Investigating the Dynamics of a Bose Einstein Condensate on an Atom Chip

Thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy.

Physics Department
Imperial College London

Iain Barr

May 2015

ABSTRACT

In this thesis I discuss work that has been carried out on the dynamics of a Bose Einstein condensate of Rb 87 produced near an atom chip.

A Bose Einstein Condensate (BEC) is a quantum state of matter where a single quantum state becomes occupied by a macroscopic number of identical Bosons. In our case this is achieved by cooling a system of trapped identical rubidium 87 atoms to its ground state.

To reach temperatures of condensation we initially laser cool atoms from room temperature, before loading them into a magnetic trap. The magnetic trap is produced through a combination of uniform magnetic fields from coils outside our vacuum chamber and currents running through wires on an atom chip. The atom chip is a microfabricated device, produced by a coating a silicon chip with a thin layer of gold and etching wires into it. Together, these fields create a magnetic field minimum $120\mu m$ from the surface of the chip which can be used to confine low field seeking hyperfine states of the atom in an elongated harmonic trap. Once the atoms are confined in the magnetic trap we used force evaporative cooling out to reach the phase space densities required for Bose Einstein condensation.

The BEC is used to investigate the relative dynamics between the fraction of the atoms in the condensate to those not in the condensate. Our atom chip provided a suitable environment to investigate this due to fragmentation of the magnetic potential close to the chip. Small imperfections in the wires on our atom chip mean that the trapping potential isn't smooth. Small regions of higher trapping frequency - or fragments - are formed. Due to the small size of these fragments it is possible to find a position where a condensate can form in the fragment, and see a potential of high frequency, whereas a non-condensed atom will see a lower frequency potential. We exploit this to set the condensed fraction moving relative to the non condensed part and investigate the subsequent damping of their motion relative to each other.

ACKNOWLEDGEMENTS

This project couldn't have been completed without the assistance of many people, too many to directly name or remember. But there are some who cannot be left unacknowledged, A huge thank you: to Ed, to Joe, Ben and Eoin, to the whole of CCM, to the CQD CDT, to Jon and Steve and Val, to Sanja and finally to Caitlin.

DECLARATION

I do declare that the work in this thesis is my own, except where otherwise acknowledged.

CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Beginnings

In 1995, a long standing goal of experimental physics was achieved: the creation of a Bose Einstein condensate from a dilute atomic gas [1][2][3]. These experiments succeeded in trapping and cooling a gas of alkali atoms down to temperatures of a few hundred nano kelvin - temperatures far lower than any naturally occurring phenomenon.

At these low temperatures and densities a remarkable regime is reached where the wavefunction of each atom extends far beyond the average inter particle spacing, and quantum effects begin to dominate. For Bosons these effects are manifest as a macroscopic occupation of the ground state of the system and the production of a truly quantum state of matter: the Bose Einstein condensate (BEC).

Bose Einstein statistics were first described by Bose in 1924 [4] and extended to massive particles by Einstein [5] and marked the first discussion of quantum statistics. Under Bose Einstein statistics particles can no longer be considered individual or uniquely identifiable, the presence of one particle intrinsically affects the statistics of another. To be precise this means that the wave function of a group of identical bosons has to be symmetric under exchange of any two particles.

A first consequence of this is that, when calculating the equilibrium distribution of a system of bosons by 'counting microstates', one needs to take into account that a configuration of identical particles in the same state count as one microstate. Following this line of reasoning produces the Bose Einstein distribution where identical particles seem to bunch together. At low temperatures and high densities this results in a macroscopic occupation of a single state - the BEC.

The occupation of this single state has profound consequences. We arrive at a system in which many particles occupy a single wavefunction of the ground state of the system. The result is a many body quantum system whose properties and dynamics can be tested on an unprecedented scale.

Dilute atomic gases are not the only systems in which Bose Einstein condensation has been achieved. Super-fluid liquid helium at temperatures below its lambda point has been described as a BEC [6][7]. The superconducting electrons in metals below a certain temperature have also been described using the idea of a BEC formed from pairs acting as bosons [8].

Other systems of quasi-particles in which condensation has been observed include photons in a microcavity [9], exciton polaritons in semiconductor microcavities [10] and magnons in a magnetic insulator [11].

Interest in atomic gas systems arise because they allow access to regimes that are inaccessible to other experiments with BECs. Unlike liquid helium, the interactions between atoms in a gas are

normally weak at the densities in which BECs are produced, which allows microscopic description. Unlike quasipartical systems the BEC in dilute gases are isolated from their environment and almost free of defects.

A further advantage of BECs in ultracold gases is the level of control and measurement attainable in such systems. The timescales on which the dynamics of such systems evolve are defined by the frequencies of the traps in which the atoms are confined, and are typically of the order of microseconds to seconds, well within experimental resolution. The ability to rapidly change the confining potential or internal states at the atoms on these timescales allows precise control.

Extracting information from such systems typically involves measuring the densities of the atoms, which is possible on timescales faster than the dynamics of the system. From these measured real space densities, information about both the real space distributions, momentum space distributions (by imaging after a time of flight), the systems can be probed at almost a single atom level.

1.2 The Place of our Experiment

Since the first creation of a BEC in dilute atomic vapours, people have been interested in their dynamic properties. Early experiments characterised the low lying collective excitations of the condensate using time varying potentials [12][13]. These excitations caused time varying density modulations of the condensate that could be probed by directly measuring the atomic density of the condensate over time. The modes excited were the centre of mass mode, the monopole mode and the quadrupole mode.

At the same time as these early experiments, theoretical models were developed to explain and characterise the properties of the system. Overviews of the properties of trapped condensates can be found in [14] [15] [16]. In the limits of zero temperature and three-dimensional systems, the entire atom cloud exists as a BEC, and the system is well described by mean field theory.

In this mean field theory each atom exists in the same wavefunction, and the total wavefunction of the system is described by the product state of these wavefunctions. Here the dynamics of the system is described by the Gross-Pitaevskii equation. This gives the shapes and frequencies of the first few collective excitations, which are found to deviate from those of a non-interacting gas and match experiments well [17].

As well as collective excitations, early experiments have managed to probe other properties of the condensate, including the speed of sound in the condensate by using non destructive imaging to following the propagation of density disturbances introduced using a blue detuned laser [18]. A focused blue detuned laser provides a repulsive potential in addition to the magnetic trap, and was also used to prepare two independent condensates in close proximity, allowing the observation of interference fringes when they were allowed to overlap [19], demonstrating the coherence of a condensate.

For a system at zero temperature, low lying collective excitations are not expected to decay, as there is no mode for them to dissipate energy into. For finite temperature systems this is no longer the case. Energy can be redistributed among the excited modes of the system by collisions between thermal atoms.

Early experiments measured this damping and found that both the rate of damping and frequency of the collective excitations varied as functions of temperature [20][21]. Here the damping was explained by a process known as Landau damping [22][23][24][25][26]. Landau damping is a process whereby a collision between a thermal excitation and the collective excitation results in an atom in a

ground state and a thermal excitation.

In order to calculate the rate at which this damping process occurs, excitations of the condensate can be modelled using the Bogoliubov approximation. Under this approximation the transition rates can be calculated using Fermi's golden rule. For three-dimensional gases, the predictions of these models matched theory well. Higher order damping mechanisms have also been observed for 'scissor' mode oscillations [27].

Measurements of centre of mass damping, made by displacing then releasing the condensate from the thermal cloud in a harmonic trap, demonstrated that while in the collisonless regime the results of Landau damping correspond well to experiment [28]. However systems in the hydrodynamic regime damped faster than expected. The collisionless and hydrodynamic regimes corresponds to systems where the number of collisions between thermal excitations per axial oscillation are small and large respectively. The authors also demonstrated that when the condensate moved through the thermal cloud at a speed faster than the critical speed of sound in the condensate, the damping was enhanced and the condensate destroyed.

The exquisite control offered by cold atom experiments means it is possible to increase the confining potential in certain directions so that the spacing of modes is much larger that the both the thermal energy and chemical potential. Here the condensate can no longer be considered three-dimensional in its quantum mechanical description - while excitations in the other modes exist, they cannot be reached.

In both one and two dimensions a system can never reach a full BEC in the thermodynamic limit due to thermal fluctuations in the system [14]. However for finite sized systems such as trapped gases it is possible to reach a point where the coherence length of the system is larger than its spatial extent, and a quasi condensate is formed. This has been achieved in both one [29] and two dimensions[30].

The system of a one-dimensional Bose gas can be solved exactly [31], and the reduced dimensionality affects the dynamics. How low dimensional system relax towards their thermal equilibrium raises a number of interesting questions. In [32] and [33] an out of equilibrium gas in an optical lattice was prepared and then measured as it evolved. In [34] and [35] it was shown that the these one-dimensional systems don't relax immediately to their equilibrium state, instead the integrability of the system means that the systems relax to a prethermal state. In extreme cases this can result in excitations that never damp [36].

Between one-dimensional systems, where damping is not expected and three-dimensional systems, where it is well explained, it is not yet understood how to describe the damping, either experimentally or computationally. Measurements in this space provide a guide to theoretical attempts to understand this system. It is in this space that our experiment takes place.

In [37] the breathing mode of a quasi-one-dimensional system was investigated. By varying the atom number and measuring the frequencies of both real space and momentum space, the change in frequency of the oscillations was mapped out as the system crossed between the ideal Bose gas and the quasi-one-dimensional BEC regimes.

1.3 Our Experiment

Our own experiment deals with a BEC of atomic rubidium 87, trapped and cooled in an ultra high vacuum environment. A brief discussion of our experiment is given in this section, before returning to an overview of the physics involved in the next chapter.

Our experiment begins in a secondary vacuum chamber with a gas of atomic rubidium being produced by a set of SAES getters dispensers [38] into our vacuum system. At this point the vapour reaches room temperature through collisions with the vacuum chamber walls.

A first round of magneto optical trapping is used to cool and trap the atoms down to 140μK. From here the atoms are transferred into our main vacuum chamber where they are collected and cooled by a second round of magneto optical trapping. At this point the atoms are 1mm below our atom chip. After subdoppler cooling and optical pumping the atoms are captured in a magnetic trap created through a combination of externally produced (relative to the vacuum chamber) magnetic fields and magnetic fields produced using the wires on our atom chip.

The magnetic trap sits 120μm beneath the surface of our atom chip, and produces a highly elongated potential (aspect ratios of the order 100 - 140), somewhere between the one and three-dimensional cases. Evaporative cooling is then used to lower the temperature further, and beyond the transition temperature to produced a Bose Einstein Condensate.

One advantage of using an atom chip is that the potential created isn't smooth, but exhibits roughness caused by imperfections in the wire. This non uniform potential allows us to excite centre of mass oscillations of the condensate (which sits at the lowest point of the magnetic trap, and sees a higher frequency harmonic potential) while keeping the thermal cloud stationary (which explore a higher energy part of the trap, and sees a lower frequency) by displacing the system at a frequency resonant only with the potential explored by the condensate. The end result is relative oscillations between the condensate and the thermal cloud that experiences damping. The relative oscillation can be read out by measuring the time dependant density of the system as it evolves.

By varying the final frequency of our radio frequency evaporation ramp we change the temperature of the system (along with other parameters). The ability to vary temperature allows us to investigate the damping of the centre of mass mode as a function of temperature. The result is a system which seems to damp more slowly than three-dimensional systems the more one-dimensional it is.

1.4 Organisation of this Thesis

This thesis is organised into three chapters following this one. In the second chapter we review the theory necessary to understand the our experiment. In chapter three we describe the design of the experiment, and provide an overview of the hardware used. In the final chapter we present the results and discuss them with regard to current theory and other experimental results.

CHAPTER 2

BACKGROUND CONCEPTS

2.1 Overview

A number of distinct physical principles are required to understand the experiments discussed here. In this chapter we aim to give an overview of the most relevant ones. In the first section 2.2 we introduce the rubidium 87 atom and its relevant properties. In the next section 2.3 we cover how the atoms are cooled from room temperature to close to their ground state. We then cover our trapping, imaging and state manipulation techniques in sections 2.4-2.5. We conclude with a discussion of the phenomenon of Bose Einstein Condensation itself 2.7.

The full details of how these techniques and ideas are implemented are deferred to the next chapter.

2.2 The Rubidium Atom

The atom used in our experiment is rubidium 87 (Rb87). Due to the accessibility of the wavelength of its cycling transition with diode lasers it is a popular choice in laser cooling experiments and one of the first atoms used to produce a BEC in atomic vapours [1]. In this section I review a few properties of the atom.

2.2.1 Level Structure

One reason for the popularity of alkali metals in laser cooling experiments is the simplicity of their level structure. All but one of the electrons form complete shells and in a first approximation this leaves a single electron to interact with light.

In this section I use standard spectroscopic notation when labelling atom levels and transitions. For a review of this notation see [39]. For a comprehensive overview of the Rb87 D2 lines, see [40].

Rubidium 87 has a nuclear spin of $I = \frac{3}{2}$. When combined with the spin $\frac{1}{2}$ of the single unpaired electron, the atom has total integer spin, making it a Boson.

From the point of view of our experiment the interesting internal structures of rubidium 87 are those of the ground state and the first excited state. The ground state is a S orbital, with combined spin orbit angular momentum coming only from the electron, giving $J = \frac{1}{2}$. When combined with the nuclear spin, we have two hyperfine ground states, one with $F = 1$ and one with $F = 2$, separated in energy by roughly 6.8GHz.

The first excited state has orbital angular moment $L = 1$. The fine structure then splits this state into two components, having total electronic angular momentum $J = \frac{1}{2}$ and $J = \frac{3}{2}$. Transitions

between the ground state and these states produce the D1 and D2 lines of rubidium respectively. Here we focus on the D2 line due to its use in laser cooling.

The $P_{3/2}$ state is split into 4 different total angular momentum states due to the hyperfine interaction. These states are labelled as $F = 0, 1, 2, 3$. The reason for focusing on the D2 line in rubidium is that selection rules allow us to form a closed transition between the $F = 2$ ground state and the $F = 3$ excited state.

A level diagram, complete with the relevant properties can be found in Figure 2.1

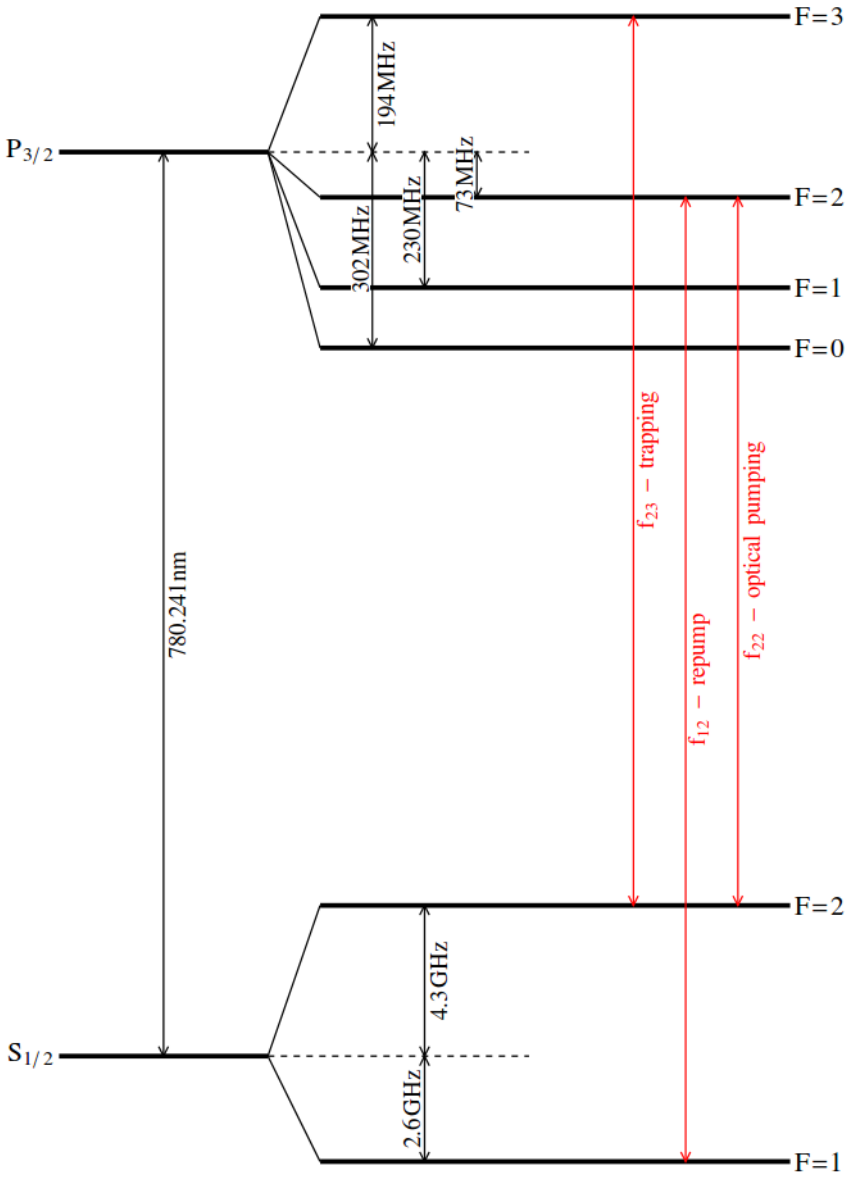

Figure 2.1: Level Scheme for the D2 levels of rubidium 87. Transition frequencies are taken from [40]. The transitions used for cooling, repumping and optical pumping have been labelled.

2.2.2 Physical Properties

Along with favourable internal structure, the rubidium 87 atom is an attractive choice for BEC experiments because the interactions between atoms are both repulsive and weak. The advantage of weak interactions is that a collection of atoms can often be considered as an ideal gas, simplifying our theoretical treatment of it, when compared to systems such as liquid helium, where strong interactions alter the physics of the condensed phase [41].

The advantage of a repulsive interaction is that it makes the system more stable. Energetically, individual rubidium atoms are not stable - a diatomic molecule is more stable. The instability means that there is a limit to densities of atoms that can exist before collisions between atoms cause the formation of rubidium molecules. When in the condensed phase the repulsive interactions act to limit the peak density (as discussed in section 2.7, the interaction energy dominates the kinetic energy). For attractive interactions, as more atoms condense the attractive interactions cause the density to increase - fundamentally limiting the size of BECs with attractive interactions [14].

2.3 Cooling Techniques

For a thermal gas the average rms wavefunction spread of each particle varies as the inverse of the square root of the temperature. To reach to the point where quantum effects due to the overlap of particle wavefunctions become relevant the system needs to be cooled significantly. To do this we use a number of well-explored cooling techniques.

2.3.1 Magneto Optical Trapping

The magneto optical trap (MOT) has become one of the mainstays of cold atomic physics. It works by harnessing two effects: the scattering force of light and level shifts due to the Zeeman effect. The result is both cooling and trapping. The MOT has been made possible through the arrival of laser technology which gives coherent, collimated light with a frequency spread smaller than the natural linewidth of atoms.

The details of this section are taken mainly from [42] and [39].

2.3.1.1 Scattering Force from a Two Level System

Given the fast spontaneous decay time of atoms from allowed optical transitions (the rubidium 87 $5^2P_{3/2}$ to $5^2S_{1/2}$ state has a natural lifetime of 30ns), an atom interacting with near resonant light can (to a first approximation) be assumed to have reached its steady state distribution. In steady state the population of the excited state is given by

$$\rho_{22} = \frac{1}{2}\frac{s}{1+s+4\delta^2}\,, \tag{2.1}$$

where $\delta = \frac{\omega_l - \omega_a}{\Gamma}$, is the difference in frequency between the atomic transition ω_a and the frequency of light incident on the atom ω_l, normalised by the spontaneous decay rate Γ, and $s = \frac{I}{I_{sat}}$ is the intensity of the incident light I, normalised by the saturation intensity of the transition $I_{sat} = \frac{c\epsilon_0\hbar^2\Gamma^2}{4|\hat{\epsilon}\cdot d|^2}$, where $|\hat{\epsilon}\cdot d|$ is the dipole matrix element between the coupled states. For the transition of interest in rubidium 87, $\Gamma = 6.0$MHz and $I_{sat} = 3.58\text{mW/cm}^2$ (for atoms that aren't being optically pumped)[40].

Because of the spherical symmetry of the system, a photon is generally considered to be emitted in a definite direction; all directions are equally likely. The force due to the spontaneously emitted photons will therefore average to zero. However, there is a force on the the atom due to the momentum transferred from absorbed photons. Because for every spontaneously emitted photon there must be exactly one photon absorbed, we can write the net force on the two level system as the average number of spontaneously emitted photons per unit time, $\Gamma\rho_{22}$ multiplied by the momentum per photon, $\hbar\mathbf{k}$, where $\mathbf{k}$ is the wave vector of the light. The resulting force is

$$\mathbf{F}_{scatter} = \frac{\mathbf{k}\hbar\Gamma}{2}\frac{s}{1+s+4\delta^2}\,. \tag{2.2}$$

The scattering force is maximised for resonant light. To understand cooling, we need to include the effect of the Doppler shift in frequencies. For light of frequency ω_l in the lab frame, an atom moving at velocity $v << c$ will see light at frequency $\omega_l' = \omega_l + \mathbf{k}\cdot\mathbf{v}$. The velocity-dependent force on the atom due to a single laser, in the direction of the laser's propagation, is given by

$$\mathbf{F}_{scatter}(\mathbf{v}) = \frac{\mathbf{k}\hbar\Gamma}{2}\frac{s}{1+s+4\left(\frac{\omega_l-\omega_a+\mathbf{k}\cdot\mathbf{v}}{\Gamma}\right)^2} \tag{2.3}$$

If we now consider the force from two counter propagating beams on an atom moving in one dimension (see Figure 2.2), we get a total force of

$$\begin{aligned} F_{total}(v) &= F_{scatter}(-v) - F_{scatter}(v) && (2.4)\\ &\simeq (F_{scatter}(0) - v\alpha) - (F_{scatter}(0) + v\alpha) && (2.5)\\ &= -2v\alpha\,, && (2.6) \end{aligned}$$

where $\alpha = \frac{\partial F_{scatter}}{\partial v}\big|_{v=0}$ and in the second line we have assumed the Doppler shift to be small, and Taylor expanded around zero velocity.

The coefficient α has the form

$$\alpha = 4k^2\hbar\frac{s\delta}{(1+s+4\delta^2)^2}\,. \tag{2.7}$$

This only holds for atoms where $\mathbf{k}\cdot\mathbf{v} << \Gamma$ - those travelling too fast will see the light as far detuned and will not scatter any photons. Roughly, atoms with velocities much bigger than $\frac{\Gamma}{k}$ will not be cooled. For rubidium 87, this is $v_c \approx 5$m/s.

Because the sign of α depends on the sign of the detuning of the laser, red detuned lasers will produce a damping force that cools the atoms, which blue detuning does not. In the low intensity limit (where the intensity is well below saturation) the effects of multiple laser beams will add, and with six beams arranged into three counter propagating perpendicular pairs we can cool the atom in every dimension.

It should be clear that the force above isn't a conservative force - this means that the entropy is lost from the system. It is lost to the light field. Low entropy laser light is converted into high entropy randomly scattered light by the atoms, allowing the subsystem consisting of the atoms to cool.

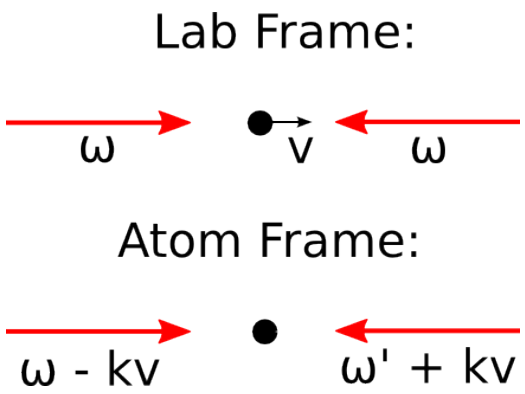

Figure 2.2: The Effect of the Doppler shift on the frequencies seen by a moving atom due to two counter propegating beams.

2.3.1.2 Zeeman Effect

In the previous section we considered the effect of light on two level systems. In reality atoms have a rich internal structure that needs to be considered. In order to get close to the idealised version of a two level system, a closed cycling transition is usually used. In our case we use the transition between the $F = 2$ and $F = 3$ hyperfine states of the D2 line in rubidum, as described in section 2.2.1.

The use of hyperfine states with non zero angular momentum means that a range of transitions between Zeeman sub levels can be accessed, the rates of which depend on the local magnetic field and polarisation of the light relative to the atom.

The effect of the magnetic field is to shift the energy of a sub level in proportion to the state's m_f number. At low magnetic fields the Hamiltonian that describes the effect of a magnetic field on the atom is given by

$$H_{Zeeman} = \mu_B g_F \mathbf{F} \cdot \mathbf{B}\,, \tag{2.8}$$

where μ_B is the Bohr magneton and g_F is the Landè g factor for the given level. For an inhomogeneous magnetic field this provides a position dependent change in transition frequency.

The effect of the light polarisation is described by the selection rules for hyperfine transitions. σ^+ light can only drive a $m_f \rightarrow m_{f'} + 1$ transition, σ^- light can only a drive $m_f \rightarrow m_{f'} - 1$ transition and π light will drive a $m_f \rightarrow m_{f'}$ transition.

If polarisations of the incoming beams and magnetic fields are combined correctly we can achieve both a cooling force and a trapping force on the atoms. To see how this is done, we consider the one-dimensional set up in Figure 2.3. Here for simplicity we are considering a $F = 0$ to $F = 1$ transition. How rubidium deviates from this will be discussed below.

Two red-detuned lasers are set counter propagating, one σ^+ polarised, the other σ^- polarised, aligned with a linear magnetic field gradient. If the atom is displaced from the zero of the magnetic field, the Zeeman shift of its hyperfine levels brings the frequency of one beam closer to resonance, while moving the frequency of the other away from resonance. The polarisation is chosen so that the light closer to resonance provides a restoring force, towards the magnetic field zero. In three dimensions, six crossed beams and a magnetic quadrupole field are used to provide the cooling and confinement.

To estimate the size of the force, we repeat the analysis from above using the fact that the detuning is now position-dependant due to the Zeeman effect and the polarisation of the beams. For one beam

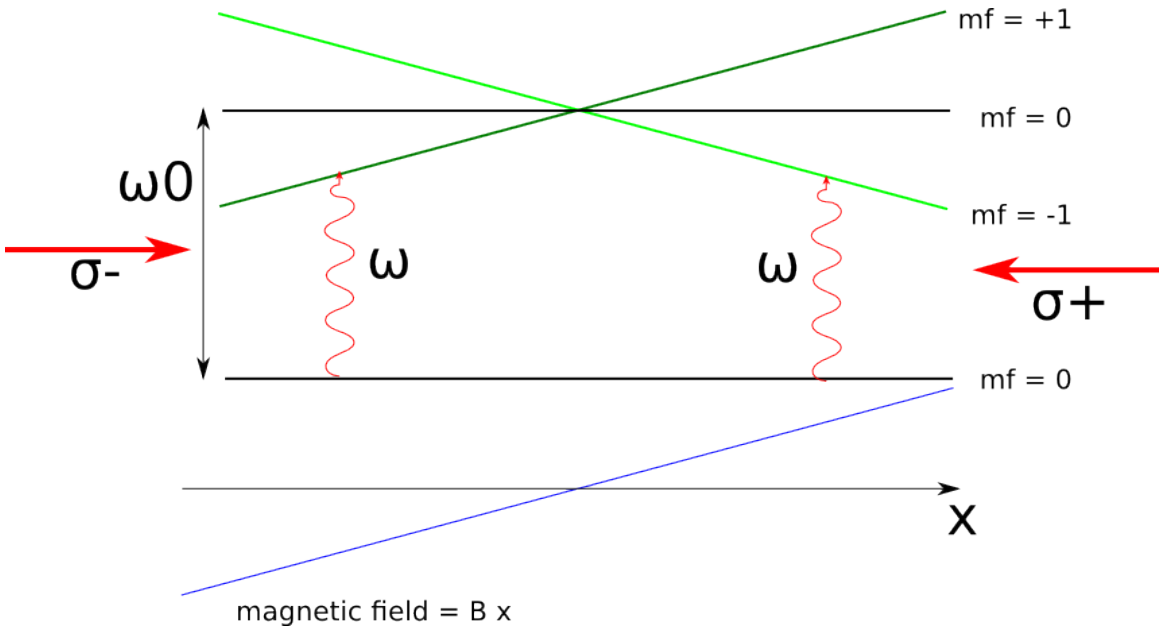

Figure 2.3: A schematic of a one-dimensional magnetic trap. The diagram shows how the energy levels of an atom change spatially due to a linear magnetic field in x (represented by the blue line). When the atom is displaced to the right, the $m_f = 1$ level is shifted closer to resonance, and addressed by the σ^+ which moves from right to left, applying a restoring force to the atom.

of polarisation $p = \pm 1$ corresponding to $\sigma^{\pm 1}$, and assuming that $g > 0$ for the levels selected, the force is given by

$$F_p(v,x) = \frac{\hbar\Gamma}{2}\frac{s}{1+s+4\left(\frac{\omega_l-\omega_a+kv+pg\mu_0 B(x)}{\Gamma}\right)^2}\,. \tag{2.9}$$

It is worth noting that the form of this force assumes that the atom is still in equilibrium: the total population of the excited states are still given by equation 2.1. Given that the decay rate is the same for all excited states, including extra excited states shouldn't change this assumption.

For two counter propagating beams of different polarisations, and a linear magnetic field gradient, the total force on an atom is given by

$$\begin{aligned} F(v,x) &= F_{+1}(v,x) + F_{-1}(v,x) \\ &= \frac{\hbar\Gamma}{2}\frac{s}{1+s+4\left(\frac{\omega_l-\omega_a+kv+g\mu_0 B(x)}{\Gamma}\right)^2} - \frac{\hbar\Gamma}{2}\frac{s}{1+s+4\left(\frac{\omega_l-\omega_a-kv-g\mu_0 B(x)}{\Gamma}\right)^2} \\ &\simeq (F(0,0) - \alpha v - \beta x) - (F(0,0) + \alpha v + \beta x) \\ &= -2\left(\alpha v + \beta x\right)\,, \end{aligned} \tag{2.10}$$

where $\alpha = \left.\frac{\partial F(v,x)}{\partial v}\right|_{v=0,x=0}$ and $\beta = \left.\frac{\partial F(v,x)}{\partial x}\right|_{v=0,x=0}$ and we have Taylor expanded the force to first order in both x and v. Here again we recover the same cooling force as before, but also a harmonic trapping force.

As with cooling, this can be extended into three dimensions using three sets of mutually orthogonal counter propagating beams and a magnetic quadrupole field, creating a magneto-optical trap.

The states we use in rubidium 87 for trapping and cooling have five hyperfine ground states and seven hyperfine excited states; however, because the g factor for the ground state is close to that of

the excited state ($g = \frac{1}{2}$ vs $g = \frac{2}{3}$), the anaysis above provides a reasonable approximation.

The analysis above also ignores a number of effects, such as collisions between atoms, re-absorption of scattered photons, coherences between internal states, however it does capture the basic physics.

2.3.1.3 The Limits of Doppler Cooling

In the previous section we considered the trapping and cooling forces. However, in this analysis we have looked at the average force on the atoms, not the fluctuations. There are two effects: the random direction the spontaneously-emitted photons are emitted, and fluctuations in the number of photons absorbed at each time step. These fluctuations increase the entropy of the atoms, heating them and introducing a fundamental limit to Doppler cooling.

To find this fundamental limit, we consider each effect on the atoms. Firstly the cooling. The average force on an atom Doppler cooled gives the equation of motion in one dimension to be

$$\frac{\partial p(t)}{\partial t} = -\alpha \frac{p}{m}\,, \tag{2.11}$$

where α is the cooling coefficient given above. Here, energy is being removed from the system at a rate of

$$\frac{\partial E}{\partial t} = \frac{\partial \frac{p^2}{2m}}{\partial t} = \frac{p}{m}\frac{\partial p}{\partial t} = -2\frac{\alpha}{m}E\,. \tag{2.12}$$

To estimate the heating rate from the fluctuations, we use the fact each spontaneously emitted photon gives the atom a kick in momentum space of magnitude q_r, the recoil momentum of the atom. Because the direction of each of these steps is random, the distribution of the final momentum follows that of a random walk: for N steps the mean momentum is given by $\hbar k\sqrt{N}$. In what follows we will be looking at the heating in only one dimension.

For two counter propagating beams the atom will on average absorb the same number of photons from each beam, however there will be fluctuations. Writing the average scattering rate for the atom due to one beam as R (and working in the regime where $I << I_{sat}$), then in time T the average number of photons absorbed from each beam will be RT. The momentum fluctuations due to both beams will be of the size $q_r\sqrt{2RT}$ (where we have assumed Poissonian fluctuations in the number of photons absorbed). The heating rate is then given by $\frac{Rq_r^2}{m}$.

Similarly, in time T the number of spontaneously emitted photons due to two beams is given by $2RT$. The average momentum gained is then $\sqrt{2RT}q_R$ [1]. The increase in energy per unit time is then given by $\frac{Rq_r^2}{m}$.

Combining the three results above, writing $E_r = \frac{q_r^2}{2m}$ as the recoil energy, we get the steady state average energy of

$$E_{doppler} = \frac{RE_r m}{\alpha}\,. \tag{2.13}$$

Using the definition for α and R given in previous sections, and ignoring saturation effects, it is possible to show that the minimum energy occurs at a detuning of $\omega_0 - \omega = \frac{\Gamma}{2}$. Inserting this into

[1] In the full three-dimensional picture an absorbed photon can scatter into any direction, not just the one direction considered here. The effect of this is to decrease the momentum in the direction considered here by a factor of 3. It also doesn't consider the beams propagating along the other dimensions. These increase the scattering rate (in the low intensity limit) by a factor of 3. Overall these effects cancel each other and we get the result presented in the text.

the equation and taking the temperature to be equal to $\frac{1}{2}k_BT = E$ (by the equipartition theorem for the energy in one dimension) we get a minimum temperature of

$$k_BT_D = \frac{\hbar\Gamma}{2}\,, \tag{2.14}$$

where k_B is Boltzmann's constant and T_D is the Doppler temperature[2]. For the cooling transition in rubidium 87 that we use this gives a minimum temperature of 146μK.

2.3.1.4 A note on relative magnitudes

Before continuing it is worth making a note of a few values.

The average velocity of a rubidium 87 atom at 146μK is roughly 12cm/s.

The recoil velocity of a rubidium 87 atom is 6mm/s, which corresponds to a temperature of 0.4μK.

2.3.2 Polarisation Gradient Cooling

Historically, when cooling using scattered light was first developed it was found that the temperatures reached were lower than the Doppler limit discussed in the previous section, and a new theory had to be developed. Effects such as light shifts of the energy levels and non equilibrium dynamics turned out to play an important role. The resulting technique became known as Polarisation Gradient Cooling and is described in this section. For a more detailed treatment see [42] and [43] and for a historical overview of its development see [44][45][46].

2.3.2.1 Standing Waves

Polarisation gradient cooling has two main flavours: lin-lin cooling and $\sigma^+\sigma^-$ cooling depending on the relative polarisation of the counter propagating beams involved. Here we discussed only the $\sigma^+\sigma^-$ as it is the one present in our experiment.

When two counter propagating circularly polarised beams are overlapped, the result is a standing wave of linearly polarised light, whose direction of polarisation rotates one full revolution every wavelength of the light. It is this spatial dependence of polarisation that gives the scheme the name 'polarisation gradient'.

2.3.2.2 Optical Pumping

In the absence of magnetic fields we can take the polarisation of the light field as the quantisation axis for the atoms. Linearly polarised light drives π transitions in the atom, and these π transitions drive the atom to a steady state population where the $m_F = 0$ ground state is more populated than other states. The redistribution of populations among the internal states of an atom is known as optical pumping.

The reason for the atom being pumped to a state in which $m_F = 0$ is more strongly populated is due to the fact that once an atom is excited, the level it spontaneously decays to is independent

[2]There might be some question of what temperature means in this situation. Temperature is usually defined for a system in equilibrium with a much larger environment. In the case of atoms interacting with light, the light is in a highly non-equilibrium state, and so is this concept of temperature is not a valid concept. Using the Fokker-Planck equation for the system, fluctuations and all, the stable distribution for the system is the Boltzmann distribution, and the temperature of this distribution is the one mentioned here [42]

of the level it was excited from - it depends only on the dipole matrix elements, which redistribute population towards the $m_F = 0$ state.

2.3.2.3 Non Equilibrium Motion and Light Shifts

If an atom moves through the laser beam sufficiently quickly, the effects of the rotating polarisation will average out and we can ignore them. If it travels sufficiently slowly the ground state population will adiabatically follow the equilibrium distribution. By 'sufficiently' fast or slow, we mean with respect to the time it takes for the atom to redistribute its internal states. It is between these two limits that polarisation gradient cooling has an effect.

If we move into a frame where the quantisation axis is parallel to the polarisation of the light, the one-dimensional motion of the atom through the polarisation lattice causes an effective rotation of the atom. For an atom moving with velocity v along the direction of the beams this equates to a rotation rate of $\dot{\theta} = kv$ where k is the wavevector of the light. Transforming from the lab frame to this rotating frame corresponds to a unitary rotation of $U = \exp(-itkvJ_z/\hbar)$ where J_z is the angular momentum operator in the direction of the beam. The Hamiltonian in the rotating frame then acquires an additional term given by

$$H_{rot} = kvJ_z\,. \tag{2.15}$$

This term couples the hyperfine levels of the atom, causing the population to move out of equilibrium as the atom moves through the polarisation lattice. In the atom frame this causes an symmetric transfer of populations among m_f states, however,in a frame aligned with the k vector of the light, this imbalance is asymmetric in m_f states.

On average the atom will absorb and scatter photons in such a way as to return to its equilibrium distribution - this means that this imbalance in m_f states causes a preferential absorption of photons from one of the beams in the lab frame. As shown below, the size of the population imbalance is proportional to the velocity of the atom, and the direction is such that the absorbed photons act against the motion, creating a velocity dependant cooling force.

To get a rough estimate of the size of the damping coefficient for polarisation gradient cooling we note that in the atom frame, to first order, the effect of 2.15 is to change the population of the levels by an amount of order $\frac{kv}{\Delta_L}$, where $\hbar\Delta_L$ is the separation in energy between the Zeeman sublevels due to the light shift[3] Ignoring numerical pre-factors of order one coming from the Clebsch-Gordan coefficients this creates a difference in Zeeman sublevel populations in the lab frame of the same order. The force on an atom is then given by this difference multiplied by the momentum transferred from each photon $\hbar k$ and the scattering rate R,

$$F \approx -v\frac{\hbar k^2 R}{\Delta_L}\,. \tag{2.16}$$

The cooling force is balanced in turn by heating due to the random kicks from the spontaneously emitted photons in the same way as with Doppler cooling. Following the arguments given in section 2.3.1.3 we can derive a rate due to just the scattering rate. Equating this to the cooling rate from equation 2.16 we get a temperature limit of

[3]This comes from treating 2.15 as a perturbation and applying perturbation theory. To first order there is no energy shift, but the levels are perturbed.

$$k_B T \approx \hbar \Delta_L \approx \frac{\hbar \Gamma^2 s}{|\delta|}\,, \tag{2.17}$$

where s is proportional to the intensity of light, as defined in section 2.3.1 and δ is the detuning of the light from resonance, and we have approximated the light shift in the limit $\delta >> \Gamma$. Comparing this result to 2.14, for laser detunings greater than Γ and $s < 1$, both readily achievable, we can reach lower temperatures than the Doppler limit.

From the above discussion we can find the rough range of atom velocities over which we would expect polarisation gradient cooling to take place. If it takes time τ to redistribute the atom's internal state to equilibrium, then velocities faster than $\frac{\lambda}{2\tau}$ will not be affected.

It also shows that this type of cooling can only happen in systems where we have three or more levels in the ground state manifold - any fewer and we don't see the required imbalance in populations due to optical pumping. For two levels it can be shown that a similar form of polarisation gradient cooling can come into effect [43].

The above only sketches the details of sub Doppler cooling. For a more quantitative treatment see [42] and [43].

A fundamental limit to polarisation gradient cooling comes from the recoil velocity of the atom from the light. Polarisation gradient cooling still relies on spontaneously emitted photons, so this provides a hard lower limit of 0.4μK. In this experiment we typically only reach temperatures two orders of magnitude above this temperature due to imperfections in the polarisation and intensity profile of the light, as it is reflected of the atom chips surface.

2.3.3 Evaporative Cooling

While the temperatures reached via laser cooling are low, they are not low enough to reach Bose Einstein condensation. To decrease the temperature further, we use evaporative cooling. The basic idea is that if we remove only the energetic atoms from a trapped gas, the average temperature of remaining trapped gas will be lower after equilibration. In our set up this is achieved by magnetically trapping the atoms (see section 2.4), and then applying a radio frequency magnetic field to couple the Zeeman sublevels (see section 2.5).

The atoms close to resonance with the rf radiation are rapidly transferred from the low field seeking $m_f = 2$ state to an untrapped sublevel, which causes these atoms to be accelerated away from the main trap; by making the resonant point far from the trap centre, such that only atoms with high energy can reach it, we can selectively address only the energetic atoms - evaporatively cooling the system.

