and the cloud to coincide, the maximum radial depth is 1.0 mT, or 670 μ K. This is only a few times larger than the temperature of the cloud. In addition, B_0 must be large enough to match the radial frequency to the size of the cloud, which further reduces the depth of the trap.

4.4.2 Mode-matching in practice

In practice, such constraints mean that the “mode-matched” trap must be optimised experimentally. The shape of the potential can then be calculated using the experimentally determined optimum parameters, as shown in figures 4.11 and 4.12. Rather than trying optimising the phase-density, we optimised the number of atoms,




Figure 4.11: Axial section of the “mode-matched” potential at the centre of the trap. The position of each wire is indicated at the bottom. The solid line is with all four end wires on, and the dotted line with just the centre two wires. The current is 15 A in both cases. In each case, the field at the centre of the trap B_0 has been subtracted.

as this is the most important parameter in determining the initial collision rate for evaporative cooling since $\gamma_{el} \propto N^{2/3} PSD^{1/3}$. The parameter that made the biggest difference to the number of atoms was the trap depth, which was made as large as possible. The axial depth was increased by running current through all four wires instead of just the centre two, as shown in figure 4.11. The current in the outer pair is in the opposite direction to the current in the inner pair. As well as reducing the axial field in the centre of the trap, the y component of the magnetic field produced by the outer end wires adds constructively to that produced by the inner wires. This gives a large increase in the axial trap depth. The z -cancellation coils were also used to reduce the field in the centre of the trap B_0 and thereby increase the radial trap depth.

To obtain good loading efficiency, it was also important to minimise the centre of mass motion (“sloshing”) of the atom cloud after loading. Sloshing causes heating, and in a shallow trap it can lead to atoms spilling out of the trap. Minimising the sloshing means optimising the relative position of the magnetic trap and the optically pumped cloud. The axial sloshing (along z) was minimised by adjusting the alignment of the MOT beams to shift the CMOT position to the centre of the magnetic trap. This can also be adjusted by applying a small magnetic field in the z direction using the earth’s field cancelling coils. The height of the magnetic trap (y direction) was adjusted by making small changes to the bias field. Due to the geometry of the imaging system, the sloshing in the x direction could not




Figure 4.12: Radial sections of the “mode-matched” potential at the centre of the trap, including gravity. The field at the bottom of the trap $B_0 = 0.37$ mT has been subtracted.

be measured directly; it is along the line of sight. Instead, the number of atoms was measured as the x position of the magnetic trap was adjusted by applying a magnetic field in the y direction. The optimum loading efficiency was obtained when no external y field was applied. The residual sloshing observed was $\sim 50 \mu\text{m}$ (p-p) in the y direction and $\sim 200 \mu\text{m}$ (p-p) in the z direction. In the y direction, this was primarily due to the anisotropic shape of the trapping potential (figure 4.12(b)), and the small amount of kinetic energy that the cloud acquires while falling during optical pumping. The residual motion in the z direction occurs because the cloud is pushed slightly during optical pumping.

The magnetic trap parameters that gave the highest number of atoms are given in table 4.3. The measured size, temperature, and phase space density of the resulting trapped cloud are given in table 4.4.

Table 4.4 shows that despite extensive experimental optimisation, an order of magnitude in number of atoms and phase space density was lost between the compressed MOT and the magnetic trap. The overall transfer efficiency depends not only on how well the magnetic trap is mode-matched, but also on how well the optical pumping works. As described in the previous section, optical pumping was reasonably effective. The main cause of the poor transfer efficiency is the low radial trap depth. In the x direction, the trap is only $470 \mu\text{K}$ deep. Only a small amount of sloshing or breathing in such a shallow potential can cause many of the atoms to be lost. As it is along the line of sight of our imaging system, we have no information

End wire current	I (A)	15
Guide current	I (A)	7.94
Bias field	B_{bias} (mT)	1.0
Radial frequency	ω_r (s^{-1})	$2\pi \times 45$
Axial frequency	ω_z (s^{-1})	$2\pi \times 17$
Residual axial field	B_0 (mT)	0.37
Radial trap depth	(mT)	0.7
Axial trap depth	(mT)	1.1
Radial gradient	(Tm^{-1})	0.8

Table 4.3: Parameters of the “mode-matched” magnetic trap.

Number of atoms	N	1.5×10^7
Radial size	σ_y (μm)	330
Axial size	σ_z (μm)	830
Peak density	n_0 (m^{-3})	1.0×10^{16}
Temperature	T (μK)	83
Phase space density	PSD	0.90×10^{-7}
Elastic collision rate	γ_{el} (s^{-1})	1.0

Table 4.4: Typical parameters of a cloud in the “mode-matched” trap.




Figure 4.13: Timing diagram for the compression phase

about the amount of oscillations in this direction. Most of the atoms are lost in the first 50 ms, corresponding to several periods of the breathing and sloshing excitations, which supports sloshing as the main loss mechanism. At times > 500 ms, the loss rate is completely determined by background gas collisions, and the trap lifetime is > 20 s.

The strategy for improving the loading efficiency is clear - more radial depth. This means increasing the current through the guide wire during the loading process. There is a risk that the increase in current could cause significant heating of the wire, which could lead to the glue holding down the wire failing under the heat and mechanical stress from thermal expansion. However, if the guide current was increased for only a short amount of time during the loading phase, this would keep the heating to a minimum. In this case, the limitation is power dissipation in the FET switch once the current has been reduced to the value for the compressed trap. Despite possible technical problems, increasing the guide current to improve the loading efficiency is worth investigating in the future.

4.5 The compressed trap

To boost the collision rate so that effective evaporative cooling is possible, the trap must be compressed to increase the density. Figure 4.13 shows how the compression was implemented. The trap is compressed in 500 ms, by increasing the bias field B_{bias} and z cancelling field, and decreasing the guide wire current.




Figure 4.14: Measurements of the radial (a) and axial (b) trap frequencies. Solid lines are fits to the measured cloud positions (diamonds).

4.5.1 Properties of the compressed trap

At the end of the compression phase, the trap was located 225 microns above the surface of the guide wire. The parameters of the compressed trap can be calculated from the fields produced by all the wires and the coils. However, uncertainties in the position of the end wires mean that it is difficult to make accurate calculations of the trap frequencies and the magnetic field at the bottom of the trap B_0 . Instead, these parameters were measured experimentally.

The trap frequency can be measured by displacing a cloud from the centre of the trap and measuring the oscillation frequency of the centre of mass of the cloud. These measurements must be performed using small displacements and cold clouds to ensure that the cloud remains in the harmonic region of the potential.

Such measurements are presented in figure 4.14. In both cases, evaporative cooling was used to prepare cold clouds at $\sim 5 \mu\text{K}$. The position of the centre of mass of the cloud is measured using Gaussian fits to the radial and axial profiles. The radial fit is an exponentially damped sine wave with a linearly varying offset. This gives an oscillation frequency of $\omega_r = 2\pi \times 838 \text{ s}^{-1}$. The damping rate is 27 s^{-1} . This is consistent with the expected damping due to elastic collisions. The drift in the offset was due to the settling time of the current control electronics. The axial fit is a simple sine wave, and yields an axial frequency of $\omega_z = 2\pi \times 26 \text{ s}^{-1}$. Another important parameter of the magnetic trap is the residual axial field at the

Radial frequency	ω_r (s^{-1})	$2\pi \times 838$
Axial frequency	ω_z (s^{-1})	$2\pi \times 26$
Residual axial field	B_0 (mT)	0.08
Radial trap depth	(mT)	2.3
Axial trap depth	(mT)	1.9
Radial gradient	(Tm^{-1})	5.9

Table 4.5: Parameters of the compressed magnetic trap

centre of the trap B_0 . This can be estimated from the RF frequency at the end of evaporative cooling to BEC (see section 4.6.1). This gives a value of 0.080(2) mT. By combining the trap frequencies, residual field and trap position, the remaining parameters of the trap such as the depth and linear field gradient can be calculated using the formulae presented in section 2.2. The parameters of the compressed trap are summarised in table 4.5.

4.5.2 The compressed cloud - initial conditions for evaporation.

After compression from the mode-matched trap, the cloud is extremely elongated, and extends beyond the field of view of our imaging system in the axial direction as shown in figure 1.5. The properties of the cloud after 500 ms in the compressed trap are summarised in table 4.6. By reversing the compression and returning the cloud to the mode matched trap, we found that no atoms are lost during the compression. The density was found by fitting Gaussians to the axial and radial distributions.

Measuring the temperature of the compressed cloud accurately is more difficult. Ballistic expansion cannot be employed radially, as the atoms rapidly fall onto the guide wire. As the cloud already extends beyond the field of view of the imaging system, measuring the temperature by releasing the cloud axially and allowing it to expand along the waveguide is very inaccurate. The temperature quoted in table 4.6 was calculated from the width radial profile of the cloud assuming that the trap is harmonic with the trap frequencies measured above. For large, hot clouds such as these, this is only approximate. Only part of the cloud is in the central harmonic region of the trap. Another error arises because the atoms are in a highly non-uniform magnetic field. At a temperature of $690 \mu\text{K}$, the atoms are sampling magnetic fields in excess of 0.1 mT. For atoms in the $F = 2, m_F = 2$ state, this leads to a Zeeman shift of $\sim 2\Gamma$ for the transition to the $F' = 3, m_{F'} = 3$ excited state. This means that at the edges of the cloud, where the field is highest, the atoms will absorb much less light than at the centre, which leads to distortion of the cloud image. This leads to a systematic error in the cloud widths, affecting both

Number of atoms	N	1.5×10^7
Radial size	σ_y (μm)	50
Axial size	σ_z (μm)	1260
Peak density	n_0 (m^{-3})	3.2×10^{17}
Temperature	T (μK)	690
Phase space density	PSD	1.1×10^{-7}
Elastic collision rate	γ_{el} (s^{-1})	54

Table 4.6: Typical parameters for a cloud in the compressed trap.

the density and the temperature measurements.

The magnetic trap is conservative and so the optimum compression is adiabatic, i.e. the initial phase space density is preserved. Compressing adiabatically also gives the highest possible collision rate for a given compression ratio. The criterion for adiabaticity is that the change in the trap frequency should be slow compared to the trap frequency:

$$\frac{d\omega_j}{dt} \ll \omega_j^2 \quad (4.5)$$

While adiabaticity is very desirable, the need for a slow compression must be balanced against the finite lifetime of the magnetic trap. Examination of the trap frequencies for the uncompressed and compressed trap (tables 4.3 and 4.5) shows that by compressing in 500 ms we are well within the adiabatic regime axially, but close to the limit radially.

The phase space density quoted in table 4.6 is close to that measured in the uncompressed trap³ which suggests that the compression is indeed adiabatic. To confirm whether the measurements in table 4.6, are consistent with an adiabatic compression, we can calculate the theoretical temperature rise in the case of a perfectly adiabatic compression. In the magnetic trap, the volume of phase space occupied is proportional to

$$\underbrace{\left(\frac{T^{3/2}}{\bar{\omega}^3}\right)}_{\text{spatial}} \times \underbrace{(T^{3/2})}_{\text{momentum}} = \left(\frac{T}{\bar{\omega}}\right)^3 \quad (4.6)$$

where $\bar{\omega}$ is the geometric mean trap frequency. Adiabatic processes conserve phase space volume, and so the ratio of initial and final temperatures is given by

$$\frac{T_f}{T_i} = \frac{\bar{\omega}_f}{\bar{\omega}_i} \quad (4.7)$$

³The phase space density quoted for the compressed cloud (table 4.6), is actually slightly *higher* than that for the uncompressed cloud (table 4.4). This discrepancy is well within what can be accounted for by the statistical and systematic errors, and does not provide evidence for an increase in phase space density on compression.




Figure 4.15: The decay of the compressed cloud. Points show the experimental data. The error bars describe a 5% fluctuation in the initial atom number. The solid line is a fit to an exponential decay with lifetimes of 6.8s and 16.7s at short and long times respectively.

For our compression, equation 4.7 predicts that the temperature should increase by a factor of 8, giving a predicted final temperature of $660 \mu\text{K}$. This is close to the temperature calculated from the width of the cloud, which suggests that our compression is close to adiabatic.

In the mode-matched trap, the collision rate is very low, and so the compression is not slow compared to the collision rate. The cloud parameters presented in table 4.6 were measured 500 ms after the end of the compression phase, when rethermalisation has taken place. If the cloud is imaged immediately after the compression, it is much wider radially and shorter axially, which reflects the greater heating in the radial direction due to the anisotropic compression.

During the compression, the collision rate has increased by a factor of 90, reaching 54s^{-1} . The decay of this cloud in the compressed trap is shown in figure 4.15. The faster decay at short times is probably due to residual background rubidium from the dispenser pulse. Using this initial lifetime, our value for $\gamma_{el}\tau_{loss}$ is therefore ~ 370 , which is within the regime where runaway evaporation can be achieved. As the evaporation progresses, the decay rate drops, which will improve the efficiency. The lifetime at long times is shorter than the 30 s observed in the “mode-matched trap”. However, for clouds that have been evaporatively cooled to a few microKelvin, the lifetime in the compressed trap can be as long as 30 s, which indicates that there is another loss mechanism other than background collisions operating for these hot clouds.

4.6 Evaporative cooling

The next step along the path to BEC is to apply forced RF evaporative cooling to the compressed atom cloud.

The RF antenna is a 5 cm diameter coil made from 2 turns of copper wire, and is taped to the outside of the top viewport. This puts the antenna as close as possible to the atoms while remaining outside the vacuum. The glass viewport reduces the effect of RF shielding by the steel chamber. In this orientation, the magnetic field produced by the antenna is perpendicular to the trap axis, so that the required $\Delta m_F = \pm 1$ transitions can be driven. A $2.2\ \Omega$ resistor in series with the coil allows the current in the antenna to be measured. The antenna is connected to the RF generator by a short coaxial cable.

The RF ramp is generated by a Stanford Research Systems DS345 30 MHz synthesised function generator. With the output of the generator set to its maximum of 10 V pk-pk, the current in the antenna at 13 MHz is 117 mA RMS. This corresponds to a magnetic field of 1×10^{-6} T at the centre of the trap, giving a Rabi frequency Ω of $2\pi \times 6$ kHz.

For these experiments, the function generator was used in sweep mode. When the generator is triggered, the output frequency sweeps between a specified start and end frequency. The frequency sweep we used was logarithmic, and is specified by three parameters; the start frequency, the end frequency, and the duration of the sweep, which are programmed via the front panel. At the end of the sweep, the output frequency jumps back to the start frequency. In this mode, the output of the generator is continuously on, and so an external RF relay is used to switch on and off the antenna.

4.6.1 Looking for BEC

The RF ramp was optimised experimentally, by measuring the final phase-space density and collision rate as the parameters were adjusted. Unfortunately, the value of the initial truncation parameter is set by the finite trap depth. This corresponds to a truncation parameter $\eta \approx 2$ (much lower than the optimum of 8). The RF frequency corresponding to the trap depth is 13 MHz, and so this is the start frequency for the ramp. From figure 2.10 this means that we do not start in the runaway regime, and therefore the ramp shape must be such that the truncation parameter increases at first, taking us into the runaway regime.

The stop frequency and ramp duration were chosen by optimising the phase-space density at the end of the ramp. By measuring the cloud size and the number of atoms we calculated the density, temperature and phase-space density. The number




Figure 4.16: The correction factor for the absorption cross-section when imaging in the compressed trap, for clouds at $7 \mu\text{K}$. At the usual imaging beam detuning of $+1.5 \text{ MHz}$, the correction factor is $1.6(1)$.

was measured by summing over the whole cloud and using formula 3.11, with the corrected cross-section for in-trap images σ'_L . The correction factor σ_L / σ'_L for cold clouds in the magnetic trap as a function of detuning is shown in figure 4.16. The imaging beam detuning was normally set to $+1.5 \text{ MHz}$; here the correction factor is $1.6(1)$. The cloud size is obtained by Gaussian fits to the radial and axial profiles of the cloud. Using these parameters, the collision rate could also be monitored. A reduction in the collision rate shows that the trajectory is no longer in the runaway regime.

Once the ramp parameters were producing cold clouds ($< 1 \mu\text{K}$) and the evaporation was clearly running away, we searched for evidence of Bose-Einstein condensation. At the end of the ramp, the axial confinement was removed and the cloud was allowed to expand along the waveguide for 8 ms . The signature for BEC was the appearance of a bimodal density distribution as the end frequency (and hence final temperature) is lowered. The experimental signature of a BEC forming in our experiment is shown in figure 4.17. The ramp used to reach BEC had a duration of 12.5 seconds , and an end frequency of $580 - 610 \text{ kHz}$. The number of atoms in the cloud when the BEC started to form at 600 kHz was $\sim 3.3 \times 10^4$.

These parameters can be contrasted with the predictions of the constant truncation parameter model presented in section 2.3. This model predicts that with a starting truncation parameter of 8, BEC would be reached in roughly the same time, but with 2.5×10^5 atoms. Clearly, our evaporative cooling ramp is not optimal, mainly because of the low value of the truncation parameter (due to trap depth) at the beginning of the ramp. A more flexible way of controlling the RF




Figure 4.17: The signature for Bose-Einstein condensation on our atom chip. As the RF frequency is lowered and the cloud cools below the BEC transition, a sharp peak appears in the absorption image of the cloud. The condensate begins to form at 600 kHz.

generator would allow us to have a ramp shape that is more complex than a single exponential. By optimising the trajectory in stages, a significant improvement in atom number may be possible. Nevertheless, it should be remarked that reaching BEC with 0.2% of the atoms that we start with compares well with the other atom chip BEC experiments [27, 28].