Evaporative cooling of trapped atomic gases was first suggested in [47] for spin polarised hydrogen and later used to achieve the first BECs in Alkali gases [1]. The overview here mostly follows from [48] and [49].

2.3.3.1 The Basics

For a gas of average energy per particle $\bar{\epsilon}$ and particle number N, if we remove dN particles each of average energy $(1+\beta)\bar{\epsilon}$ we get a relationship between number and average energy of

$$\frac{\bar{\epsilon}}{\bar{\epsilon}(0)} = \left(\frac{N}{N(0)}\right)^{\beta}\,, \tag{2.18}$$

where $N(0)$ and $\bar{\epsilon}(0)$ show the initial atom number and average energy per atom.

In order to relate average energy to temperature we assume that at all times the atoms within the trap are thermally distributed[4]. This assumption has be shown to be approximately correct via numerical simulation in [50].

Because average energy is proportional to temperature in this system, we get

$$\frac{T}{T(0)} = \left(\frac{N}{N(0)}\right)^{\beta}, \tag{2.19}$$

where $T(0)$ and $N(0)$ are the initial temperature and number respectively.

This equation suggests that in order to lower the temperature the most for a given atom loss, we should make the energy at which we remove particles β as high as possible. This shouldn't be surprising - if we set the energy cut off at the total energy of the atom cloud, eventually a single atom will be ejected, leaving the rest of the system in the ground state. The problem is that this would take a prohibitively long time, and in real conditions other mechanisms cause heating or atom loss, making such a scheme untenable. We need to look at the rate at which these effects happen.

2.3.3.2 Rate of Evaporation

Once confined in the magnetic trap, there are two main mechanisms for atom loss in our experiment: collisions with residual particles in our vacuum chamber, and three body recombination. As discussed above, collisions between atoms tend to be elastic at low temperature. This is because although a rubidium molecule of two atoms is chemically more stable than single atoms, in a two body collision there is no mechanism to remove the excess binding energy. In a three body collision this can happen - two atoms form a molecule and the third gains the excess energy as kinetic energy. The rate at which this will occur is proportional to the density of the atoms squared, with constant of proportionality $K_3 = 1.810^{-41} m^6 s^{-1}$ [51]. In our experiment we don't reach densities at which this effect becomes relevant, and in this overview I will ignore these effects from here on.

The loss due to background collisions depends on the background pressure and the type of atoms present in the chamber. It will be constant in this experiment.

The loss rate fundamentally limits how slowly we can proceed in evaporation, however there is also a limit to how quickly we can proceed: to evaporate optimally, atoms need to 'rethermalise' as they evaporate. The time scale for rethermalisation is governed by the rate of elastic collisions between atoms in the cloud. The average time between collisions is $\tau_{col} = nv\sigma$, where n is the number density, v is the average relative velocity between atoms and σ is the scattering cross section of the atoms (we assume that the atoms are cold enough that the cross section is independent of atom energy).

Because the density of the atoms enters into the rate of rethermalisation, the effect of evaporation is non linear: cooling leads to an increase in density, which in turn speeds up the thermalisation. The goal of our evaporation ramp is to get the speed up in cooling before we run out of atoms, to reach the temperatures necessary for BEC. A rate equation approach to modelling the system can be found in Appendix A.

In practice when carrying out evaporation we measure the change in N and T as we change the position of the evaporation ramp over time. Here the equations in Appendix A provide a target for

[4]For a trap in which all the atoms above a certain energy are removed true thermal equilibrium is only reached when the trap empties. Here we assume the timescales on which we operate are short enough this equilibrium isn't reached and thermal distribution is approximated by a truncated Boltzmann distribution.

us to aim for, although significant experimental optimisation is required to reach BEC.

2.4 Magnetic Trapping

After having been trapped and cooled in a MOT and a round of polarisation gradient cooling the atoms are optically pumped to the $F = 2$, $m_F = 2$ hyperfine state. This maximally stretched state is minimal field seeking, a requirement for static trapping. For sufficiently slow particles with a magnetic dipole moment, as the particle moves through the magnetic field its spin adiabatically follows the direction of the field. In this case, the energy of the particle doesn't depend on the direction of the magnetic field, but only its magnitude[52]. The Zeeman effect gives an energy shift of $\Delta E = \mu_B g_F m_F |\mathbf{B}|$.

For the adiabatic condition to hold, the rate of Larmor precession has to be much faster than the rate at which the magnetic field changes. For rubidium 87 in the $m_F = 2$ state the Larmor precession frequency is of the order 1.4MHz/Guass. For a field at 1Guass this suggests we need a velocity much smaller than 140m/s for this condition to hold - far smaller than the speed of our cooled atoms.

There are however two problems: the magnetic dipole of an atom is weak, it requires field gradients of 15G/cm to support the atoms against gravity and the atom should not find itself in zero field. The finite depth of the trap means we need to precool the atoms - even the large field gradients we produce close to our atom chip have a finite depth, and atoms travelling too fast will not be trapped by the conservative potential of the magnetic trap. The second point is because the Larmor precession rate is proportional to the magnetic field strength - if the field ever becomes small enough the adiabatic condition cannot be met and the atoms can transition to untrapped states.

To avoid having a field minimum, we use a Ioffe-Prichard trap. Along two axes a quadrupole field is set up of the form

$$B(x, y, z) = B_1 \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix} . \tag{2.20}$$

On top of this field we apply a field that is harmonic along the z axis, and a constant background field. The result is the full magnetic field

$$B(x, y, z) = B_1 \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix} + B_2 \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{x^2+y^2}{2} \end{pmatrix} + B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} , \tag{2.21}$$

where the final terms in the equation are required by Maxwell's equations to make the divergence and curl equal to zero. The magnitude of the magnetic field is then given by

$$|B(x, y, z)| = \sqrt{(B_1 y + B_2 xz)^2 + (B_1 x - B_2 yz)^2 + (B_0 + B_2(z^2 - \frac{x^2 + y^2}{2}))^2} . \tag{2.22}$$

The magnitude of the field has cylindrical symmetry. We work in a regime where the trapping in the z direction is weak compared to the radial trapping, $B_1 >> B_2$. In this case, for small displacements around the origin, the above equation simplifies to

$$|B(x, y, z)| \approx B_0 + \frac{B_1^2}{2B_0}(x^2 + y^2) + \frac{B_z}{2} z^2 , \tag{2.23}$$

which is a harmonic trap. However, for large values of x or y the trap can start become linear in x or

y, similar to the quadrupole that forms it. Details of how we form this trap are deferred to the next chapter.

2.5 Radio Frequency Radiation

Another tool we have to control the atoms is radio frequency (RF) radiation. The magnetic component of this radiation couples the hyperfine Zeeman sublevel of the atom. Restricting our attention to only the Zeeman sublevels of the $F = 2$ hyperfine state, the Hamiltonian of the system is given by

$$\begin{aligned} H &= \mu_B g_F \hat{\mathbf{F}} \cdot \mathbf{B} \\ &= \mu_B g_F \left(\hat{\mathbf{F}} \cdot \mathbf{B}_{static} + \hat{\mathbf{F}} \cdot \mathbf{B}_{RF} \right) , \end{aligned} \tag{2.24}$$

where F is the total angular momentum quantum number of the system, $\mathbf{B}_{static}$ is the static magnetic field, $\mathbf{B}_{RF}$ is the (time-varying) magnetic field due to the RF radiation, μ_B is the Bohr Magneton and g_F is the Landè g-factor for the level. If we work in a system where the quantisation axis is taken to be the direction of the local static magnetic field, and take the radiation to be of frequency ω_{RF} the Hamiltonian becomes

$$H = \mu_B g_F \left(F_z B_{static} + \left(F_z B_z + F_y B_y + F_x B_x \right) \cos(\omega_{RF} t) \right) , \tag{2.25}$$

where F_i is the angular momentum operator along axis i.

To proceed, we can now move into a rotating frame, rotating about the quantisation axis of the atom with frequency ω_0. This is the equivalent of applying the unitary transformation $U = \exp\left(-i\hat{F}_z \omega t\right)$ to the system. Using the identity $H' = U^\dagger H U - i\hbar U^\dagger \frac{\partial}{\partial t} U$ to find the Hamiltonian that acts on the transformed system we get

$$\begin{aligned} H' = \mu_B g_F \Bigg(& F_z (B_{static} - \frac{\hbar\omega_0}{\mu_B g_F}) + \\ & \left(F_z B_z + \left(F_x \cos(\omega_0 t) + F_y \sin(\omega_0 t) \right) B_x + \right. \\ & \left. \left(-F_x \sin(\omega_0 t) + F_y \cos(\omega_0 t) \right) B_y \right) \cos(\omega_{RF} t) \Bigg) . \end{aligned}$$

The above can be simplified by making the rotating wave approximation. If we set $\omega_0 = \omega_{RF}$ several of the trigonometric terms simplify to give constants and several combine to give terms oscillating at twice the RF. The rotating wave approximation amounts to neglecting those terms that oscillate at twice the frequency, assuming they average out to zero. We also ignore the terms oscillating at the RF frequency, which are those parallel to the static magnetic field - in our experiment they will be small relative to the other fields. The result of these approximations is

$$H'_{RWA} = \mu_B g_F \left(F_z (B_{static} - \frac{\hbar\omega_{RF}}{\mu_B g_F}) + \frac{1}{2} \left(B_x F_x + B_y F_y \right) \right) , \tag{2.26}$$

which can be written succinctly as

$$H'_{RWA} = \hbar \Delta F_z + \frac{\hbar}{2} \left(F_+ \Omega + F_- \Omega^* \right) , \tag{2.27}$$

where $\Delta = \left(\frac{\mu_B g_F B_{static}}{\hbar} - \omega_{RF}\right)$, $\Omega = \frac{\mu_B g_F}{2}(B_x - iB_y)$ and $F_\pm = F_x \pm iF_y$ are the raising and lowering operators.

In this rotating wave picture the Hamiltonian can be diagonalised to get a new set of states with eigenvalues

$$E_m = m\hbar\sqrt{\Delta^2 + \Omega^2}\,, \tag{2.28}$$

where m is an integer that runs over the same range as the projected total angular moment of the atom.

The result is similar to that found in a coupled two level system, but extended to a multi level system. As with a coupled two level system we see the population of the atom oscillate between the Zeeman sublevels with a frequency proportional to $\sqrt{\Delta^2 + \Omega^2}$, and the minimum population of the initial state is given by $\frac{\Delta^4}{(\Delta^2+\Omega^2)^4}$ (assuming that the atom begins in one of the extreme states).

Directly applying single frequency radiation at a fixed power to an atom in free space allows us to drive transitions between the levels - the populations of these levels can then be measured using a Stern-Gerlach type set up. Applying a magnetic field gradient causes a differential acceleration of the hyperfine states, separating them in space and allowing then to be imaged and counted.

When combined with non uniform magnetic fields the application of RF can open up a number of possibilities. As the term Δ depends on the strength of the magnetic field, it can be arranged that the RF radiation only interacts with atoms in a certain spatial region. This is the idea behind forced evaporative cooling.

A related idea is that of creating RF adiabatic potentials. By turning up the RF power slowly, the atoms don't oscillate between hyperfine states, but adiabatically evolve into an eigenstate of the combined Hamiltonian 2.27. In this dressed state picture the atoms no longer just see the bare potential, they see a potential given by 2.28. Complex potentials can be developed as Δ depends on the magnitude of the magnetic field and Ω depends on both the magnitude of the RF radiation and the relative polarisation of the radiation and the static field. The initial idea for using RF adiabatic potentials in this way was developed in [53] and has been used to produce a double well geometry for splitting BECs [54]. This technique has also been extended to include time averaged adiabatic potentials [55].

2.6 Imaging Techniques

After all is said and done, there is only one way we get information about our atoms out of the experiment, and that is by taking a picture of them. There are two main ways to do that: fluorescence imaging and absorption imaging.

2.6.1 Fluorescence Imaging

During the MOT and polarisation gradient cooling stage of the experiments the atoms scatter photons. By imaging these photons onto a CCD camera we gain information about the number of atoms and their spatial extent. To determine the number of atoms we measure the scattered intensity and compare it with the scattering rate per atom, given by 2.1 multiplied by the spontaneous decay rate.

Because each photon is scattered in a random direction we only collect a fraction of these photons, given by the solid angle fraction that our imaging system can collect - in our case about 0.2%. To

collect even one photon per atom requires many scattering events, over which time the atom has undergone a random walk, 'blurring' their positions and heating them. During loading of the MOT this is not a problem because the MOT cools them. When the atoms are in a BEC, the blurring of the image means that we do not gain as much information as we would like - it is at this point we use absorption imaging. It is however worth noting that techniques have been developed to get information from ultracold atoms using fluorescence imaging [56].

2.6.2 Absorption Imaging

The idea behind absorption imaging is that when illuminated with a beam of on resonant laser light, the atoms scatter photons in random directions, leaving a shadow on the beam. By comparing the intensity distribution with and without atoms, we gain information about the density of the atoms.

When the laser light is weak (compared to the saturation intensity of the atoms) the resulting intensity is given by Beer's law,

$$I(x, y) = I_0 \exp(-\sigma n(x, y)) \,, \tag{2.29}$$

where $n(x, y)$ is the column density of atoms and σ is the absorption cross section of the transition.

For higher intensities σ is reduced because of the response of the atoms saturate. By taking the logarithm of $\frac{I_0}{I(x,y)}$ we obtain the two-dimensional column density of the atoms. This technique has the advantage being insensitive to slow drifts in beam intensity. More details on absorption imaging near an atom chip can be found in [57].

We typically use an imaging beam of intensity 0.25mW/cm^2, compared with the 1.67mW/cm^2 saturation intensity of atoms in the $m_F = 2$ ground state driven by σ^+ light.

The quantum efficiency of our camera is 0.18 at 780nm, meaning that over 100μs, each pixel (size 3.45μm by 3.45μm) receives an average of 2000 photons when no atoms are present. For Poissonian statistics, this puts the shot noise at ± 45 counts per pixel. Propagating this error in the standard way we obtain a fundamental noise level for atom counts of ± 1.3 atoms. During imaging of the BEC, typical densities of our atom cloud are of the order 10^{17} atoms/m^3, meaning that we typically image four atoms per pixel.

2.6.3 Post Processing Images

While photon shot noise places a fundamental limit on the accuracy of our imaging, often the images are limited by other issues. Due to the large number of optical elements, fringes in intensity can arise on the profile of the imaging beam. Because of the way absorption images are processed, this would not be a problem if the fringes were stationary, however wavelength scale vibrations of optical elements in the time between taking the absorption image and the background image can cause these fringes can move, creating unwanted noise.

One way around this is to use the fact that while these fringes move in time, they look the same most of the time. This means that if we look at a lot of background images taken over a short timescale, the chances are that one of those images is going to have the same fringe pattern as our absorption image, allowing us to remove the fringes. This idea is formalised in [58] and is known as the Eigenface method.

Whenever we 'take' an absorption image we really capture two images: the absorption image, A and the background image B, usually taken 50ms later with the magnetic fields turned off but

otherwise under the same conditions. In every sequence of images we end up with a number of background images all taken under similar conditions. To remove any fringe patterns present in the images, we look for a linear combination of these background images that corresponds as closely as possible to an area in the absorption image where no atoms are present. To do this we minimise the quantity

$$\sum_i \left(A_i - \sum_j c_j B_{j,i} \right)^2 , \tag{2.30}$$

where A_i is the value of a pixel i in absorption image A and the sum runs over an area where no atoms are present. $B_{j,i}$ is the value of a pixel i in background image j and c_j are the linear coefficients we hope to find.

Differentiating with respect to the linear coefficients and solving to find the values that minimise the equation above, after a few lines of algebra we get

$$\underline{c} = (\underline{\underline{B}}\,\underline{\underline{B^T}})^{-1}\underline{\underline{B}}\,\underline{A}\,, \tag{2.31}$$

where A and B have been written as a vector and a matrix.

An example of two absorption images, one cleaned and one uncleaned can be found in Figure 2.4. It is worth noting that the entire region of this image corresponds to the 'foreground' image - the part of the image we are trying to clean. The background image used for the comparison comes from a region of similar size, but offset from the image shown in Figure 2.4.

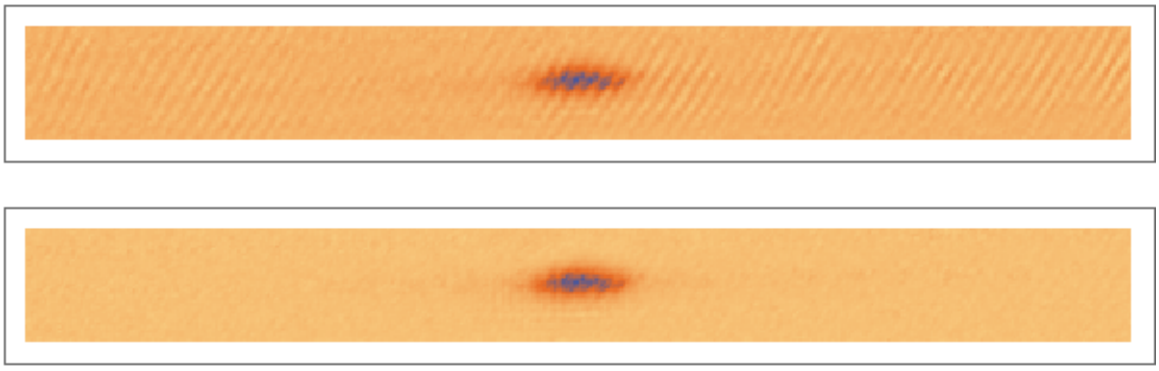

Figure 2.4: A comparison of two absorption images of a BEC. The top image has just been processed normally, the lower image has had the eigenface technique applied to it to remove noise due to optical fringes.

If we look at a region where no atoms are present and count the number of 'atoms' detected per pixel we get a measure of the error in our detection system. A histogram of both the cleaned and uncleaned image pixel counts, taken over many images and several runs can be found in Figure 2.5.

The standard deviations of the distributions are 1.8 atoms/pixel for the uncleaned image, and 1.4 atoms/pixel for the cleaned images. In what follows the value 1.4 atoms will be used as our error per pixel for absorption imaging.

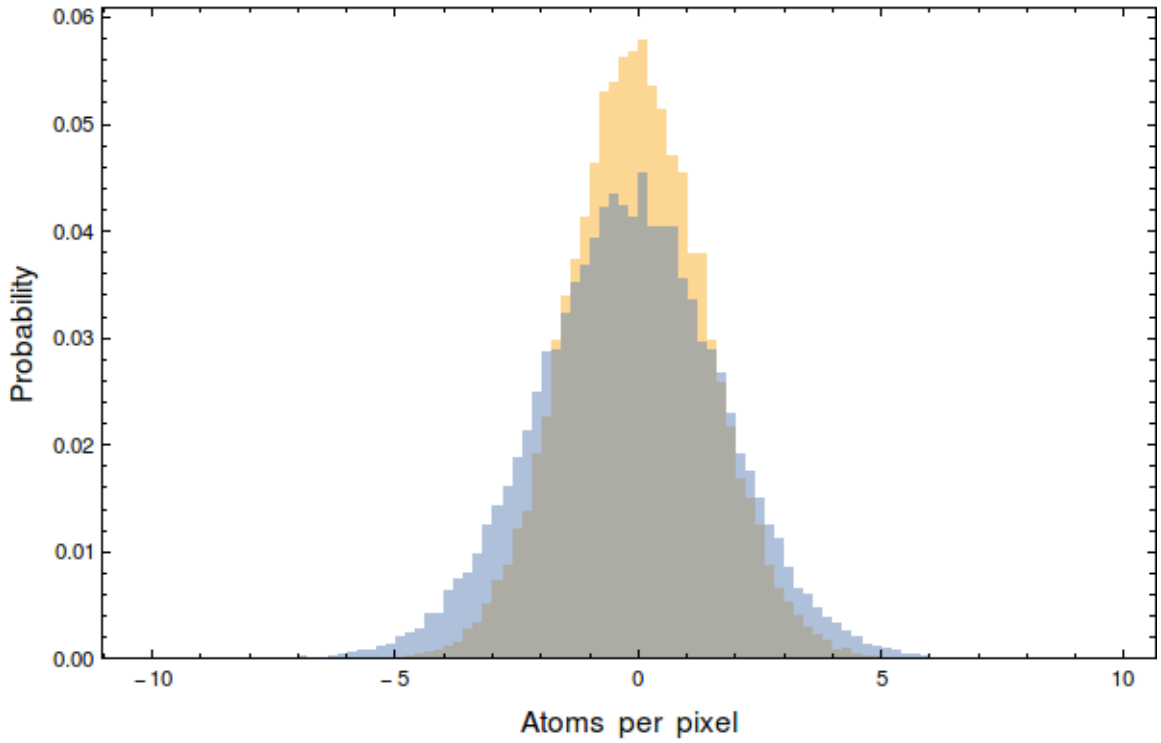

Figure 2.5: A histogram of background pixel counts taken from many absorption images. The blue bars show the counts from uncleaned images, the orange bars show the counts from the cleaned images.

2.7 Bose Einstein Condensation

In this section we describe what a Bose Einstein Condensate is, and the maths used to describe it. None of the working here is my own, only an interpretation of references [14] and [15].

2.7.1 Atom Bunching

When indistinguishable particles of integer spin (Bosons) are cooled, they have a habit of 'bunching' together. Such bunching is a result of the boson statistics - symmetry requirements mean that the wavefunction of multi-particle states have to be symmetric under interchange of any two indistinguishable Bosons. To demonstrate this difference, consider two coins that can either show heads (H) or tails (T). Where the coins are distinguishable there are four possible combinations, HH,HT, TH and HH. In the absence of any other information, we assign equal probability to each outcome, giving probability 0.5 of finding the coins in that same state. Here distinguishably of the coins means that the outcomes HT and TH are measurably different, and therefore separate outcomes. For indistinguishable coins, there are only three possible outcomes: HH, TT and HT. Each of these has equal probability of occurring, and we have probability 2/3 of the coins being in the same state. Without interacting, the coins have bunched together in their outcome due to symmetry requirements.

In systems where the occupancy of a given energy level in much less than one - the typical case for an atomic gas at room temperature - the distribution of particles among energy levels is close to that expected of a classical, distinguishable gas, the Maxwell Boltzmann distribution. However, as we cool the system the occupancy will increase and quantum statistics become increasingly relevant. For Bosons, the distribution of atoms among energy levels obey the Bose-Einstein distribution:

$$P(E_i) \propto \frac{g(E_i)}{\exp \frac{E_i - \mu}{kT} - 1}\,, \tag{2.32}$$

where $P(E_i)$ is the probability of finding the system in a state with energy E_i, $g(E_i)$ is the density of states of energy level E_i, μ is the chemical potential, k is Boltzmann's constant and T is the temperature.

To derive this distribution we follow the route of maximising the number of microstates with a fixed total energy and particle number, subject to the condition of indistinguishability.

If we continue to lower temperature further still this 'bunching' manifests itself as the macroscopic occupation of the ground state of the system. If we consider the occupation of ground state as an order parameter, this macroscopic occupation of the ground state of a system represents a phase transition.

The transition to a macroscopically occupied state is known as Bose-Einstein Condensation (BEC). For dilute atomic gases, this has been happened since 1995 [1] [2].

What makes this state interesting is the fact that it acts as a single quantum object: many atoms occupy the same quantum state. The ability to experimentally access such a state opens up a path to test predictions about many body quantum mechanics.

2.7.2 Transition Temperature

To find the temperature below which atoms begin to occupy the ground state, we consider the role the chemical potential plays in the Bose Einstein distribution. From the form of the distribution 2.32, if we take the energy of the lowest state to be zero the chemical potential must be less than zero - a positive chemical potential would result in a negative probability of occupation, which is non-physical. Also from this equation we note that the more negative the chemical potential, the lower the probability of any given state being occupied. A consequence of this is the fact that at a given temperature, there is a maximum number of atoms that can occupy the excited states and this occurs for a chemical potential of zero.

If we consider the spectrum of excitations to be continuous except for the ground state of the system, we can write the maximum number of particles in the excited states as

$$N_{EX} = \int_0^\infty g(\epsilon) \frac{1}{e^{\frac{\epsilon}{k_B T}} - 1}\, d\epsilon\,, \tag{2.33}$$

where $g(\epsilon) = \frac{\epsilon^2}{2\hbar^3 \omega_x \omega_y \omega_z}$ is the density of state for a harmonic trap. We can evaluate this integral using the identity

$$\int_0^\infty \frac{x^{\alpha-1}}{e^x - 1}\, dx = \Gamma(\alpha)\zeta(\alpha)\,, \tag{2.34}$$

where Γ is the gamma function and ζ is the Riemann zeta function.

For our three-dimensional harmonic trap this evaluated to about 2.4. The temperature below which the ground state starts to become occupied is then given by

$$k_B T_C \approx 0.94 N^{1/3} \hbar \bar{\omega}\,, \tag{2.35}$$

where $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$.

To get an idea for the numbers, a gas of 20000 atoms confined in a potential of $\omega_x = \omega_y =$

$2\pi\times1.4\times10^3 Hz$, $\omega_z = 2\pi\times10 Hz$ (typical parameters for our experiment) has a transition temperature of roughly 310nK. We can also use the above equation to calculate how the number of condensed atoms changes as a function of temperature. Below the critical temperature we can set the chemical potential to zero. The number of atoms in the excited state is proportional to T^3, meaning the number in the condensed fraction is given by

$$N_C = N\left(1 - \left(\frac{T}{T_C}\right)^3\right). \tag{2.36}$$

Another measure of condensation can be given by phase space density. Condensation requires a macroscopic number of particles to occupy a single quantum state - this can happen if you increase the density of particles without increasing the temperature, or decease the temperature without changing the atom number. A useful metric that combines these values into one number is phase space density. For a thermal system it is defined as the number of particles inside a cube with side length one thermal de Broglie wavelength, $\lambda_B = \sqrt{\frac{2\pi\hbar^2}{mk_BT}}$. If we define the volume of the trap as a cube of the harmonic oscillator lengths $\sqrt{\frac{2k_BT}{m}}\frac{1}{\omega_i}$, the phase space density of the trap is

$$\rho = \frac{N\bar{\omega}^3\hbar^3\pi^{3/2}}{(k_BT)^3}. \tag{2.37}$$

Using the transition temperature found above, we find that condensation occurs for phase space densities greater than $\rho_C \approx 3.7$. This provides a useful figure of merit when developing a route to condensation as number, temperature and trapping frequencies are readily deduced from absorption images of the condensate.

In deriving the transition temperature for condensation it may not seem clear what the difference between bosons and distinguishable particles is. If we follow the above route naively we can arrive at a transition temperature for distinguishable particles, after which a macroscopic occupation of the ground state seems to occur. The difference lies in our treatment of the chemical potential, μ.

In both the distinguishable and indistinguishable particle case, μ arises from requirement to conserve total particle number. For bosons, this number is bounded from below by zero due to the nature of the distribution, for distinguishable particles this is not the case, meaning the assumption we made above - that the chemical potential is zero, is not valid. Following this line of reasoning we can provide some justification for treating the chemical potential as zero when calculating the critical temperature: If N_0 particles occupy the ground state of the system, the chemical potential for bosons is given by

$$\frac{\mu}{k_BT} = \log\left(1 + \frac{1}{N_0}\right) \approx \frac{1}{N_0}. \tag{2.38}$$

For any significant population of the ground state, the chemical potential is close to zero. This is a consequence of the fact that for non interacting distinguishable particles their distribution is independent of each other - whereas for indistinguishable boson, even in the non interacting case their distribution depends on the states of other particles.

All of this only applies for non-interacting particles. In our experiment interactions mean that we have a chemical potential of around 1-2kHz, but the above reasoning can still be used as a guide.

2.7.3 The State at the Bottom

In the previous section, we discussed how a BEC can exist as the ground state of a Bosonic system. In the absence of interactions this ground state will just be the product state of single particle ground state wavefunctions. When interactions are included we need to work a little harder. A more detailed approach can be found in [14].

In order to find the ground state wavefunction of many interacting atoms, we have to minimise the energy of such a state. To do this we take a mean field approach and assume that we can treat the ground state wavefunction as a product of single particle wavefunctions:

$$\phi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_n) = \prod_i \psi(\mathbf{r}_i)\,. \tag{2.39}$$

The Hamiltonian is given by

$$H = \sum_i \left(\mathbf{p}_i^2/2m + V(\mathbf{r}_i)\right) + \sum_{i,j<i} U(\mathbf{r}_i - \mathbf{r}_j)\,, \tag{2.40}$$

where $\mathbf{r}_i$ are the position operators of the atoms, $\mathbf{p}_i$ are their momentum operators and V and U are the external potential and inter-atom potential respectively. The sums run over the number of atoms.

For an atomic system at low temperatures we can treat the interaction term as a delta function, $U_0\delta(\mathbf{r}_i - \mathbf{r}_j)$, due to larger momentum scattering paths being 'frozen out' [14]. The energy of the system then becomes

$$E = N\int d^3\mathbf{r}\left(-\frac{\hbar^2}{2m}\psi^*(\mathbf{r})\nabla^2\psi(\mathbf{r}) + V(\mathbf{r})|\psi(\mathbf{r})|^2 + \frac{(N-1)}{2}U_0|\psi(\mathbf{r})|^4\right), \tag{2.41}$$

where we have used the assumption that the wavefunction for each atom is identical.

When we minimise this energy, the total number of atoms has to be kept constant. To do this we introduce a Lagrange multiplier μ, and minimise $E - \mu N$. The result is

$$-\frac{\hbar^2}{2m}\nabla^2\psi(\mathbf{r}) + V(\mathbf{r})\psi(\mathbf{r}) + NU_0|\psi(\mathbf{r})|^2\psi(\mathbf{r}) = \mu\psi(\mathbf{r})\,, \tag{2.42}$$

We have also taken $N \approx N - 1$. The equation is known as the Gross-Pitaevskii equation.

The Gross-Pitaevskii equation allows us to calculate the ground state for the system, however it is possible to make one more approximation. Before making this approximation, let us compare the size of the energies from each term in equation 2.41.

For a single particle, in a harmonic potential $V = \frac{m\omega^2 r^2}{2}$, the potential energy is roughly $\frac{\hbar m\omega^2 R^2}{2}$ where R is average spatial extent of the cloud. The kinetic energy is roughly $\frac{\hbar^2}{2mR^2}$, as by Heisenberg's uncertainty principle the momentum is roughly $\frac{\hbar}{R}$ for a particle confined to distance R. The interaction energy is roughly $\frac{NU_0}{R^3}$.

In the absence of interactions the kinetic and potential energy would be equal. However, once we include interactions, which in the case of rubidium are repulsive, the particles become more spread out. The effect is to increase the potential energy while reducing the kinetic energy. In our crude approximations, and taking the typical length scale as the harmonic oscillator length $R = \sqrt{\frac{\hbar}{m\omega}}$, the ratio of interaction energy to kinetic energy is $\frac{2NU_0 m^{3/2}\omega^{1/2}}{\hbar^{5/2}}$. Putting in typical values for our condensate, ($\bar{\omega} = 2\pi 200$Hz, $N = 40000$), we get 10^4, showing the interaction energy dominates the kinetic energy.

The small kinetic energy allows us to neglect it, an approximation know as the Thomas Fermi approximation. This makes the Gross-Pitaevskii equation simple to solve:

$$N|\psi(\mathbf{r})|^2 = \mathrm{Max}\left(0, \frac{\mu - V(\mathbf{r})}{U_0}\right). \tag{2.43}$$

The term of the left hand side is just the density profile. The term on the right hand side is the chemical potential minus the trapping potential, showing that in a harmonic trap the condensate takes on the form of a unturned parabola. Without the interaction term the density profile would be a Gaussian. The appearance of a parabolic density profile from a thermal Gaussian one is a clear experimental sign of Bose Einstein Condensation in a dilute atomic gas.

Given that there is a fixed number of atoms in the condensate, integrating the above equation over all space we will get the total number of atoms. This can then be inverted to get the chemical potential of the condensate

$$\mu = \left(\frac{15N\bar{\omega}U_0}{8\pi}\right)^{2/5}\left(\frac{m}{2}\right)^{3/5}. \tag{2.44}$$

Using the fact that the chemical potential $\mu = \frac{\partial E}{\partial N}$ we can use the above equation to write the total energy of the condensate as

$$E = \frac{5}{7}N\mu\,. \tag{2.45}$$

So the energy per particle in the condensate is of order μ. For our experiment we find that $\mu \approx 2\hbar\omega_r$.

2.7.4 Basic Condensate Dynamics

In the previous section we described what the ground state of a system looked like. In this section we investigate its time dependence.

2.7.4.1 The Bogoliubov Approximation

In the previous section, we treated the wavefunction as just a complex field, but if we want to take into account quantum fluctuations, we need to describe the system in terms of bosonic creation and annihilation operators. In this form, the Hamiltonian takes the from

$$H = N\int d\mathbf{r}\left(-\hat{\psi}^\dagger(\mathbf{r})\frac{\hbar^2}{2m}\nabla^2\hat{\psi}(\mathbf{r}) + \hat{\psi}^\dagger(\mathbf{r})\hat{\psi}(\mathbf{r})V(\mathbf{r}) + \hat{\psi}^\dagger(\mathbf{r})\hat{\psi}^\dagger(\mathbf{r})\hat{\psi}(\mathbf{r})\hat{\psi}(\mathbf{r})\frac{U_0N}{2}\right), \tag{2.46}$$

where $\hat{\psi}(\mathbf{r})$ and $\hat{\psi}^\dagger(\mathbf{r})$ are the creation and annihilation operators for a boson at position $\mathbf{r}$ respectively.

We want to know what happens to a BEC perturbed around its ground state. To do this we introduce the Bogoliubov approximation: we treat the ground state as a classical field $\psi(\mathbf{r})$ and fluctuations of this as operators corresponding to small perturbations. This amounts to taking the number of atoms in the ground state as being large enough that removing one doesn't affect the properties of the system, similar to a coherent state in quantum optics.

$$\hat{\psi}(\mathbf{r}) = \psi(\mathbf{r}) + \delta\hat{\psi}(\mathbf{r})\,. \tag{2.47}$$

For more details of this description, see [15]. Substituting this into the Hamiltonian of the system

and minimising the energy subject to constant atom number (equivalent to minimising $E-\mu N$) , the solution for the ground state wavefunction is just the solution to the Gross Pitaevski equation 2.42, and by definition the contributions from terms linear in $\delta\hat{\psi}$ are zero. The remaining terms are

$$\begin{aligned} H-\mu N =& E_0-\mu N+N\int d\mathbf{r}\left(-\delta\hat{\psi}^\dagger(\mathbf{r})\frac{\hbar^2}{2m}\nabla^2\delta\hat{\psi}(\mathbf{r})\right. \\ &+\left[V(\mathbf{r})+2U_0|\psi(\mathbf{r})|^2-\mu\right]\delta\hat{\psi}^\dagger(\mathbf{r})\delta\hat{\psi}(\mathbf{r}) \\ &+\frac{U_0}{2}\left(\psi(\mathbf{r})^2(\delta\hat{\psi}^\dagger)^2+\psi^*(\mathbf{r})^2(\delta\hat{\psi})^2\right) \\ &+\frac{U_0}{2}\left(2\psi(\mathbf{r})(\delta\hat{\psi}^\dagger)^2\delta\hat{\psi}(\mathbf{r})+2\psi^*(\mathbf{r})(\delta\hat{\psi})^2\delta\hat{\psi}^\dagger(\mathbf{r})\right. \\ &\left.\left.+\delta\hat{\psi}^\dagger(\mathbf{r})\delta\hat{\psi}^\dagger(\mathbf{r})\delta\hat{\psi}(\mathbf{r})\delta\hat{\psi}(\mathbf{r})\right)\right) , \end{aligned} \tag{2.48}$$

where E_0 is the energy of the ground state given by the solution to the Gross-Pitaevskii equation. We are left with terms quadratic in the fluctuation operators, and higher order terms. The higher order terms on the final two lines correspond to interactions directly between excitations. These will be ignored for now, but we will return to them later. Keeping only the terms quadratic in the fluctuation operators leaves a Hamiltonian that does not commute with the fluctuation operators $\delta\psi$ and $\delta\psi^\dagger$, so these fluctuations are not constant in time.

The next step is to try and find operators that diagonalize the Hamiltonian. To proceed we find the time dependence of the fluctuations, using the following

$$i\hbar\frac{\partial\delta\hat{\psi}}{\partial t}=\left[\delta\hat{\psi},H\right] \tag{2.49}$$

$$i\hbar\frac{\partial\delta\hat{\psi}^\dagger}{\partial t}=\left[\delta\hat{\psi}^\dagger,H\right] . \tag{2.50}$$

Using the Bosonic commutation relations, $\left[\delta\hat{\psi}(\mathbf{r}),\delta\hat{\psi}^\dagger(\mathbf{r}')\right]=\delta(\mathbf{r}-\mathbf{r}')$, $\left[\delta\hat{\psi}(\mathbf{r}),\delta\hat{\psi}(\mathbf{r}')\right]=0$ we get

$$i\hbar\frac{\partial\delta\hat{\psi}}{\partial t}=\left(-\frac{\hbar^2}{2m}\nabla^2+V(\mathbf{r})+2U_0|\psi(\mathbf{r})|^2-\mu\right)\delta\hat{\psi}+U_0\psi(\mathbf{r})^2\delta\hat{\psi}^\dagger \tag{2.51}$$

$$-i\hbar\frac{\partial\delta\hat{\psi}^\dagger}{\partial t}=\left(-\frac{\hbar^2}{2m}\nabla^2+V(\mathbf{r})+2U_0|\psi(\mathbf{r})|^2-\mu\right)\delta\hat{\psi}^\dagger+U_0\psi^*(\mathbf{r})^2\delta\hat{\psi} . \tag{2.52}$$

Comparing the previous equations with equation 2.48, keeping terms only second order in the fluctuations, we can recast the remaining part of equation 2.48 in the form

$$H_2=N\frac{i\hbar}{2}\int d\mathbf{r}\left(\delta\hat{\psi}^\dagger(\mathbf{r},t)\frac{\partial\delta\hat{\psi}(\mathbf{r},t)}{\partial t}-\frac{\partial\delta\hat{\psi}^\dagger(\mathbf{r},t)}{\partial t}\delta\psi(\mathbf{r},t)\right) , \tag{2.53}$$

where H_2 has been used to label the part of equation 2.48 that is quadratic in the fluctuation operators.

We can now make the transformation

$$\delta\hat{\psi}(\mathbf{r},t)=\sum_j\left(u_j(\mathbf{r})\hat{a}_je^{-i\epsilon_jt/\hbar}-v_j^*(\mathbf{r})\hat{a}_j^\dagger e^{+i\epsilon_jt/\hbar}\right) , \tag{2.54}$$

where $\hat{a}_j$ and $\hat{a}_j^\dagger$ are the creation and annihilation operators for bosonic mode j, and u and v are

functions to be determined.

After some manipulation, we get the diagonalised Hamiltonian

$$H_2 = \sum_j \epsilon_j \hat{a}^\dagger \hat{a}\,, \tag{2.55}$$

where we have made use of the fact that in order for the transformation 2.54 to obey the bosonic commutation relations, u and v must obey the orthogonality relation

$$\int d\mathbf{r}\left(u_i u_j^* - v_i^* v_j\right) = \delta_{ij}\,. \tag{2.56}$$

The preceding might not seems as if it has improved out understanding much, but what we've done is transform the description of the condensate from one where we looked only at the mean atom wavefunction ψ, to that one where we can consider the excitations of the state.

To find the u and v, we substitute the transformation 2.54 into equations 2.51 and collect terms of each mode together. The result is a set of coupled equations, known as Bogoliubov's equations

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + 2U_0 n(\mathbf{r}) - \mu - \epsilon_i\right) u_i(\mathbf{r}) - U_0 n(\mathbf{r}) v_i(\mathbf{r}) = 0 \tag{2.57}$$

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + 2U_0 n(\mathbf{r}) - \mu + \epsilon_i\right) v_i(\mathbf{r}) - U_0 n(\mathbf{r}) u_i(\mathbf{r}) = 0\,, \tag{2.58}$$

where the ground state wavefunction has been assumed to be real, been replaced with the ground state density $n(\mathbf{r})$. Solving these equations for a trapped condensate is far from trivial. Instead we will try and gain some understanding by looking at a uniform Bose gas.

2.7.4.2 A Uniform Gas

The advantage of thinking about a uniform Bose gas rather than a trapped one is that it makes the problem translationally invariant. Noting this we write u and v as momentum eigenstates

$$u(\mathbf{r}) = u_q \exp(i\mathbf{q}\cdot\mathbf{r}) \tag{2.59}$$

$$v(\mathbf{r}) = v_q \exp(i\mathbf{q}\cdot\mathbf{r}) \tag{2.60}$$

which gives the Bogoliubov equations the simple form

$$\left(\frac{\hbar^2 q^2}{2m} + 2U_0 n - \mu - \epsilon_q\right) u_q = U_0 n v_q \tag{2.61}$$

$$\left(\frac{\hbar^2 q^2}{2m} + 2U_0 n - \mu + \epsilon_q\right) v_q = U_0 n u_q\,. \tag{2.62}$$

If we recast the equations in the from of a matrix, the only way we can get non-zero u_q and v_q is if its determinant is zero. As the chemical potential of a uniform system is just nU_0, we get

$$\epsilon_q = \sqrt{\frac{\hbar^2 q^2}{2m}\left(\frac{\hbar^2 q^2}{2m} + 2U_0 n\right)}\,. \tag{2.63}$$

This is the dispersion relationship for a Bose gas. In the long wavelength, small q limit, the energy of these excitations scales linearly with the momentum q, like waves, whereas in the short wavelength limit the energy scales as q^2, like free particles.