4.7 Evaluation

The path to BEC described here was a success - we gain 17 orders of magnitude in phase space density in just 35 seconds. The experiment has produced one of the first BECs on an atom chip in the world. We cross the threshold for BEC with 3×10^4 atoms remaining in the trap. However, it is clear from the discussion above that there remains room for optimisation.

The weakest point is the transfer of atoms from the compressed MOT to the initial magnetic trap. A gain of a factor of two in atom number here could make a large difference to the number of atoms remaining in the condensate. In the future, improving this step by using a deeper trap and compressing the MOT further would be worthwhile. The problem of a low trap depth should also be kept in mind for the design of future atom chips. Another area where there remains room for optimisation is in the RF trajectory. A more sophisticated experiment control system would allow direct control of the RF generator in real time, so that the trajectory would no longer be constrained to a single exponential.

Chapter 5

Properties of the BEC

In this chapter, images of the condensate forming such as those in figure 4.17 are analysed in detail. I will show that these partially condensed clouds are well described by a purely thermal uncondensed component and a Thomas-Fermi BEC component that do not interact with each other. By comparing the model and the experimental data, the static properties of the BEC in the magnetic trap are extracted.

5.1 A preliminary analysis

The experiments on condensate formation analysed here were performed as follows. At the end of the evaporative cooling ramp, the cloud was held in the magnetic trap for a further 30 ms. The axial confinement was then removed by turning off both the end wires and the z -cancelling coils. On its own, this causes the axial field B_0 to vanish as well, changing the radial shape of the potential. To avoid this, the z -earth coils were simultaneously switched on to apply a uniform axial field closely matching B_0 . This kept the radial frequency approximately the same during the expansion and helped reduce heating. The cloud expanded along the guide, and was imaged 8 ms after the fields were switched.

The images were processed as described in section 3.6.2. Each image is normalised using an area far from the atom cloud to take account of shot-to-shot variations in intensity and divided by a normalised background image. The logarithm is then taken. The data analysed here consists of a set of seven such images at end frequencies from 610 kHz to 580 kHz, which were all taken on the same day under the same experimental conditions. An example is shown in figure 5.1.

The image shows that the cloud is only a few pixels wide in the radial direction. A




Figure 5.1: Image of a partially condensed cloud at $f_{end} = 585$ kHz, after 8 ms of axial expansion in the guide. False colour indicates the amount of absorption, varying from blue to white for areas of low to high absorption respectively. A black line indicates the position of the surface.

typical radial cross-section is shown in figure 5.11, where the width can be compared to the pixel size. For all the images of clouds close to BEC, the radial profile can be fitted by a Gaussian with a width σ of 1.25(3) pixels (6.9(2) microns). This width is reached at a temperature of $\sim 10 \mu\text{K}$, below which it does not change. In this regime, the radial width is determined by the imaging process, and contains little information on the shape of the cloud. The origin of this resolution limit is discussed in more detail in section 5.3.

Axially, the cloud is much larger. Figure 5.2 shows how the axial profile of the cloud changes as the final RF frequency is lowered. Each profile is taken from a single image, and the profiles are obtained by averaging over three pixels around the cloud centre in the radial direction. This helps reduce the noise, but makes no significant difference to the shape of the profile. The change in the profile shape as the frequency is lowered is striking. A sharp central feature corresponding to BEC starts to appear at 595 kHz; at an end frequency of 580 kHz it completely dominates the cloud shape.

In the rest of this chapter, these experimental axial profiles are compared with theory. This requires a theory for the density distribution of the cloud, and a theory for how the corresponding absorption image is obtained. The density distribution is treated in section 5.2 and the imaging in 5.3. The result is theoretical axial profiles that are fitted to the experimental data with just two free parameters: the temperature and the amount by which the cloud is blurred due to recoil from spontaneous emission during imaging. Theory and experiment are compared in section 5.5.




Figure 5.2: Axial profiles after 8ms of expansion.

5.2 The density distribution

5.2.1 The cloud in the trap

The density distribution of a thermal cloud and a pure BEC in a harmonic trap were both considered in Chapter 2. The thermal cloud was shown to have a Gaussian spatial distribution, while the ground-state of the condensate was shown to be the parabolic Thomas-Fermi distribution. To analyse the data on condensate formation, the discussion must be extended to include the density distribution of partially condensed clouds.

The simplest approximation is to consider the cloud as two non-interacting parts; a thermal component described by Boltzmann statistics with a Gaussian spatial distribution, and a condensate component, described by the Thomas-Fermi distribution. In this case the combined density distribution can be written as:

$$n(\mathbf{r}) = (1 - f_c)n_{th}(\mathbf{r}) + f_cn_{TF}(\mathbf{r}) \quad (5.1)$$

where f_c is the condensed fraction.

For comparison with the images, it is the column density distribution n_{2D} that we are interested in (see section 5.3.1). This is obtained from equation 5.1 by integrating in the x direction. The resulting column density distribution can be written as

$$n_{2D}(y, z) = N \left\{ (1 - f_c) \left[\frac{1}{2\pi\sigma_y\sigma_z} \exp\left(-\left(\frac{y^2}{2\sigma_y^2} + \frac{z^2}{2\sigma_z^2}\right)\right) \right] + f_c \left[\frac{5}{2\pi R_y R_z} \max\left(1 - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2}, 0\right)^{3/2} \right] \right\}. \quad (5.2)$$

By combining equations 2.56 and 2.57 the condensate fraction f_c can be written as

$$f_c(\omega_r, \omega_z, N, T) = \frac{N_0}{N} = 1 - \left(\frac{k_B T}{0.94\hbar\bar{\omega}}\right)^3 \frac{T^3}{N}. \quad (5.3)$$

The widths of the thermal cloud σ_j depend on the temperature and the trap frequencies according to equation 2.43:

$$\sigma_j(\omega_j, T) = \sqrt{\frac{k_B T}{m\omega_j^2}} \quad (5.4)$$

Finally, by combining equations 2.65 and 2.64 and noting that the number of atoms in the condensate $N_0 = f_c N$ the Thomas-Fermi radii R_j of the cloud can be written as:

$$R_j(\omega_j, N, T) = \frac{1}{\omega_j} \sqrt{\frac{\hbar\bar{\omega}}{m}} \left(15f_c(T, N_{Tc}) N a \sqrt{\frac{m\bar{\omega}}{\hbar}} \right)^{1/5} \quad (5.5)$$

BEC	λ_r	0.87
	λ_z	1.73
Thermal	$\lambda_{th,r}$	1
	$\lambda_{th,z}$	1.65

Table 5.1: Scaling factors at $t = 8$ ms

From equations 5.3, 5.4 and 5.5 it follows that the column density distribution 5.2 depends on four experimental parameters: the radial and axial trap frequencies ω_r ($\omega_r = \omega_x = \omega_y$) and ω_z , the number of atoms N and the temperature T . The first three of these are known. Measurements of the axial and radial frequencies of the compressed trap were presented in section 4.5.1. The number of atoms in each cloud is obtained by summing over the image of the cloud and using equation 3.11 with the absorption cross-section for cold clouds in the magnetic trap σ'_L measured in section 4.6.1. The temperature cannot be obtained from the images without fitting something to the data. Therefore it is treated as a fitting parameter for each cloud.

5.2.2 Axial expansion

The axial expansion of the cloud along the guide must also be taken into account. As shown in sections 2.4.3 and 2.5.4, the expansion of both the thermal and BEC components can be described by a scaling factor $\lambda(t)$.

To obtain the scaling parameters for the condensate, we solve the nonlinear differential equations 2.73 that describe the evolution in the Thomas-Fermi limit with the initial conditions $\beta_r = 1$, $\beta_z = 0$ at $t = 0$, which correspond to turning off the axial confinement and leaving the radial frequency unchanged. The axial scaling factor for the thermal component $\lambda_{th,z}$ is given by equation 2.46. The radial width does not change, as the radial frequency remains constant. The scaling factors are plotted versus time in figure 5.3, and their values at $t = 8$ ms are given table 5.1. For both the condensate and the thermal component, the scaling factors depend only on the trap frequencies, and so no additional parameters are introduced. The column density distribution after expansion is obtained by using the scaled widths in equation 5.2.

5.3 Imaging

Now that we have a model for the density distribution of the partially condensed clouds, the process of image formation must be considered before we compare theory with experiment. The imaging process is modelled in two steps. The interaction




Figure 5.3: The scaling factors for axial expansion from our trap. The solid lines show the axial and radial Castin-Dum scaling factors for the BEC. The dashed line shows the axial scaling factor for the thermal component.

of the atoms with the probe light is considered first. This determines the spatial intensity distribution of the light after it has propagated through the atom cloud. The second step is to consider the resolution of the imaging system.

5.3.1 The atom-light interaction

In section 3.6.2, a simple model was introduced for the interaction of the atoms with the light during the probe pulse. The interaction is represented by a constant cross-section σ'_L for light absorption by atoms in the magnetic trap. The result is that the light is absorbed exponentially in the cloud, with the amount of absorption depending on the optical density $D(y, z)$:

$$D(y, z) = \sigma'_L n_{2D} \quad (5.6)$$

A measurement of the absorption cross-section σ'_L for cold atoms in the magnetic trap was presented in section 4.6.1. The optical density of the $7\ \mu\text{K}$ clouds is ~ 3 , and so the low-density theory applies, allowing us to relate the number of atoms to the cross-section according to equation 3.11. The resulting cross-section was $\sigma'_L = 8.2 \times 10^{-14}\ \text{m}^2$.

5.3.2 Recoil blurring

During the imaging process, the recoil from spontaneous emission will cause the atom cloud to be heated, blurring out the original density distribution. This has two effects. The first is to reduce the spatial resolution of the image. The second is to reduce the density, and therefore the optical density, of the cloud.

If an atom scatters N_p photons in the duration of the imaging pulse Δt , then recoil from the spontaneous emission leads to a random, approximately isotropic velocity of $v_{rms} = \sqrt{N_p} v_{rec}$, where v_{rec} is the recoil velocity [33]. This random velocity leads to a random position at the end of the pulse given by

$$r_{rms} = \sqrt{\frac{N_p}{3}} v_{rec} \Delta t \quad (5.7)$$

For the parameters of our imaging pulse ($N_p = 195$, $\Delta t = 30\ \mu\text{s}$), $r_{rms} = 1.41\ \mu\text{m}$. As well as heating during the probe pulse, the cloud will also be heated by light that leaks through the imaging AOM in the time between the opening of the shutter and the firing of the imaging pulse. This is important because although the AOM attenuates the light by a factor of $\sim 10^{-4}$, the atoms are exposed for much longer ($\sim 1\ \text{ms}$) than during the pulse. The resulting blurring is not known because the exact amount of leakage and the exposure time were not measured.

To take into account the effect of recoil heating, the column density distribution calculated in section 5.2 is convolved with a two-dimensional Gaussian blurring function with single width s_{rec} . Rather than trying to calculate s_{rec} , I treat it as a fit parameter for the axial profiles. Its value is the same for each profile. An additional check is performed by comparing the theoretical and experimental *radial* profile. As the cloud itself is very narrow radially, its apparent radial width in the image is set by the amount of recoil blurring and the optical resolution. The optical resolution is known as discussed below, and so the comparison of theory and experiment in the radial direction provides a good check on the amount of recoil blurring.

The most important consequence of the recoil heating is a big reduction in the optical density of the cloud, as shown in figure 5.4. This is because radially, the amount of blurring is much larger than the initial cloud size ($\sim 0.8 \mu\text{m}$). The shape also changes slightly; the wings become more important compared to the peak. Figure 5.5 shows that without the recoil blurring, the amount of absorption seen on the camera would also be much lower. This is because the cloud diameter is much less than 1 pixel before blurring, which allows most of the light to go by the cloud, even on the central pixel. Figure 5.5 also shows that without the recoil blurring, the spatial information about the condensate would be washed out. This occurs because for very high optical densities, the transmitted intensity is close to zero and does not change very much.

Once the optical density $D(x, y)$ of the blurred cloud has been calculated, the transmission $T(y, z)$ in the object plane is obtained:

$$T(y, z) = e^{-D(y, z)} \quad (5.8)$$

The next step is to calculate the corresponding intensity distribution in the image plane at the camera.

5.3.3 The optical resolution

The process of image formation is most easily considered using Fourier methods. It can be shown that for the case of an extended object illuminated with coherent light, such as our atom cloud, the image $U_1(x_1, y_1)$ of an intensity distribution in the object plane $U_0(x_0, y_0)$ is given by the convolution of U_0 with the transmission function $K(x_0, y_0)$ of the imaging system [84]:

$$U_1(x_1, y_1) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} U_0(x_0, y_0) K(x_1 - x_0, y_1 - y_0) dx_0 dy_0 \quad (5.9)$$




Figure 5.4: An axial profile of the optical density before (dashed) and after (solid) the recoil blurring is taken into account. The amount of motional blurring used here is $s_{rec} = 3.2 \mu\text{m}$, which is the value obtained from the fit to the axial profiles.




Figure 5.5: The cloud profile obtained from the model with (solid) and without (dashed) recoil blurring. The optical resolution has been included in both cases. The transmitted intensity is higher *without* the recoil blurring because more light gets around the narrower cloud.




Figure 5.6: The diffraction limited Airy transmission function (solid) and the Gaussian resolution function from the fit shown in fig. 5.7 (dashed).




Figure 5.7: Section through a raw image, showing the edge of the shadow of the wire. Points are the counts on each pixel, and the solid line is an error function fit.

In a “diffraction-limited system”, $K(x_0, y_0)$ is simply the diffraction pattern of the aperture of the imaging system. For a circular aperture this is an Airy pattern centred on (x_0, y_0) . The Airy pattern for our imaging system is shown in figure 5.6.

Although diffraction sets the ultimate limit on the optical resolution of our imaging system, aberrations from the viewports and slight errors in focussing the imaging system mean that the experimental optical resolution is lower. The experimental optical resolution can be estimated using the blurring of the edge of the guide wire. To a good approximation, the atoms are trapped directly above the centre of the wire, the edge of which therefore is also in the object plane of the imaging system. Figure 5.7 shows a section through an unprocessed image of the edge of the wire. In principle, under the assumption of a perfectly sharp edge, deconvolution would yield the exact experimental transmission function. However, the information is spread over a small number of pixels, and the pixellation makes deconvolution unreliable. I have chosen to use a simpler approach. Here, the transmission function is represented by a simple Gaussian “blurring” function

$$K(x - x_0, y - y_0) = \frac{1}{2\pi s_{opt}^2} \exp \left[\frac{-((x - x_0)^2 + (y - y_0)^2)}{2s_{opt}^2} \right]. \quad (5.10)$$

The width s_{opt} is obtained by fitting an error function to the image of the edge of

the wire as shown in figure 5.7, which gives $s_{opt} = 4.1 \mu\text{m}$. This blurring function is compared with the diffraction limited Airy pattern for our lens in figure 5.6. Comparison of the two suggests that the experimental resolution is only slightly worse than the diffraction limit.

The intensity in the image plane is obtained by convolving $T(x, y)$ with this Gaussian blur function. To obtain the axial profiles corresponding to those in figure 5.2 it remains to take the log of the intensity distribution and integrate over a width corresponding to 3 pixels in the y direction.

5.4 Calculating the theoretical axial profiles

To summarise the above, it is useful to describe step-by-step how the calculations of the theoretical profiles were performed. All the calculations were performed using the *Mathematica* package from Wolfram Research. In principle each step can be performed analytically. However the final convolution stage involves integrals of the form $\int \exp(-\exp(-x^2))$, which *Mathematica* was unable to evaluate analytically. Therefore I have used a partly numerical model. The calculations were performed as follows:

- The expression for the column density (equation 5.2) is evaluated on a two-dimensional grid of points $20 \mu\text{m}$ wide in the y direction and $600 \mu\text{m}$ wide in the z direction, for a particular T . The resolution of the grid is $0.1 \mu\text{m}$ in the y -direction, and $2 \mu\text{m}$ in the z direction.
- The recoil blurring function is evaluated on the same grid and numerically convolved with the column density distribution.
- To obtain the transmitted intensity, each point is multiplied by the cross-section and the exponential is calculated.
- This transmitted intensity is numerically convolved with the optical resolution function.
- The axial profile corresponding to those in figure 5.2 is calculated by taking the natural log and integrating over the size of three pixels in the object plane ($16.7 \mu\text{m}$) around the centre of the array in the y direction.

The slowest step in this procedure is the evaluation of the column density on the grid, which takes ~ 60 s on a fast PC. As this must be repeated each time the temperature is changed, least-squares fitting of the temperature and s_{rec} is

impractical. Instead, the fitting is performed “by eye”. These two parameters are adjusted until reasonable agreement is obtained for each image, with s_{rec} constrained to be the same for each image.

5.5 Comparison with the data

The experimental and theoretical axial profiles are compared in figures 5.8, 5.9 and 5.10. The agreement is good in all cases, from the purely thermal cloud at 610 kHz to clouds with a very high condensate fraction at lower frequencies (583 kHz, 580 kHz). Both the amount of absorption and the cloud shapes are well reproduced. The pattern of residuals in each of the fits shows no significant trend with decreasing RF frequency and temperature. Such a trend could indicate a problem with the way that the imaging process has been modelled, or a systematic problem with a parameter of the density distribution, such as a systematic error in the number of atoms.