Now that we have the dispersion relationship for the gas we can look at another of its interesting properties - superfluidity. A fluid will act like a super-fluid if it is not energetically possible for an object moving through the fluid to create an excitation. With no excitations, there can be no damping.

Consider a condensate of energy E, total momentum $\mathbf{P}$ and total mass M in the lab frame. Performing a Galilean transformation to a frame moving at $\mathbf{v}$ to the lab frame transforms the momentum as

$$\mathbf{P}' = \mathbf{P} - M\mathbf{v}\,. \tag{2.64}$$

In this new frame the energy of the system changes by

$$\Delta E = \frac{|\mathbf{P}'|^2}{2M} - \frac{|\mathbf{P}|^2}{2M} = -\mathbf{v}\cdot\mathbf{P} + \frac{1}{2}M\mathbf{v}^2\,. \tag{2.65}$$

Now imagine a condensate stationary in the lab frame, with an object moving through it with constant velocity $\mathbf{v}$. For the object to dissipate energy to the condensate, it needs to create and excitation. If there are no excitations, in a frame where the object is stationary the total energy of the system is

$$E_0 = E + \frac{1}{2}Mv^2\,. \tag{2.66}$$

Consider the object now creating an excitation with energy ϵ_p and momentum $\mathbf{p}$. The total energy of the condensate system (ignore the moving object) in the lab frame is now $E_0 + \epsilon_p$, where E_0 is the energy of the condensate at rest. Performing a Galilean transformation to a frame where the object is stationary we get a total energy of

$$E_1 = E + \epsilon_p - \mathbf{p}\cdot\mathbf{v} + \frac{1}{2}Mv^2\,. \tag{2.67}$$

Subtracting 2.66 from 2.67 we find that the energy required to create an excitation is $\epsilon_p - \mathbf{p}\cdot\mathbf{v}$. In the object's frame it is stationary and cannot transfer energy. Therefore the minimum velocity required to create an excitation is $\frac{\epsilon_p}{p}$. This is known as Landau's critical velocity.

Combined with the dispersion relationship we previously found, the critical velocity for an object moving through our condensate is given by

$$v_c = \sqrt{\frac{U_0 n}{m}}\,. \tag{2.68}$$

2.7.4.3 Dynamics of a Trapped Gas

To describe small excitations of a trapped BEC about its ground state, we need to solve Bogoliubov's equations, 2.61, for our trapping potential, $V(\mathbf{r}) = \frac{m\omega^2}{2}\left(x^2 + y^2 + \lambda z^2\right)$, a harmonic trap elongated along the z axis. The elongation is described by the parameter λ. This has been done in [59] for low lying excitations. In thermal equilibrium, these states will be populated according to the Bose-Einstein distribution, 2.32. The exact form of these solutions is complex, and will not be discussed here, except to say that solutions exist.

Up to this point, we have ignored interactions between excitations, given by the final two lines in equation 2.48. When studying low lying excitations of BECs, we need to move beyond these approximations - the interactions can redistribute energy between oscillation modes, causing damping of a macroscopically occupied oscillation.

2.7.5 Coherence and Dimensionality in BECs

So far we have considered the condensation in three dimensions - in highly elongated systems this may no longer be the case. Our system can be consider cylindrically symmetric to first order: a harmonic trap with axial trapping frequency ω_z and radial trapping frequency ω_ρ where $\frac{\omega_\rho}{\omega_z}$ is of order 140. For a system to be considered truly one-dimensional, the energy required for excitations in the radial direction needs to be much higher then the average energy of the system. Radial degrees of freedom are then 'frozen out'.

The average energy per particle in our experiment is $\mu \approx 2\hbar\omega_\rho$, which puts us in the regime between one and three dimensions. Given that our axial frequency is much lower than our radial frequency many axial states will be populated.

Because of our elongated geometry, our condensate might still not achieve complete coherence after dropping below the transition temperature, as described in [60] [61]. Thermal excitations of low-lying Bogoliubov states mean that the phase of the ground state can fluctuate along the length of the condensate.

These fluctuations in phase don't affect the density of the condensate in trap, but upon release they do have an effect. In the mean field approximation the velocity of the condensate is described by $\mathbf{v} = \hbar\frac{\nabla\phi}{m}$. If the phase fluctuates along the length of the condensate then after release and some time in free fall this will be manifest as a variation of velocity along the length of the condensate, resulting in density modulations. This behaviour is not expected from a thermal gas and can be used as an indicator of the onset of condensation.

As the phase of the condensate is random and varies from shot to shot, so does the density modulation pattern. Because the fluctuations are caused by thermal excitations, the correlation length of the density modulation pattern after free fall depends on the temperature of the condensate. In [60] it was shown that the mean squared amplitude of the phase between two points z, z' close to the centre of a condensate is given by

$$\left\langle \big(\phi(z) - \phi(z')\big)^2 \right\rangle = \delta_L \frac{|z - z'|}{L}\,, \tag{2.69}$$

where L is the length of the condensate and

$$\delta_L = \frac{32\mu T}{15N(\hbar\omega_z)^2}\,. \tag{2.70}$$

Here δ_L can be considered a measure of the fluctuations over the length of the condensate. When $\delta_L << 1$ the phase is approximately constant over the length of the condensate.

CHAPTER 3

THE EXPERIMENT

3.1 The More Things Change

In this chapter I detail the hardware used for the experiments detailed in this thesis.

Building the hardware to run an experiment of this nature is more than a one man job. Most of the work to build the infrastructure to produce a Bose Einstein condensate had be completed by previous students, before I arrived on the scene. As such, this chapter has substantial overlap with the theses of previous students, Rob Sewell [62], Florian Baumgartner [63] and Ben Yuen [64].

3.2 Overview of the Machine

The hardware required to preform the experiment can be split into several areas: the atom chip 3.2.1, the magnetic fields 3.2.2, the light 3.2.3, the vacuum system 3.2.4, the imaging system 3.2.5 and the computer control 3.2.6. For each of these areas I will give a description of our set up.

3.2.1 The Atom Chip

Atom chips are micro-fabricated devices designed to simplify and allow the compact creation of electrical, magnetic and light fields. By reducing the scale of the devices, higher magnetic field gradients can be produced with less current, and precise positioning of wires on the atom chip means that the relative orientation of the fields can be controlled. For an overview of atoms chips, see [65] [66].

The production and manipulation of our BEC takes place just below the chip. In this section I describe the atom chip used in this experiment.

3.2.1.1 Our Atom Chip

Our atom chip is a micro-fabricated device consisting of current carrying wires etched on a reflective surface. It consists of a 20mm by 20mm square chip of silicon, coated with a 100nm layer of SiO_2 to act as an insulating barrier, then a 50nm layer of chromium to improve adhesion, finally coated with a 3μm layer of gold. The wires have been created by etching through the gold and chromium.

The entire chip is mounted on a substructure, discussed below. Electrical connections are made to the chip by gold wire bonding. For a detailed description of the atom chip's fabrication, see [67].

A photograph of the atom chip can be found in Figure 3.1, along with a schematic of the various wires on the chip. There are four central wires on our chip that run parallel to each other and form a 'Z' shape. The two central wires, called the 'small Z wires' are 50μm wide and have a centre-to-centre

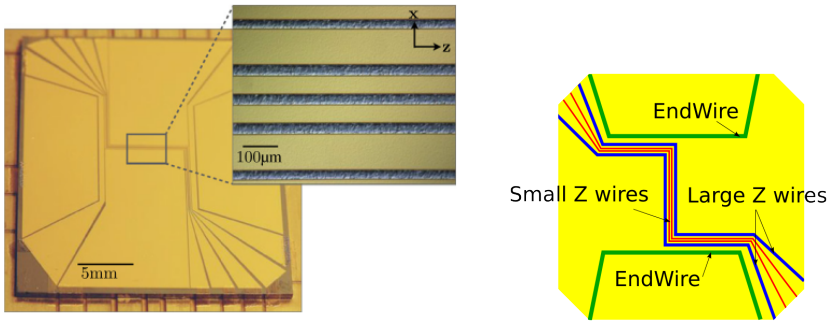

(a) Photo of the atom chip

(b) Schematic of the atom chip

Figure 3.1: Figure 3.1a shows a photo of the atom chip we use, the insert show a zoomed in section of the central part of the 'Z' wires. Figure 3.1b shows an idealised version of the chip, not drawn to scale. The red lines show two central wires on the chip, each $50\mu m$ wide, refered to as the 'small Z wires', the blue lines show two outer wires, each $100\mu m$ wide, refered to as the 'large Z wires' and the green lines show two end wires, refered to in the text as the 'end wires'.

spacing of 85μm. The outer two wires wire of the 'Z', called the 'large Z wires', are 100μm wide each with centre-to-centre of 300μm. Each is etched into the gold surface, with a gap between the wires of width $15\mu m$. As well as the 'Z' wires, we have two 'end wires', each $200\mu m$ wide which run along the top of the 'Z's.

3.2.1.2 Note on Orientation

The atom chip forms a natural basis to define what I will call the 'lab coordinates'. The chip is mounted on the top flange of the main vacuum chamber, with its surface perpendicular to the direction of gravity. We label the direction defined by gravity as the Y axis, the direction parallel to the central parts of the Z wires the Z axis and the direction perpendicular to these wires as the X axis.

3.2.1.3 Chip Substructure

The chip sits directly on a ceramic base plate, as illustrated in Figure 3.2, which provides electrical insulation, but good thermal conduction for the atom chip. The ceramic base plate is mounted on a copper structure that provides electrical connection between the atom chip and wires going to the vacuum feed throughs.

Directly behind the centre of the baseplate is copper in an 'H' shape. This structure is used to apply RF fields to the atoms during forced evaporation. The H is aligned with the central 'Z' shape of the atom chip wires, and sits 1mm below the surface of the chip.

3.2.2 Magnetic Fields

Magnetic fields play a central role in both trapping the atoms and manipulating their internal states. We use a combination of static fields produced by currents in external coils, static fields produced by the atom chip wires, and radio frequency currents on the atom chip and its substructure.

3.2.2.1 Magnetic Field Coils

Outside our vacuum chamber we use six pairs of wound coils to produce bias fields for the experiment. These coils are arranged into roughly Helmholtz configurations to provide uniform fields at the centre

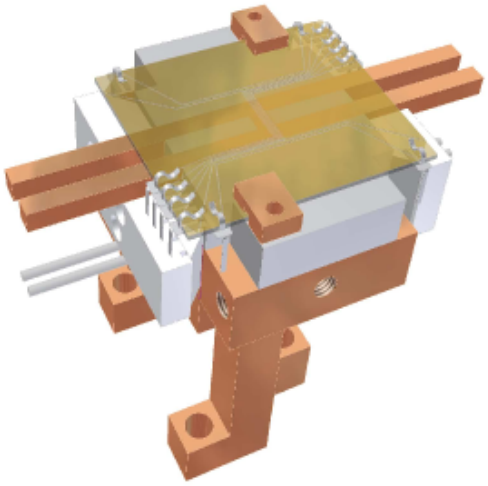

(a) CAD drawing of the atom chip's mount.

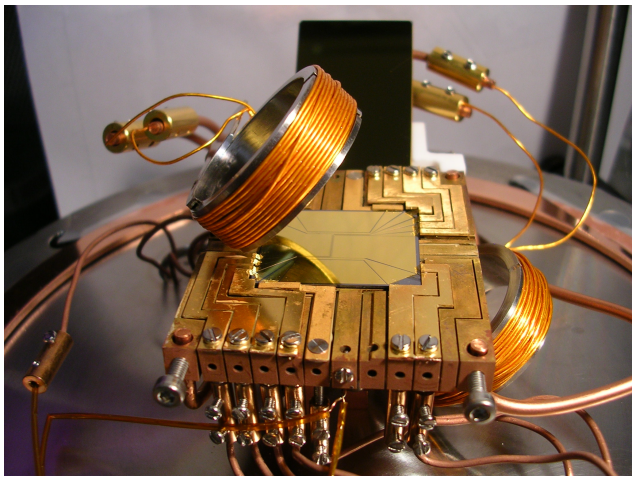

(b) photo of the assembled atom chip mount.

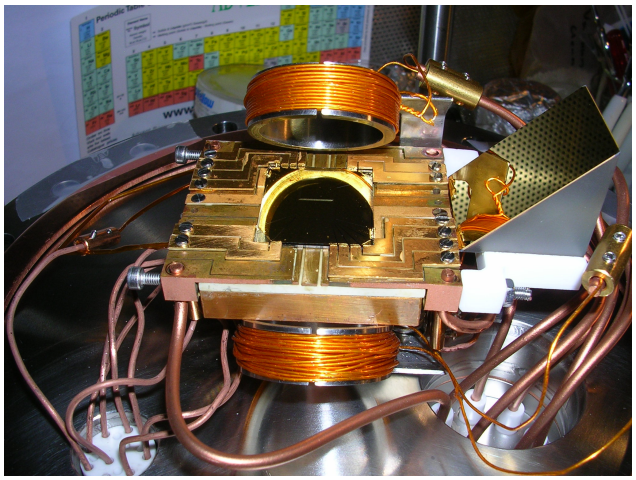

(c) Another Photo of the assembled atom chip mount.

Figure 3.2: Figure 3.2a shows a schematic of the structure used to hold the atom chip. Figure 3.2b and Figure 3.2c show photos of the assembled atom chip's substructure. Figure 3.2a shows an electrical connection between the substructure and the atom chip being made by metal clips. These were used in a previous version of the experiment, but for the work presented here we use bonded gold wires to make the connection.

of the vacuum chamber.

Three pairs of coils are arranged to provide constant magnetic fields for atom manipulation. These pairs are capable of creating magnetic field strengths of 30-40G the centre of the vacuum chamber. In this thesis we will refer to these as the X bias, Y bias and Z bias fields depending on the axis along which their fields are aligned, following the convention in 3.2.1.2.

The Z bias and Y bias pairs of coils are mounted directly on the vacuum chamber. This constrains their alignment to be precise relative to the atom chip. The X bias however is not fixed relative to the vacuum chamber and is held by stands mounted to the optical table. When these coils were set up, they were misaligned and there is a 3 degree tilt causing the field from these coils to have a small component along the Z axis. These coils were set up in the distant past of the experiment and realigning them would mean disassembling the vacuum chamber, something which when the error was discovered was impractical. Instead we work around this fact.

Each pair of coils is wired in series to ensure the same current runs in each coil. The current in the coils is controlled with a set of home built current drivers based around the high current op-amp OPA549. The current is directly controlled by a variable analogue voltage line from the computer. The magnetic field produced by each coil was measured roughly using a magnetometer before the coils were installed, and then improved upon using the atoms themselves.

As well as these three main bias coil pairs we have three pairs of 'cancellation coils', which are designed to provide field strengths of up to 500mG to cancel small scale drifts in magnetic fields over the course of the day. Between experimental runs the ambient magnetic field is measured along each axis using magnetometers mounted on the top of the chamber. The current through the cancellation fields is then adjusted to reduce all three components to zero. The current is then held at this point for the duration of the experiment, after which it is released. It has been found that these cancellation coils stabilise the magnetic field to within 10mG[1].

Finally we have a set of magnetic field coils mounted inside the vacuum chamber. The orientation of these coils can be seen in Figure 3.2b. They are used to produce the quadrupole field we use for making the MOT in the main chamber. These coils are arranged to produce a quadrupole field and the current in them is controlled using a similar home built current driver to the bias coils.

3.2.2.2 Magnetic Fields, the Atom Chip and Magnetic Trapping

The area where the atoms are trapped is in the centre of the atom chip, above the small wires. When far from the chip (relative to the separation of the small wires), it is possible to treat the magnetic field produced as that of a single wire with double the current (in what follows I will call this current I. In reality we are putting current $I/2$ through two wires). In this case the field produced from the wire alone is

$$\mathbf{B}_{wire}(x,y,z) = \frac{I\mu_0}{2\pi(x^2+y^2)} \begin{pmatrix} y \\ -x \\ 0 \end{pmatrix}, \tag{3.1}$$

where μ_0 is the permeability of free space, and the coordinate system is taken to be the one detailed in 3.2.1.2.

[1]This works well enough to cancel an 8T magnetic field of a condensed matter experiment being ramped two floors away.

If we apply a bias field B_x along the X axis, the resulting field is

$$\mathbf{B}_{wire+bias}(x,y,z) = \frac{I\mu_0}{2\pi(x^2+y^2)} \begin{pmatrix} \frac{2\pi(x^2+y^2)B_x}{I\mu_0} + y \\ -x \\ 0 \end{pmatrix}, \tag{3.2}$$

creating a magnetic field minimum at $x = 0$, $y_0 = -\frac{I\mu_0}{2\pi B_x}$. Changing the centre of the coordinate system to be centred around the magnetic field minimum, and looking at small displacements about this point ($\Delta y, \Delta x << y_0$), the magnetic field is given by:

$$\mathbf{B}_{wire+bias,recentered}(x,y,z) = \frac{2\pi B_x^2}{I\mu_0} \begin{pmatrix} y \\ -x \\ 0 \end{pmatrix}, \tag{3.3}$$

this field provides confinement in the X and Y directions for atoms in a weak field seeking Zeeman sub level, as discussed in section 2.4.

The confinement in the axial direction comes from the current in the ends of the small Z wires. Approximating the ends of the Z wires as infinitely long and thin wires, separated by $2L$, the magnetic field is given by

$$\begin{aligned} \mathbf{B}_{ends}(x,y,z) &= \frac{I\mu_0}{2\pi} \left(\frac{1}{y^2+(L-z)^2} \begin{pmatrix} 0 \\ -(L-z) \\ y \end{pmatrix} + \frac{1}{y^2+(L+z)^2} \begin{pmatrix} 0 \\ -(L+z) \\ y \end{pmatrix} \right) \\ &\approx \frac{I\mu_0}{2\pi} \left(\begin{pmatrix} 0 \\ -\frac{2}{L} \\ \frac{2y}{L^2} + \frac{6y_0z^2}{L^4} \end{pmatrix} \right), \end{aligned}$$

where we have Taylor expanded and used the fact that $z, y << L$ and neglected some terms in B_Y as the high trapping frequency from the quadrupole field will make these term negligible. y_0 is the distance from the chip of the quadrupole centre. The terms linear in y causes the trap to be displaced slightly in the y direction from the quadrupole centre, however this term is small and will be neglected from here on.

We are now in a position to calculate the trapping potential the atoms see in the low temperature limit. Following the discussion in section 2.4 the potential energy of the atoms is proportional to the strength of the static field, and at low temperatures the trap is approximately harmonic with frequencies

$$\omega_r = \sqrt{\frac{\mu_B m_F g_F B_1^2}{B_0 m}}, \tag{3.4}$$

$$\omega_z = \sqrt{\frac{\mu_B m_F g_F B_2}{m}}, \tag{3.5}$$

where ω_r and ω_z are the radial and axial angular frequencies respectively. B_0 is the field applied from the Z bias, $B_1 = \frac{2\pi B_x}{I\mu_0}$ and $B_2 = \frac{6}{B_x L^4}\left(\frac{I\mu_0}{2\pi}\right)^2$.

Standard parameters for our experiment (I = 2A, B_0 = 2G, L = 3.5mm, B_X = 30G) give

frequencies $f_r = 2\text{kHz}$, $f_z = 3\text{Hz}$ and a distance to the chip of $120\mu\text{m}$. The radial frequency isn't far from the measured frequency of 1.4kHz, but at low temperatures the axial frequency is typically much higher due to fragmentation, discussed in section 3.2.2.4.

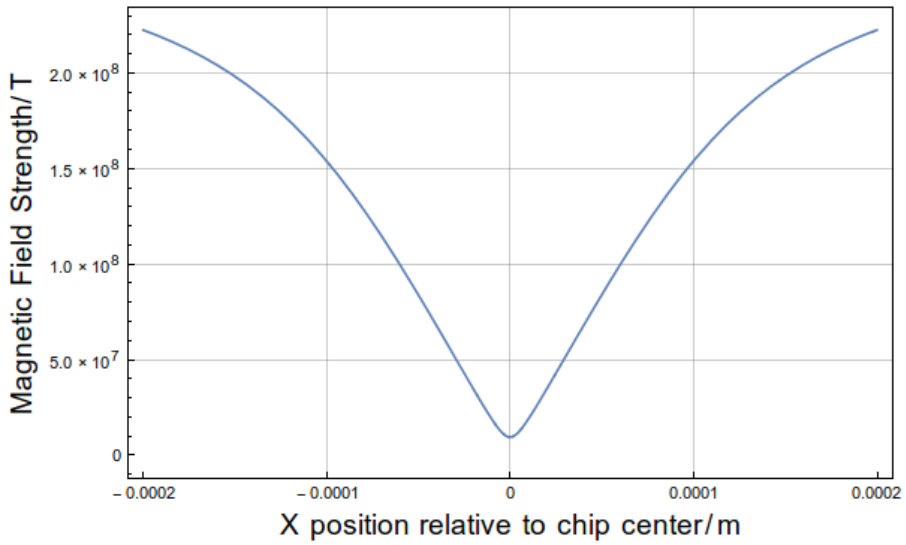

Figure 3.3: Simulations of the magnetic field strength in the X direction taking into account the finite sizes of the chip wires, under standard conditions.

To get a feel for the shape of the magnetic field magnitudes, numerical models taking into account the finite wire sizes have been produced. Around the trap minimum the potential energy of the atoms in the X direction is shown in Figure 3.3. Here you can see a small region where the potential is harmonic, followed by a large region where the trap is approximately linear, before it decays off. At BEC temperatures the atoms see only the harmonic part of the trap, but atoms loaded from a MOT will explore the linear part of the trap.

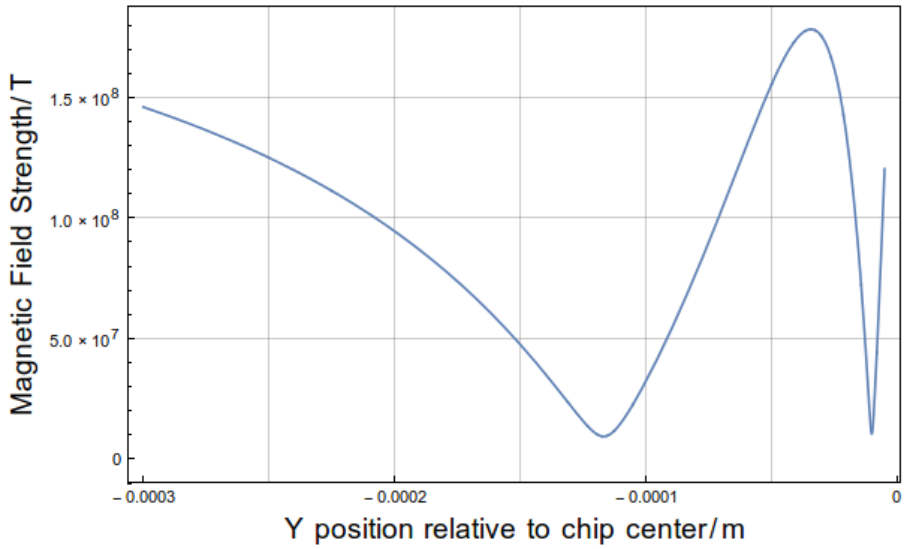

Figure 3.4: Simulations of the magnetic field strength in the Y direction taking into account the finite sizes of the chip wires, under standard conditions.

Around the trap minimum the potential energy of the atoms in the Y direction is shown in Figure 3.4. Here we see two things which aren't taken into account by our idealised model above. The first is the fact that we see two minima, which is a consequence of using a two wire configuration. Under normal conditions this second minima makes no difference to our experiment, but it does limit

how close we can move the atoms to the atom chip. As you increase the X bias (moving the quadrupole centre closer to the chip) at some point these two minima will cross and move horizontally apart [68]. This happens when the distance of the minimum to the chip is the same as the separation of the two wires, 85μm.

Along the y axis the trap has a finite depth due to the roll off of the magnetic field strength and the action of gravity of the atoms. From our numerical model, taking into account the effect of gravity, we can calculate the depth of the magnetic trap to be of the order 1500μK. At low temperatures this has negligible effect on our cold cloud, however when we initially load the atoms into the magnetic trap it provides a mechanism for passive evaporation.

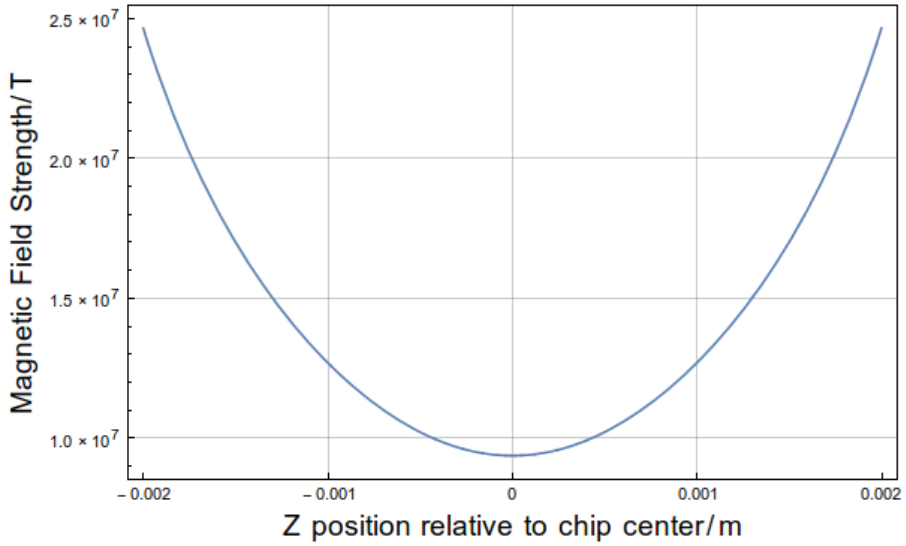

Figure 3.5: Simulations of the magnetic field strength in the Z direction taking into account the finite sizes of the chip wires, under standard conditions.

Around the trap minimum the potential energy of the atoms in the Z direction is shown in Figure 3.5. Notice the difference in scale between this plot and the previous two: the trap formed is harmonic, but much weaker that the other axis. This only holds near the centre of the 'Z', at the edges the ends of the 'Z' wire cause a much steeper potential.

Because the end wires are part of a Z they break the symmetry of the system and the trap minimum doesn't lie exactly parallel to the central part of the Z, instead it is rotated slightly. Numerical calculations of the trap show that this rotation is typically of the order 0.2 degrees, which is small enough to ignore for most purposes.

3.2.2.3 End Wires

As well as the the small Z wires we have end wires on the chip (as seen in Figure 3.1) that can influence the magnetic trap. These wires act in a similar way to the ends of the small Z wires, but can be controlled independently. They allow us to move the potential minimum of the trap in the Z direction - something that is useful for both centring the atoms over a high frequency fragment during evaporation and for exciting centre of mass oscillations in the condensate.

3.2.2.4 Fragmentation

It is common in atom chip experiments that magnetic fields close to wires don't follow the pattern expected of ideal current carrying elements [69] [70] [71] [72]. In elongated traps like the one here,

where the radial confinement comes from the quadrupole field above a wire this non-ideal behaviour manifests itself as a 'roughness' along the axial direction. For this reason we see additional structure appearing in the density of the cloud when it is colder than 1μK. The cloud can even fragment into separated regions, hence the name.

The roughness of the potential comes from a variation in the Z-component of the magnetic field. The net current in the axial direction must be constant by conservation of current, but roughness on the edges of the wire can create currents perpendicular to the axis of the wire, and these are the cause of the roughness in the potential.

While this fragmentation may seem to be a disadvantage, we use it to our advantage in this experiment. Firstly it creates regions where axial trapping frequencies are higher than for the bare trap, allowing fast forced evaporation. By adjusting the current in the end wires we centre the cold cloud over a region of higher axial frequency during part of the evaporation ramp.

A second advantage of the fragmentation potential is that it creates a region in which the condensed part of our atom cloud experiences a different trapping frequency to the thermal part of the atom cloud. We use this to excite a relative motion between the two components. More details of this will be covered in the next chapter. A measurement of the fragmentation potential used for measurements of the centre of mass potential is discussed in 3.4.4.

3.2.2.5 Radio Frequencies

In addition to the static magnetic fields applied to the system we also apply currents oscillating at radio frequencies to force evaporation of the atoms from the magnetic trap, and to create the adiabatic potentials. For evaporation we use the copper H described in Figure 3.2 and for adiabatic potentials we usually pass current along the large Z wires on the atom chip, described in Figure 3.1.

To produce the radio frequency currents, we use two computer-controlled VFG frequency synthesisers. These are connected to fast TTL controlled RF switches, then to 24dB RF amplifiers to produce the required 20dBm of power before being fed into the vacuum chamber. Both the RF switches and amplifiers are manufactured by Mini-Circuits.

3.2.3 Light

As discussed in section 2.2.1, we use the rubidium 87 atom and interact with it using the D2 lines. To do this we need several stable frequencies of light around 780.241nm. In this section I describe how we generate this light.

3.2.3.1 The Reference Laser

To provide a consistent frequency reference for subsequent light, we start with an extended cavity diode laser (ECDL) locked to resonance to the $F = 2$ to $F' = 3$ transition. The laser is a commercial diode designed for operation at 785nm run on a home build current driver and temperature stabilised using a Peltier cooler.

The extended cavity forces the diode to the desired frequency and provides single mode operation through optical feedback. This is achieved using a diffraction grating mounted in front of the diode and aligned so that the first order diffracted mode is fed back into the diode, in effect creating an external cavity around the diode. As the cavity narrows the range of allowed modes, we can achieve

single mode operation. By adjusting the cavity length using a piezoelectric actuator we can tune the laser frequency over a narrow range. For more details about the ECDL configuration, see [73].

To lock this laser to the desired atomic transition frequency, polarisation spectroscopy of a rubidium vapour cell is used. Polarisation spectroscopy is a modified version of saturated absorption spectroscopy that uses induced birefringence of atoms in a vapour cell to provide an error signal about the desired atomic transition. A schematic of the set up can be found in Figure 3.6.

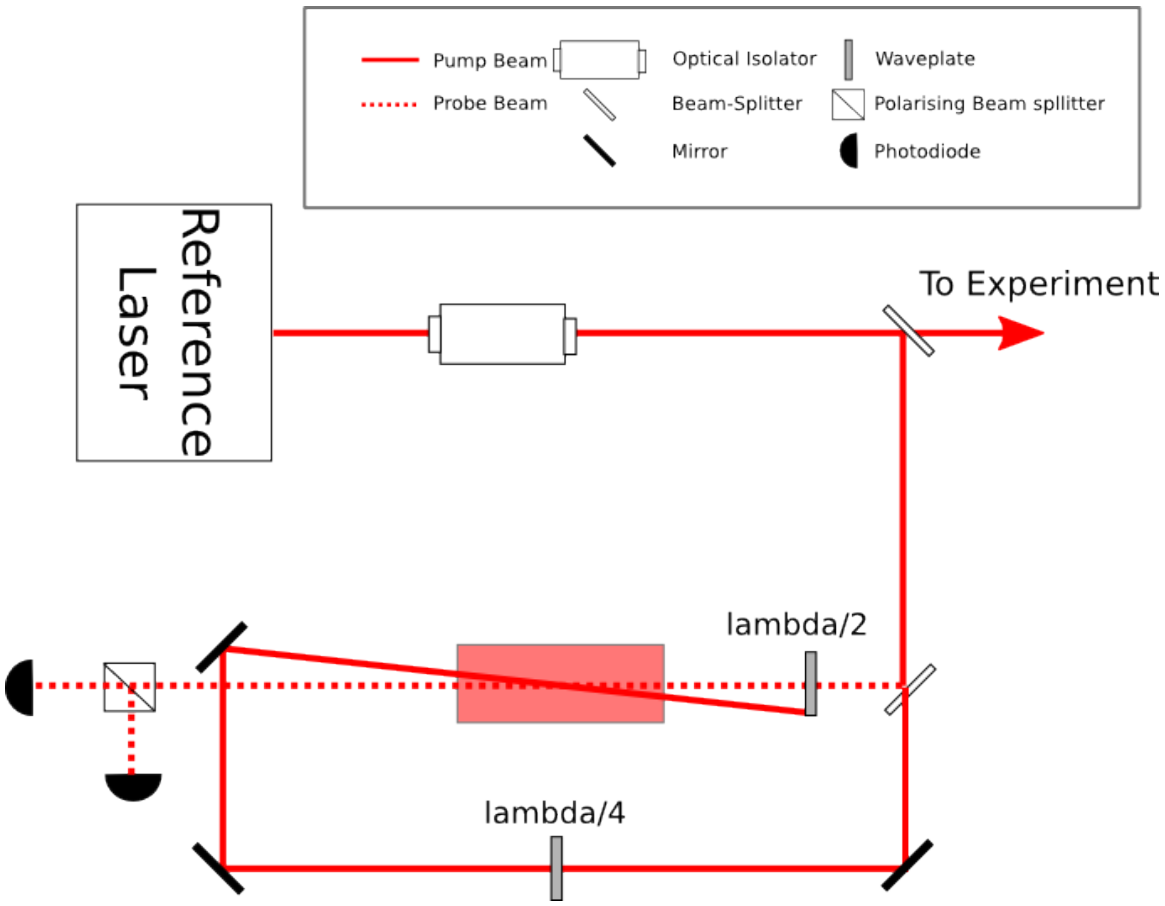

Figure 3.6: A schematic of the polarisation spectroscopy set up used to lock our reference laser to an atomic transition in rubidium. The polarisations of the pump and probe beams are controlled through $\lambda/4$ and $\lambda/2$ waveplates respectively. For more details of polarisation spectroscopy, see [74]

Two counter propagating beams are incident on the vapour cell. The first, known as the pump beam, is circularly polarised and when close to the transition in frequency, optically pumps the rubidium atoms in to an extremal m_f state. The second beam, known as the probe beam, is derived from the same beam as the pump, and so has the same frequency. The polarisation of the probe beam is linear and so can be thought of as an equal combination of left and right circularly polarised light.

Atoms in an extreme m_f state interact with circularly polarised light of different handedness differently. The difference arises from the different Clebsch-Gordan coefficients, causing different absorption cross sections. If we write difference in cross section on resonance as $\Delta\alpha_0$, its dependence on frequency will have the form $\frac{\Delta\alpha_0}{1+\delta^2}$ where δ is the detuning of the laser from resonance in units of natural linewidth. Via the Kramer Kronig relations, the birefringence of the material with this absorption is then proportional to $\frac{\delta\Delta\alpha_0}{1+\delta^2}$, a dispersive shape with a zero crossing when the laser is on resonance with the transition. This birefringence introduces a phase shift between the different components of the probe light.

This phase shift can be read off by passing the probe beam through a polarising beam splitter (PBS) and measuring the intensity of each component with a photodiode. By taking the difference in signal of the two beams we directly measure the error signal and use it to drive the piezo actuators to stabilise the laser frequency.

In a similar way to saturated absorption spectroscopy, the counter propagating nature of the beams means that only those atoms of zero velocity relative to the direction of beam propagation are addressed, so the polarisation feature is narrow in frequency.

A weakly magnetised mu metal shield covers the rubidium vapour cell. This shield has the advantage of limiting the influence stray magnetic fields can have on spectroscopy through Zeeman shifting the levels.

For more details on the polarisation spectroscopy see [74].

3.2.3.2 The Cooling Laser

To cool the atoms we use a second laser, currently a Toptica TA pro. This TA pro is an ECLD diode laser with a tapered amplifier to get 1W of single frequency light. The frequency of the ECDL can be controlled using a piezo actuator, allowing us to scan the frequency. The light out of the laser is immediately coupled into a single mode fibre to allow control over where the light goes, as well as cleaning up the shape of the beam.

Note that although we currently use a TA pro to provide the cooling light, for part of the experiment reported here a Coherent MBR-110 Ti:sapphire laser, pumped by an 8W Coherent Verdi V-8 was used. The modularity of our laser systems meant that we could exchange the two without any real change[2] to the rest of our setup or locking schemes and this difference will be ignored from here on.

To lock the cooling laser to a specific frequency, we split off a small amount of the light and mix it with the reference laser on a fast photodiode. The intensity measured on the photodiode will include a component at a frequency equal to the difference in frequency between the reference laser and the cooling laser, known as the beat note signal. This signal is passed through a high pass filter to remove any DC components, then mixed with the output of a voltage controlled oscillator oscillating in the range 200-390MHz. A low pass filter then removes the component proportional to the sum of their frequencies, leaving only a signal at the difference frequency.

To convert this signal into an error signal, the signal is passed through a high pass filter with a 3dB point at 120MHz. The signal is then averaged to produce a voltage proportional to the amplitude of the signal, which is then offset so that the 3dB point of the filter is centred around zero. This creates an error signal from the frequency response of the high pass filter. Locking to this feature allows us to fix the difference in frequency between the VCO and the beat note of the frequency to 120MHz. This scheme is known as side of filter locking [75] and a schematic can be found in Figure 3.7.

The advantage of using this scheme is that it allows us to change the frequency of the cooling laser by adjusting the voltage input to the VCO, giving us fast computer control. It does mean that the cooling laser is locked at least 80MHz red detuned from the cooling transition in rubidium 87, whereas the ideal cooling frequency is around 12MHz. This is a deliberate choice as it allows us to bring the frequency closer to resonance we use a sequence of Acousto Optical Modulators (AOMs), as discussed in a section **??**. These AOMs allow us to control the power in the beams precisely.

[2]The main impact of changing the laser system has been to substantially reduce the time spend maintaining the laser.

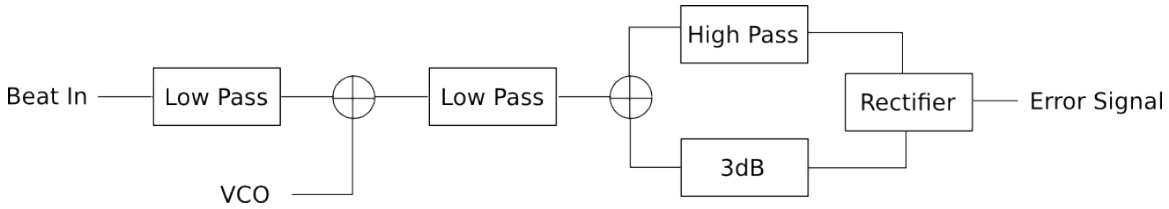

Figure 3.7: A schematic of the side of filter lock technique. The incoming beat signal between two lasers is low pass filtered, before being mixed with the signal from a VCO and low pass filtered again. The resulting signal is split equally in two, one part passing through a high pass filter, the other through an attenuator. The two paths are then rectified and the difference results in an error signal. The design is taken from [75].

3.2.3.3 The Repump Laser

The transition we cool on isn't completely closed. The finite linewidth of our laser means that with low probability it is possible for our laser to excite a transition between the $F = 2$ ground state and the $F' = 2$ excited state. Such an excitation is only expected to happen once every 9000 excitation events, but when it does it is possible for the atom to decay into the $F = 1$ ground state, effectively removing it from the cooling scheme. To move these atoms back into the $F = 2$ ground state we need a repump laser.

To achieve this we use a second home built diode laser, configured and locked in a similar way to our reference laser (section 3.2.3.1), to drive a $F = 1$ to $F = 2$ transition in rubidium 87. This laser, known as the repump, allows us to return atoms from the $F = 1$ ground state to the $F = 2$ ground state, as the $F = 2$ excited state can decay to both ground states.

3.2.3.4 The AOMs

As well as requiring the MOT light, we need several other frequencies in order to cool, optically pump and image the atoms. We use acousto-optic modulators (AOMs) to shift the light to the required frequencies. AOMs are made from crystals that exhibit the acousto-optic effect, a change in refractive index of the crystals due to a mechanical stress. A wave of sound travelling in the crystal creates a periodic change in refractive index, causing light incident on the crystal to undergo diffraction.

Because the sound wave is travelling through the crystal, the diffracted light undergoes a frequency shift depending on the order it is diffracted into. This effect is the same as light diffracted from a moving grating: each order of the light is Doppler shifted by different amounts and exits the AOM at a different angle. Another way of thinking about AOMs is to consider that each photon can absorb or emit a fixed number of phonons from the crystal. Because the phonons are travelling perpendicular to the photons, their momentum causes a deflection of the photon, and the energy change creates a frequency shift. Both these effects are quantized due to the fact the photon can only absorb or emit an integer number of phonons, resulting in separate diffraction orders of light, each with their own frequency. By turning the sound waves on and off, it is possible to achieve both fast switching of the light (switching times of the order 100ns) and frequency control.

In this experiment we use a combination of home build and commercial AOMs shifting the light by frequencies of roughly 100MHz. To produce the sound waves at these frequencies, we use VCOs followed by RF amplifiers. RF switches allow fast switching on and off of the AOMs.

3.2.3.5 Shutters

While the AOMs allow us to switch the light quickly, they do scatter a small number of photons even when switched off. Because any stray, resonant light can create unwanted scatter which can heat or optically pump the atoms in the magnetic trap, unwanted light must be minimised as much as possible. To achieve this we use mechanical shutters. The shutters are controlled by solenoids, driven by home built circuits that achieve switching times of the order 2ms.