The fits to the axial profiles give a value of $s_{rec} = 3.2 \mu\text{m}$ for the amount of recoil blurring. This value can be cross-checked by comparing the experimental and theoretical radial profiles. As shown in figure 5.11, the radial shape of the cloud is also well reproduced. This shows that the model is consistent, and that the axial fits have provided a good estimate for the amount of recoil blurring.

Figure 5.8(a) shows that at 610 kHz, we have a purely thermal cloud. As the end frequency of the RF ramp is lowered, the cloud cools. The cloud at 600 kHz (figure 5.8(b)) is close to the BEC transition. The axial profile is best fit by the Thomas-Fermi plus thermal theory with a temperature of 381 nK. The dashed line shows a purely thermal cloud at the same temperature. The agreement is only slightly worse, showing that the cloud is very near the temperature at which the condensate starts to form. As the temperature is lowered further, the condensate fraction increases and the narrow central peak in the cloud profile becomes clear (figures 5.8(c) and 5.8(d)). In the plot of the profile at 585 kHz (figure 5.9(a)), the Thomas-Fermi plus thermal cloud model is compared with a purely thermal cloud at the same temperature. The purely thermal model does not reproduce the amount of absorption or the shape of the cloud. Finally, in figures 5.9(b) and 5.10 almost two thirds of the atoms are in the condensate, and the axial profile is dominated by the condensate peak.

Figure 5.10 gives an indication of how accurately we can determine the temperature by comparing theory and experiment “by eye”. The black line in the figure is the best fit, which gives a temperature of 215 nK. The blue and red curves are

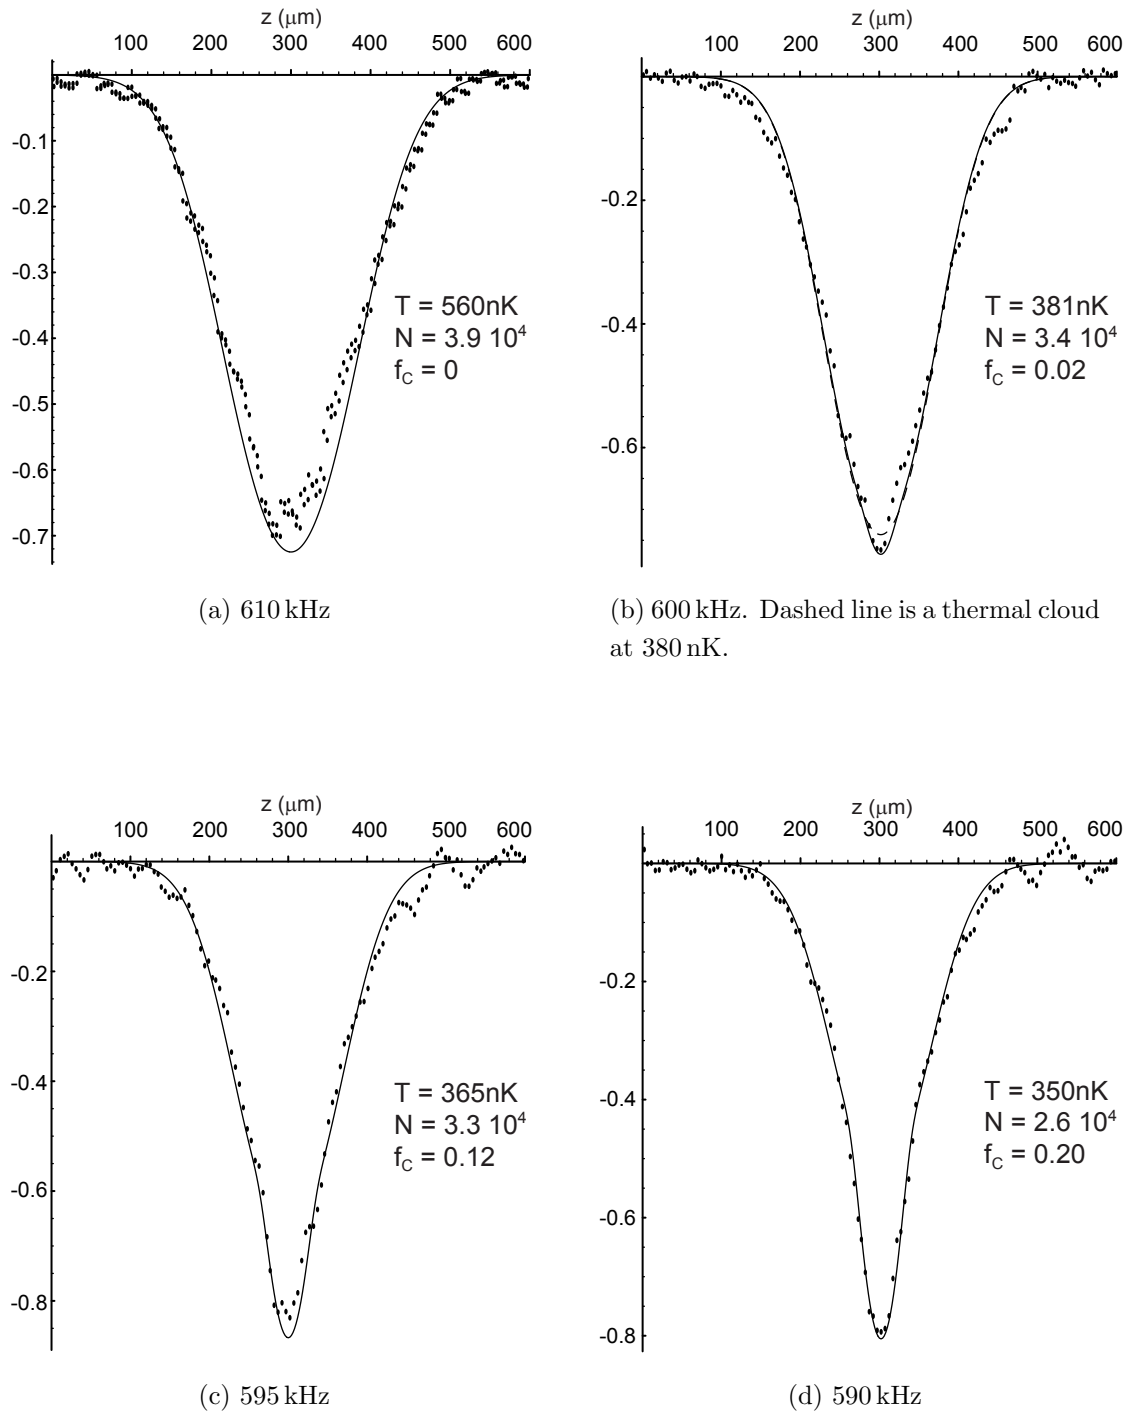


Figure 5.8: Comparison of model (solid line) and experimental (dots) axial profiles. Each plot is labelled with the number of atoms in the whole cloud N , the temperature T and the condensed fraction f_c . The only fit parameters were the temperature and the amount of recoil blurring.




Figure 5.9: Comparison of model (solid line) and experimental (dots) axial profiles at various final RF frequencies. Each plot is labelled with the number of atoms in the whole cloud N , the temperature T and the condensate fraction f .




Figure 5.10: Comparison of the data and the model at 580 kHz for different temperatures. The black line is a “best fit” by eye, which gives a temperature of 215 nK. The blue and red curves are at 190 nK and 240 nK. Both are in significantly worse agreement with the data, which shows that by eye it is straightforward to determine T to $\sim 10\%$.




Figure 5.11: Comparison of experimental (bars) and theoretical (solid line) radial profiles at 590 kHz.

theoretical axial profiles for clouds that are 25 nK colder and hotter respectively. They do not fit the data as well. The blue curve at 190 nK fits the condensate peak very well, but it cannot reproduce the shape of the wings. In contrast, the red cloud at 240 nK fits reasonably well in the wings, but does not get the height of the condensate peak right. This shows that the temperature can be estimated to at least 10 %.

The agreement of theory and experiment over such a wide range of temperatures supports the use of the two-component, non-interacting model for the density distribution. Recent work studying the quadrupole oscillations of a trapped BEC in the presence of a thermal cloud in a magnetic trap has shown that the BEC and the thermal cloud are coupled via a mean field interaction. Although such an interaction could modify the dynamics of the axial expansion of the cloud, no evidence of that is seen here. This may well be because the expansion time of 8 ms is not long enough to allow the coupling to have any influence. The model fits the width of the condensate peak very well, which means that the axial expansion of the condensate is well described by the Castin-Dum model discussed in section 2.5.4. The experiment demonstrates that the condensate is behaving as a coherent matter wave as it expands into the waveguide.

Number of condensate atoms	N_0	1.0×10^4
Peak density	n_0 (m^{-3})	2.8×10^{20}
Thomas-Fermi radius	$R_{x,y,TF}$ (μm)	0.9
Thomas-Fermi length	$2R_{z,TF}$ (μm)	56
Chemical potential	μ (nK)	110

Table 5.2: Properties of the BEC in the trap, calculated from the fit to the 580 kHz image.

5.6 Summary

In summary, these fits confirm that this sequence of images does indeed correspond to the formation of a condensate on our atom chip. The data is well reproduced by a model for the density distribution based on a Thomas-Fermi condensate fraction and a classical thermal uncondensed fraction that do not interact. In particular, the expansion of the condensate is well described by a model based on the Thomas-Fermi approximation, which demonstrates that the condensate is propagating as a coherent matter wave in our magnetic waveguide. The agreement between theory and experiment is sufficiently good that quantitative information about each cloud can be obtained, in particular the temperature and condensate fraction. This allows the static properties of the condensate in the magnetic trap to be calculated. The values for the cloud at 580 kHz (figure 5.10) are given in table 5.2.

Chapter 6

Conclusions

6.1 Summary

In this thesis I have described the preparation of one of the first Bose-Einstein condensates to be made on an atom chip. The condensate contained 1×10^4 ^{87}Rb atoms, and formed at temperatures below 380 nK. It was prepared just 220 μm from a room temperature (290 K) substrate. By improving the coupling from the MOT to the magnetic trap and optimising the evaporative cooling RF trajectory, it should be possible to increase the number of atoms in the condensate by a factor of 10.

As a first demonstration of the potential for manipulating condensates with atom chips, the cloud was allowed to expand axially along a magnetic waveguide before it was imaged. In chapter 5, these images were compared to a theoretical model. This showed that the expanding, partially condensed clouds are well described by a theory based on a classical thermal component and a Thomas-Fermi BEC component that do not interact with each other. In particular, these results showed that the BEC propagates along the waveguide as a coherent matter wave. Comparison of theory and experiment also allowed the temperature of the partially condensed clouds to be accurately determined, as well as the static properties of the BEC in the magnetic trap.

6.2 Future directions

The combination of Bose-Einstein condensates and atom chips opens up a huge range of possibilities. Here I will briefly describe the latest work we are doing, and the directions that the experiment may take in the future.

Guiding and interferometry Perhaps the most obvious future experiment is to

study the propagation of the condensate along the magnetic waveguide in more detail. The propagation of a BEC in a magnetic waveguide above an atom chip was briefly investigated in [30], where they saw single-mode propagation except where deformations of their micro-fabricated wires caused excitations of the condensate. In recent experiments [85] the Tübingen group have shown that the propagation of their condensate in a magnetic waveguide can be reasonably well described by the Castin-Dum scaling factor model even after ~ 100 ms of expansion. In this thesis, I have studied the expansion of a partially condensed cloud rather than a pure condensate. It would be interesting to see whether after longer expansion times we observe any disagreement with the predictions of the Castin-Dum theory due to interactions between the thermal cloud and the condensate.

A major motivation for atom chip BEC experiments is the prospect of building chip-based atom interferometers. This is difficult with the single-wire geometry of our current atom chip. A new atom chip based on a multiple-wire “Z” trap is currently being fabricated in collaboration with the microelectronics group at Southampton University. This new atom chip will use lithographically patterned gold wires, and will allow us to realise the the kind of on-chip interferometer proposed by our group in [24].

Atom-surface interactions A unique feature of atom-chip BEC experiments is that ultra-cold ($< 1 \times 10^{-6}$ K) clouds of atoms are trapped very close to a room temperature (~ 290 K) surface. Interactions between the trapped atoms and the surface are the subject of much current interest. The Tübingen group observed two effects [86]. Firstly, when an ultra-cold cloud was brought close ($< 70 \mu\text{m}$) to the surface of the conductor, the cloud was observed to fragment due to corrugations of the potential. The same effect was observed above the much larger conductor used in our experiments. Figure 6.1 shows a fragmented thermal cloud on our atom chip. The cloud has a temperature of $\sim 4 \mu\text{K}$, and is about $15 \mu\text{m}$ above the surface of the wire. Fragmentation was also observed by the MIT group [30]. Recently, the Tübingen group have shown that the corrugation of the potential is due to spatially varying axial component of the magnetic field produced by the current in the wire [87]. As yet the origin of this modulated axial field remains unknown.

The second effect is that the lifetime decreases as the atoms are brought closer to the surface. The dependence of lifetime on height is still an open question. In [86] the Tübingen group observed a linear dependence on height above




Figure 6.1: Fragmentation of a thermal cloud near the surface of the guide wire. The atoms are at $\sim 4 \mu\text{K}$ and are about $15 \mu\text{m}$ above the surface, which is indicated by the white line.

the *centre* of the conductor, whereas recently the MIT group [88] report no height dependence at all. Both these measurements disagree with theoretical predictions based on losses due to spin-flips caused by thermally induced currents in the surface [89]. We are currently performing a detailed study of the height dependence of the lifetime of ultra-cold clouds on our chip, including a comparison with this theory.

Light on the atom chip In the longer term, we hope to incorporate optical elements on to the atom chip using optical fibres. The combination of atom chips with micro-optics promises to be very powerful. For example, optical fibres with small lenses formed on the end can make waist sizes of $\sim 1 \mu\text{m}$. Such fibre lenses could be used for detecting a single atom [90]. They can also produce microscopic optical dipole traps which can be loaded from the magnetically trapped cloud.

The ultimate goal is to use the combination of atom chips and micro-optics to do quantum information processing. In a very beautiful recent experiment [91], I. Bloch and co-workers at the MPQ in Garching, Germany have shown that by putting a BEC into an optical lattice, it is possible to go from the phase coherent superfluid state to a “Mott insulator” state, where the atoms are completely localised in individual wells of the lattice. The result is a grid of atoms, one per lattice site, all of which are in the ground state of their local potential well. By making the lattice state-dependent, they have demonstrated that the atoms can be made to interact coherently with their nearest neighbours by collisions. Using optical fibres, such an optical lattice could be integrated onto an atom chip. We are also investigating making microscopic optical cavities so that individual atoms can be coherently coupled to a light field.

In short, the combination of Bose-Einstein condensates and atom chips holds great promise for the future.

Bibliography

- [1] J. A. Retter *Cold Atom Microtraps above a Videotape Surface*, D. Phil. thesis, University of Sussex, (2002)
- [2] E. A. Hinds and I. G. Hughes, *Magnetic atom optics: guides, traps and chips for atoms*, J. Phys. D, **32**, R119, (1999)
- [3] J. Schmiedmayer and R. Folman, *Miniaturizing atom optics: from wires to atom chips*, C. R. Acad. Sci. Paris, t. 2, Série 4, p551 (2001)
- [4] R. Folman et al., *Controlling cold atoms using nanofabricated surfaces: Atom Chips*, Phys. Rev. Lett., **84**, 4749, (2000)
- [5] D. Müller et al., *Guiding Neutral Atoms around curves with lithographically patterned current-carrying wires*, Phys. Rev. Lett., **83**, 5194, (1999)
- [6] N. H. Dekker et al., *Guiding Neutral Atoms on a Chip*, Phys. Rev. Lett., **84**, 1124, (2000)
- [7] J. Reichel, W. Hänsel and T. W. Hänsch, *Atomic Manipulation with Magnetic Surface Traps*, Phys. Rev. Lett., **83**, 3398, (1999)
- [8] D. Cassettari et al., *Beam splitter for guided atoms*, Phys. Rev. Lett., **85**, 5483, (2000)
- [9] W. Hänsel et al., *Magnetic Conveyor Belt for Transporting and Merging Trapped Atom Clouds*, Phys. Rev. Lett., **86**, 609, 2001
- [10] M. H. Anderson et al., *Observation of Bose-Einstein condensation in a dilute atomic vapour*, Science, **269**, 198, (1995)
- [11] K.B. Davis et al., *Bose-Einstein condensation in a gas of sodium atoms*, Phys. Rev. Lett., **75**, 3969, (1995)
- [12] M. R. Andrews et al., *Observation of interference between two Bose condensates*, Science, **275**, 637, (1997)