3.2.3.6 Putting it All Together

Half of the optical table used in the experiment is dedicated to creating the right frequencies of light and controlling their timing. A schematic of this half of the table can be found in Figure 3.8.

We derive four frequencies of light. The cooling light comes from the TA pro which is offset locked 100MHz red detuned from the cooling transition; its frequency is then stepped up by 88MHz by an AOM, before being mixed with repumper light. Both the cooling and repump light can be independently blocked by shutters.

The imaging light is derived from the cooling light, and is shifted up in frequency by 100MHz by an AOM. The optical pumping light is derived from the reference laser light, and stepped down in frequency by double passing it through an AOM at roughly 140MHz. It is then mixed with one branch of the imaging light before being sent to the main chamber. The reference laser's light was used so that the frequency of the cooling laser could be swept without having to wait for it to be brought back to resonance when we need to optically pump the atoms. Both the imaging light and the optical pumping light can be blocked independently by shutters.

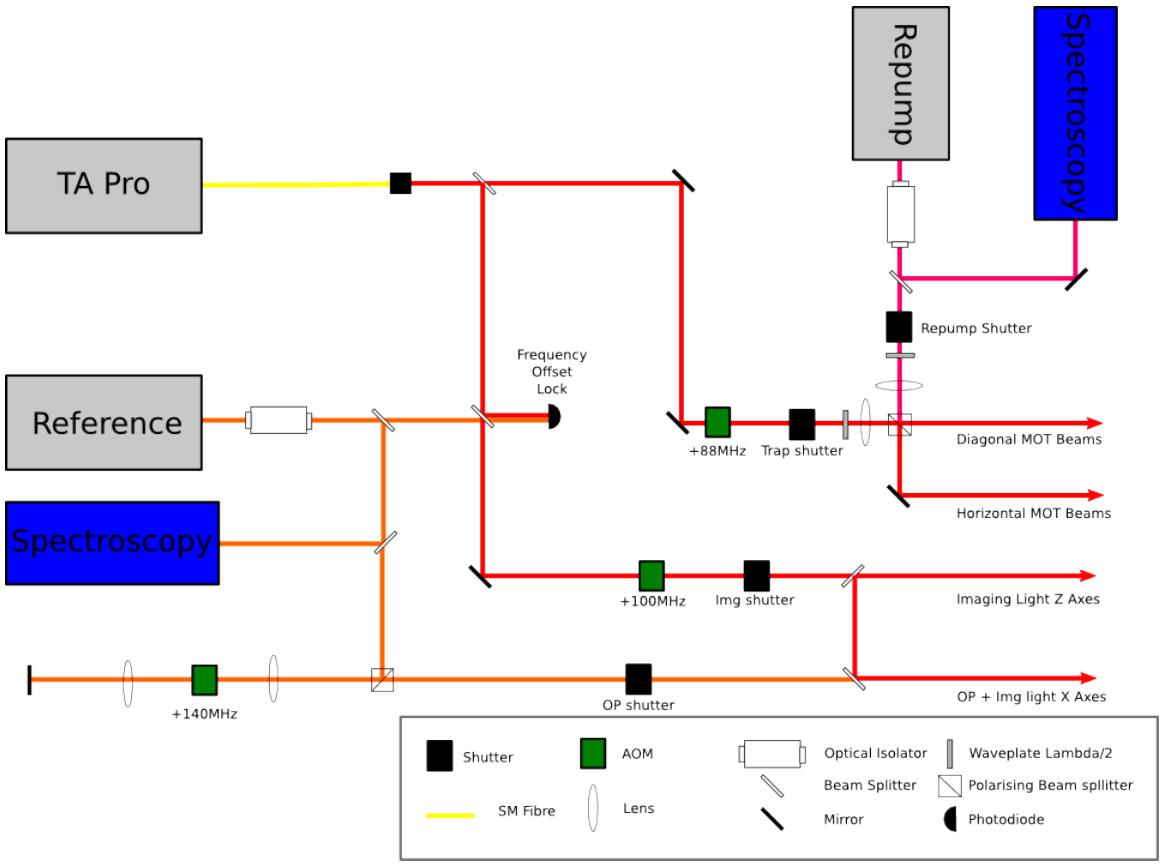

Figure 3.8: The layout of the part of our optical table used to prepare the light.

3.2.4 The Vacuum Chamber

The vacuum system consists of two chambers: the science chamber and the LVIS chamber. The two are connected by a 1mm aperture designed to allow a flux of cold atoms to pass through, while maintaining a pressure differential between the two. Their designs are discussed in this section.

3.2.4.1 The Science Chamber

The main chamber is where we carry out our experiments. A diagram of it can be found in 3.9.

It consists of a 10 port octagon chamber made from 316 stainless steel by Kimball Physics. The eight ports around the edge of the octagon are $2\frac{3}{4}$ inch Conflat and there are two 8 inch Conflat ports on the top and bottom of the chamber. The atom chip assembly is mounted on the top flange, which also contains feedthroughs that carry the wires that carry current to the atom chip. The bottom port, and four on the edge of the chamber are sealed viewports to provide optical access.

Connected to the other four port are an ion pump (Varian UHV 25l/s, with a soft iron magnetic shielding), the LVIS chamber, a non evaporable getter and a T piece which in turn is connected to an angle valve (to provide a connection for a turbo pump) and ion gauge (a UHV-24p Nude Bayard-Alpert ionisation gauge).

3.2.4.2 The LVIS Chamber

LVIS stands for low velocity intense source [76], and is a magneto optical trap with a hole in it. The LVIS is formed using six beams, arranged into three sets of retro reflected beams. One of the mirrors used for retro reflection contains an aperture. The aperture creates a 'hole' in one of the cooling beams into which cold atoms are pushed. It is from this beam that we get atoms in the main chamber. The design of the LVIS is identical to the one described in [77].

Because the LVIS loads atoms from the thermal tail of the background gas, the pressure needs to be high (at least relative to the main chamber), typically of the order 10^{-8}mbar. These pressures are too high to produce a BEC, which is why we use the two chamber setup.

The LVIS chamber consists of two six way crosses connected by a T piece, along with extra sections. See the diagram 3.10 for full details. Magnetic fields for the LVIS are created by coils wound round the view ports.

The connection between the LVIS chamber and the science chamber consists of a 6mm thick $\frac{\lambda}{4}$ waveplate with a 1mm diameter hole drilled in the centre. The back surface of the waveplate is coated in gold. It is mounted in a modified $2\frac{3}{4}$ inch flange.

3.2.4.3 Getting Light Into the Vacuum Chamber

Having prepared the light, the final step is to launch it into the vacuum chamber (details of the vacuum chamber are discussed in the previous section). The cooling light and repump light are kept as free space beams.

Following the labelling in Figure 3.8, the diagonal MOT beam is split into two separate beams which are launched into the vacuum chamber at 45 degrees to the horizontal, entering through the large port on the base of our vacuum chamber. Each of these beams reflects off the atom chip, and then exits the vacuum chamber. The beams are aligned such that they are counter propagating and overlap. Together these two beams and their reflections off the atom chip make up four of our six MOT beams.

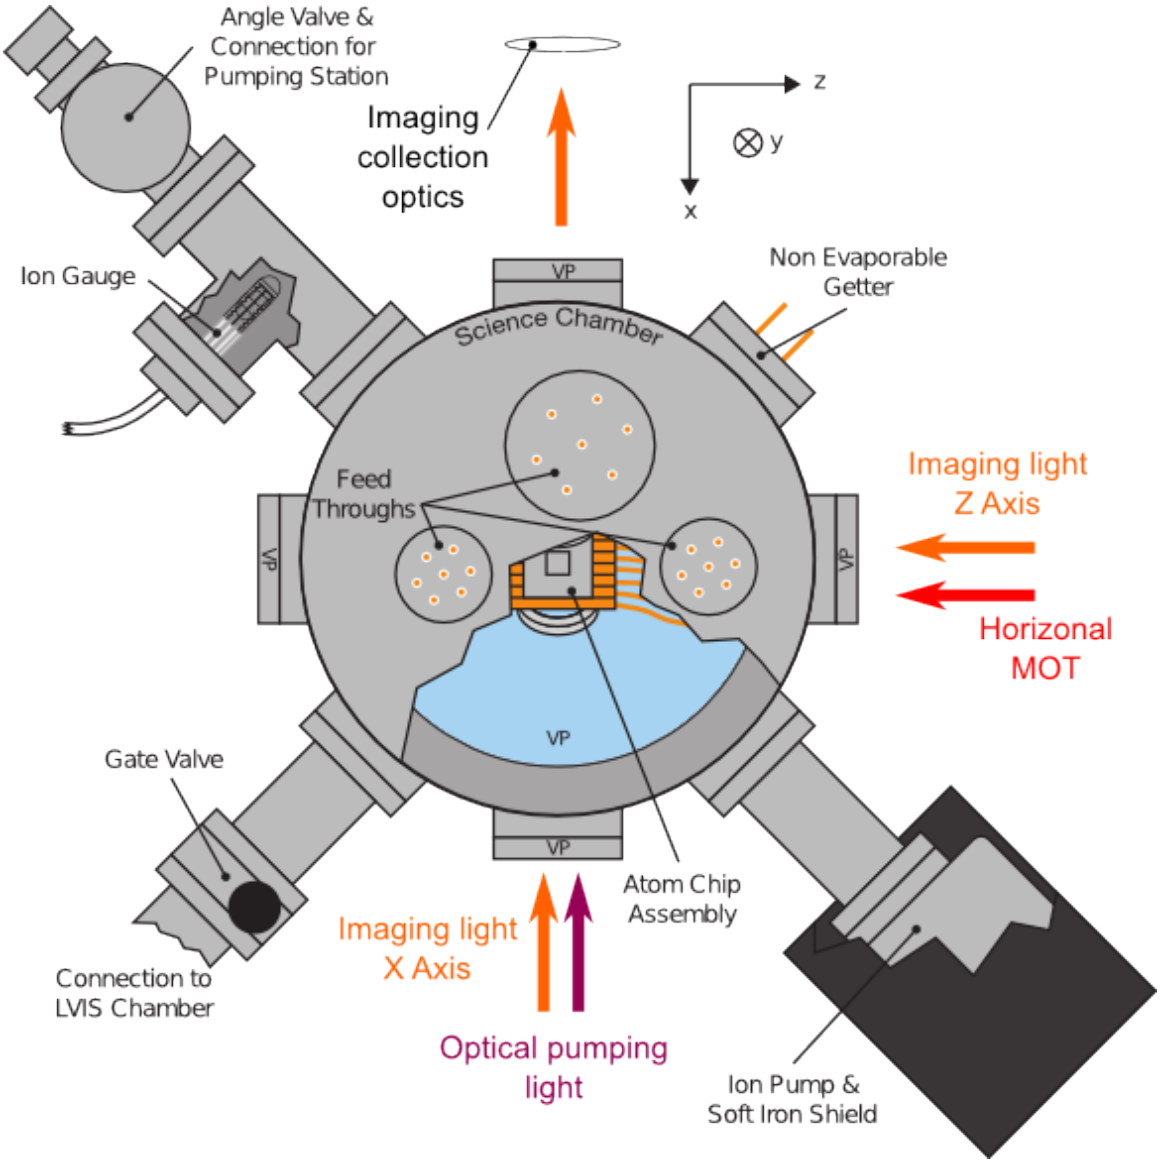

Figure 3.9: The science chamber. Parts labelled VP are view ports. The section labelled ̀connection to LVIS chamber' joins with the connection in the LVIS chamber diagram in 3.10. The direction of entry for one set of horizontal MOT beams, the imaging light and the optical pumping light is shown. The MOT light and the imaging light on the Z axis exit the chamber along the y axis, and out of the plane of this diagram. In this diagram the light for the MOT and imaging beams along the Z axis and the optical pumping and imaging light along the X axis are shown offset from each other for clarity, in reality they overlap. This figure is a modified version of one taken from [64] with permission.

The other two MOT beams are derived from the horizontal MOT beams, as labelled in Figure 3.8. One of these beams enters through one of the small view ports on our vacuum chamber, the other enters vertically from the view port below the vacuum chamber. The two are overlapped on a gold coated prism mounted at the edge of our atom chip and visible in Figure 3.2c. The polarisation and intensities of the MOT beams are controlled using wave plates and PBSs prior to entering the vacuum chamber.

Optical pumping light and imaging light are fibre coupled using single mode polarisation maintaining fibres prior to being launched into the camber, making aligning them easier.

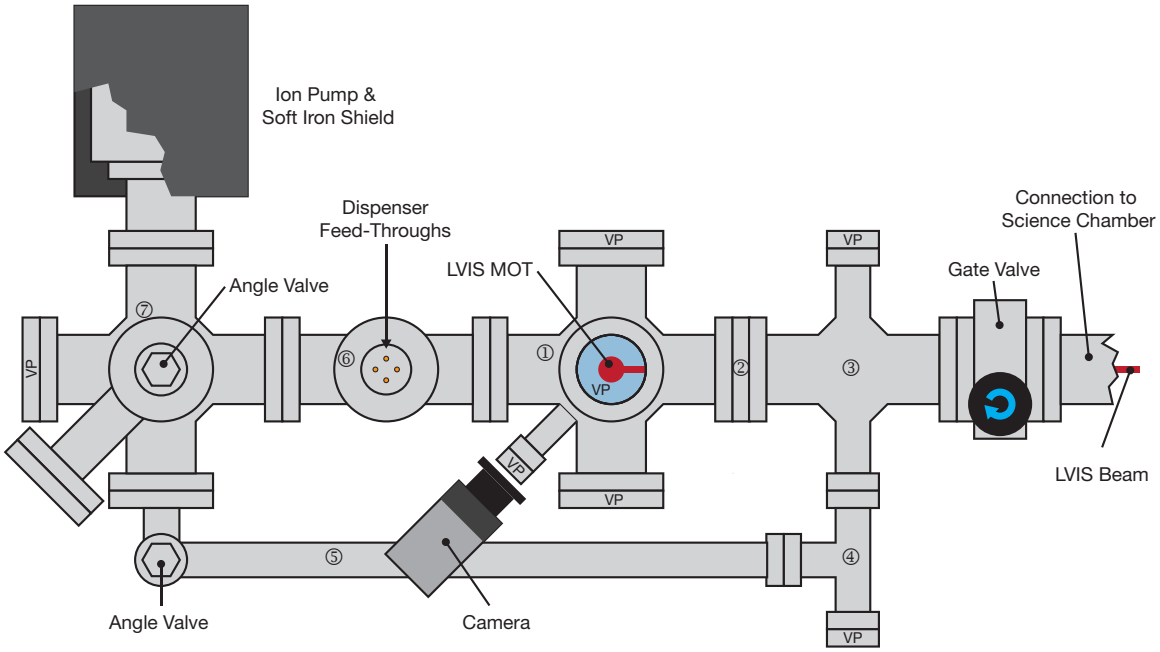

Figure 3.10: The LVIS vacuum chamber. The sections are numbered and listed following the sequence of the main text. 1 LVIS MOT six-way cross. 2 Extraction mirror flange. 3 Four-way cross (UHV). 4 By-pass T-piece (UHV). 5 By-pass tube (UHV) and angle valve. 6 Dispenser T-piece. 7 Six-way cross. Sections 3 to 5 are under UHV. Other sections are under HV, maintained by a 25l/s ion pump. Flanges labelled VP are view ports. The by-pass tube is only used during the initial pump down of the system. This figure is taken from [64] with permission.

3.2.5 Imaging System

To view the atoms we use two different cameras and imaging systems. In the coordinate frame described in 3.2.1.2 they are aligned with the Z axis and X axis. Both cameras are centred on the point in the vacuum chamber where the atoms are formed.

In this thesis, we us the X most commonly. It provides a side view of the cigar shaped trap, allowing us to image in the Y-Z plane using the orientation described in 3.2.1.2, and can be used for both absorption and fluorescence imaging. The direction in which this light enters and leaves our chamber can be found in Figure 3.9.

After leaving the chamber, the light is collected out of the view port of the vacuum chamber by a pair of lenses arranged in a 1f-2f-1f configuration, focal length 15cm, before being imaged on a CCD camera, AVT F505B. The lens arrangement is designed to produce 1:1 magnification of the atom cloud. For absorption imaging the light is launched into the vacuum chamber from the opposing view port (along with the optical pumping light), having been expanded and collimated to a $1/e$ size of 7.5mm. This imaging system is diffraction limited by the size of the view port on the vacuum system, Rayleigh limited to be able to distinguish features separated by 3μm. The pixel size of the camera used is 3.45μm by 3.45μm. The beam has a 1mm 1/e radius.

We are limited in how far it is possible to view using the X axis imaging system, as the coil formers for the MOT coils block our view of the atoms. We can only view 2mm below the atom chip, so the cloud is only visible for 20ms of free fall.

The imaging system along the Z axis is aligned with the horizontal MOT beams in the vacuum chamber (visible in Figure 3.9). We combine and separate the beams using polarising beam splitters. Because the polarisation isn't perfect in the MOT beams there is always some leak through into the

Z axis imaging system, meaning it cannot be used for florescence imaging. It is useful however for absorption imaging. In the chamber this imaging light is reflected off the gold coated prism mounted by the atom chip and visible in Figure 3.2c.

After leaving the chamber the light is collected below the vacuum chamber and passed through a 1f-2f-1f imaging system, focal length 12cm. The resulting 1:1 magnified image is imaged on a AVT 303B camera. If required a third lens can be added to increase the magnification. The resolution along this axis is diffraction limited by the size of the polarisation optic used for the MOT beam, and is similar to the other axis.

3.2.6 Computer Control

The entire experiment is controlled using software known as Tyche, written by Rob Nyman and Manuel Succo [78]. It provides a platform from which we can create precisely timed sequences of voltage pulses to control the experiment.

Most of the control sequence is managed by two national instruments digital and analogue output cards. We use a NI-DAQ 6713 card to provide 8 analogue outputs in the range -10 to 10V, and a NI-DAQ 6543 card to provide 32 digital (0 or 5V) TTL signals. The control sequence is loaded onto these cards, prior to the experiment. A second PC controls the magnetic feedback coils and monitors the conditions of the experiment for diagnostic purposes.

Most of the hardware (current drivers for the magnetic fields, mechanical shutters, AOM frequency drivers etc.) is controlled directly by the voltages these cards output. Devices not directly controlled by these cards such as the RF frequency generators and the cameras, are configured prior to the sequence and are activated when a digital trigger line goes high - the line comes from the digital output card noted above.

3.3 A Typical Sequence

Having outlined the various components of the experiment, I now describe a typical experimental sequence that produces a Bose Einstein Condensate.

3.3.1 Getting the Atoms

The rubidium atoms themselves come from a SAES Getter dispenser - a resistive element coated in rubidium chromate and a reducing agent of zinc and aluminium. When a current is passed through the element a reduction reaction occurs producing free rubidium atoms in the LVIS chamber. Typically we run the dispenser at a constant current in the 4-8A range, adjusting it occasionally to keep the amount of rubidium fixed.

These atoms are then trapped and cooled in the LVIS chamber by a MOT. By applying a constant field on top of the LVIS's quadrupole field the MOT forms in a displaced position so that its centre is not in front of the aperture to the main chamber.

When the sequence begins the offset field on the LVIS MOT is changed to align the LVIS MOT with the aperture in one of its retro reflecting mirrors so that atoms are pushed out of the LVIS and into the main chamber.

3.3.2 The First MOT in the Science Chamber

At the same time as the magnetic fields are adjusted in the LVIS, the science chamber's MOT coils and beams are turned on, along with the bias fields necessary to form a MOT from the incoming atoms. These have been adjusted empirically to maximise the number of atoms loaded into the initial MOT. Initially the MOT forms 5mm from the chip's surface and is loaded for 10s until it typically contains 1×10^9 atoms, as measured by fluorescence imaging. At this point no current flows in the chip wires.

3.3.3 The Second MOT in the science Chamber: the UMOT

Having loaded enough atoms into the initial science chamber MOT, the current in the MOT coils is ramped down and the currents on the atom chip are ramped up. During this time, the external bias fields are adjusted to bring the entire MOT closer to the atom chip and alter its shape from being potato shaped (quadrupole bias fields mean that one of the three axes has twice the field gradient, therefore twice the trapping frequency of the other) to cigar shaped (the quadrupole field is now created by the small Z wires on the atom chip, making it highly elongated along this axis). We call this the UMOT. The name comes from the fact that the quadrupole field now derives from current running in a U on the atom chip: 1A is passed through each of the small Z wires, and 2A is passed through one of the end wires, effectively creating a U shape of current of the chip surface.

The transfer of atoms between MOTs has been modelled and empirically optimised. For details of how exactly it is done, see [64]. The end result is that we typically transfer half of the atoms in the original MOT to the UMOT, situated about 400μm from the surface of the atom chip at a temperature of 140μK. This transfer happens over 80ms.

Having moved the MOT close to the chip, the UMOT magnetic field gradient is increased and the frequency of the cooling laser is red detuned by a further 70MHz from resonance over 15ms. We do this to position and increase the density of the atom cloud prior to loading the magnetic trap. At this point a significant loss factor is thermal atoms colliding with the atom chip (which sits at just above room temperature) and being either absorbed or heated. At the end of this sequence we typically measure 5×10^7 atoms through absorption imaging.

3.3.4 Sub Doppler Cooling and Optical Pumping

Following the UMOT, the magnetic fields are turned off in such a way that the atoms are given a slight upwards kick - enough to stop them from falling away due to gravity, but not enough that a significant fraction is lost in collisions with the chip. A small sideways kick is also applied to compensate for the momentum about to be imparted by the optical pumping step (see below). The kick is created by differentially adjusting the rate at which the chip wire currents and the external bias fields that make up the trapping potential are turned off. These rates are adjusted empirically to optimise loading into the magnetic trap.

The cooling light is kept on for 2ms to induce subdoppler cooling, before being quickly turned off using the AOMs. This cooling typically lowers the temperature of the cloud to 50μK.

An external bias field is then turned on along the Z axis to provide a well defined quantisation axis, and a 1ms pulse of optical pumping light is applied. The light is σ^+ polarised relative to the magnetic field and is used to drive a transition between the $F = 2$ ground state and the $F = 2$ excited state level, increasing the m_F state of the atom with each absorbed photon. Combined with

spontaneous decay of the atom, the light pumps the atoms in to the $F = 2$, $m_F = 2$ hyperfine ground state. The repump light is also applied during the optical pumping step. At this point all light is turned off quickly by AOMs and then blocked completely using mechanical shutters.

The optical pumping is important because it increases the number of atoms loaded into the magnetic trap (see below) by a factor of three. The origin of the number three is currently unknown, but attempts to improve this further by adjusting factors such as the alignment, length, intensity and orientation of the optical pumping pulse have proved futile.

3.3.5 Into the Magnetic Trap

Having pumped the atoms into the $F = 2$, $m_F = 2$ ground state for trapping, we turn on the magnetic trap, by ramping up the current in the small Z wire of the atom chip and turning on the X and Z bias fields to create the trap discussed in section 2.4. The trap is ramped on over 15ms, using a large Z bias to keep the radial trapping frequency low so that the trap as is broad as possible to collect the atoms. Over the next 50ms the trap is compressed by reducing the Z bias to the point where evaporation will take place. The atoms are held 125μm below the atom chip.

Typically we load 30 - 35 million atoms into the magnetic trap, at an elevated temperature of 200μK. At this point there are some small centre of mass and quadrupole oscillations due to the imperfect overlap between the UMOT and the magnetic trap. The increase in temperature arises from our inability to load the magnetic trap adiabatically. The thermal cloud of atoms has a finite size after sub Doppler cooling and when we turn on the magnetic trap outer parts of it gain potential energy in addition to the thermal kinetic energy.

Now we wait for 600ms. Because our trap has a finite depth (see section 3.2.2), unforced evaporative cooling takes place - reducing the number of atoms in the trap, but also lowering the temperature. We end up with 20 - 30 million atoms in the trap at a temperature of 140μK.

3.3.6 Evaporation

We apply rf radiation to the atoms by driving rf currents into the copper H situated 1mm behind the atom chip. As discussed in section 2.3.3, this has the effect of 'dressing' the magnetic trap, lowering its depth according to its frequency. The more energetic atoms evaporate, lowering the temperature. We change the frequency and intensity of the RF over time using a mix of modelled and empirically optimised linear and exponential frequency ramps. Most of the work on finding our ramps was done by Ben Yuen, and is discussed in great depth in his thesis [64].

The full evaporation sequence typically lasts 8.5s, and results in a BEC of 10000 - 50000 atoms at a temperature of 150nK. We use this as the starting point for the experiments discussed in the next chapter.

3.3.7 Run to Run Changes

In this experiment the system is sensitive to small changes. Drifts in static magnetic field can affect the position of the bottom of the magnetic trap, affecting the final temperature and atom number. RF noise anywhere in the system can cause atom loss from the magnetic trap. Changes in laser frequency can alter the number of atoms initially trapped and cooled. Changes in temperature can result in small changes in alignment of both the laser systems, affecting the efficiency of cooling and trapping and shifting the position of the imaging system.

Where possible we counteract these changes. We monitor slow changes in the local magnetic field and attempt to cancel them by changing our bias fields. The magnetic field coils (the main source of heating in the experiment) are kept running between runs to maintain the temperature of the system. Laser frequency, polarisation and power are monitored and logged to make sure they stay within acceptable bounds.

Despite these attempts to keep everything as stable as possible there are always small changes between each run. These changes result in changes in properties of our atomic system that we would otherwise like to keep constant. As our imaging techniques are destructive, to collect data over many cycles of the experiment we must deal with the changes. Where they cannot be removed through improvements in our design we try to remove any systematic effect they have on our result by randomising the order in which the data is taken.

For otherwise identical conditions we find that the atom number in our initial magnetic trap typically fluctuated by 10%. The temperature at this point in the cycle typically varies by 10μK. After evaporatively cooling this results in atom number changes of around 30% and temperature variations of the order 20nK.

3.4 Characterising the Atom Cloud

Throughout the experiment we extract information about the atoms we are trapping and cooling from images of the atoms, mostly absorption images. Our imaging system is described in 3.2.5 and the theory of absorption imaging is discussed in 2.6. Here we present some measurements of properties of the system, and our techniques for measuring them.

3.4.1 Measuring Temperature

Throughout the experiment, especially during evaporation, knowing the temperature of the atomic cloud is important. To measure the temperature we measure the expansion of the cloud after being released from the trap. A classical gas in thermal equilibrium follows the Boltzmann distribution. For a non interacting gas this is given in the x direction by

$$\rho(x, p_x) \propto \exp\left(-\frac{p_x^2}{2mk_BT} - \frac{m\omega_x^2 x^2}{2k_BT}\right) , \tag{3.6}$$

where $\rho(x, p_x)$ is the phase space density at position x and momentum p_x.

To find how the cloud expands once the potential is turned off we assume that the atoms follow classical trajectories. The position of an atom beginning at x_0 with momentum p_x is given by $x(t) = x_0 + \frac{p_x}{m}t$. The distribution at time t is then

$$\begin{aligned}\rho(x,t) &= \iint dx_0 dp_x\, \rho(x_0, p_x)\delta(x - x_0 - \frac{p_x}{m}t) \\ &\propto \exp\left(-\frac{m\omega_x^2 x^2}{2k_BT(1+\omega_x^2 t^2)}\right) ,\end{aligned} \tag{3.7}$$

showing that the cloud expands quadratically in time once the width is sufficiently larger than its initial width, and the rate of this expansion is proportional to the square root of temperature.

To measure this we release the cloud and then image it with absorption imaging along the Z axis.

We use imaging system along the Z axis for temperature measurement because we have imaging access over a larger range of heights than the other axis, allowing a larger range of expansion times, and because the initial trap has much higher trapping frequencies in the X and Y directions, removing the effects of the initial width. By taking a sequence of images at different expansion times and comparing the width to the time we extract temperature.

An example of this process can be found in Figure 3.11. Here we extract how the width of the cloud varies in time. From equation 3.7 for an ideal gas the relationship between the two is $\sigma^2 = \sigma_0^2 + t^2 \frac{Tk_B}{m}$, so by fitting a straight line to the points using least squares regression we can measure the temperature.

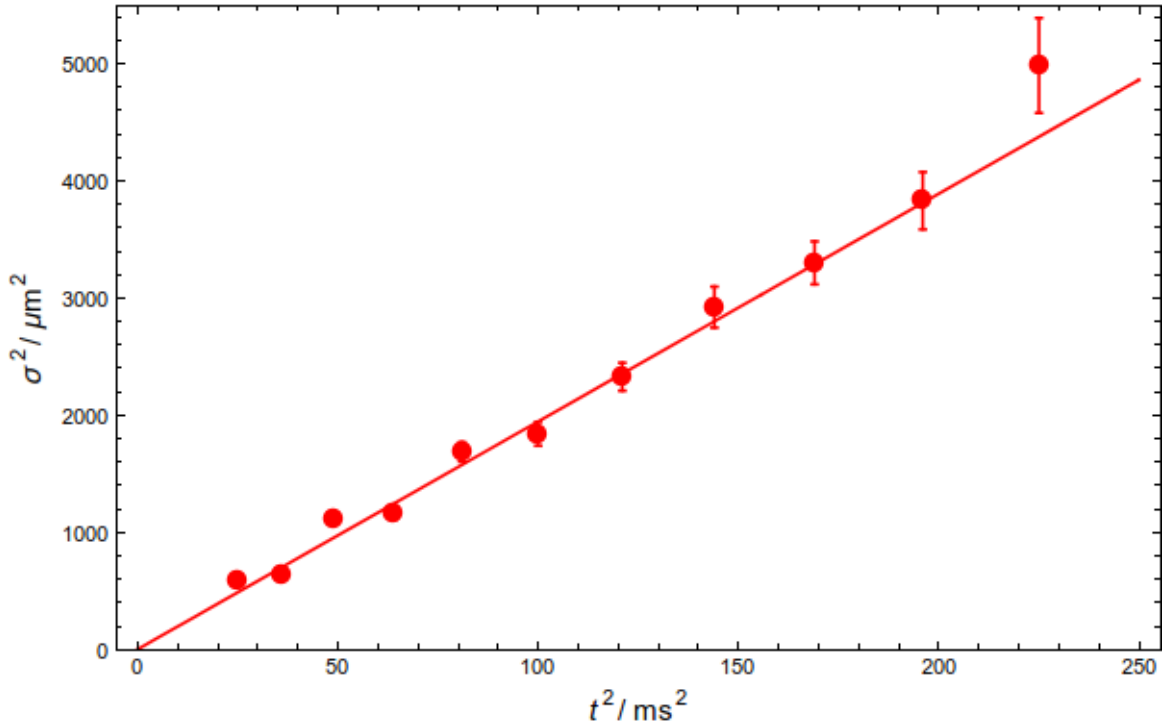

Figure 3.11: A sequence of atomic clouds are created and released and allowed to fall for a varied length of time before being imaged along its length to measure the temperature of the cloud. We measure the width of the cloud for each time, and plot width squared against time squared. Here we plot the width of the could along the X axis. Fitting a straight line to this allows us to measure the temperature, as discussed in the text. The cloud in this measurement is at $0.8 \pm 0.05\mu$K.

The above method only works for thermal ideal clouds. For cold clouds close to or below condensation temperatures the effects of repulsive atom atom interactions can drive expansion faster than the rate given above, making it an unreliable measure of temperature. In this regime we use another technique described below.

3.4.2 Characterising the Trap

3.4.2.1 Trapping Frequencies

To measure the frequencies of the trap in the harmonic regime, we displace the condensate slightly from its centre and then release it, measuring how the centre-of-mass of the condensate evolves.

To measure the frequency in the radial direction we adiabatically lower the small Z wire current by 0.1A (from its typical current of 1A) over 50ms and then quickly ($< 400\mu$s) return it to its original value. We then allow the condensate to evolve in this trap for a varied length of time, before releasing

it, letting it fall for 2ms and taking an absorption image. We fit to this image to extract the centre of mass of the cloud. The centre-of-mass position at imaging time is proportional to the momentum of the condensate at its time of release due to the time of flight. Fitting a sinusoid to the cloud allows us to extract the trap frequency.

We make measurements of this form for several different Z bias currents. Plots of the results can be found in Figure 3.12. We vary the hold time such that we capture one cycle of the oscillations after the excitation and one cycle at a later time. We do this to improve the accuracy of the measurement of frequency. The frequencies of these plots can be found in Figure 3.13.

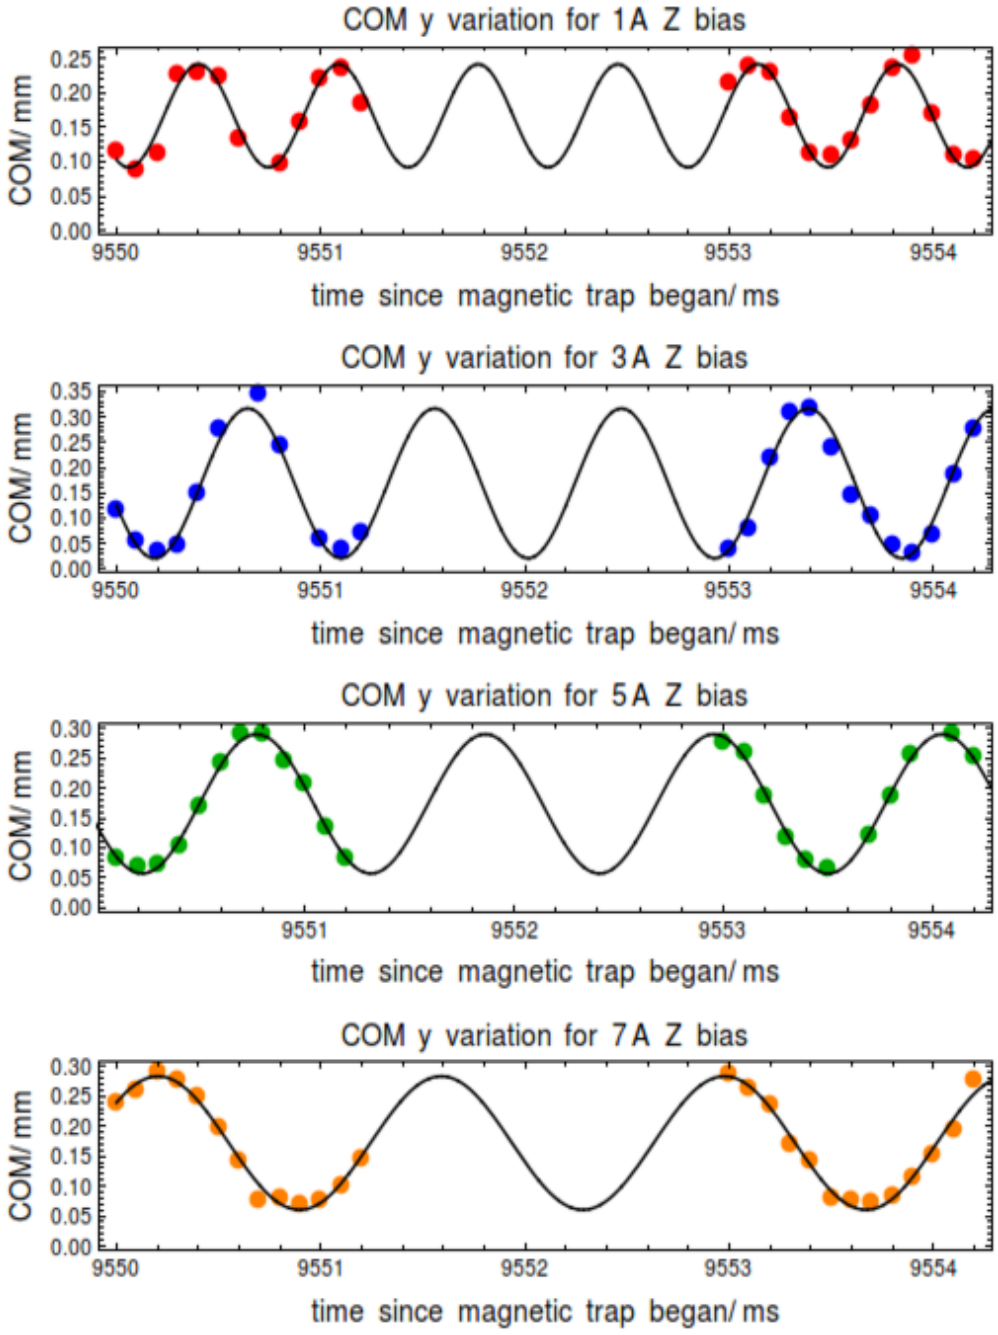

Figure 3.12: Measurement of the radial oscillation frequencies. By displacing the Y position of the condensate and releasing it suddenly we produce radial oscillations. Here we plot the centre of mass (COM) positions as a function of hold time. Above are the results of these measurements for four different traps at four different Z bias magnetic fields. Coloured points show the COM positions, the black lines show a sinosoid fitted to these points to extract a frequency.

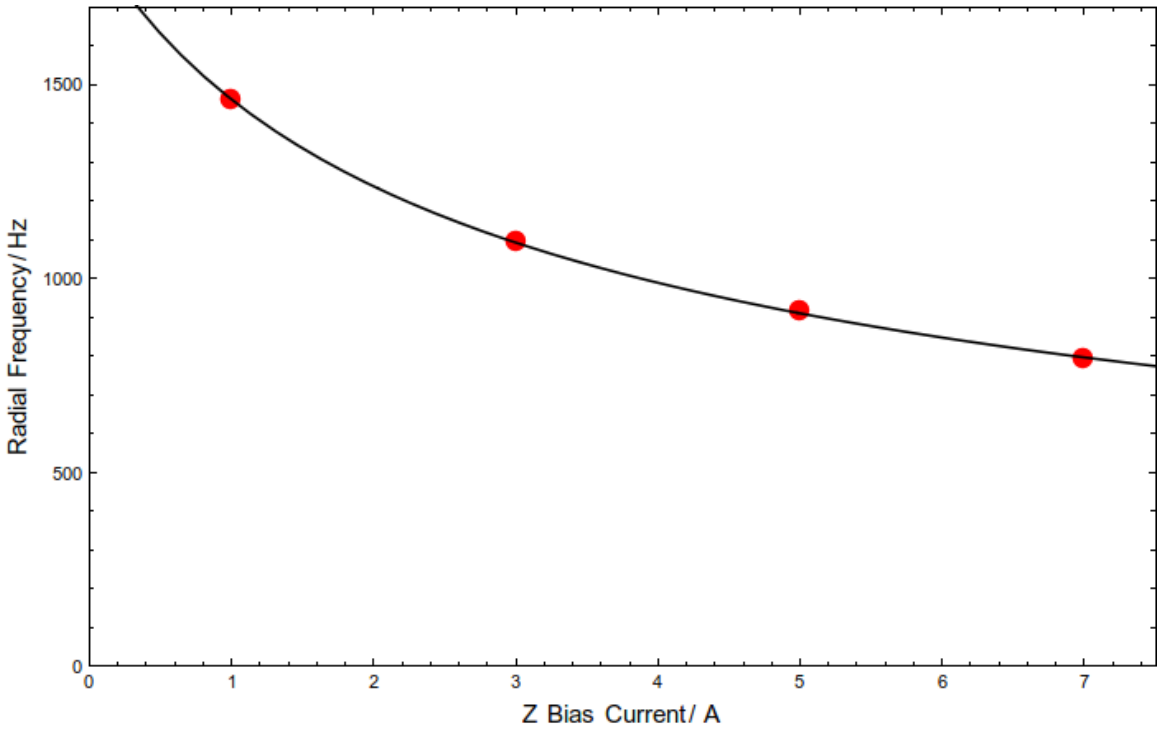

Figure 3.13: Comparing the radial frequencies of our magnetic trap at different Z bias fields. The data points correspond to frequencies found from Figure 3.12. The current is proportional to the Z bias magnetic field. The fitted line is dicussed in the text.

From the discussion in section 3.2.2 we would expect the frequency to be proportional to the inverse square root of the Z bias field. However in moving from a 1A trap to the 5A trap we have only decreased the frequency of the trap by 35%. The resolution to this is the fact that the Z bias field isn't the only source of constant magnetic field along this axis. As mentioned in 3.2.2, a slight misalignment in the X bias field means that there is contribution along the Z axis that causes the apparent offset. In Figure 3.13 we fit a line to the data where $f \propto \frac{1}{\sqrt{aI+c}}$, where f is the measured frequency of the trap, I is the current in the coils, a is a calibration factor and c is an offset field. The fit returns the offset value of 1.3 ± 0.1G, comparable with what we have measured elsewhere.

Despite this offset we continue to label the different trap widths using the current in the Z bias coils, as this is the quantity varied in the lab.

3.4.2.2 The Trap Bottom

Along with the trapping frequencies, an important figure to measure is the magnetic field at the centre of the trap, which we call the trap bottom. We specify this by giving the Zeeman splitting between the $F = 2$ magnetic sublevels. This frequency is useful when applying radio frequency magnetic fields to the system for evaporation: if we ramp the frequency below this value there will be no atoms left.

To measure this value we use the RF radiation itself. As discussed in section 2.5, the effect of applying nearly resonant RF radiation is to couple the hyperfine states of the atom, causing them to undergo Rabi oscillations between the states at a rate $\propto \sqrt{\Delta^2 + \Omega^2}$.