- [13] M. Kozuma et al., *Coherent splitting of Bose-Einstein condensed atoms with optically induced Bragg diffraction*, Phys. Rev. Lett., **82**, 871, (1999)
- [14] E. W. Hagley et al., *A well-collimated quasi-continuous atom laser*, Science, **283**, 1706, (1999)
- [15] S. Inouye et al., *Observation of phase-coherent amplification of atomic matter waves*, Nature, **402**, 641-644 (1999).
- [16] M. Kozuma et al., *Phase-Coherent Amplification of Matter Waves*, Science, **286**, 2309, (1999)
- [17] L. Deng et al., *Four-wave mixing with matter waves*, Nature, **398**, 218, (1999)
- [18] Y. Tori et al., *Mach-Zender Bragg interferometer for a Bose-Einstein condensate*, Phys. Rev. A, **61**, 041602(R), (2000)
- [19] S. Gupta et al., *Contrast Interferometry using Bose-Einstein Condensates to Measure h/m and α* , preprint, arXiv:cond-mat/0202452, (2002)
- [20] K. Bongs et al., *Waveguide for Bose-Einstein condensates*, Phys. Rev. A, **63**, 031602(R), (2001)
- [21] K. Bongs et al., *Coherent evolution of bouncing Bose-Einstein condensates*, Phys. Rev. Lett., **83**, 3577, (1999)
- [22] A.S. Arnold., C. MacCormick, and M. G. Boshier, *Adaptive inelastic magnetic mirror for Bose-Einstein condensates*, Phys. Rev. A., **65**, 031601(R), (2002)
- [23] C. MacCormick, *Coherent atom optics with a BEC*, D. Phil. thesis, University of Sussex, 2002
- [24] E. A. Hinds, C. J. Vale, and M. G. Boshier, *Two-Wire waveguide and interferometer for cold atoms*, Phys. Rev. Lett., **86**, 1462 (2001)
- [25] W. Hänsel et al., *Trapped-atom interferometer in a magnetic microtrap*, Phys. Rev. A, **64**, 063607, (2001)
- [26] E. Andersson, *Multimode Interferometer for Guided Matter Waves*, Phys. Rev. Lett., **88**, 100401, (2002)
- [27] H. Ott et al., *Bose-Einstein Condensation in a Surface Micro Trap*, Phys. Rev. Lett., **87**, 230401 (2001)

- [28] W. Hänsel, P. Hommelhoff, T. W. Hänsch and J. Reichel, *Bose-Einstein condensation on a microelectronic chip*, Nature, **413**, 498, (2001)
- [29] The BEC news page at <http://amo.phy.gasou.edu/bec.html/>
- [30] A. E. Leanhardt et al., *Propagation of Bose-Einstein condensates in a magnetic waveguide*, Phys. Rev. Lett., **89**, 040401, (2002)
- [31] J. E. Simsarian et al., *Magneto-Optic Trapping of ^{210}Fr* , Phys. Rev. Lett. **76**, 3522 (1996).
- [32] J. R. Anglin and W. Ketterle, *Bose-Einstein condensation of atomic gases*, Nature, **416**, 211, (2002)
- [33] W. Ketterle, D.S. Durfee and D. M. Stamper-Kurn, *Making, probing and understanding Bose-Einstein condensates*, in *Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics “Enrico Fermi”*, Course CXL, IOS Press, 1999, p96
- [34] M. D. Barrett, J. A. Sauer, and M. S. Chapman, *All-Optical Formation of an Atomic Bose-Einstein Condensate*, Phys. Rev. Lett., **87**, 010404, (2001).
- [35] C.S. Adams and E. Riis, *Laser Cooling and Trapping of Neutral Atoms*, Prog. Quant. Electr., **21**, 1, (1997)
- [36] H.J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Springer-Verlag, 1999
- [37] T. W. Hänsch and A.L. Schawlow, *Cooling of gases by laser radiation*, Opt. Comm., **13**, 68, (1975)
- [38] J. Dalibard et al., in *Proceedings of the 11th International Conference on Atomic Physics*, edited by S. Haroche, J. C. Gay, and G. Grynberg, World Scientific, 1989
- [39] S. Chu et al., in *Proceedings of the 11th International Conference on Atomic Physics*, edited by S. Haroche, J. C. Gay, and G. Grynberg, World Scientific, 1989
- [40] J. Dalibard and C. Cohen-Tannoudji, *Laser cooling below the Doppler limit by polarisation gradients: simple theoretical models*, J. Opt. Soc. Am. B, **6**, 2023, (1989)

- [41] E. L. Raab et al., *Trapping of neutral sodium atoms with radiation pressure*, Phys. Rev. Lett., **59**, (1987)
- [42] G. K. Woodgate, *Elementary Atomic Structure*, OUP, 2nd edn., 1983, pp191-194
- [43] L.I. Schiff, *Quantum Mechanics*, McGraw-Hill, 3rd edn., 1968, p289
- [44] E. Majorana, *Nuovo Cimento*, **9**, 43, (1932)
- [45] E. A. Hinds and C. Eberlein, *Quantum propagation of neutral atoms in a magnetic quadrupole guide*, Phys. Rev. A, **61**, 033614, (2000), and erratum by the same authors Phys. Rev. A, **64**, 039902[E], (2001)
- [46] A. Goepfert et al., *Efficient magnetic guiding and deflection of atomic beams with moderate velocities*, Appl. Phys. B, **69**, 217, (1999)
- [47] J. A. Richmond et al., *Magnetic guide for neutral atoms*, **65**, 033422, (2002)
- [48] J. Denschlag, D. Cassettari and J. Schmiedmayer, *Guiding Neutral Atoms with a Wire*, Phys. Rev. Lett., **82**, 2014, (1999)
- [49] M. Key et al., *Propagation of Cold Atoms along a Miniature Magnetic Guide*, Phys. Rev. Lett., **84**, 1371, 2000
- [50] M. G. Key, *A Miniature Magnetic Waveguide for Cold Atoms*, D. Phil. thesis, University of Sussex, 2000
- [51] J. Fortagh et al., *Miniaturized Wire Trap for Neutral Atoms*, Phys. Rev. Lett., **81**, 5310, (2000)
- [52] D. E. Pritchard, *Cooling neutral atoms in a magnetic trap for precision spectroscopy*, Phys. Rev. Lett., **51**, 1336, (1983)
- [53] O.J. Luiten et al., *Kinetic theory of the evaporative cooling of a trapped gas*, Phys. Rev. A, **53**, 381, (1996)
- [54] K. B. Davis, M. O. Mewes and W. Ketterle, *An analytical model for the evaporative cooling of a trapped gas*, Appl. Phys. B, **60**, 155, (1995)
- [55] A. S. Arnold, *Preparation and Manipulation of an ^{87}Rb Bose-Einstein Condensate*, D. Phil. thesis, University of Sussex, 1999

- [56] D. W. Snoke and J. P. Wolfe, *Population dynamics of a Bose gas near saturation*, Phys. Rev. B, **39**, 4030, (1989)
- [57] J. P. Burke, Jr. et al., *Prospects for Mixed-Isotope Bose-Einstein Condensates in Rubidium*, Phys. Rev. Lett., **80**, 2097, (1998)
- [58] S. L. Cornish et al., *Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions*, Phys. Rev. Lett., **85**, 1795, (2000)
- [59] C.A. Sackett et al., *Optimization of evaporative cooling*, Phys. Rev. A, **55**, 3797, (1997)
- [60] F. Dalfovo et al., *Theory of Bose-Einstein condensation in trapped gases*, Rev. Mod. Phys., **71**, 463, (1999)
- [61] F. Mandl, *Statistical Physics*, Wiley, 2nd edn., 1988
- [62] V. Bagnato, D. E. Pritchard and D. Kleppner, *Bose-Einstein condensation in an external potential*, Phys. Rev. A, **35**, 4354, (1987)
- [63] J. R. Ensher et al., *Bose-Einstein Condensation in a Dilute Gas: Measurement of Energy and Ground-State Occupation*, Phys. Rev. Lett., **77**, 4984, (1996)
- [64] R. Loudon, *The Quantum Theory of Light*, 2nd edn., OUP, 1997, p120-157
- [65] G. Baym and C.J. Pethick, *Ground-State Properties of Magnetically Trapped Bose-Condensed Rubidium Gas*, Phys. Rev. Lett., **76**, 6, (1996)
- [66] Y. Castin and R.Dum, *Bose-Einstein Condensates in Time Dependent Traps*, Phys. Rev. Lett., **77**, 5315, (1996)
- [67] J. Fortagh et al., *Fast loading of a magneto-optical trap from a pulsed thermal source*, J. Appl. Phys., **84**, 6499, (1989)
- [68] Carl E. Wieman and Leo Holberg, *Using diode lasers for atomic physics*, Rev. Sci. Instrum., **62**, 1, (1991)
- [69] K.B. MacAdam, A. Steinbach and C. Wieman, *A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb*, Am. J. Phys., **60**, 1098, (1992)
- [70] A. S. Arnold, J.S. Wilson and M.G. Boshier, *A simple extended cavity diode laser*, Rev. Sci. Instrum., **69**, 1236, (1998)