Because the atoms occupy a finite area in the trap, apply rf radiation exactly on resonance with the lowest point of the trap doesn't remove the largest number of atom, this happens at a slightly higher frequency, and it is this point we call the trap bottom. We do this because this point follows the the true trap minimum, and for evaporation is the more useful parameter to know. Changes in atom

number or chemical potential may shift this point relative to the true trap minimum, however this changes are expected to be small relative to our resolution. To find this point we apply rf radiation for a fixed length of time and measure the remaining number of atoms. Scanning the frequency between runs we can find the trap bottom by finding the point at which the atom number is most reduced.

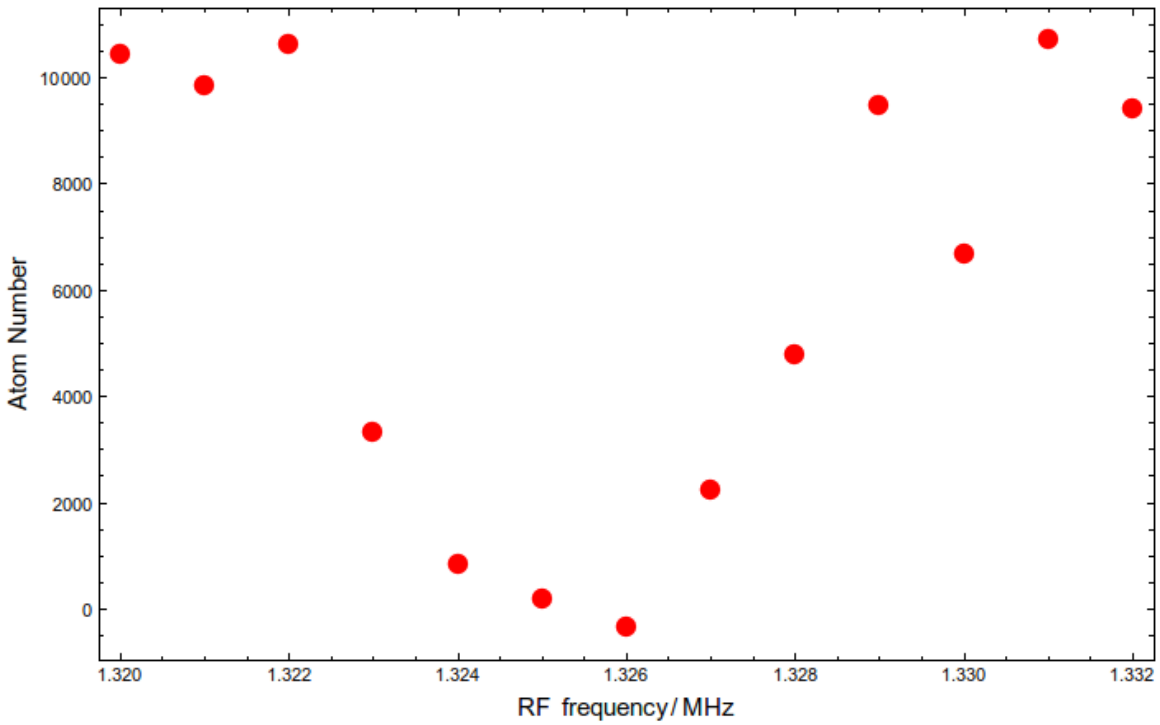

Figure 3.14: An example of how we measure the position of the trap bottom. A condensate is created and then an RF pulse is applied for 10ms. When the RF is resonant with the splitting between the Zeeman sublevels the atoms are transferred to an untrapped state and lost.

An example of a measurement of trap bottom can be found in Figure 3.14. From the shape of the atom loss curve we can calculate the trap bottom to be at 1.325 ± 0.001MHz. To get this in terms of magnetic field write the splitting between Zeeman sublevels as $\Delta f = \frac{g_F \mu_B B}{h}$ to get a field of 1.895 ± 0.001G.

It is using this technique that we confirmed that the X bias coils were misaligned relative to the axes defined by the atom chip. We measured the trap bottom as a function of X bias current and noticed the trap bottom changed. The results can be found in Figure 3.15. Fitting a line to the data points gives a contribution to the Z bias field from the X bias coils of 0.19G/A. Given the current to field calibration of the X bias coils of 3.45G/A this suggests a misalignment of 3.1 degrees.

Because our condensate occupies a finite space in the trap, it 'sees' a variation in magnetic field. This broadens the peak in atom number removed, limiting the resolution of this technique to the chemical potential.

3.4.2.3 Calibrating Fields

Although the magnetic field coils had been calibrated prior to being mounted on the chamber, the alignment may have shifted and the chamber or optical table may affect the calibration, so it is useful to use the atoms themselves to calibrate the fields.

To calibrate the bias fields we follow a similar technique to the one used in the previous section. We release the atoms from the magnetic trap and turn off the magnetic fields except the one we want

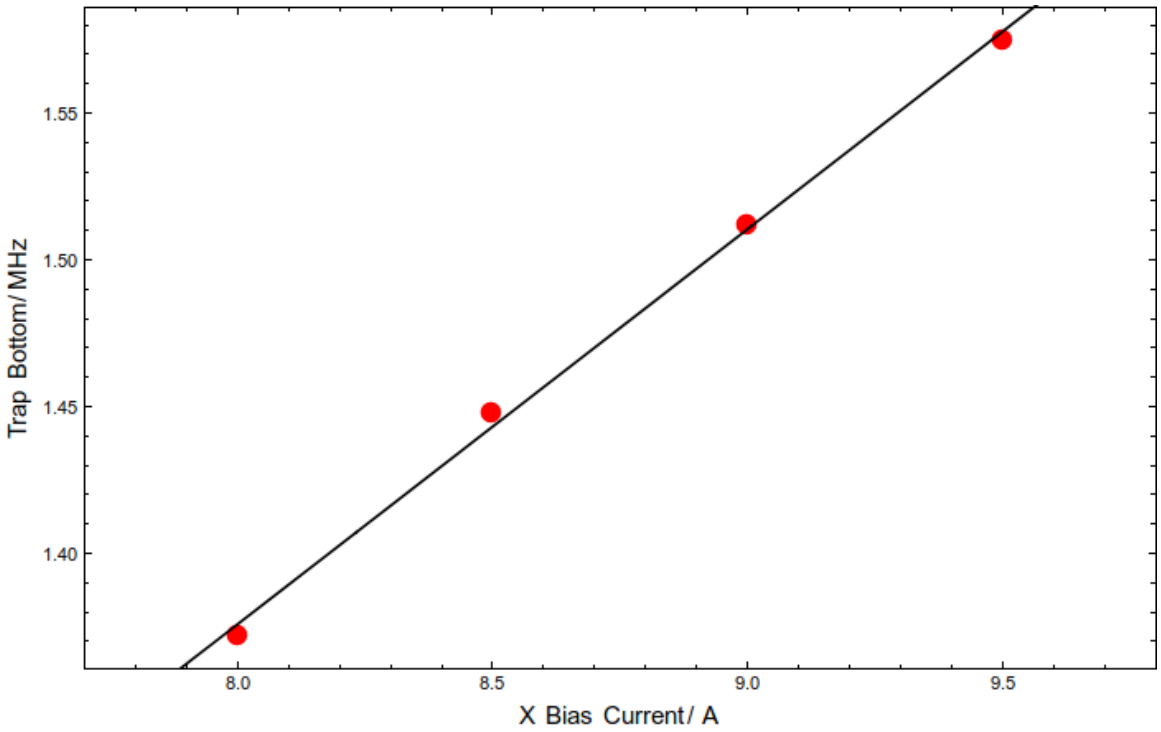

Figure 3.15: Measuring how the trap bottom changes as a function of applied X bias field. If the coils were prefectly aligned there would be no change. The change is due to a contribution to the Z magnetic field due to the X bias coils.

to measure. We then apply a pulse of RF radiation, followed by a current pulse through the small z wires of the chip to create a field gradient that through the Stern-Gerlach effect applies a different force to each m_F state. After a few ms the different m_F states separate in space and we image the atoms and measure the number in each state.

By scanning the frequency of the rf over this range and looking for maximum transfers from the $m_F = 2$ state we can find the point at which $\Delta = 0$. Measuring how the resonant point varies as a function of the current applied to the bias coils we measure a calibration curve for the coils.

When applying this technique, we need to adjust it slightly depending on which bias field we are calibrating. The rf field produced using the large Z wires on the chip can have various polarisations depending on the relative phase delay between the currents in each wire. When in phase the rf field produced is polarised parallel to the Y axis. When π out of phase it is polarised parallel to the X axis. In between these values we get elliptical polarisation. From the discussion in section 2.5 the rf only strongly couples Zeeman sublevels when its polarisation is perpendicular to the static field used as the quantisation axis. For this reason we use in phase rf currents when measuring the X bias and out of phase radiation when measuring the Y bias. Either can be used to measure the Z bias field.

An example of using this process to find the calibration of the Z bias can be found in Figure 3.16. The first set of images, Figure 3.16a shows a sequence of absorption images for the procedure above at different rf frequencies. The atom clouds have a curved shape as the magnetic field is being applied by the small wires of the atom chip (situated at the centre top of each image). The band at the bottom of most images is the $m_F = 2$ state, displaced most because it is weak field seeking. Above that some images show a small cloud of a similar pattern. This is the $m_F = 1$ state. For the image taken at 0.81MHz you can see the appearance of the $m_F = 0$ state, showing that this is closest to resonance.

Because we are coupling the Zeeman sublevels the population will oscillate between the levels if a

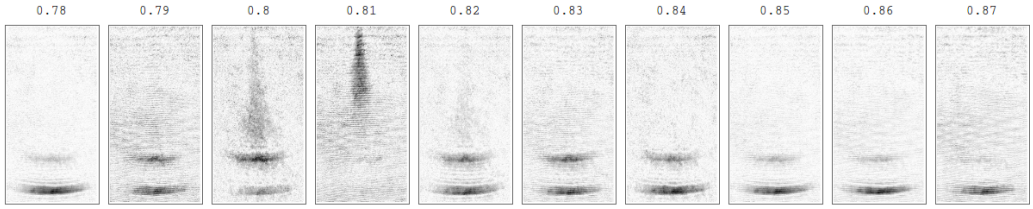

(a) Finding the rf resonance for the Z bias field

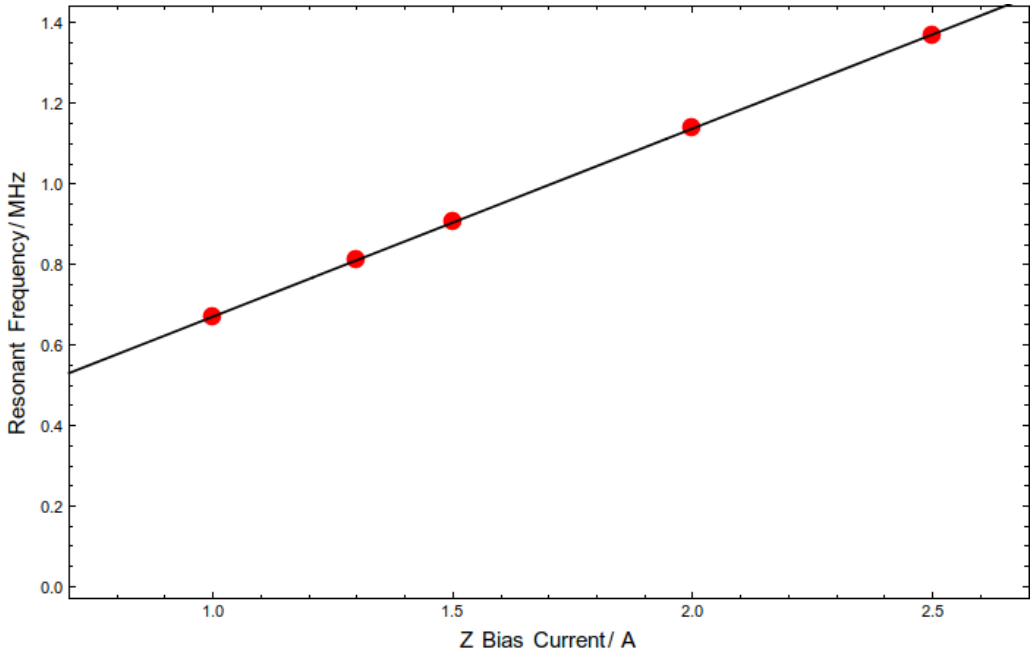

(b) Calibrating the Z bias

Figure 3.16: The proceedure we use to calibrate the Z bias. Figure 3.16a shows a sequence of absorption images taken using the imaging system along the Z axis after having applied an RF pulse followed by a magnetic field gradient to seperate out the sublevels. The numbers above each image show the frequency of the applied rf radiation. Figure 3.16b shows the calibration curve of resonant rf frequency as a function of current in the Z bias coils. The same technique is used to characterise the bias coils on the other axes.

field is applied. To ensure that we are only seeing transfer when the system is on resonance we need to know the rf power and adjust the frequencies accordingly.

To measure the direct strength of applied RF fields we use a similar technique. On resonance ($\Delta = 0$) the oscillation rate of the atoms between different hyperfine states is directly proportional to the strength of the applied RF fields. Measuring this rate is a more reliable indicator of RF field strength than calculating the rf fields from the rf currents in the wires entering the chamber. We believe this is due to capacitive coupling between wires and the ground plane of the chip itself causing current to leak out.

An example of this process can be found in Figure 3.17. The images in Figure 3.17a show a sequence of absorption images, taken on axis 1, after following the above process. Each image is taken at the same place, but for rf pulse durations separated by 0.001ms each. The coloured rectangles on these images show the locations of the atoms in different Zeeman sublevels after applying a magnetic field gradient to separate them. Red, Blue and Green correspond to the $m_F = 0, 1, 2$ levels respectively.

Counting the number of atoms in each Zeeman sublevel as a function of rf pulse duration we get Figure 3.17b. On top of these points we have fitted the analytic solution to how we expect these populations to change, treating Δ and Ω as free parameters (all three lines were fitted to the same Δ

and Ω).

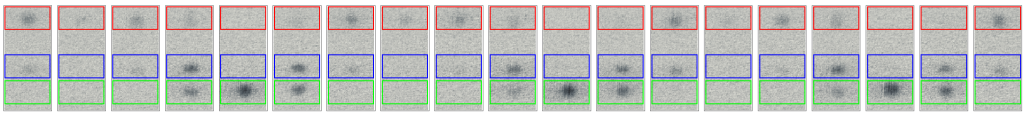

(a) Measuring the rf power, raw images

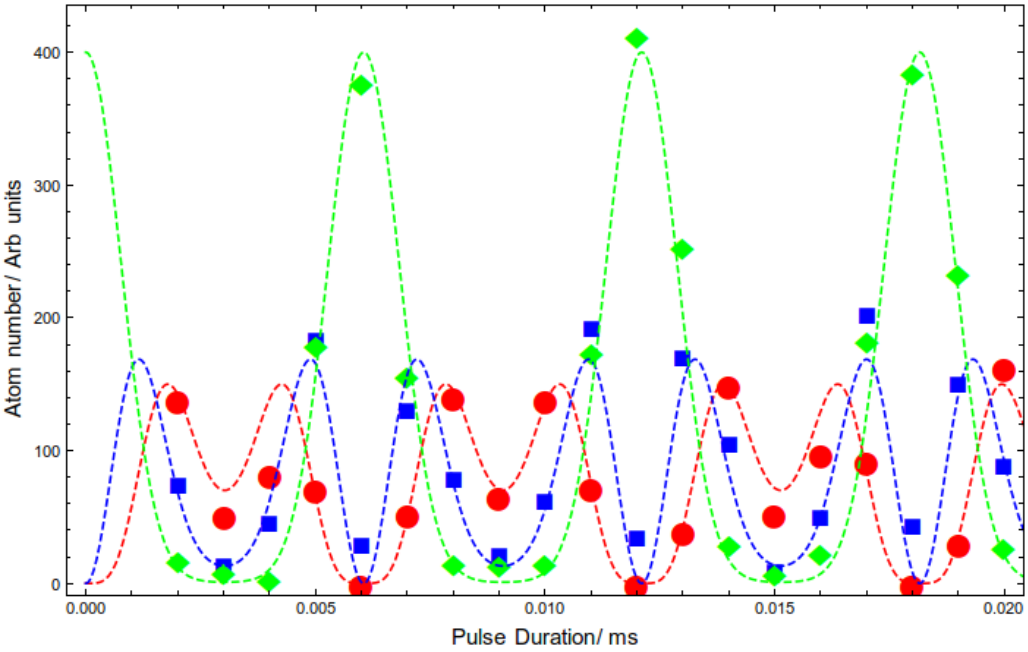

(b) Measuring the rf power, processed data

Figure 3.17: An example of how we calibrate the rf power. The images in Figure 3.17a show a sequence of absorption images, taken along the Z axis, after applying rf radiation at a fixed frequency for a varibale length of time. Each image is taken at the same place, but for rf pulse durations separate by 0.001ms each. The coloured rectangles on these images show the locations of the atoms in different Zeeman sublevels after applying a magnetic field gradient to separate them. Red, Blue and Green correspond to the $m_F = 0, 1, 2$ levels respectively. In Figure 3.17b we plot the populations of each m_F using the same colours as the previous image to label the states. The dashed lines show the predicted population for fitted values of Δ and Ω.

We can repeat this process at different amplitudes of the rf power. The result is a calibration between rf power controlled using the settings on our rf frequency generator and the power experienced by the atoms.

3.4.3 Recognising a BEC

In our experiment there are three tell-tail signs that we have produced a Bose Einstein Condensate.

3.4.3.1 Density Increase

The first sign that we have a Bose Einstein condensate is the appearance of a high-density core when we take an absorption image of the condensate. This indicates a high population of the ground state. In a gas where interactions can be neglected, the gas can be thought of as purely two modes: a macroscopically occupied ground state, and a thermal part.

As discussed in section 2.7, the shape of the ground state of the system in the case of interactions is described by a parabola for a harmonic trap, in contrast with a Gaussian distribution for a thermal

cloud.

Images and integrated density profiles of a cloud of atoms as they cross the transition temperature can be found in Figure 3.18. The left hand side of this figure shows the absorption images, the right hand side shows the integrated profiles of these clouds. As you move vertically down the temperature between shots decreases.

The first image shows the condensate above the transition temperature, where the profile remains Gaussian. In the second image a small area of high density is visible in the image and the profile is showing a small, narrow peak sticking out of the Gaussian tails. In the third image this bi-model density distribution becomes clearer as the condensed fraction of atoms increases until we reach the final image, where the majority of the atoms are in the condensate.

In this sequence of images we can see another property of our magnetic trap becoming visible. As the temperature decreases the trap becomes less symmetric. To the right of the centre a low tail of atoms can be seen along the X axis. This is due to the fragmentation of our atom chip discussed in section 3.2.2.4.

3.4.3.2 Asymmetric Expansion

An increase in the density of the atoms isn't a definitive proof of a BEC, as it could be caused by a difference in the potential only becoming apparent as the atoms cool, so we look for other signs. Another characteristic of a BEC a in non-spherical trap is a symmetric expansion once released from the magnetic trap.

A demonstration of this can be found in Figure 3.19, which shows a sequence of absorption images of a BEC after different release times from the magnetic trap.

For a distinguishable non interacting thermal gas in a harmonic trap, expansion after release from a trap is expected to be isotropic: the separability of the Boltzmann distribution means that for any position in the trap, each particle has the same momentum probability distribution. This momentum distribution is isotropic, and therefore so is the expansion.

For a BEC the situation is different. Even for a non interacting BEC we would expect asymmetric expansion: the Heisenberg uncertainty principle means that when we tightly confine an atom in one spatial direction we increase the uncertainty of the momentum in this direction. This seems to explain the asymmetric expansion we see, where a 'cigar' turns into a 'pancake'. The zero point energy of the system means that even in the ground state there is kinetic energy to make the condensate expand.

For an interacting gas we see similar behaviour - but for different reasons. Here, even for classical gases in the zero temperature limit we get asymmetric expansion. On turning off the trap, the interaction energy due to repulsive interactions between the atoms drive the cloud to expand. From the discussion in section 2.7, our condensate lies in the Thomas Fermi regime, meaning that the potential energy is much larger than the kinetic energy of the ground state and the asymmetric expansion we see is driven by this, rather than the zero point energy discussed above.

Both the classical and quantum case of asymmetric expansion can be described quantitatively using a scaling argument given in [79]. Here, the density distribution of a cloud in a harmonic trap is described using the Thomas-Fermi approximation at a given time, and at some later time it is described by a rescaled version of this distribution.

Following [79] we describe the widths of the cloud using scaling parameters λ_i, where $i = x, y, z$, so that the density distribution of the BEC is described as an inverted parabola with time-varying width $R_i(t) = R_i(0)\lambda_i(t)$. At $t = 0$ the cloud is stationary in the trap and $\lambda(0) = 1$. Upon release the

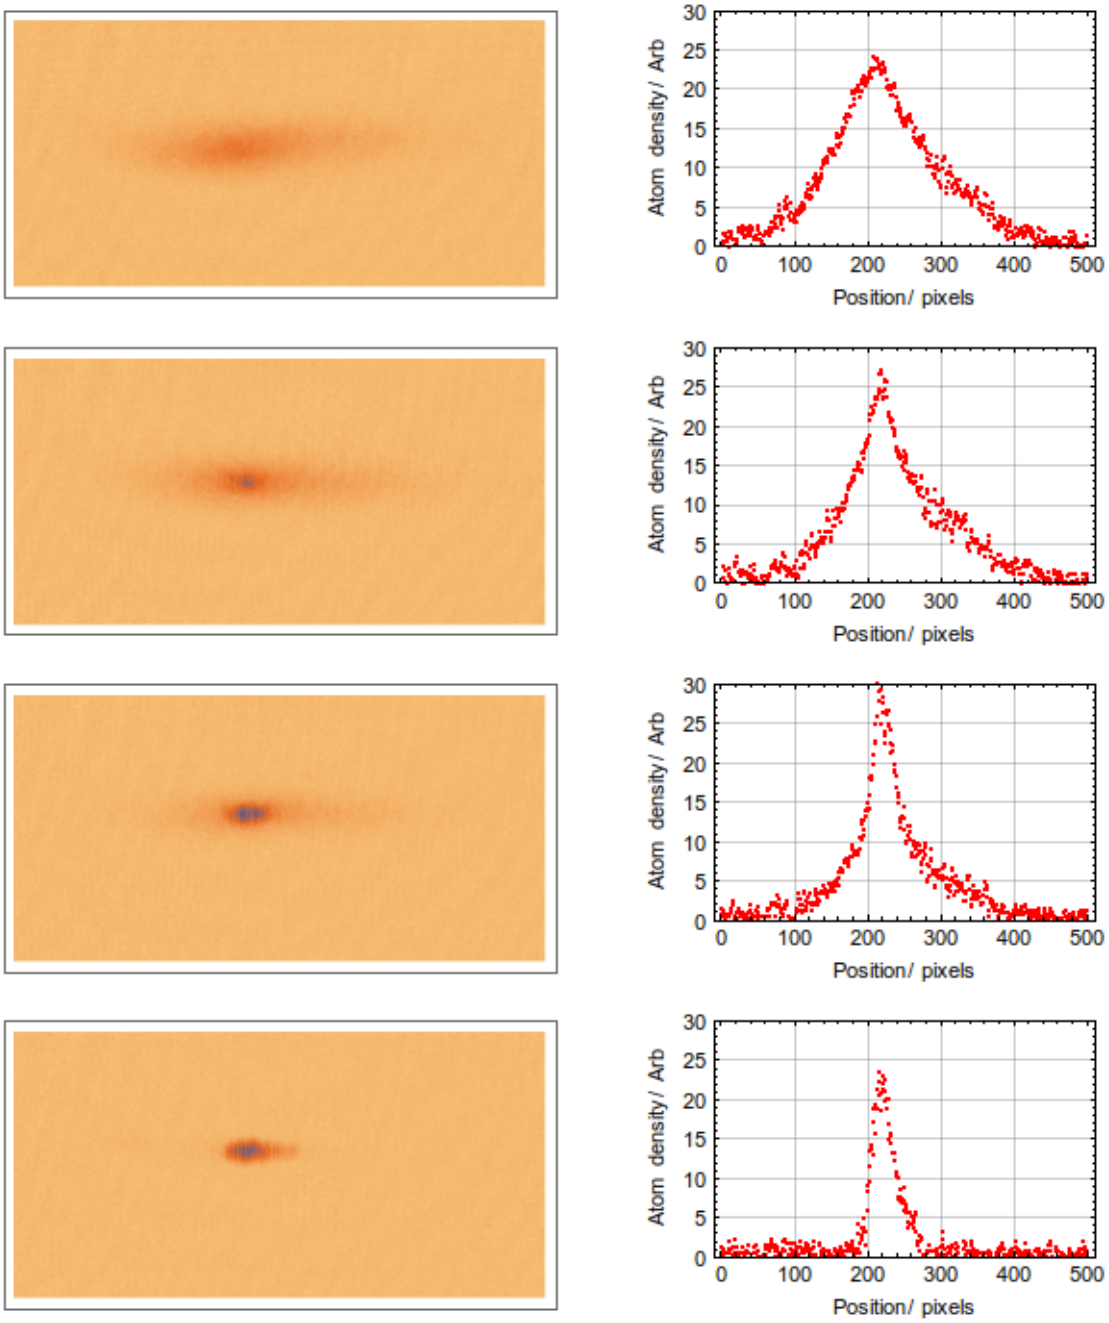

Figure 3.18: A sequence of absorption images of a BEC forming. Left hand side: absorbtion images of atoms 6ms after having been released from the magnetic trap. Right hand side: An integerated density profile of the atom cloud to the left, integrated along the Y axis of the images. X axis units correspond to the distance measured in pixels, y axis units are arbitrary, but proportional to atom density. For top to bottom: each pair of images shows the effect of lowering the end point of the RF evaporative cooling end point - this has the effect of both removing atoms and lowering the temperature. In the sequence you can see the growing condensed fraction, making the distribution of atoms narrower and more dense as the temperature falls. These images have been taken using the imaging system along the X axis of our experiment. The X axis in the images is aligned with the Z axis of our experiment and the Y axis of the images is aligned with the Y axis of our experiment.

scaling parameters evolve according to

$$\ddot{\lambda_i} = \frac{\omega_i^2}{\lambda_i \lambda_x \lambda_y \lambda_z}\,, \tag{3.8}$$

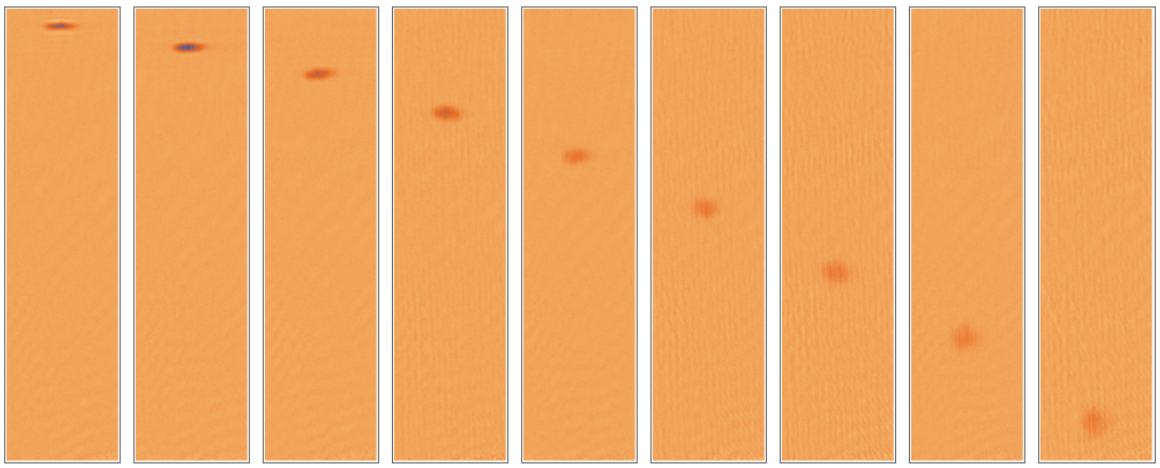

Figure 3.19: A sequence of absorption images of a BEC expanding after different times of flight. The first image on the left shows an absorption image of a BEC after 2ms time of flight, each subsequent image shows it after another 2ms time of flight. The asymmetric expansion ratio is a characteristic of a BEC. In these images the X axis corresponds to the Z axis of our experiment, and the Y axis corresponds to the Y axis of our experiment.

where the dots represent differentiation with respect to time. It is worth noting here even though this expression contains λ_x, λ_y and λ_z, the i in λ_i also corresponds to $i = x, y, z$. Working in the elongated cylindrically symmetric case, $\omega_x = \omega_y = \omega_\rho >> \omega_z$ these equations of motion simplify to

$$\ddot{\lambda_\rho} = \frac{1}{\lambda_\rho^3 \lambda_z}$$
$$\ddot{\lambda_z} = \frac{\epsilon^2}{\lambda_\rho^2 \lambda_z^2}\,,$$

where $\epsilon = \frac{\omega_z}{\omega_\rho} << 1$ and the time coordinate has been rescaled to $\tau = t\omega_\rho$.

By expanding these equations in powers of ϵ and looking at only the first order solutions we get

$$\lambda_\rho = \sqrt{1+\tau^2}$$
$$\lambda_z = 1\,,$$

showing that to first order the could expands asymmetrically. Going further than this we can use the asymmetric expansion as a measure of axial frequency of the trap. To first order, the ratio of axial width to radial width as a function of time is given by

$$\frac{W_\rho}{W_z} = \frac{\lambda_\rho \sqrt{2\mu/m\omega_\rho^2}}{\lambda_z \sqrt{2\mu/m\omega_z^2}} = \frac{\epsilon \lambda_\rho}{\lambda_z} = \sqrt{\epsilon^2 + \omega_z^2 t^2}\,. \tag{3.9}$$

So the rate of expansion of the ratio of widths is given by the axial trap frequency multiplied by time.

For the data shown in Figure 3.19 we can apply this theory. Plotting the ratio of the widths of the clouds obtained by fitting a Thomas Fermi profile to the clouds in Figure 3.19 we get Figure 3.20. Fitting a straight line to the data we can use the preceding theory to get an estimate of ω_z of

$2\pi \times 6.9 \pm 0.2$Hz, close to the value measured using other techniques[3].

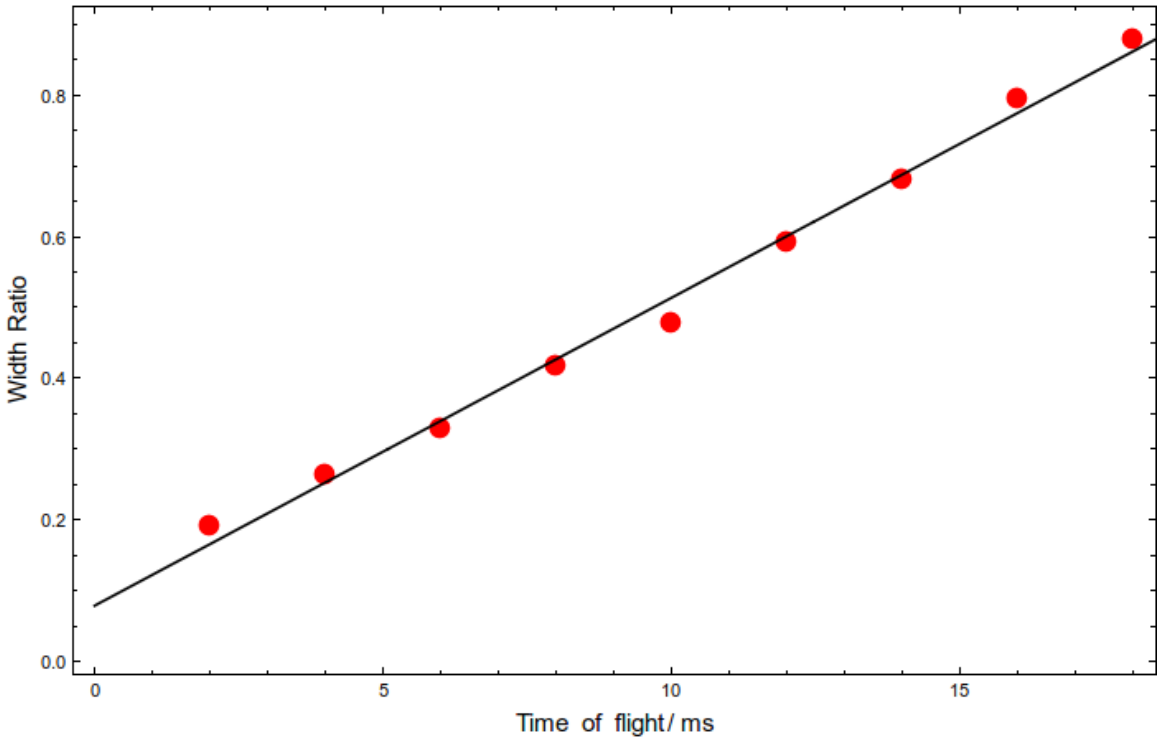

Figure 3.20: Measurements of the change in the width ratio of the data from Figure 3.19. A Thomas Fermi profile was fitted to these clouds in the horizontal and vertical directions, and the resulting ratio was plotted as a function of fall time. The line of the graph shows the result of a linear least squared fit to the data.

The above description of asymmetric expansion holds for both classical and quantum gases in the Thomas Fermi regime, where the kinetic energy is negligible. In this sense the asymmetric doesn't uniquely identify a quantum gas, rather one at close to zero temperature.

3.4.3.3 Density Fluctuations

What does demonstrate the quantum nature of the BEC is the appearance of density fluctuations along the length of the condensate. These density fluctuations appear after a time of flight expansion of the condensate, and are not just the atom bunching expected from bosons (although this has been observed [80]), indeed these density patterns are not observed while the condensate is trapped - it is something specific about the expansion that drives their appearance.

An example of clear density fringes can be found in Figure 3.21. Here we allow a BEC to expand for 14ms, then take an absorption image. Figure 3.21a shows the raw absorption image. Figure 3.21b shows the profile of the cloud, found by integrating the density onto the Z axis. Onto this profile we have fitted a bimodal distribution (discussed in the next section). It clearly shows both the thermal and condensed part of the cloud. It is also possible to observe the appearance of fringes in Figure 3.18 as the BEC critical temperature is crossed.

Figure 3.21c shows the residuals of the fitting process, showing a much higher deviation from the fit in the region of the condensate, a sign of the onset of condensation. It is worth reiterating that this isn't just a consequence of the fact that this part of the cloud has signal in it, and therefore higher noise. At early times of flight, the images in Figure 3.20 do not show fluctuations of this size.

[3]The data referred to here are from a different fragment to that used in the next chapter.

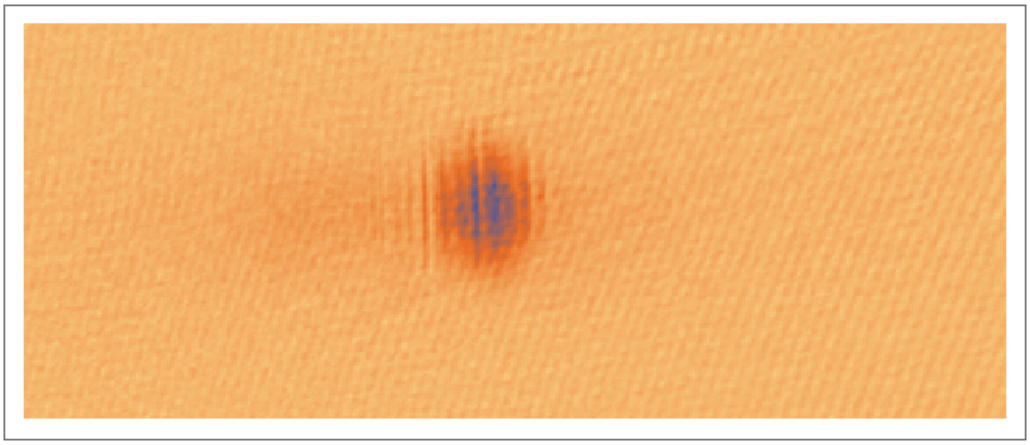

(a) Absorption image of a BEC showing density fluctuations

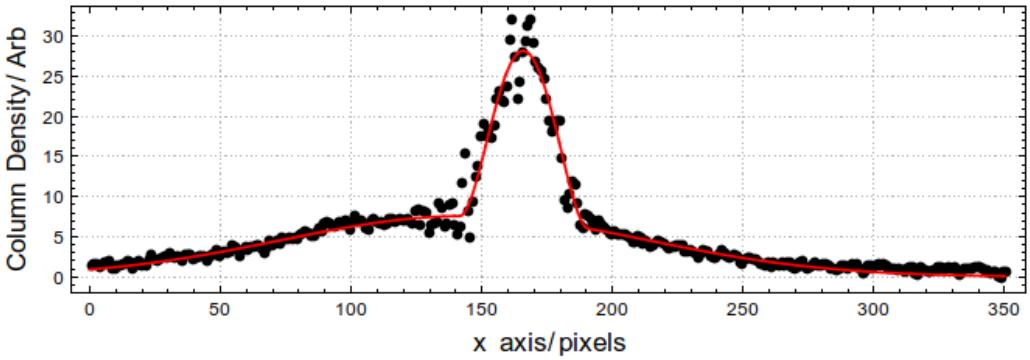

(b) Integrated profile of the same cloud

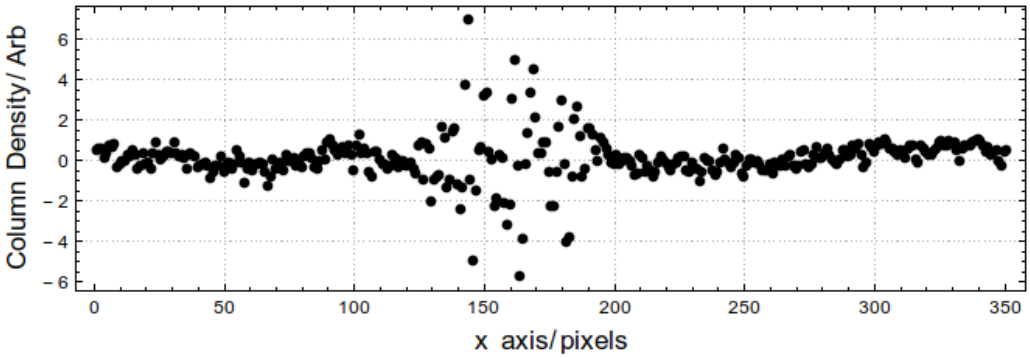

(c) Residuals of the fitting process

Figure 3.21: Examples of density fluctuations in an absorption image of a BEC allowed to expand for 14ms. Figure 3.21a shows the raw absorption image of of a BEC 14ms after being released from the trap. Figure 3.21b shows the profile of the cloud, found by integrating the density onto the Z axis. To this we have fitted a bimodal distribution, shown in red. Figure 3.21c shows the residuals of the fitting process. The deviation from the fit is much higher over the condensate. In the top image the X axis corresponds to the Z axis of the experiment and the Y axis of the image corresponds to the Y axis of the experiment.

What is interesting about these fluctuations is that they appear as deep bands of high and low density, perpendicular to the weakly confined axis of the magnetic trap. On each realisation of the experiment the positions of these modulations changes, suggesting their origin isn't related to the underlying potential.

As discussed in section 2.7.5 and [60], these density fluctuations are a result of low-lying Bogoliubov excitations in the axial direction, and can occur even in elongated three-dimensional BECs [81]. While trapped, the repulsion between atoms means that these excitations do not manifest themselves as

density modulations, but instead as phase fluctuations along the weakly confined axis of the trap. On release from the trap, these phase gradients of the BEC cause velocity gradients, resulting the observed density modulation. As noted in [81], these fluctuations will depend on atom number, temperature and the aspect ratio of the trap.

For a purely one-dimensional BEC, an analytic form for the correlation function of the axial profile after a time of flight was developed in [82]. In [83] the correlation function of such a cloud was measured and used as a probe for temperature. In this experiment they find the phase correlation length is characterised by

$$\lambda_T = \frac{2\hbar^2 n}{m k_B T}\,, \tag{3.10}$$

where n is the one-dimensional atom density in the trap, m is the mass of the atom, k_B is Boltzmann's constant and T is the temperature of the system. In the asymptotic expansion time limit, the second order correlation function of the axial density follows $g_2(x) \approx 1+\exp(-2|x|/\lambda_T)$. For our experimental parameters λ_T is in the range 5 - 30 μm, similar to the size of the density ripples we see.

It is also worth noting that similar density ripples have been observed in [84], where they are described as planar soliton defects, created by a Kibble - Zurek mechanism as the atom cloud undergoes the BEC phase transition. In this work, it is shown that the number of density modulations is related to how fast the system undergoes the phase transition, not just the final temperature.

3.4.4 Extracting Information about the Condensate from an Image

As mentioned previously, all our information about the condensate comes from absorption images of the atoms. Each of these images gives us a measure of the 2D column density of the atom distribution at a given time.

Typically we image some time after releasing the condensate from the magnetic trap. There are two reasons: (i) When in the trap, the radial width of our condensate is of the order 0.5μm, comparable with the wavelength of our imaging light and making it hard to gain any information about the radial distribution of the atoms due to diffraction of the imaging light. (ii) The magnetic fields in the trap are not aligned with the wavevector of the imaging light, meaning the atoms don't see pure σ^+ light, and the transition is no long closed. This reduces the signal to noise ratio of the measurement. To get around these issues we release the atoms from the trap and let them expand for 2ms. After this time the magnetic fields in the system have had a chance to switch fully and the cloud has had time to expand.

Currently there is no analytic expression for the density of a trapped interacting, Bose gas at close to zero temperature, however there are two limiting cases we can use for guidance. In the non interacting limit the gas will follow the standard thermodynamic distribution. In the zero temperature limit an interacting gas will follow the Thomas Fermi distribution discussed in section 2.7.

To extract information from our cloud, which is at non-zero temperature and interacting we must make some approximations. The first is that the central part of the cloud can be modelled using the Thomas Fermi distribution - this allows use to extract information about the centre of mass and the width of the condensed fraction of the atoms. The second approximation we make is that the thermal part of the cloud is expected to follow the distribution of non interacting atoms. This allows us to extract information about the temperature of the cloud. We also assume that because the thermal energy is much larger than the energy spacing of the trap, we can replace sums over states with integrals over phase space. See [85] for a discussion of the issues involved in these approximations.

These approximations come into conflict in the region of the condensate where the repulsive interaction between the condensate and the thermal atoms can have an effect. We follow [64] to consider the effect of the repulsion from the condensate on the thermal atoms. As in [64] we neglect the action of the thermal cloud on the condensate because the density of the thermal atoms is low compared with the density of the condensate.

From the discussion in section 2.7, the effective potential felt by the thermal atoms is then given by

$$V_{eff}(r) = V(r) + 2U_0 n_c(r)\,, \tag{3.11}$$

where V is the external potential felt by the atoms, n_c is the condensate density and U_0 is the interaction energy of the atoms. The factor of two comes from equation 2.48 where we treat the thermal atoms as Bogoliubov excitations. In this approximation we have assumed that the condensate can be treated as a mean field.

From section 2.7 we expect the Thomas Fermi distribution to take the form

$$n_c(r) = \mathrm{Max}\left(0, \frac{\mu - V(r)}{U_0}\right)\,, \tag{3.12}$$

where μ is the chemical potential.

Integrating over x and y we get the one-dimensional Thomas Fermi distribution

$$n_c(z) = \frac{\pi}{U_0 m\omega_r^2}\left(\mu - V(z)\right)^2\,, \tag{3.13}$$

showing that under the above approximations the condensate extends only to the regions $V(z) \leq \mu$.

The phase space distribution of a non interacting thermal atom cloud is given by

$$n_{th}(\mathbf{r},\mathbf{p}) = \frac{1}{(2\pi\hbar)^3}\frac{1}{e^{\frac{E(\mathbf{r},\mathbf{p})-\mu}{k_B T}} - 1}\,. \tag{3.14}$$

Integrating over momentum space to get the real space density we obtain

$$n_{th}(\mathbf{r}) = \frac{(2\pi m k_B T)^{3/2}}{(2\pi\hbar)^3} g_{\frac{3}{2}}\left(\exp\left(\frac{\mu - V(\mathbf{r})}{k_B T}\right)\right)\,, \tag{3.15}$$

where $g_j(a) = \sum_{n=0}^{\infty} a^n/n^j$, is the poly logarithm function.

To get the projected one-dimensional density of the atom cloud, we again have to integrate over x and y of the atom. In the region $V(z) > \mu$ (region 1) this is given by

$$n1_{th}(\mathbf{r}) = \frac{2\pi k_B T}{m\omega_r^2}\frac{(2\pi m k_B T)^{3/2}}{(2\pi\hbar)^3} g_{\frac{5}{2}}\left(\exp\left(\frac{\mu - V(z)}{k_B T}\right)\right)\,. \tag{3.16}$$

In the region $V(z) \leq \mu$, where the condensate is (region 2), the density of the cloud is given by

$$n2_{th}(\mathbf{r}) = \frac{2\pi k_B T}{m\omega_r^2}\frac{(2\pi m k_B T)^{3/2}}{(2\pi\hbar)^3}\sum_{n=1}^{\infty}\frac{1}{n^{5/2}}\left(2 - \exp\left(n\frac{\mu - V(z)}{k_B T}\right)\right)\,. \tag{3.17}$$

We assume that the cloud takes the form

$$n_{total}(z) = n_{th}(z) + n_c(z)\,, \tag{3.18}$$

where n_c is given by equation 3.13 and $n_{th} = n1_{th}$ for $V(z) > \mu$ or by $n_{th} = n2_{th}$ for $V(z) \leq \mu$.

The model takes the same chemical potential for the thermal gas and the condensate.

In order to proceed we need to know the potential, $V(z)$. As discussed in 3.2.2.4, this is not simply harmonic due to imperfections of the current flow in the chip wire. Instead we must measure the potential directly.

We determine $V(z)$ by imaging the density distribution $n(z)$ of a thermal cloud at 550nK. In thermal equilibrium the Boltzmann factor gives $n(z) \propto \exp(-V(z)/k_B T)$. Inverting this we get $V(z) = -k_B T \log(n(z)/n_0)$ where n_0 is a normalisation factor.

The result of this measurement can be found in Figure 3.22. Here the black points show the measured potential shape in units of energy. To get a form of the potential we can use we apply a Fourier filter to remove high frequency components of the potential given by the scatter in atom number. To the filtered profile we fit an interpolating function, plotted in blue. It is this interpolating function we use as V.

On top of this potential we have fitted two harmonic potentials. In red we show the harmonic fit to just the lowest part of the potential where the condensate sits. The measured frequency of the fit is 9.7Hz, similar to the condensate centre-of-mass frequency measured in the next chapter. The green line shows the harmonic fit to the potential ignoring the region where the condensate sits. The measured frequency of this fit is 3.2Hz.

We also show the regions of the trap where the condensate and the thermal fraction occupy, indicated by the vertical lines. The indicated lines for the condensate show its average end to end length. The lines for the thermal cloud show the average Gaussian width.

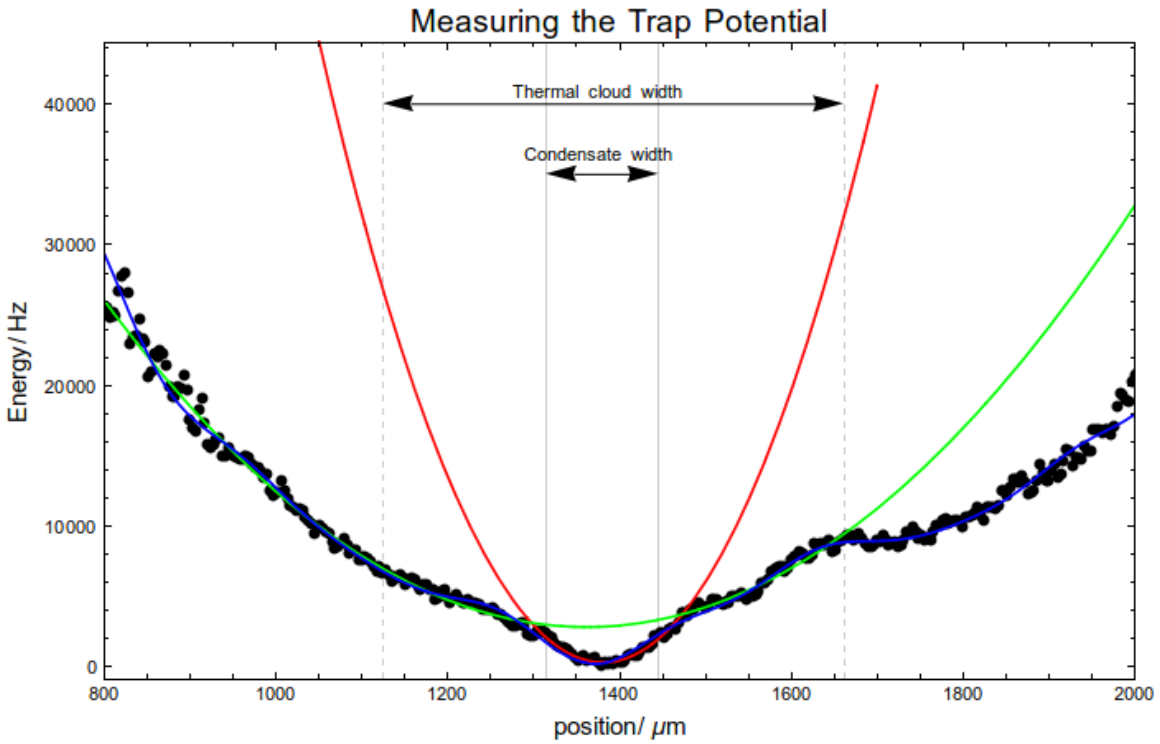

Figure 3.22: An example of the measurement of the axial potential around its minimum. The black dots show the inverted density measurement, the red lines shows a harmonic fit to the region at the trap minimum where the condensate forms. The green line shows a harmonic fit to the region around the condensate where we would expect the wings of the thermal cloud to sit. From this image we extract a frequency of 3.2Hz for the region the thermal atoms would see. Here the an image taken along the experiments X axis has projected onto the experiment's Z axis.

It is worth noting the our analysis applies to the distribution of thermal atoms in the trap, whereas our images were taken after 2ms of expansion. However, we find that in the z direction this brief expansion has negligible effect. Using equation 3.7, in the region occupied by thermal atoms, the trap frequency is found to be 3.2Hz, so we would expect a change in width of only 0.5% due to the expansion.

It is also worth noting that in using this model we have assumed that the system is in thermal equilibrium. For measurements of oscillations in centre of mass and width this assumption is not true, however the energy of the oscillations is very small relative to the thermal energy, and therefore we consider thermal equilibrium to be valid. Smaller scale properties, such as centre of mass position and width of the condensate are found by allowing the position and width of the condensate in the model above to vary slightly relative to the equilibrium position.

To fit to the clouds we use Mathematica's NonlinearModelFit function. Examples of fits to the cloud and extracted information can be found in Figure 3.23. Here we take three clouds that roughly cover the range of our measurements and fit to them. Displayed on each of the images are both the extracted profile and the fit. The red lines show the results of the full fits, and the green and blue lines show only the thermal and condensed parts of the fit respectively. The temperatures extracted from the fits are displayed above each image. The fits return the temperature with 68% confidence intervals of about 4nK.

To check the validity of the method presented here we use a second method to measure the temperature to compare the two. At the edge of the thermal cloud, where the density is lowest, we would expect the effect of the interactions between atoms to be weakest. Here, the density distribution is expected to follow that of an ideal Bose gas. By removing the central section of the profile and fitting a poly logarithm Gaussian to the wings of the cloud we can get a second measure of temperature. For a harmonic trap of frequency ω we would expect the width of the Gaussian to follow

$$\sigma^2 = \frac{k_B T}{m\omega^2}\,. \tag{3.19}$$

If we then take the potential experienced by the atoms at the wings to be harmonic with a frequency of 3.2Hz, from the width of the fit we can find the temperature. Plotting the temperature obtained by fitting to the wings against that obtained using the full semi ideal fit we get Figure 3.24.

From this figure we can see that at high temperatures (wide clouds) the two temperatures agree. At low temperatures where the width of the cloud is narrow it appears fitting to the wings underestimates the temperature. We believe this is due to the fact that when the atoms are colder they lie in the range of the trap where the frequency is higher.

3.4.5 BEC Lifetime

There are three things which limit our ability to hold the condensate for arbitrary lengths of time: chip wire heating, atom loss and atom heating.

3.4.5.1 Chip Wire Heating

The first of these, the heating of the chip wires is a technical limitation that has in the past damaged atom chips. Passing current in the wires causes them to heat, which can only flow away through the chip and into the substructure. To monitor the heating of the chip wires, we monitor the changes in their resistance. The increase in resistance of gold with temperature is well characterised.

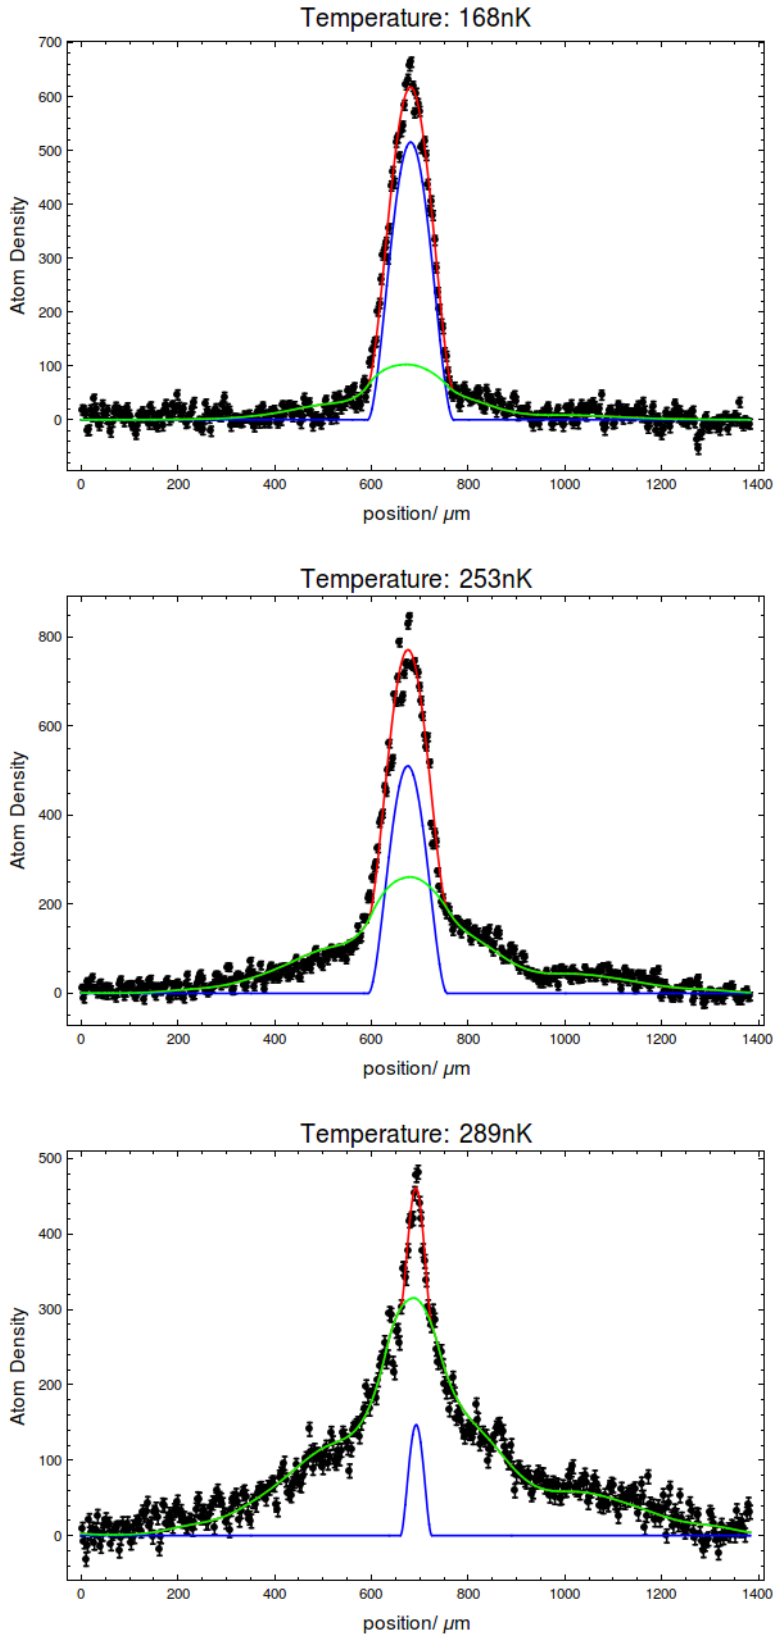

Figure 3.23: Examples of how we fit to the one-dimensional projected density of a atom cloud. The black points show the measured density, complete with errors. The red line shows the full fitted profile, the green line shows just the thermal part of the fitted profile and the blue line shows just the condensed part of the profile. The profiles come from absorbtion images of clounds prepared in the 1A magnetic trap, imaged along the X axis of the experiment, then with their density profile projected onto the Z axis of the experiment.

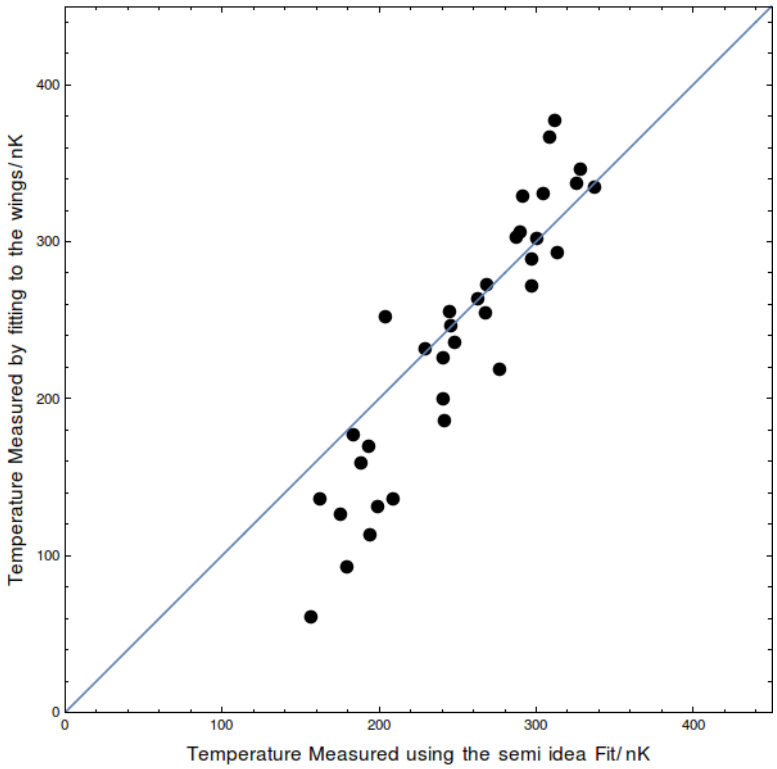

Figure 3.24: Comparing the temperatures obtained from a semi idea fit with that obtained by fitting a poly logarithm Gaussion to just the wings of the distribution, and assuming the trap is harmonic with frequency 3.2Hz. The blue line on the graph shows the line $y = x$.

We observe two types of behaviour of resistance in time for a fixed applied current. For low currents the resistance increases up to a certain point then plateaus when the rate at which heat is conducted away matches the power dissipated into the wire. At higher currents we see a thermal runaway. Here the increase in resistance with temperature causes more heat to be dissipated in the wires, in turn increasing the resistance in a self reinforcing way that ultimately destroys the wire.

Running at 1A in each of the small z wires we observe the second kind of behaviour and what looks like part of an uncontrolled runaway. For this reason we are wary of running current in the chip wires for more then 12 seconds, and this places an upper limit on how long we can hold a BEC. Given that our loading and evaporation sequence run for approximately 9 seconds this leaves 3 seconds for an experiment. However there are other limits.

3.4.5.2 Atom Loss

A second limitation is atom loss. The atoms in our magnetic trap can be hit by background gas molecules in our vacuum chamber causing both spin flips of the atoms and momentum transfer. Both eject atoms from the trap. This loss affects all our trapped atoms - both thermal and condensed equally, making it is easy to measure. At the pressures in our main chamber we see a background loss lifetime of approximately 13 seconds - long enough that it usually doesn't affect the experiments on BECs.

3.4.5.3 Heating

A more serious issue for BEC experiments is heating. In our experiment a BEC held in a seemingly stationary magnetic trap will heat at a rate of roughly 100nK/s. Given that a typical BEC transition

temperature in our experiment is around 300nK this does limit how long we can hold it for. It is believed that this heating rate comes from acoustic vibrations of the experiment itself, however attempts to measure these vibrations, or to reduce them, have so far been futile.

It is worth noting that the while this heating rate limits how long we can run an experiment for, our heating is comparable with those in similar experiments.

CHAPTER 4

RESULTS

The previous chapter ends with a description of how we produce a BEC, and how we measure some of it's properties. In this chapter we discuss what we have done with the BEC after that point.

4.1 Overview

As discussed in section 2.7, in the absence of interaction we would expect the BEC to act as a superfluid, making the observation of damping of it's centre-of-mass motion a surprise. However in deriving the property of superfluidity of the condensate we ignored interactions between the condensed fraction of the atoms and the non-condensed fraction. It is these interactions that cause the damping.

Under the Bogoliubov approximation the lowest order interaction that can cause damping is due to the scattering between a condensate excitation and a thermal excitation resulting in a higher energy thermal excitation. This process is known as Landau Damping. Because the scattering involves thermally distributed atoms we expect the damping rate to depend on the temperature. It is this temperature dependence that we measure.

4.2 The Experiment

4.2.1 Overview

Each experiment begins with the creation of a condensate. Here we have some freedom in the parameters of the magnetic trap. As discussed in section 3.2.2, by changing the strength of the Z bias field we can tune the axial trapping frequency. This is done after the atoms have been collected into the magnetic trap, but before evaporation begins. The experiments discussed here take place at one of two values: 1A and 5A currents in the Zbias, corresponding to radial trapping frequencies of 1.4kHz and 0.9kHz respectively. In what follows we will label these the 1A and 5A datasets. Other than this difference the procedure remains the same.

We then set the temperature of the condensate to a predetermined value and set the condensate oscillating relative to the thermal cloud. Taking a sequence of images of the condensate at various times after these oscillations have been excited we can map out how the centre-of-mass (COM) of the condensate damps as a function of the properties of the atomic cloud.

In this section we begin by describing how a single COM data point is collected. From there we describe how a 'run' of data points is collected, analysed and aggregated. We then discuss the results of these aggregated data runs, and how the changing of various experimental parameters affects the

properties of the atom cloud. Finally we present the overall results of the damping experiments as a whole.

4.2.2 A Single Data Point

4.2.2.1 How we take a Data Point

The collection of a single data point begins with the creation of a BEC. To do this we follow the procedure in section 3.3, which ends with an evaporation ramp and a BEC with rf radiation applied 10kHz above the trap bottom. Over the next 100ms we raise this rf frequency to a fixed value Δf above the trap bottom. In this time the atomic cloud thermalises to a temperature that depends on the value of Δf. The dependence of T on Δf is discussed in section 4.2.4.1.

We then turn off the rf radiation to reduce unwanted atoms loss while we excite the centre-of-mass of the condensate by modulating the current sinusoidally in one of the end wires. We use an amplitude of 0.1A to produce a magnetic field gradient of 10mG/cm in the axial direction at the position of the atom cloud, which displaces the position of the trap minimum by 3μm in the axial direction. Due to the high trapping frequencies in the radial directions, the displacement due to the current in the end wires is less than a nanometre and negligible. We oscillate this current at 10Hz for two complete cycles. We find that driving the condensate for longer than two cycles does not increase the oscillation amplitude. Modelling the system as a single particle in a harmonic potential shows that after two cycles, the oscillation amplitude should be 20μm. Displacing the particle by more than this amount moves it into a region where the trapping potential is no longer harmonic and therefore the drive is no longer on resonance.

After driving the condensate the rf radiation is applied again at Δf above the trap bottom. We do this to keep the temperature constant over the course of a run, at the expense of losing atoms over the run. The constant rf is high enough that it doesn't interfere with the COM oscillations, but does provide an evaporation mechanism for high energy atoms.

We then allow the condensate to oscillate freely for a fixed period of time Δt, before turning off the magnetic trap, letting it expand for 2ms, and then we take an absorption image. The magnetic trap is extinguished by turning off the small chip wires over 3μs, which is significantly faster than other time scales.

4.2.2.2 Examples of the Data Points

At the end of this procedure we end up with two images: an absorption image where atoms are present and a background image taken after the atoms have fallen away. We then process these images in a method described in section 2.6 and analyse the resulting profile of the images using the method described in section 3.4.4.

To give an overview of the process, we repeat the procedure for 3 clouds taken in three different runs at three different temperatures. In Figure 4.1 we show the raw foreground images, the raw background images and the processed absorption image. In the foreground and background images the intensities at each pixel are proportional to the intensity of light falling on each pixel. The shadow cast by the atoms is clear, but sits on a large background. After processing these images to produce the absorption image, the intensity is proportional to the column density of the atoms.

In these images you can begin to see the differences between the one amp trap and the five amp trap, where the width of the condensate along y is noticeably bigger. This is due to the radial

frequency of the trap being higher in the one amp trap. The differences in backgrounds between the one amp and five amp traps are due to the images being taken on different days. It appears as if a piece of dust ended up on the imaging optic on the day the five amp data were taken. Small motions of the camera between days (probably owing to us moving things in the experiment) account for the slight differences in positions.

Given the cleaned and processed absorption images we aim to extract information about the atom cloud. We do this by fitting to the profile of the cloud: the density of the cloud projected onto the z axis, and equation 3.18 is fitted to the points. An example of the fits to the same three clouds can be found in Figure 4.2. Above each of these images is the extracted temperature of the cloud. In each image the red curve corresponds to the total profile fit, the green and blue curves correspond to the thermal and condensed parts of the fit respectively.

Each of these profiles gives us information about the COM of the condensate and thermal clouds, the number of atoms, the temperature and the chemical potential of the atom cloud.

4.2.2.3 Errors in extracted information

When we fit to a profile of an image to extract information, we also assign a confidence interval to the parameter. The method for doing this is outlined in Appendix B. Unless otherwise noted the errors quoted are taken to be the 68% confidence interval of the parameter found using these techniques.

A special case in our experiment is the COM position of the condensate. Typically the error in this type of fitting is of the order 0.3μm for a single shot. It has been found that when we take a sequence of images of a condensate where no excitation has taken place, there is some scatter in its reported position larger than this value.

An example of this can be found in Figure 4.3a where we have taken images of 120 condensates created one after the other. The position is extracted using the procedure above. The scatter on the points is much larger than the error predicted just from fitting.

For a condensate that should be stationary the positions we measure are normally distributed with width $\sigma \approx 1.4\mu$m. An example of a histogram of the points previously collected can be found in Figure 4.3b.

From Figure 4.3a it looks like there is some correlation between subsequent runs, although there is no overall pattern. We can quantify this by looking at the autocorrelation of the positions. A measure of the autocorrelation of the position shows that it takes about five samples before the correlation between shots drops to zero (this corresponds to about 2 minutes). A plot of the autocorrelation can be found in Figure 4.3c.

The drift is thought to be due to slow drifts in the imaging system because the same sort of scatter is present even if we vary the time over which the condensate is held in the magnetic trap. If the scatter in position were due to some sort of shift in the magnetic fields or an excitation of some mode of the condensate it would be expected to oscillate with hold time.

Attempts to remove the drift through stabilising the imaging system seemed to have no effect on the scatter. We were also unable to compare the position of the cloud with the position of the chip itself as there are no features that we could focus on.

As we were unable to remove this drift we instead included it as the dominant error in our knowledge of position and quote our COM position as having a one sigma confidence interval of 1.4μm. We randomise the order in which we take our images, allowing us to remove correlations in shot to shot position that would otherwise add an unwanted slope to our data. We will further justify

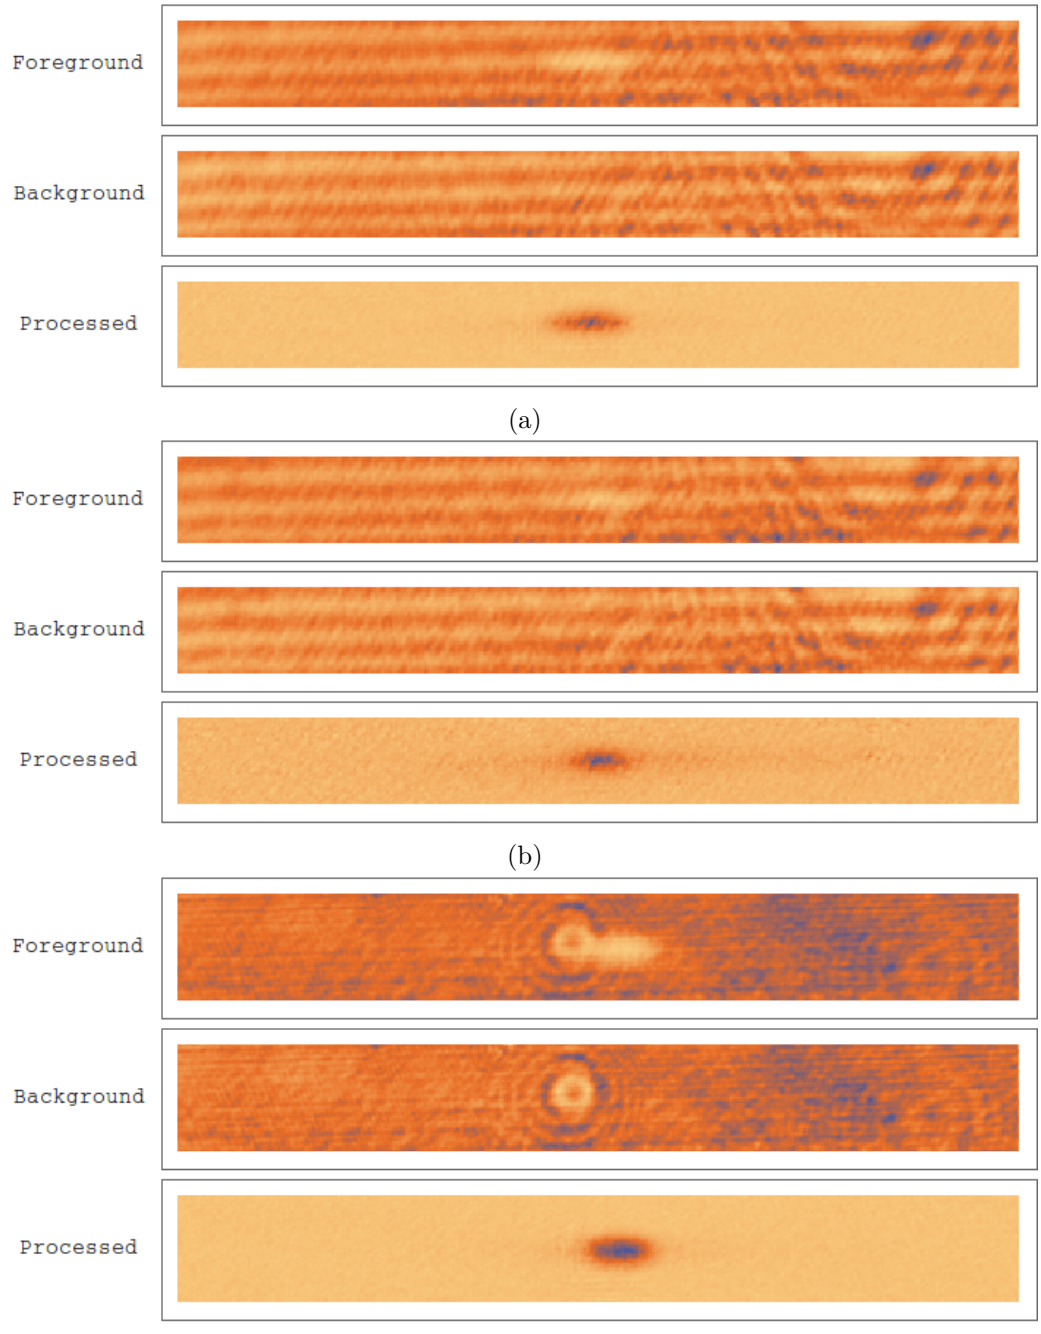

Figure 4.1: Here we show examples of how we process the foreground and background images taken during an absorption image into a measure of the column density of atoms. Figure 4.1a and Figure 4.1b show images taken from the 1A dataset, Figure 4.1c shows an image of a shot from the 5A dataset.

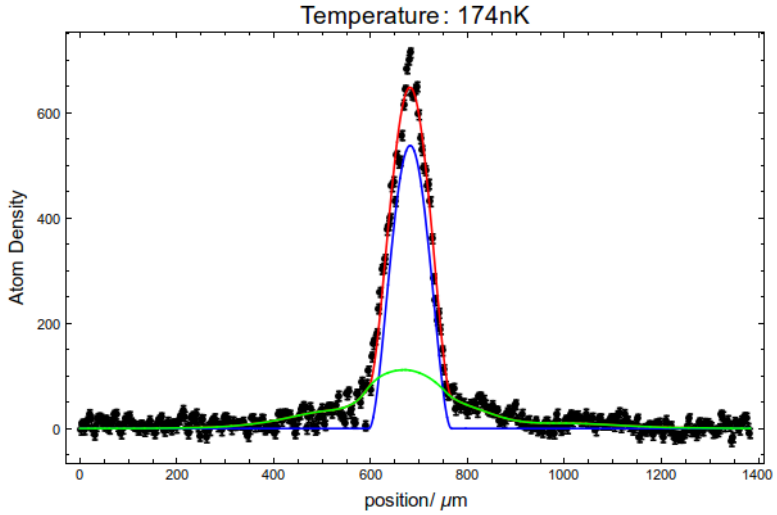

(a)

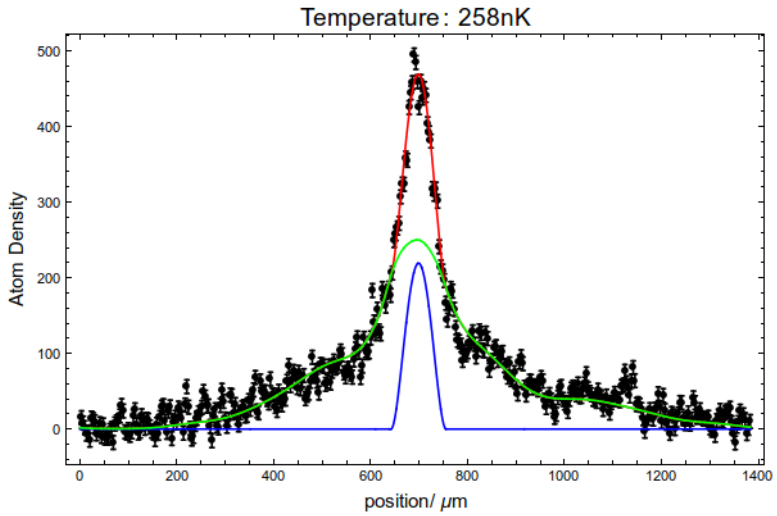

(b)

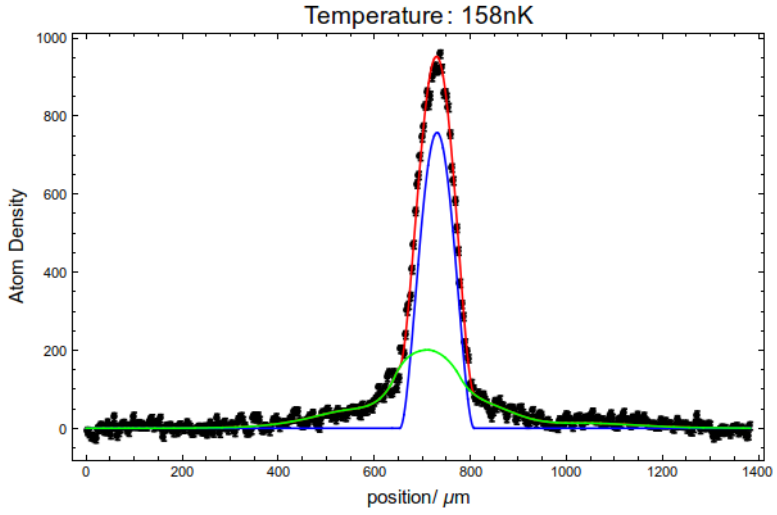

(c)

Figure 4.2: Examples of how we fit to the profile of absorption images. Figure 4.2a, Figure 4.2b and Figure 4.2c correspond to the images shown in Figure 4.1a, Figure 4.1b and Figure 4.1c respectively.

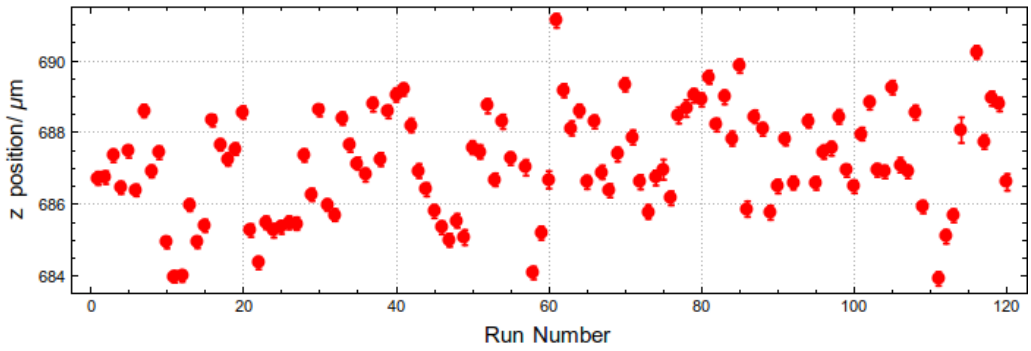

(a)

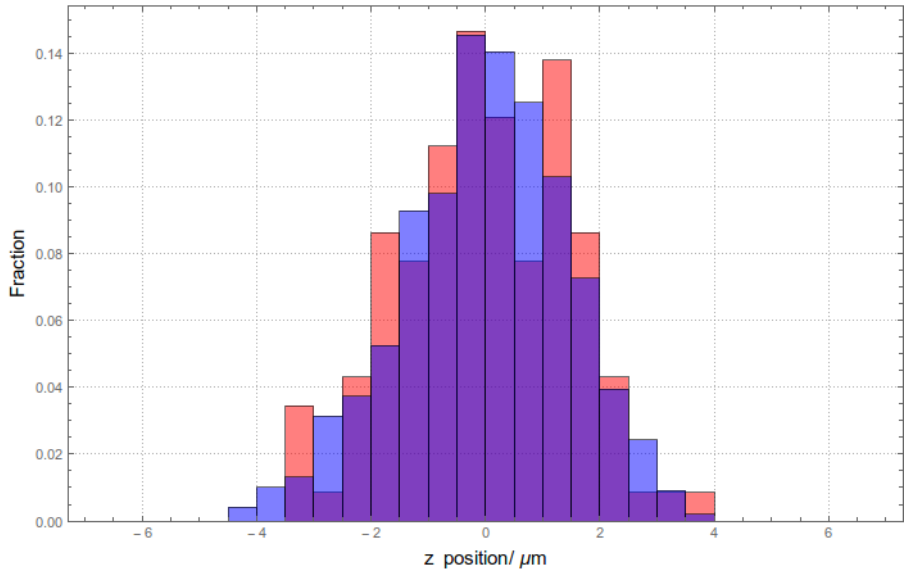

(b)

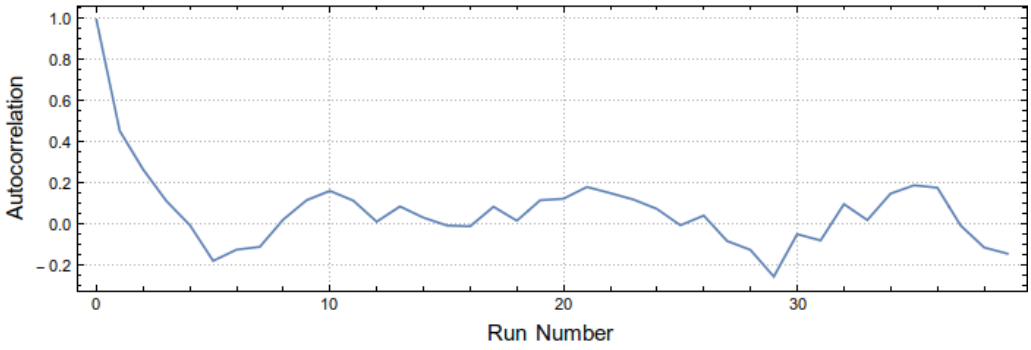

(c)

Figure 4.3: An example of drifts in our imaging system. Figure 4.3a shows 120 condensate COM positions collected consecutively. This image also shows the errors from the fits used to extract the COM. Figure 4.3b shows a histogram of the COM positions. Figure 4.3c shows the autocorrelation of these images.

the use of this error in section 4.2.4.4.

4.2.3 A Run

For one set of oscillation data, called a run, we repeat the procedure above for Δt in the range 0-400ms, spaced by 12.5ms for a total of 32 data points. Because our imaging technique is destructive, we have to create a new BEC for each data point. Each run takes about 30 minutes to complete, during which we try to hold the parameters of the experiment constant so that the only change between runs is Δt. In order to decouple any drifts in our experiment that may cause the properties of the cloud to change with time from the change of Δt we randomise the order in which we choose Δt.

We choose the time range to be long enough that we can cover several oscillations, but short enough that the heating of the atom cloud does not become significant. The number of points we take is limited by the time over which the experiment can repeatedly reproduce a condensate with similar parameters.

When taking a run, there are usually a few shots that do not turn out well. Sometimes this is due to the imaging system not firing correctly, sometimes the atom number is too low or high and sometimes some other parameter changes in an unexpected way. These data points are removed before processing a run and leave an average of 30 data points per run.