- [71] W. Demtröder, *Laser Spectroscopy*, Springer-Verlag, 2nd edn., 1993, p 444
- [72] C. Wieman and T.W. Hänsch, *Doppler-Free Laser Polarization Spectroscopy*, Phys. Rev. Lett., **30**, 1170, (1976)
- [73] Jung Bog Kim, Hong Jin Kong, and Sang Soo Lee, *Dye laser frequency locking to the hyperfine structure ($3S_{1/2}, F = 2 \rightarrow 3P_{1/2}, F = 2$) of sodium D_1 line by using polarization spectroscopy*, Appl. Phys. Lett., **52**, 417, (1988)
- [74] G. P. T. Lancaster et al., *A polarisation locked diode laser for trapping cold atoms*, Opt. Comm., **170**, 79, (1999)
- [75] U. Schünemann et al., *Simple scheme for tunable frequency offset locking of two lasers*, Rev. Sci Instrum., **70**, 242, 1999
- [76] Coherent Inc., *Model MBR-110 Operator's Manual*
- [77] C.G. Townsend et al., *Phase-space density in the magneto-optical trap*, Phys. Rev. A, **52**, 1423, (1995)
- [78] K. Lindquist, M. Stephens and C. Wieman, *Experimental and theoretical study of the vapor cell Zeeman optical trap*, Phys. Rev. A, **46**, 4082, (1992)
- [79] W. Petrich et al., *Behaviour of atoms in a compressed magneto-optical trap*, J. Opt. Soc Am. B, **11**, 1332, (1994)
- [80] A.M. Steane, M. Chowdhury and C.J. Foot, *Radiation force in the magneto-optical trap*, J. Opt. Soc. Am. B, **9**, 2142, (1992)
- [81] W. Salzmann, *Atoms in a long magneto-optical trap*, Diplome Thesis, University of Sussex, 2001
- [82] W. Ketterle et al., *High densities of cold atoms in a dark spontaneous force optical trap*, Phys. Rev. Lett., **70**, 2253, (1993)
- [83] C. Townsend et al., *High density trapping of cesium atoms in a dark magneto-optical trap*, Phys. Rev. A, **53**, 1702, (1996)
- [84] M. Born and E. Wolf, *Principles of Optics*, C.U.P., 7th edn, 1999, pp 543-553
- [85] J. Fortagh et al., *Bose-Einstein Condensates in Magnetic Waveguides*, arXiv: cond-mat/0210547, (2002), to be published in Appl. Phys. B

- [86] J. Fortagh et al., *Surface effects in magnetic microtraps*, Phys. Rev. A, **66**, 041604(R), (2002)
- [87] S. Kraft et al., *Anomalous longitudinal magnetic field near the surface of copper conductors*, J. Phys. B, **35**, L469, (2002)
- [88] A. E. Leanhardt et al., *Bose-Einstein condensates near a microfabricated surface*, arXiv:cond-mat/0211345, (2002)
- [89] C. Henkel, S. Pötting and M. Wilkens, *Loss and heating of particles in small and noisy traps*, Appl. Phys. B, **69**, 379, (1999)
- [90] P. Horak et al., *Towards single-atom detection on a chip*, arXiv:quant-ph/0210090
- [91] M. Greiner et al., *Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms*, Nature, **415**, 39, (2002)
- [92] W. Demtröder, *Laser Spectroscopy*, Springer-Verlag, 2nd edn., 1993, pp 454-466
- [93] C. P. Pearman et al., *Polarisation spectroscopy of a closed atomic transition: applications to laser frequency locking*, to be published (2002)

Appendix A

Polarisation spectroscopy lineshape

To see how we obtain a dispersion signal from our polarisation spectrometer, we need to look at the line shape in detail. This analysis follows that in [92] and [93].

If the probe beam propagates along the z axis and is linearly polarised such that its polarisation vector makes an angle θ with the x axis, then the electric field before the vapour cell ($z = 0$) is:

$$\mathbf{E} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} = E_0 e^{i\omega t} \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix} \quad (\text{A.1})$$

This can be written in terms of circular polarisation basis vectors (σ^\pm) as:

$$\mathbf{E} = E_0 e^{i\omega t} \left[\frac{e^{i\theta}}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + \frac{e^{-i\theta}}{2} \begin{pmatrix} 1 \\ -i \end{pmatrix} \right]. \quad (\text{A.2})$$

As described in section 3.5.3, optical pumping by the pump beam causes changes in the absorption coefficient α and refractive index n for the two polarisation components. After a travelling a distance L through the cell, the electric field of the probe beam becomes

$$\mathbf{E} = E_0 e^{i\omega t} \left[e^{-i(\frac{\omega L}{c} n^+ - i\alpha^+ L/2)} \frac{e^{i\theta}}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + e^{-i(\frac{\omega L}{c} n^- - i\alpha^- L/2)} \frac{e^{-i\theta}}{2} \begin{pmatrix} 1 \\ -i \end{pmatrix} \right] \quad (\text{A.3})$$

where n^\pm and α^\pm are the refractive indices and absorption coefficients respectively for the σ^\pm polarisation components. This expression can be rewritten as

$$\mathbf{E} = E_0 e^{i\omega t} e^{-i(\frac{\omega L}{c} n - i\alpha L/2)} \left[e^{i\Delta} \frac{e^{i\theta}}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + e^{-i\Delta} \frac{e^{-i\theta}}{2} \begin{pmatrix} 1 \\ -i \end{pmatrix} \right]. \quad (\text{A.4})$$

Here $n = \frac{1}{2}(n^+ + n^-)$ and $\alpha = \frac{1}{2}(\alpha^+ + \alpha^-)$ are the average refractive index and absorption coefficient. The quantity Δ depends on the differences between the two polarisations $\Delta n = (n^+ - n^-)$ and $\Delta\alpha = (\alpha^+ - \alpha^-)$:

$$\Delta = \frac{\omega L}{2c} \Delta n - i\Delta\alpha/4 \quad (\text{A.5})$$

The polarising beam splitter cube splits the probe beam into horizontal (x) and vertical (y) components, each of which is measured by a photodiode. The intensity on each photodiode $I_{x,y}$ is obtained by multiplying A.4 by its complex conjugate. This gives

$$I_x = \frac{I_0}{2} e^{-\alpha L} \left[\frac{1}{2} (e^{\Delta\alpha/2} + e^{-\Delta\alpha/2}) + \cos \left(2\theta + \frac{\omega L}{c} \Delta n \right) \right] \quad (\text{A.6})$$

$$I_y = \frac{I_0}{2} e^{-\alpha L} \left[\frac{1}{2} (e^{\Delta\alpha/2} + e^{-\Delta\alpha/2}) - \cos \left(2\theta + \frac{\omega L}{c} \Delta n \right) \right] \quad (\text{A.7})$$

The dispersion signal that we lock to is the difference in intensity between the two photodiodes. Subtracting the above two expressions, we find that the contribution to the lineshape due to the difference in the absorption coefficient $\Delta\alpha$ cancels out, and we are left with

$$S = I_x - I_y = I_0 e^{-\alpha L} \cos \left(2\theta + \frac{\omega L}{c} \Delta n \right) \quad (\text{A.8})$$

$$= I_0 e^{-\alpha L} \left[\cos(2\theta) \cos \left(\frac{\omega L}{c} \Delta n \right) - \sin(2\theta) \sin \left(\frac{\omega L}{c} \Delta n \right) \right] \quad (\text{A.9})$$

To make progress, we need to know the spectral profile of Δn as the laser frequency is scanned across the atomic transition. As in saturated absorption spectroscopy, at low intensity, the spectral profile of the difference in absorption $\Delta\alpha$ is Lorentzian

$$\Delta\alpha = \frac{\Delta\alpha_0}{1 + x^2} \quad (\text{A.10})$$

where x is the detuning in linewidths $x = 2(\omega_0 - \omega)/\Gamma$. The absorption coefficient and the refractive index are related by the Kramers - Kronig relation [92]. This yields a dispersive profile for the refractive index:

$$\Delta n = \frac{c}{\omega_0} \frac{\Delta\alpha_0 x}{1 + x^2} . \quad (\text{A.11})$$

To obtain the dispersion curve that we lock the lasers to in the experiment, the angle of the probe beam polarisation is set to $\theta = 45^\circ$. The change in polarisation

due to the vapour is small, and so equation A.8 yields the following for the error signal as the laser is scanned across the transition:

$$S(x) \simeq -I_0 e^{-\alpha L} L \Delta \alpha_0 \frac{x}{1+x^2}. \quad (\text{A.12})$$

For laser stabilisation, this signal has two advantages over the usual polarisation spectroscopy signal, which is obtained using a single photodiode and a small angle θ (see [71, 93]). Unlike the latter scheme, $S(x)$ has no background and no Lorentzian component.

In this treatment, the atomic physics underlying the changes in absorption coefficient and refractive index of the vapour due to the pump beam has been swept into the parameter α_0 - the differential absorption at the line centre. For a particular transition, the amount of anisotropy will obviously depend on the angular momenta of the states involved. In the limit of low intensity and neglecting spontaneous emission, [72] presents a simple formula for the amount of anisotropy in terms of the total angular momentum quantum numbers of the upper and lower states and their lifetimes. In the case of rubidium, this formula predicts that dispersion effects should be observed on both the trap ($F = 2 \rightarrow F' = 3$) and repump ($F = 1 \rightarrow F' = 2$) transitions. Experimentally, we observe a strong signal on the trap transition as expected, but almost nothing on the repump transition (section 3.5.3). This is because the latter case does not form a closed two-level system.