To infer the damping of the system, we look at the COM of the condensate as a function of t. We expect it to follow a damped sinusoid,

$$z(t) = Ae^{-\gamma t}\sin(\omega t + \phi) + B\,, \tag{4.1}$$

the equation of motion for a damped harmonic oscillator. Here A is the amplitude of the oscillation, B is the central position, ω is the frequency of oscillation, t is time and ϕ is a phase offset. To fit to this curve we perform a numerical minimisation of the residues[1].

As an example of the data taken over a run, we look at the how the COM position of the condensate, the temperature and the atom number of each data point in three runs vary as a function of Δt. The runs shown are those from which the data from Figure 4.1 and Figure 4.2 were taken. The result can be found in Figure 4.4, Figure 4.5 and Figure 4.6. On each of the figures we also show all of the absorption images for that run to give a sense of scale. The oscillations are only around 10μm in amplitude, whereas the condensate is around 100μm in length and the thermal cloud extends 500μm.

These figures show a good representation of the data collected. There are two trends worth noting: (i) in Figure 4.4 the temperature reduces slightly as a function of time whereas in Figure 4.5 it seems to increase. (ii) The atom number in Figure 4.4 goes down as a function of time whereas for Figure 4.5 it stays constant.

Both these effects are results of where we position the rf during the run. In Figure 4.4 it is set to 10kHz above the trap bottom, whereas for Figure 4.5 it was 300kHz above the trap bottom. When the rf knife is lower it has the effect of removing energetic atoms and lowering the temperature. When it is high the atoms heat up due to the intrinsic heating rate of our trap. In Figure 4.6 the rf is set to 50kHz above the trap bottom, where neither effect dominates and atom number and temperature stay about constant.

While these effects are noticeable across a run, they are small compared to differences between runs and small enough that their effects on a measurement of damping will be negligible. In the rest

[1] Again using Mathematica's NonlinearModelFit function.

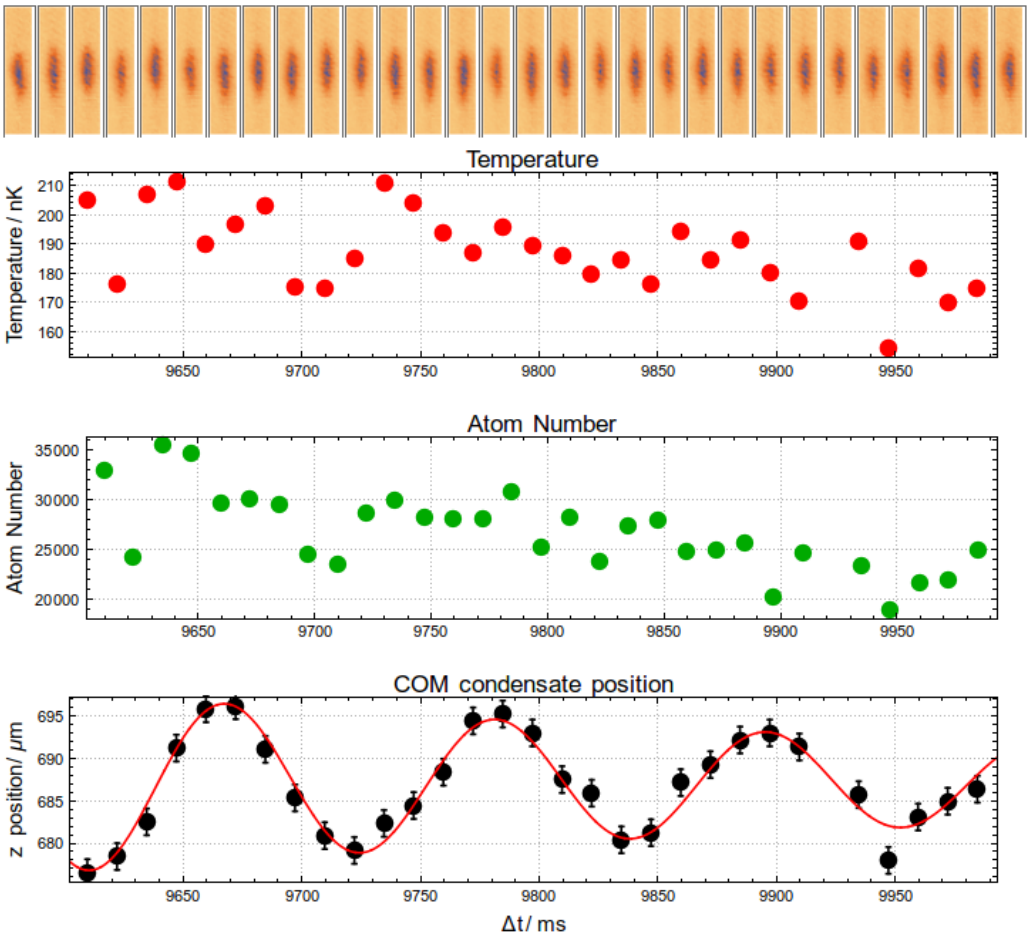

Figure 4.4: Here we show examples of the information extracted from a processed run. We show how the temperature, the atom number and the COM position of the condensate vary as a function of Δt. For the COM position, the red line shows the result of fitting equation 4.1 to the datapoints. This run contains the datapoint presented in Figure 4.2a.

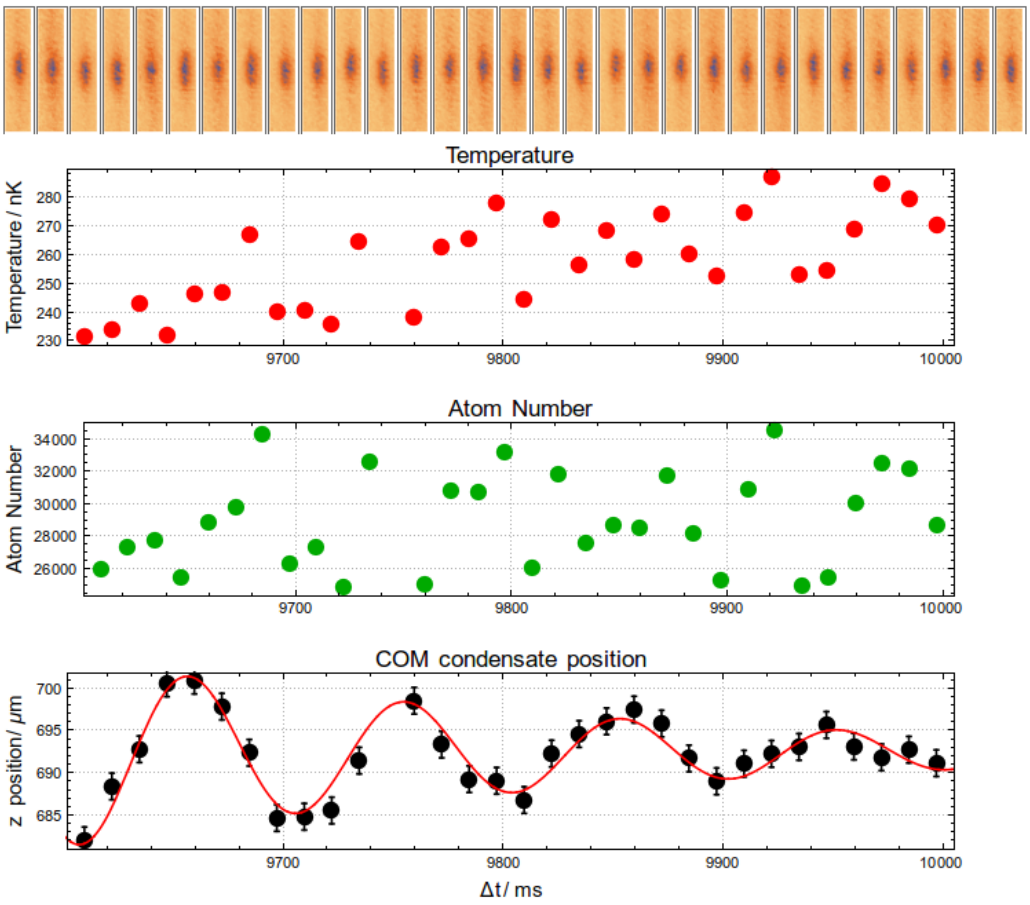

Figure 4.5: Here we show examples of the information extracted from a processed run. We show how the temperature, the atom number and the COM position of the condensate vary as a function of Δt. For the COM position, the red line shows the result of fitting equation 4.1 to the datapoints. This run contains the datapoint presented in Figure 4.2b.

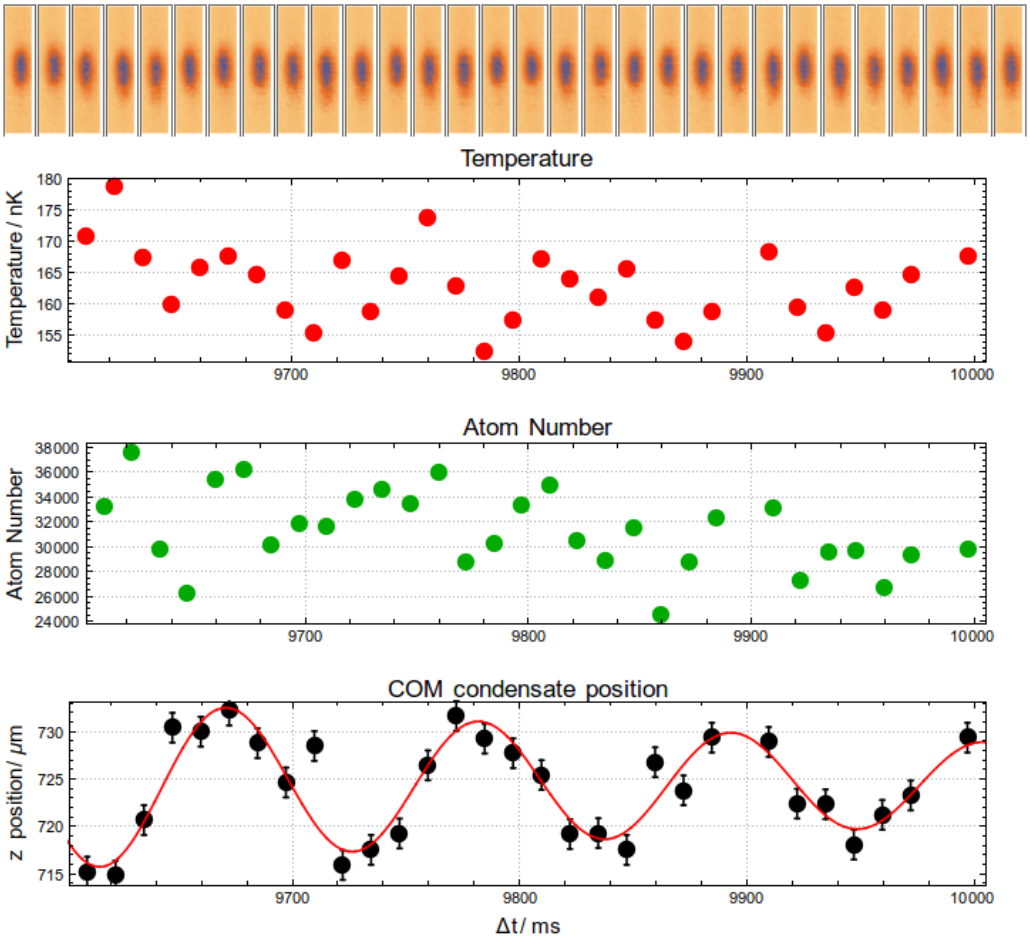

Figure 4.6: Here we show examples of the information extracted from a processed run. We show how the temperature, the atom number and the COM position of the condensate vary as a function of Δt. For the COM position, the red line shows the result of fitting equation 4.1 to the datapoints. This run contains the datapoint presented in Figure 4.2c.

of this chapter when we talk about the temperature (or some other parameter of a run) we will use the mean value of that parameter over a run, and quote the error as the standard error of this mean.

4.2.4 Comparing Runs

Having taken a series of runs, the next step is to compare them. We will defer a discussion of the COM damping until the next section; in this section we will focus of which other parameters change between runs.

Our method of aggregation of data points into runs results in us assigning eight parameters to each run: $A, B, \omega, \phi, \gamma, T, N, \mu$. The first five of these come from fitting the damped sinusoid in equation 4.1 to the centre-of-mass of each run, the remaining three come from averaging information from the fit of each data point in the run. The parameter B, which indicates the mean position of the run, rarely changes and seems uncorrelated with everything else, so we will ignore it.

We will also focus on the one amp dataset because the data were taken over a wider range of rf knife positions than the five amp dataset. The trends reported are similar.

4.2.4.1 The Effect of the rf Knife

We are aiming to measure the damping of the system as a function of temperature. To do this we need to be able to control the temperature of the system well over the course of one run, as well as being able to change the temperature between runs.

Our main method of controlling the temperature of the condensate comes from changing the position of the rf knife for each run. The effect of this rf knife is to remove atoms above a certain energy from the magnetic trap, and is discussed in section 2.3.3 as a tool for evaporative cooling. We keep the rf knife on after setting the condensate in motion to prevent heating of the condensate due to both the intrinsic heating of the trap, and the heating caused by setting the condensate in motion. Because of this we need to make sure we understand what other possible effects it can have on the system.

We expect the rf knife to have little bearing on the phase and the amplitude of the oscillations. Plotting the two as functions of frequency of the rf knife in Figure 4.7 we see that this is the case.

In a similar way we can look at how the temperature, number and chemical potential change as functions of rf knife position. The results are plotted in Figure 4.8.

We extract the chemical potential from a data point using the fitting described previously. This fitting means that the chemical potential is directly related to the number of atoms in the condensate - itself a function of temperature and total number of the atom cloud. As such the trend seen in Figure 4.8c is a consequence of the changes in atom number and temperature, rather than an independent property which is changing as a function of rf knife position. It does however give us information about how the condensed fraction of the atoms changes as a function of temperature. This will be discussed in section 4.2.4.2.

To qualitatively explain the behaviour of the number and temperature with rf knife position we use a simplified model. We imagine that we start with 45000 atoms in the magnetic trap in thermal equilibrium at 300nK. We then, in a single step, remove all the atoms above the escape energy determined by the rf knife and allow the cloud to equilibrate. The resulting atom number and temperature are those given by the black line in Figure 4.8a and Figure 4.8b.

Ignoring the points when the rf knife is at 300kHz this model has qualitatively the same behaviour

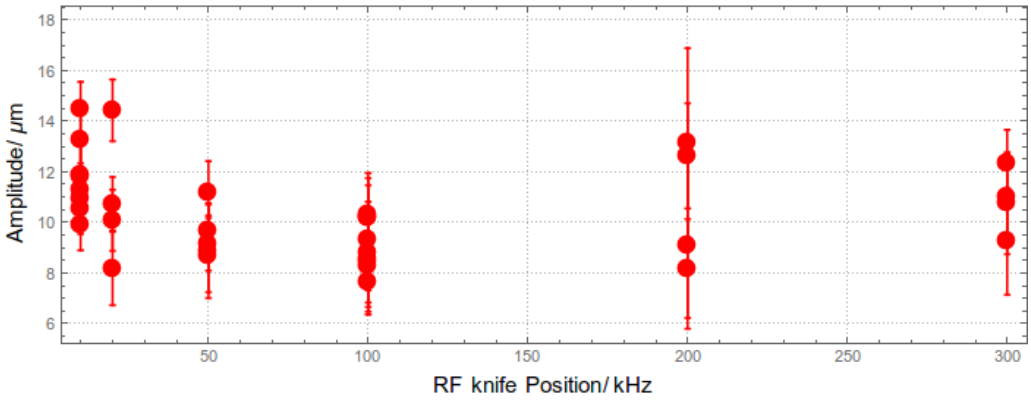

(a)

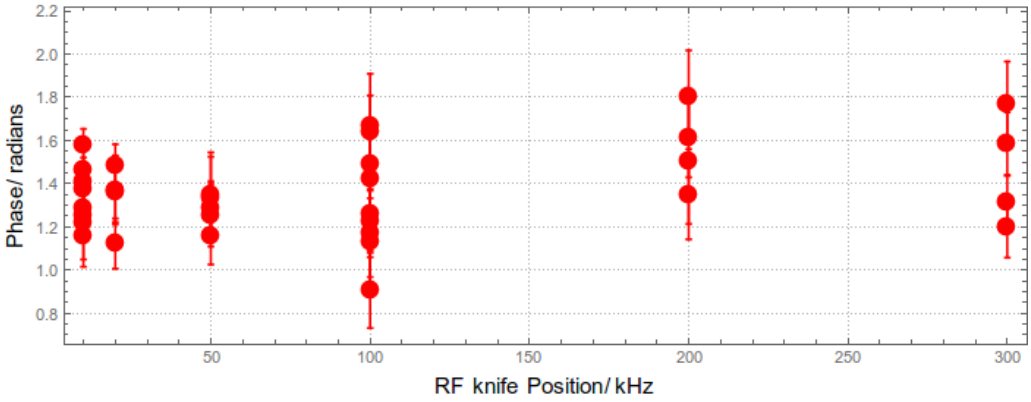

(b)

Figure 4.7: We compare how the amplitude (Figure 4.7a) and the phase (Figure 4.7b) of the condensate COM oscillations vary as functions of rf knife position.

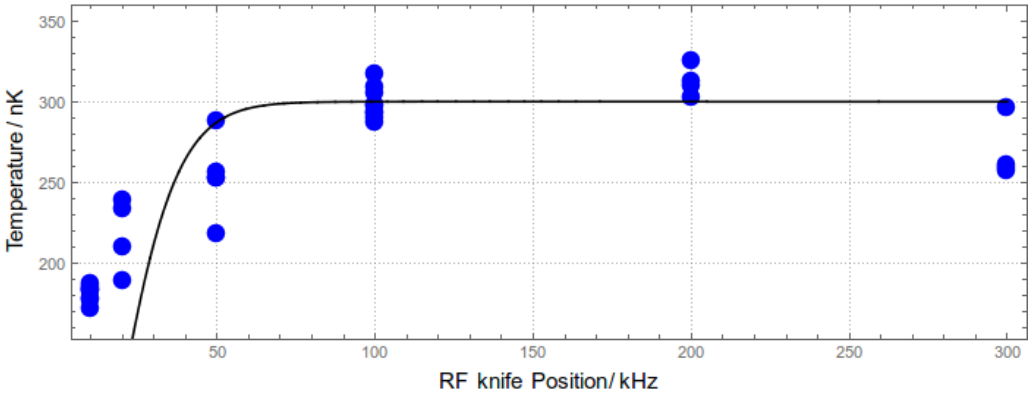

(a)

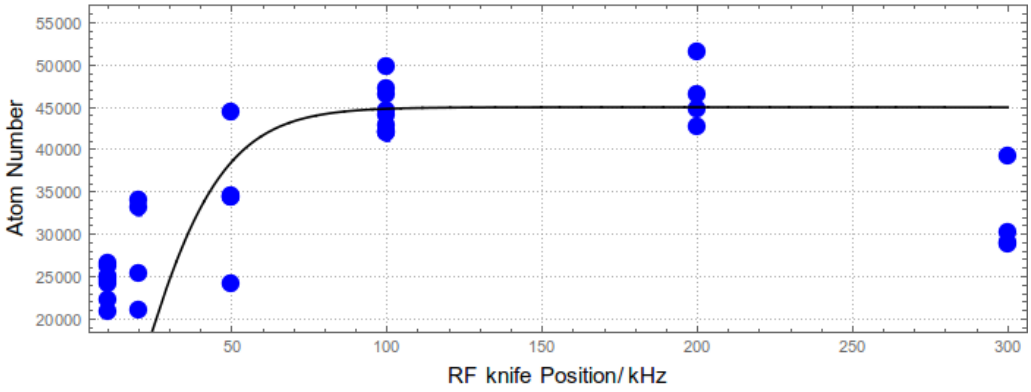

(b)

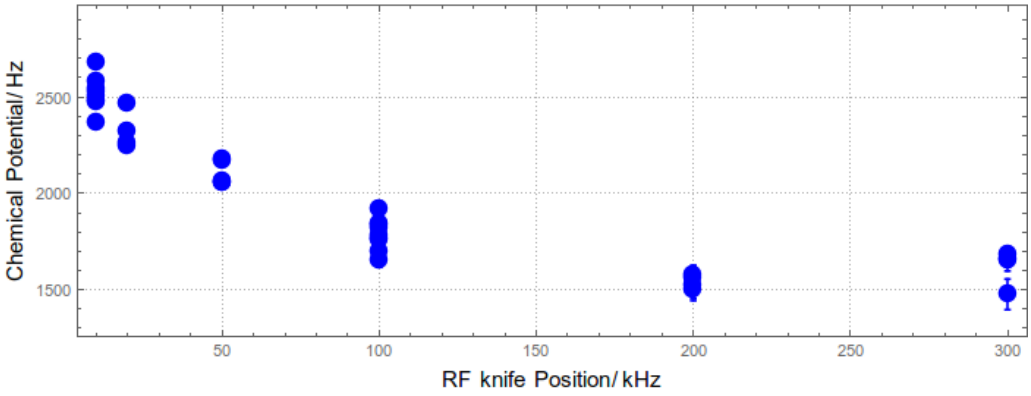

(c)

Figure 4.8: We compare how the temperature (Figure 4.8a), the atom number (Figure 4.8b) and the chemical potential (Figure 4.8c) of the atom cloud vary as functions of rf knife position. The black lines in Figure 4.8a and Figure 4.8b are based on a simplified model described in the text.

as our data, although is deviates at low positions of rf knife. The reason for this is that we don't turn on the rf knife at its final position, instead it is moved there slowly. This means that the cloud has a chance to redistribute energy as the rf knife moves, meaning more atoms remain and the temperature doesn't drop as much our idealised model suggests.

The points where the cloud is at 300kHz do not fit even this simple model. We believe this is because the data for these points were all taken on a single day, separate from the other data. It is possible that drifts in our experiment meant that the initial temperature and atom number on these days were lower than the rest of the runs.

A final property that appears to have a dependence on the position of the rf knife is the frequency of the oscillation, plotted in Figure 4.9. From this plot it seems as if the frequency of oscillation may be lower at lower rf knife frequencies. This effect is investigated further in section 4.2.4.3.

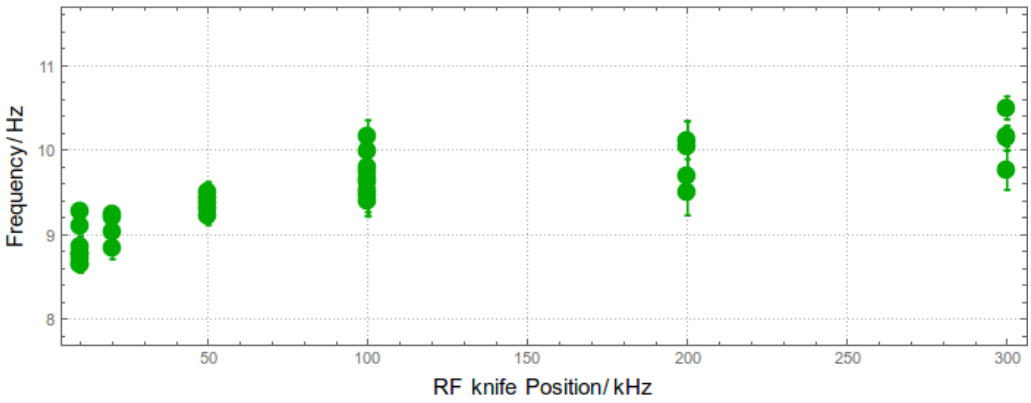

Figure 4.9: We compare how the oscillation frequency of the condensate varies as functions of rf knife position. There appears to be some dependence. This is discussed in section 4.2.4.3.

4.2.4.2 Condensed Fraction

As discussed in the previous section, our measured value of chemical potential gives us a measure of the number of atoms in the condensate. This is due to our assumption that we can treat the condensed part of the atom cloud using the Thomas Fermi description. It is useful to investigate how this dependence compares to that of the non-interacting model to check the consistency of our approach.

Using the Thomas Fermi approximation, we can look at how the fraction of the atoms in the atomic cloud can vary as a function of temperature. Collecting together all the data points we have, we plot this relationship out in Figure 4.10a. Here we show points from both the one amp dataset (red) and the five amp dataset (blue).

From our discussion of an ideal cloud in section 2.7.2 we would expect the condensed fraction to follow

$$\frac{N_c}{N} = \left(1 - \left(\frac{T}{T_C}\right)^3\right) , \tag{4.2}$$

where N_C is the number of atoms in the condensate, N is the total number of atoms in the cloud, T is the temperature of the cloud and T_C is the transition temperature. There is an offset between the two lines, since T_C depends on the density of atoms, and this is higher in the one amp trap, where the radial frequency is higher.

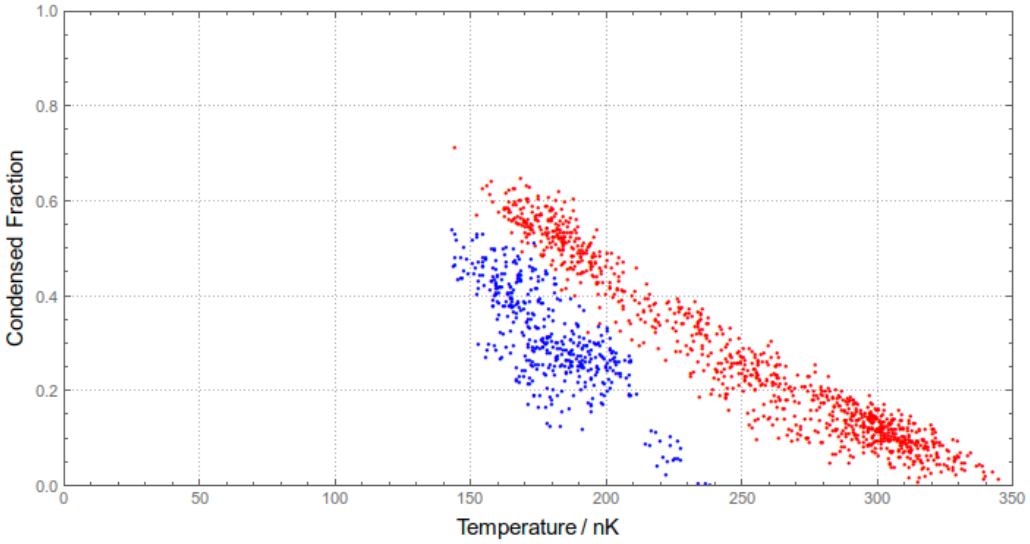

(a)

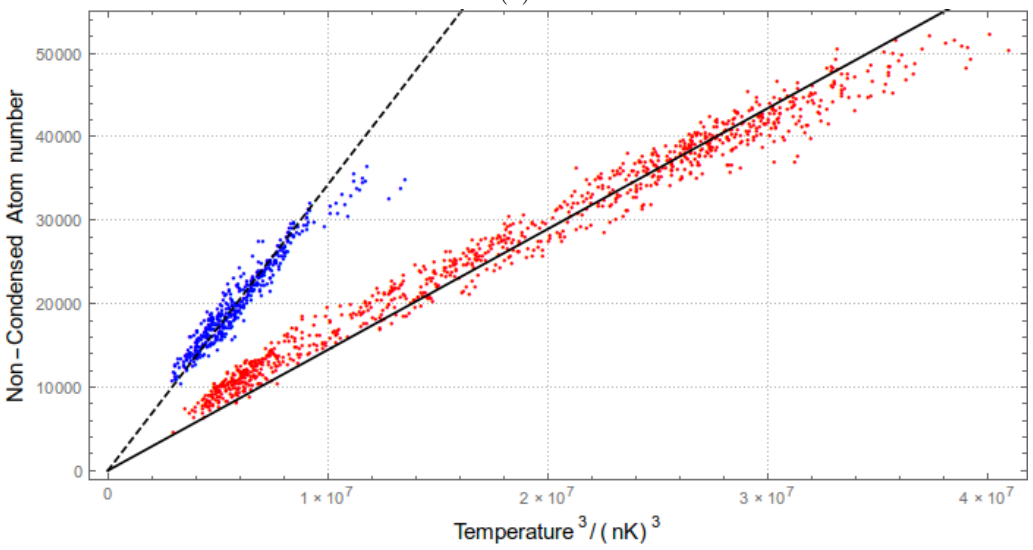

(b)

Figure 4.10: In Figure 4.10a we plot the fraction of atoms in the condensate as a function of temperature. The red points show data from the one amp dataset, the blue points show data from the five amp dataset. In Figure 4.10b we plot the number of non condensed atoms against temperature cubed. The colours are the same as in Figure 4.10a. The lines show the expected result of the non-interacting theory.

To be more quantitative we need to take into account the atom number change between runs. Noting that in the non interacting case the transition temperature is given by $k_B T_c = 0.94 N^{1/3} \bar{\omega} \hbar$, where $\bar{\omega}$ is the geometric mean of the trapping frequencies, we can write $T_C = \tilde{T}_C N^{1/3}$. We can then rearrange the equation for the condensed fraction to get

$$N - N_C = \left(\frac{T}{\tilde{T}_C} \right)^3 . \tag{4.3}$$

Plotting the number of non-condensed atoms against temperature cubed we get Figure 4.10b. Again we show points from both the one amp dataset (red) and the five amp dataset (blue). To these data points we have fitted two straight lines, describing the non-interacting BEC theory. When fitting each of these lines there was only one degree of freedom: the axial trapping frequency. For both lines the optimal fit was obtained for an axial trapping frequency of 4.5 Hz.

It is satisfying that the value of 4.5Hz is between the trapping frequency at the lowest part of the trap (9.5Hz) and the trap seen by thermal cloud (3.2Hz).

4.2.4.3 Oscillation Frequency

When we plotted the frequency of the condensate COM motion against rf frequency in Figure 4.9 it appeared that the lower rf knife positions resulted in lower frequency oscillations. We believe that this trend is due to changes in atom number. Given the strong correlation between atom number, rf knife position and temperature it is hard to disentangle which effects are due to what.

To investigate further we added a set of coils around the rubidium vapour cell used for locking the repump laser. By applying a current to the coils, we Zeeman shift the atoms used for locking, detuning the repump laser from resonance. The effect is to lower the atom number in the initial MOT without changing any of the other parameters of the system and hence without any unwanted effect on our magnetic trap.

Keeping all other parameters constant we made a standard oscillation run, but with the repump detuning randomly set to be either on or off. The result of this measurement can be found in Figure 4.11. The two oscillate at similar frequencies, but one can see that the cloud with fewer atoms (red points) oscillates at a slightly higher frequency than the cloud with more atoms (blue points). There is also a phase difference between the oscillations of the two clouds. This can be explained by the difference in frequencies if we assume a phase shift of zero when the driving oscillation began.

In Figure 4.11a you can find the COM positions of the two oscillation runs. A clear change in frequency and phase of the oscillations is present. In Figure 4.11b and Figure 4.11c we look at how both the atom number and the temperature of the clouds change during the runs. We see a mean difference in atom number between the two runs of 6000 atoms, but no statistically significant change in temperature (the measured mean values are 178 ± 5nK and 184 ± 6nK). From the fits in Figure 4.11a we extract frequencies of 9.10 ± 0.07Hz and 8.77 ± 0.07Hz for the high and low atom number runs respectively. While this difference is small, it does explain the changes in frequency at low rf knife values in Figure 4.9.

We believe the changes in frequency with atom number are due to small anharmonicities in the magnetic trap. As the atom number increases, the length of the condensate increases and it explores more of these regions, changing the oscillation frequency.

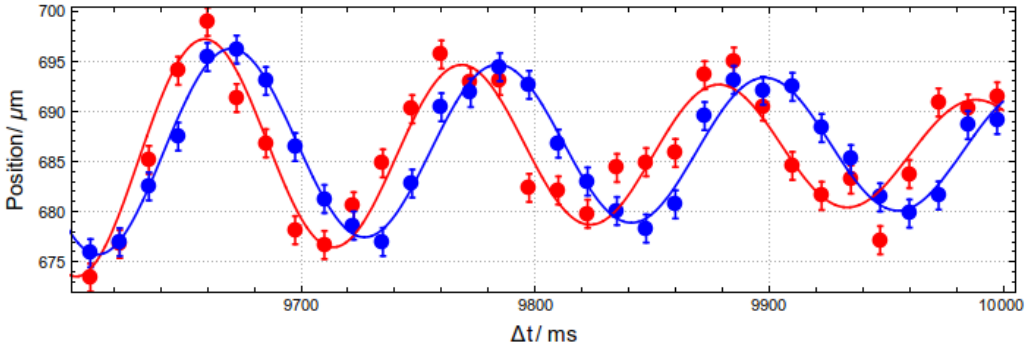

(a)

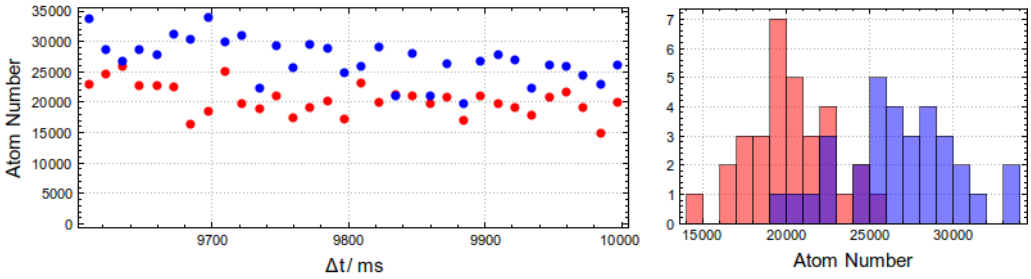

(b)

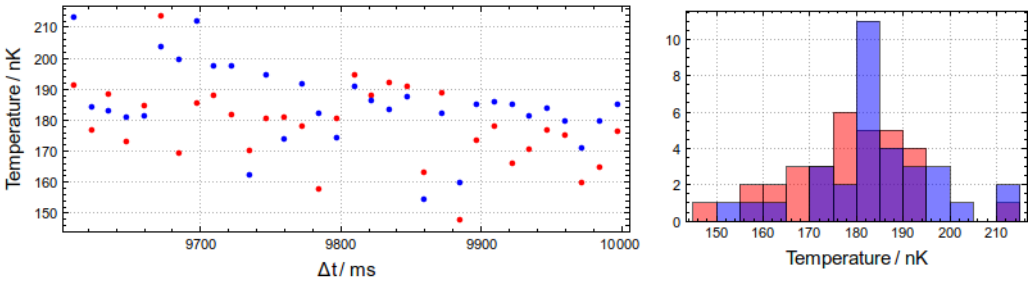

(c)

Figure 4.11: We compare the effect of atom number of oscillation frequency. The blue points show a normal run, the red points show a run with the repump laser detuned (see text). In Figure 4.11a we show how the COM positions of the condensates change as functions of Δt. In Figure 4.11b and Figure 4.11c we show how the atom number and temperature changed over that run. These results show that a cloud with fewer atoms but the same temperature oscillates as a slightly higher frequency.

4.2.4.4 COM Errors Revisited

In section 4.2.2.3 we stated that the dominate contribution to determining the COM position of each condensate was a scatter independent of the fitting procedure. Having fitted an expected damping curve to the COM positions we should check that this is accurately reflected in our data.

Having applied the fitting we can then look at the residuals of the COM positions of each condensate. The results are plotted in Figure 4.12a in blue, overlaid with this is the histogram from Figure 4.3b in red. The standard deviation of the residuals from our fit is 1.41, showing that the majority of the measured deviations from oscillations can be explained as just random motion in our imaging system.

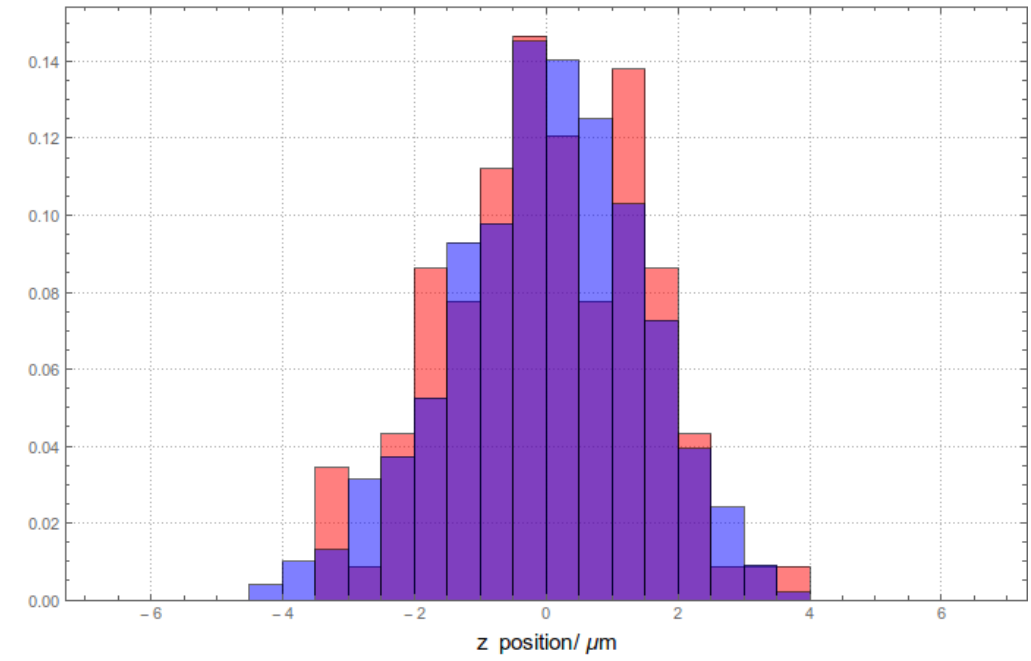

(a)

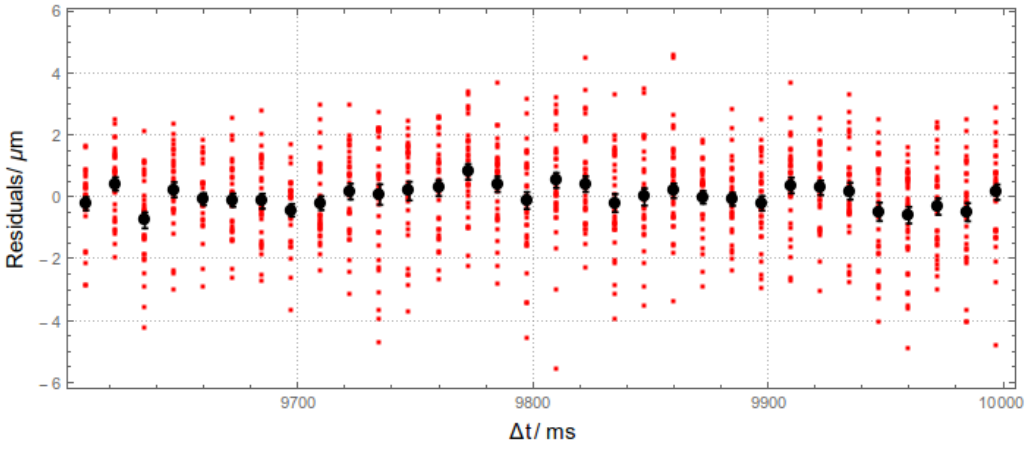

(b)

Figure 4.12: In Figure 4.12a we compare the scatter in position of an unexcited condensate in blue with the residuals of fitting a damped sinusoid to an excited condensate in red. In Figure 4.12b we look at how the residuals of our fitting vary as a function of Δt. The red points show the residuals of every fit, and the black points show the mean and standard error on these residuals.

A second check we can apply is to look at how the residuals of the fit change as a function of Δt. The results can be found in Figure 4.12b. The red points show the residuals of every fit, and the

black points show the mean and standard error on these residuals.

It has been suggested that there is some pattern to the residuals plotted in Figure 4.12b. One reasonable approach is to assume that any meaningful pattern in residuals would follow a sinusoidal pattern, similar to the underlying data. If this were the case, it may be a sign that our model isn't fully capturing the variation in the data.

We consider two hypothesises: (1) that the distribution of residuals is given by a random Gaussian variable, mean zero. (2) that the distribution of the residuals is given by a sinusoid, with a frequency between 2-50Hz with Gaussian noise, mean zero. From a fit to the residuals of a damped sinusoid with constrained frequency, we can calculate the likelihood of (2)[2]. Is is found that this value is smaller than the likelihood of (1). From this we conclude that for our purposes (1) is accurate.

4.2.4.5 Excitations of the Thermal Cloud

We believe that the relative motion between the condensate and the thermal part of the atom cloud results in damping of the condensate COM motion. We therefore need to examine the motion of the thermal cloud, as well as that of the condensate.

When we fit to the profile of a shot we extract information about the COM position of the thermal cloud. Typically the scatter of these points is much larger than that of the condensate COM. This is due to the fact that the fit to the thermal cloud is dominated by the fit to the wings of the profile, where the density is lower than for the condensate. For a single run there does not appear to be any pattern in the position of the thermal cloud.

An example of a single run, comparing the COM of the condensate and thermal cloud can be found in Figure 4.13a. The red points show the COM of the condensate which follows a damped sinusoid. The Black points show the COM of the thermal cloud for the same shots. The relative offset between the two is due to a small offset in the fragmentation potential between the centre of the the high frequency part of the trap the condensate experiences and the centre of the trap the thermal cloud experiences. It matches the offset we find from Figure 3.22.

When we look at all the points a new pattern emerges. In Figure 4.13b and Figure 4.13c we plot the positions of the COM of the condensate and the thermal cloud repectively as functions of time for all runs in the one amp dataset. The black points show the means of their positions at each Δt. For clarity the means of the positions are reproduced together in Figure 4.13d.

When looking at the COM of just the thermal cloud, averaged over many runs (Figure 4.13c), we can see the emergence of what appears to be a coherent oscillation. Fitting a damped sinusoid to Figure 4.13c we find the mean seems to oscillate at 3.2 ± 0.2Hz, which agrees well with the frequency we extract when fitting to the potential found from the data in Figure 3.22. We conclude that when we are shaking the trap we are exciting some damped motion of the thermal cloud, as well as the condensate. However the thermal cloud oscillates at a much lower frequency.

We give two reasons why this slow oscillation is expected to cause the same damping to the condensate as one that is stationary. Firstly the difference in velocity between the condensate and the thermal cloud is maximised during the first oscillation cycle of the condensate COM. If there was an effect due to the different velocities we would expect the Condensate COM motion to be asymmetric in the first cycle - we do not see this. Secondly, numerical modelling of the system as two classical masses in a harmonic trap with damping depending on their relative velocity with similar parameters

[2]It is worth pointing out that the parameters found from this fit are very dependent on the initial parameters entered into the fitting function. A range of these were tried and the one which minimised the χ^2 value of the fit were used.

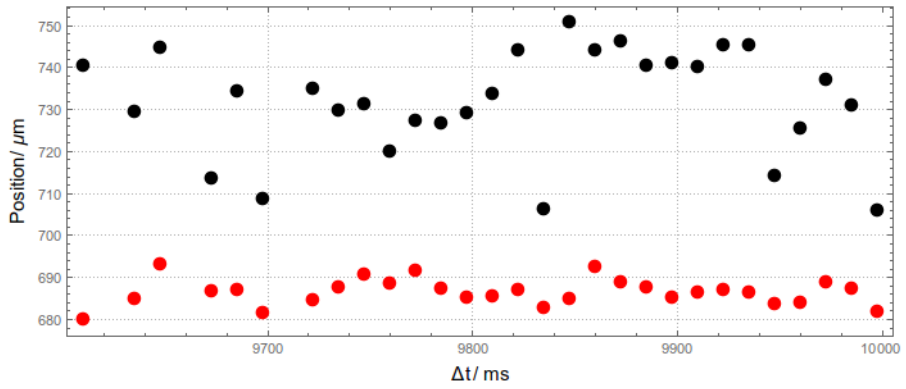

(a)

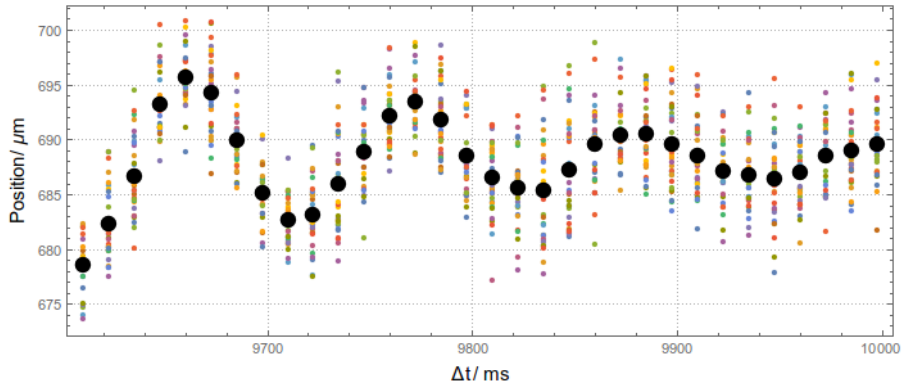

(b)

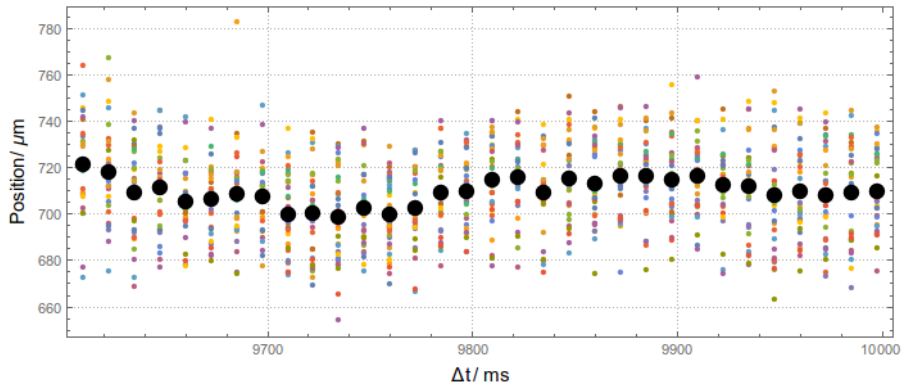

(c)

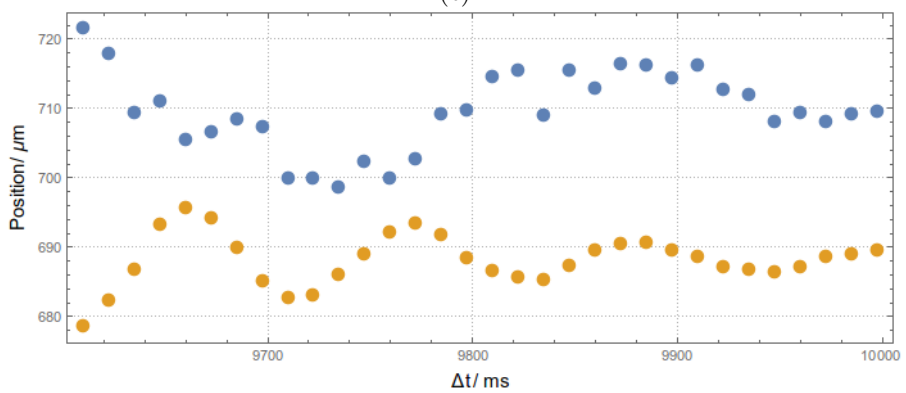

(d)

Figure 4.13: Oscillations of the cloud. Figure 4.13a A single run from the 1A dataset. Red shows the condensate COM, black the thermal cloud COM. Figure 4.13b and Figure 4.13c show the COM positions of the condensate and the thermal cloud for every run taken in the one amp trap respectively. The Black points on these graphs show the mean positions of all the runs. Figure 4.13d shows the mean positions for the COM of the condensate (orange) and the thermal cloud (blue).

to our experiment do not show a difference of a few percent from the case of a stationary thermal cloud.

4.2.5 Other Excitations

In the previous section we saw that we are exciting centre-of-mass motion of the condensate and the thermal cloud. It is worth asking whether we are exciting any other collective modes of the condensate.

To investigate this we follow the work presented in [86], which uses a procedure known as principal component analysis (PCA) to identify regions of density in the atomic cloud that vary between shots. It is also worth pointing out that a similar technique used to measure the phase of a BEC [87].

In using PCA to analyse the absorption images we treat each of the N pixels in a single image as independent random variables, and the set of M images obtained in a run as M samples of those N variables.We subtract the mean of each pixel across the M images and compute the N by N covariance matrix of M sample images. The diagonal of this matrix gives the variance of each pixel, and the off diagonal terms give the covariance between two pixels.

The eigenvectors of this matrix are known as the 'principal components' of the system, and it is possible to use them as a basis set in which to express an image. The eigenvalues of the system represent the variance of the data projected onto the directions defined by the eigenvectors. The principal components of the system are numbered in order of descending size of eigenvalue.

Given that the definition above makes no mention of the time or spatial orderings of each pixel, it is worth asking how we recover information about the modes of the system.

Following [17], we note that in the hydrodynamic model small excitations of the system can be described by

$$\rho(\mathbf{r},t) = \rho_0(\mathbf{r}) + \sum_k c_k \cos(\omega_k t + \phi_k) f_k(\mathbf{r})\,, \tag{4.4}$$

where $\rho(\mathbf{r},t)$ is the density of the cloud at position $\mathbf{r}$ and time t, and the sum runs over the normal modes of the system each with spatial shape $f_k(\mathbf{r})$ and frequency ω_k, and c_k is related to the mode occupation.

Having removed the mean and taken the covariance between each pixel, the resulting covariance matrix, S, is given by

$$S_{i,j} = E_t\left[\left(\sum_k c_k \cos(\omega_k t + \phi_k) f_k(\mathbf{r}_i)\right)\left(\sum_{k'} c_{k'} \cos(\omega_{k'} t + \phi_{k'}) f_{k'}(\mathbf{r}_j)\right)\right]\,, \tag{4.5}$$

where E_t is the expectation value taken over time. In our case the intervals between shots are discrete, however if we assume that the time interval between shots is much faster that the period of any of the excited modes the terms that contain terms of different frequencies average to zero and we get

$$S_{i,j} = \frac{1}{N}\sum_k c_k^2 f_k(\mathbf{r}_i) f_k(\mathbf{r}_j)\,. \tag{4.6}$$

Due to the orthogonality of the normal modes f, the matrix S has eigenvectors f_k, with eigenvalues c_k^2. This shows that we can get a measure of the excited mode shapes and amplitudes using PCA.

The result of applying PCA to a single run can be found in Figure 4.14, where (a) shows the average images of a run with $M = 31$ and (b) - (g) show the first six principal components.

The numbers in the left hand column in Figure 4.14 show the normalised eigenvalue of each

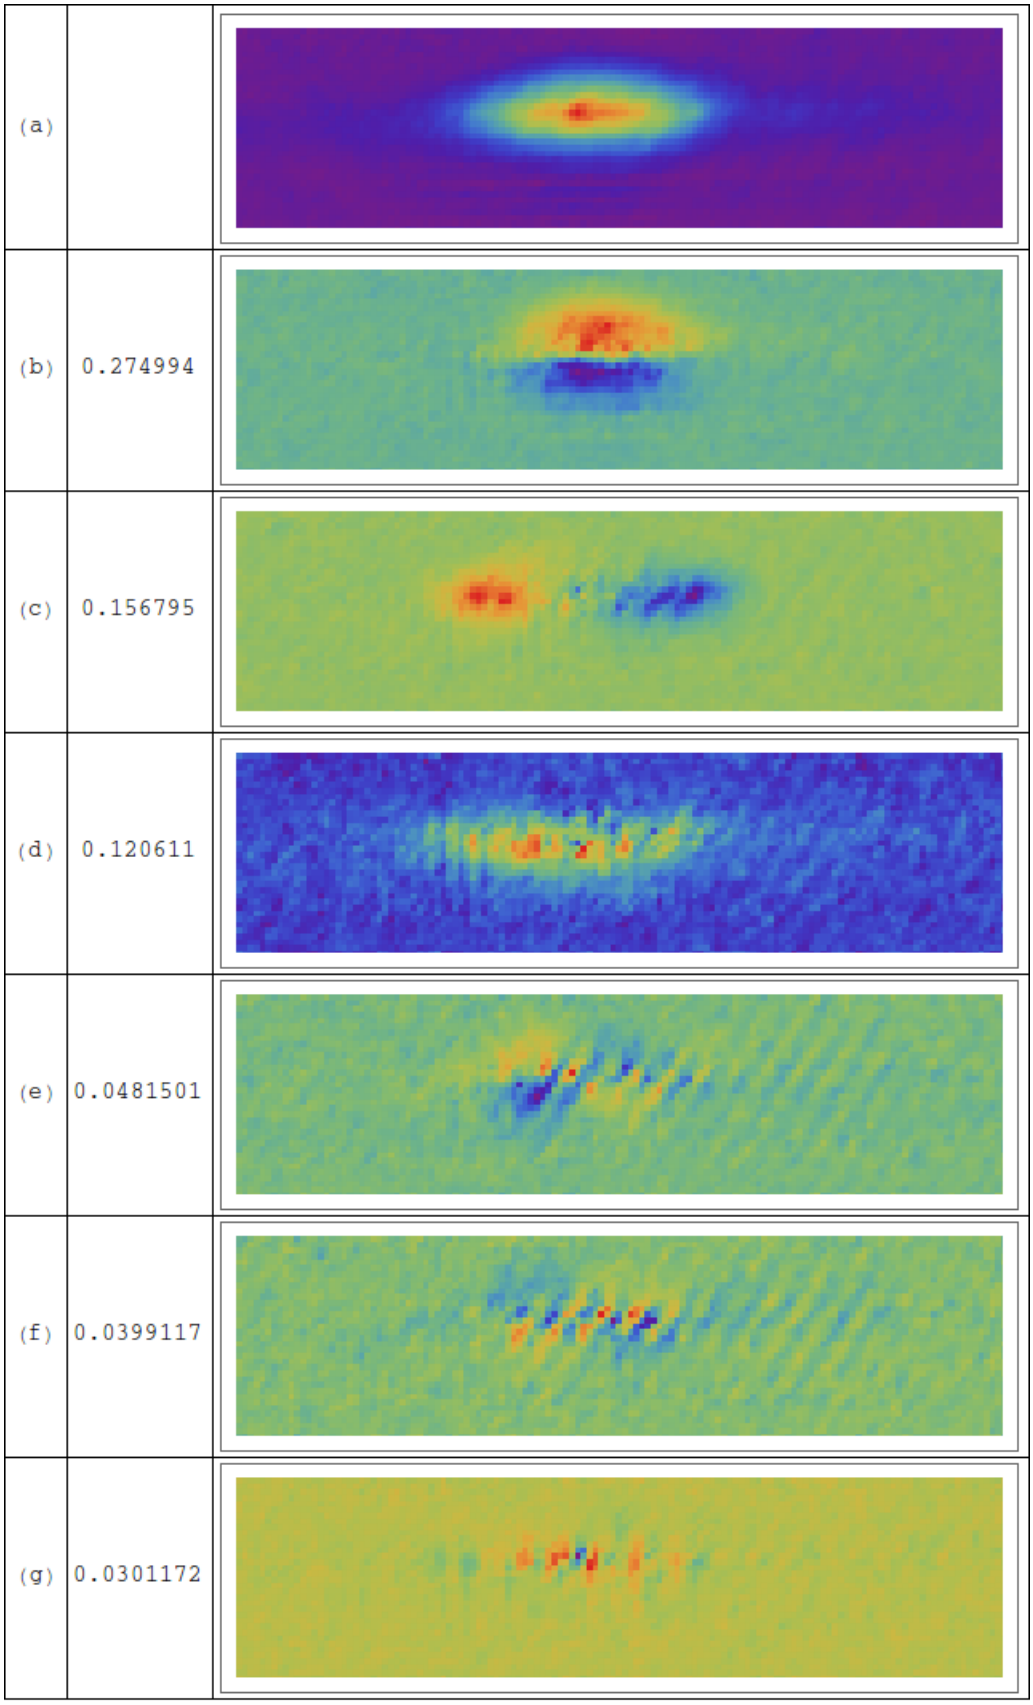

Figure 4.14: The top image shows an averaged absorption image of the data taken from a single oscillation run in the one amp trap. Each subsequent image shows a principal component of the run. The numbers in the second column show the relative amount of the variance of the dataset explained by that principal component. The colour scaling is arbritary, and varies between images.

component, showing how much of the shot to shot pixel variation can be explain by that particular principal component.

Looking at the shape the first component, (b), we learn something interest of: that there is motion along the y axis. This can be confirmed by looking at the COM position of the condensate on the y axis. It shows random scatter with a standard deviation of 1.4μm. From our discussion in section 4.2.2.3, we believed this to be due to the same slow drift we observe along z.

To check that we are not exciting motion in the y direction when we set the condensate oscillating we have taken a set of shots over a short time scale after exciting the condensate. The results, shown in Figure 4.15, do not exhibit any oscillation along y.

Given that it is random motion in the y direction, it is not surprising that is is responsible for the largest principal component: motion in the y direction causes the largest change in pixel values from the 'average image'. This is because the condensate is long in the z direction and thin in the y direction, so any change in the y direction COM causes a large change in how the condensates would overlap.

The principal components give us other information about a run. In Figure 4.14 component (c) corresponds to changes of the centre-of-mass in the z direction (the change in COM we have been primarily looking at) and component (d) corresponds to changes in atom number. While these components have been identified simply by looking at them, we can make it more quantitative by decomposing each image into its principal components and investigating their the physical properties.

We do this in Figure 4.16. Here we take each image in a run and take the scaler product between it and the first three principal components of Figure 4.14. The result is used to label the x-axis of Figure 4.16, where the labelling corresponds to that in Figure 4.14. On the y-axis we show experimental parameters determined previously. There are strong correlations between between (b) and the y position, (c) and the x position and (d) and the atom number,

One might think to remove the centre-of-mass motion before applying PCA by centring each image. This is not straightforward, because the motion of the cloud is on the micron scale, whereas the pixel size of our images is 3.5μm.

It is reassuring that the PCA has identified known changes in the atom cloud measurable by other means, but so far, apart from motion along the y axis, we haven't used it to identify anything new. We apply it as a technique because as noted in [86] it would allow us to find higher order collective excitations of the condensate.

These higher order collective excitations manifest themselves as time-varying principal component decompositions. The remaining principal components identified in Figure 4.14 show no pattern in time, and seem uncorrelated with other parts of the experiment, however it might be suggested that principal component (e) is related to a twisting motion of the condensate. These modes are expected to oscillate at a much higher frequency than our measurements can distinguish, making their identification here difficult.

As a comparison, we can ask what would be expected if higher order collective excitations were present. We can approach this using a run taken using a different excitation mechanism. When investigating methods to excite the condensate we also tried to lower the current in the small Z wires slowly, and return it to the original position quickly (similar to the technique used to measure radial frequency of the trap in section 3.4.2.1, but with a lower amplitude). The result of applying PCA to this run can be found in Figure 4.17.

Here we see components that were present in the previous PCA: (a) again shows the mean density,

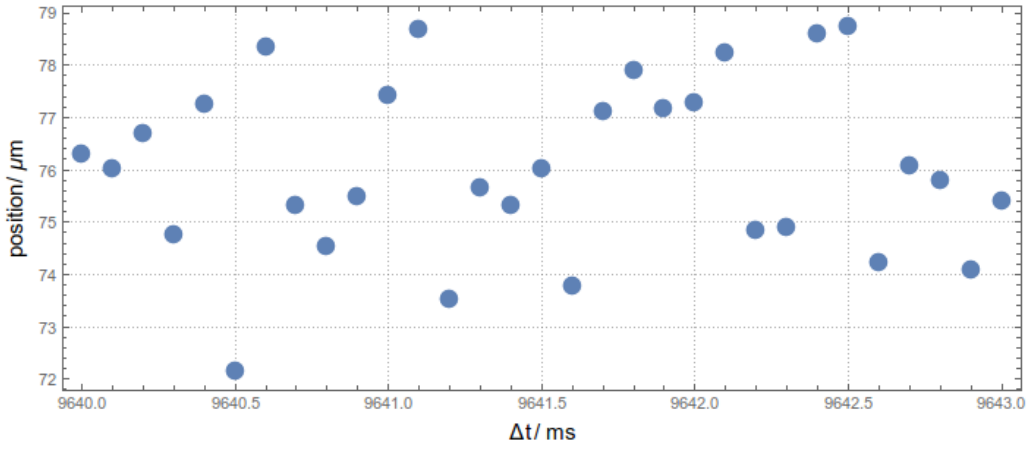

Figure 4.15: Here we look at how the y position of the condensate changes as a function of Δt over a short time after setting the condensate oscillating in the z direction.

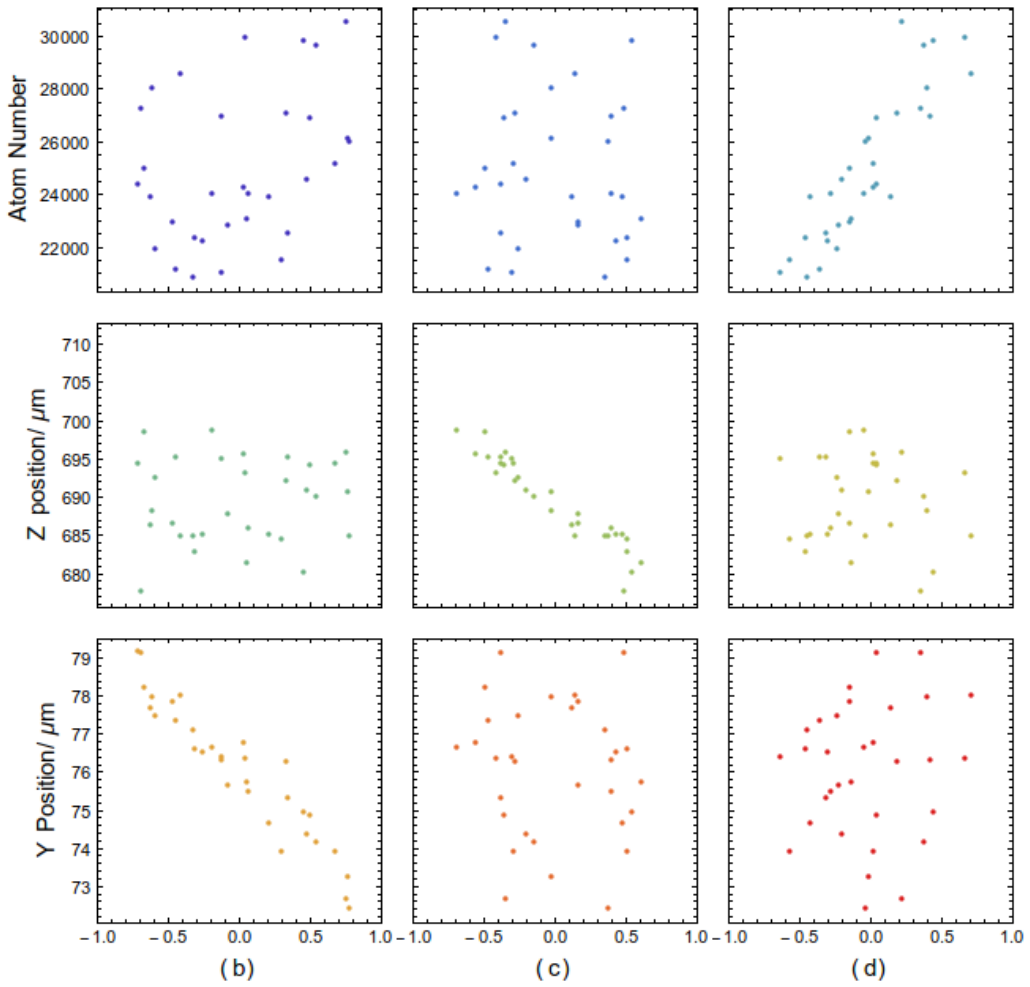

Figure 4.16: Here we compare the decomposition of the images of a run into the principal components identified in Figure 4.14 with the atom number and Z and Y positions measured using other techniques. The x axis of the plots label the components identified in Figure 4.14.

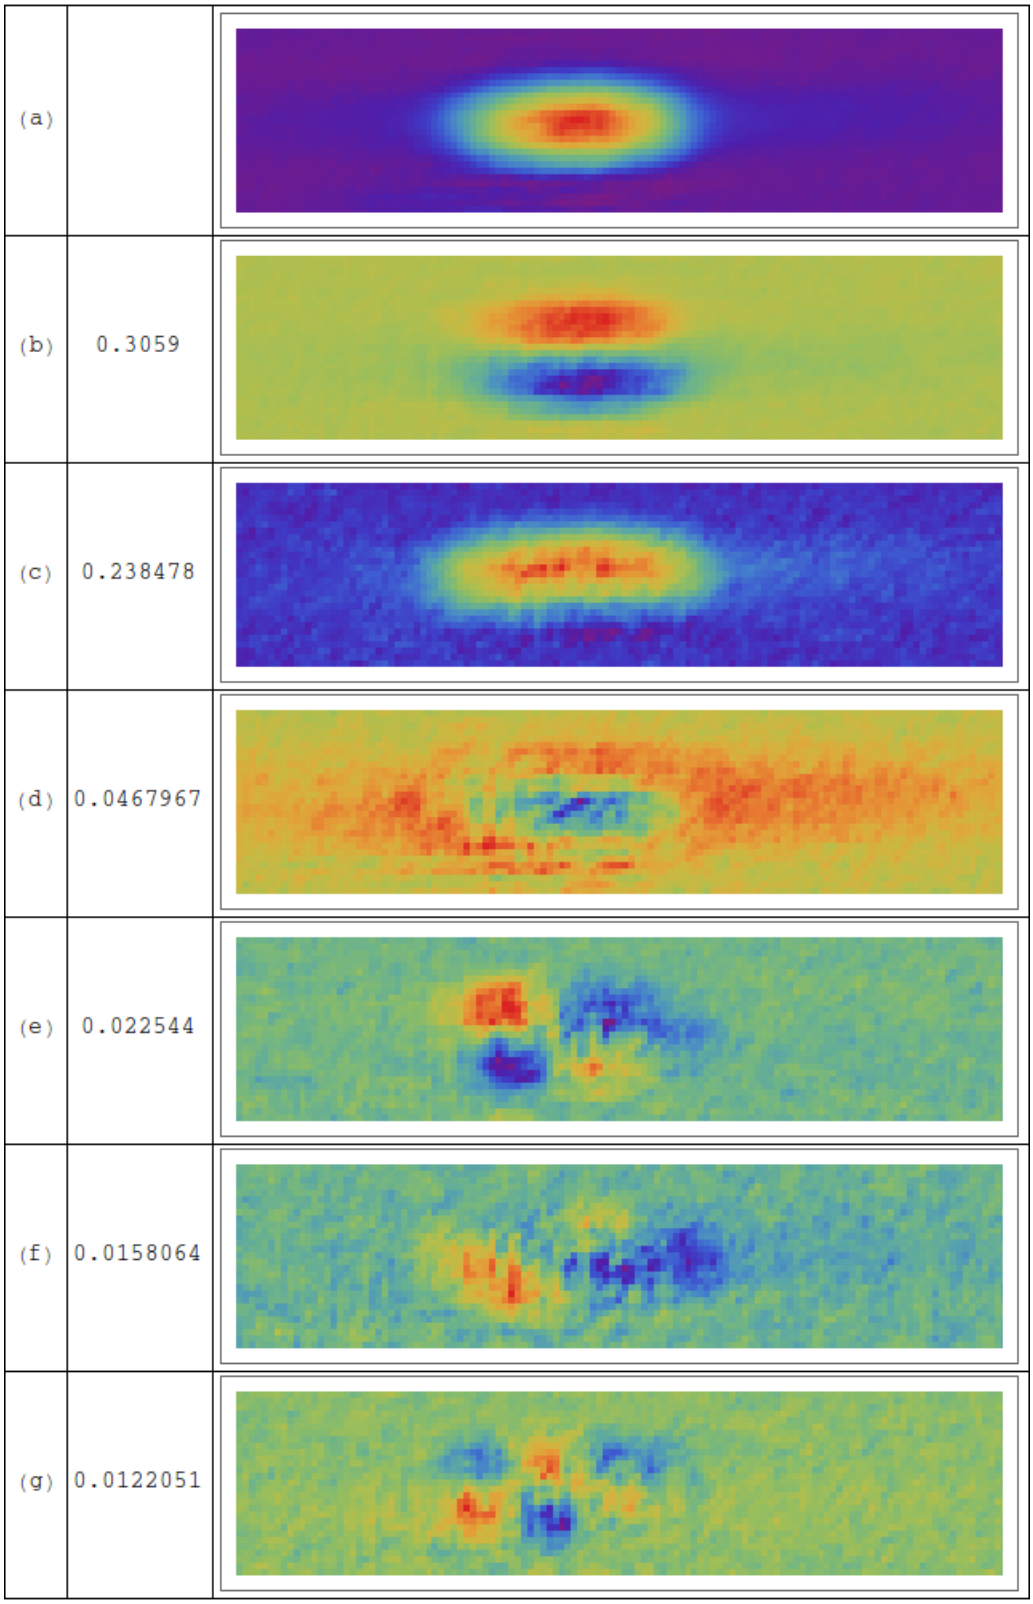

Figure 4.17: Here we show PCA applied to a run with an alternative excitation mechninism. The top image shows an averaged absorption image of the data taken over a run. Each subsequent image shows a principal component of the run. The numbers in the second column show the relative amount of the variance of the dataset explained by that principal component. The colour scaling is arbritary, and varies between images.

(b) shows motion in the y direction, (c) shows atom number changes, (d) shows temperature changes (in this run no rf knife was used, and experimental parameters were less well controlled, resulting in a wider temperature variation, not necessarily correlated with atom number). What we also see in this run are a twisting excitation in (e) and a bending excitation in (g). The fact that these principal components are not seen in Figure 4.14 indicated they are much less strongly excited there.

It is worth noting that we did not detect any time dependence of the principal components (e) and (g), perhaps because this was too fast for the time resolution of our experiment. We would expect these excitations to oscillate at a frequency related to the radial frequency of our trap, but the shots were taken over a range to look at the axial excitation frequency. The PCA was carried out well after the original experiments were run, at a point when going back and retaking the data was difficult. It does however provide an interesting approach to identifying collective excitations in future work.

4.3 Interlude

At this point, it is worth pausing and reviewing what we have covered so far. In the preceding section we have characterised the motion quite carefully, showing how we have both created and excited the centre-of-mass motion relative to the thermal cloud. In doing so we covered how the temperature of the condensate was changed using the rf knife, along with how other parameters of the experiment changed.

The goal of the preceding was to outline in detail how we move from a set of single images of the condensate to that damping curves that are going to be presented in the next section and to understand exactly what control we have over the system.

We are now ready to return to the main point, which is the damping of the axial centre-of-mass motion of the condensate. In Figure 4.18 we plot this damping rate gamma, which is defined in equation 4.1, against temperature, for all the date we have. We see that the motion is more strongly damped in the 5A trap (blue) than in the 1A trap (red). For both traps the damping increases with temperature. In the next section we discuss these results, comparing them with theory and with other experiments.

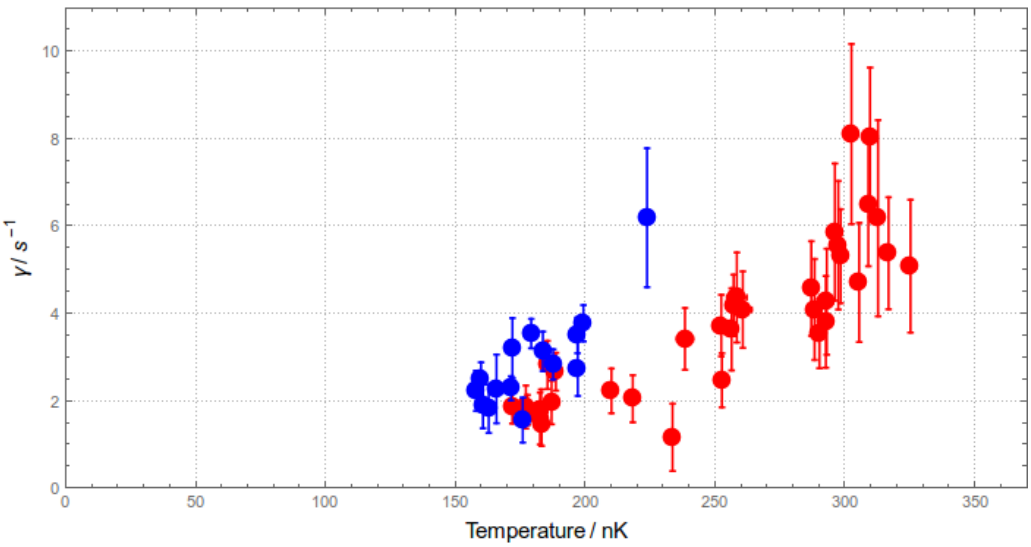

Figure 4.18: Here we look at how the damping rate of the condensate COM position varies as a function of cloud temperature. The red points belong to the one amp dataset and the blue points belong to the five amp dataset.

4.4 Damping of the Oscillations

4.4.1 Theoretical Description

Landau damping [22][23][24][25][26] provides a simple mechanism for damping the oscillations of the condensate COM. The idea is that motion of the condensate is damped via energy exchange between the coherent condensate motion ν and thermal excitations of the cloud ϵ, which scatter into ϵ'. The equation is microscopically reversible.

$$\epsilon + \nu \leftrightarrow \epsilon' \,, \tag{4.7}$$

Where ϵ, ϵ' and ν correspond to states.

In thermal equilibrium, the lower energy thermal excitations are more populated than those at high energy. This asymmetry ensures a net damping of the coherent motion. The damping rate can be calculated from Fermi's golden rule.

The rate at which energy is removed from the coherent motion is[22][23]

$$\Gamma_\nu = \frac{2\pi}{\hbar} \sum_{\epsilon,\epsilon'} \left|\left\langle \epsilon, \nu | H' | 0, \epsilon' \right\rangle\right|^2 \delta(E_\epsilon + E_\nu - E_{\epsilon'})(f_\epsilon - f_{\epsilon'}) \,, \tag{4.8}$$

where the sum runs over all thermal excitations ϵ, ϵ', with equilibrium occupation f_ϵ and $f_{\epsilon'}$ respectively being given by the Bose Einstein distribution, E_i gives the energy of the excitation i and H' is the Hamiltonian that gives the interactions between excitations, and is given in equation 2.48.

Given that the energy of the oscillation is proportional to the amplitude of the oscillation squared, we expect the amplitude to decay at half the rate of the energy. The rate of decay of the amplitude, γ is then given by

$$\gamma = \pi\omega_\nu \sum_{\epsilon,\epsilon'} |M_{\epsilon,\epsilon'}|^2 \delta(E_\epsilon + \hbar\omega_\nu - E_{\epsilon'}) \frac{\partial f_\epsilon}{\partial E_\epsilon} \,, \tag{4.9}$$

where the matrix elements from the previous equation have been written as M and where we have made the further approximation that E_ν is small relative to T and we can approximate the difference in populations of the levels as a differential.

To relate this equation to our experiment we need to know the matrix elements, M, in equation 4.9. For a three-dimensional trapped gas these have been computed in [23][26] with the result

$$\gamma = A_\nu \omega_\nu \frac{k_B T}{\sqrt{n_0}} \left(\frac{a^3}{U_0^2}\right)^{1/2} \,. \tag{4.10}$$

Here n_0 is the peak density of the condensate, $U_0 = \frac{4\pi\hbar^2 a}{m}$ is the point interaction energy of the atoms, a is the absolute scattering length of the atoms and A_ν is a constant that has to be calculated numerically and depends on the shape of the mode that is being damped. Its full form is detailed and will not be produced here, but can be found in [26].

At the beginning of a damping run, the condensate has a centre-of-mass amplitude of roughly 10μm. For our 10Hz trap this corresponds to an energy of $4.5\hbar\omega$ per atom, which is outside the regime of a single Bogoliubov excitation assumed in the theory above. However, we know of no theory treating this regime, so we compare our results with equation 4.10.

In order to compare equation 4.10 with our results we need to measure how our values ω and n_0

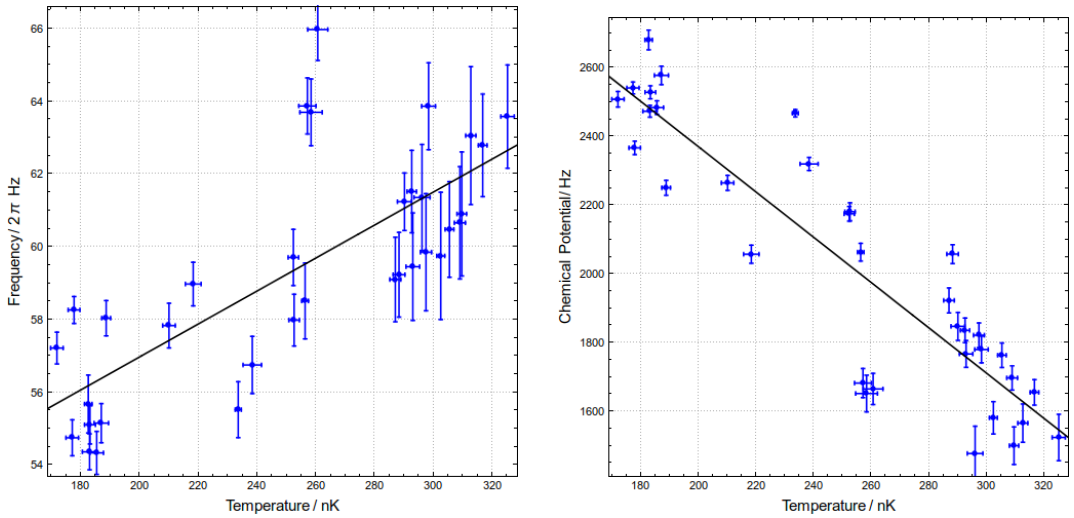

(a) one amp

(b) one amp

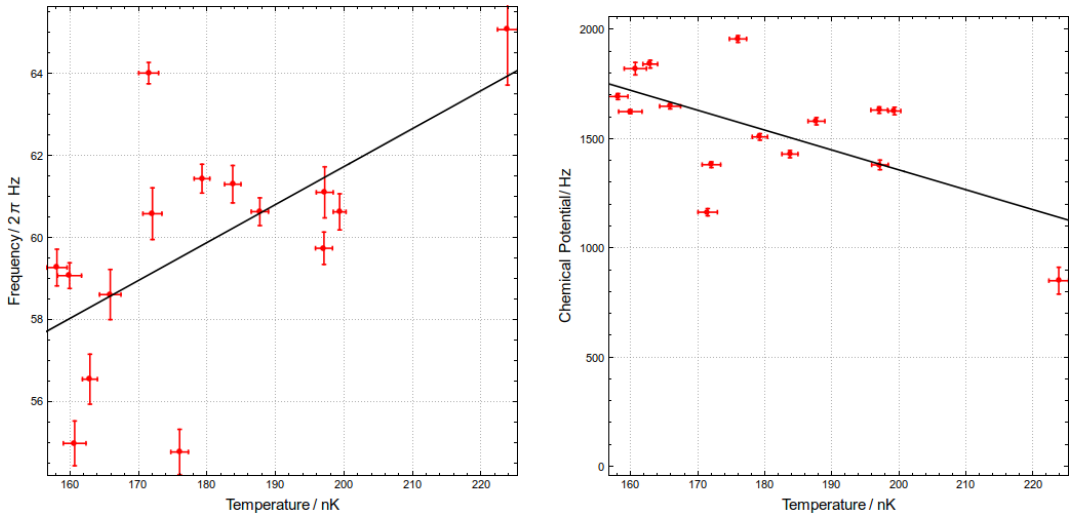

(c) five amp

(d) five amp

Figure 4.19: Here we take into account changes in the chemcial potential and frequency with temperature by approximating them linearly. Figure 4.19a and Figure 4.19c show how the oscillation frequencies change as functions of temperature in the one amp and five amp datasets respectively. Figure 4.19b and Figure 4.19d show how the chemical potentials change as functions of temperature in the one amp and five amp datasets respectively. The black lines show linear fits to the points.

change with T. Figure 4.19a and Figure 4.19c show the observed change in frequency with temperature for the one amp and five amp datasets. Figure 4.19b and Figure 4.19d show the observed change in chemical potential with temperature for the one amp and five amp datasets. For all datasets we make linear fits to model the variation. We derive the peak number density from the relation $n_0 = \frac{\mu}{U_0}$. Although this formula assumes the Thomas-Fermi approximation, which is not strictly applicable here, this nevertheless provides a reasonable estimate for the present purposes. With n_0 and ω thus determined, we now fit equation 4.10 to our data, with A_μ as the only free parameter. The results are shown in Figure 4.20.

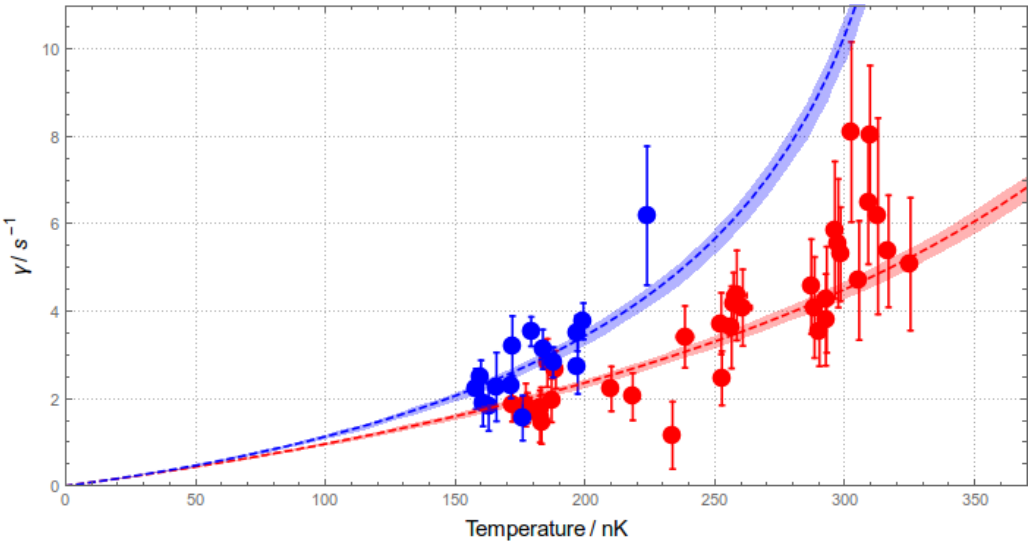

Figure 4.20: Here we show the data presented in Figure 4.18 overlaid with lines obtain by fitting to equation 4.10. Again the blue data points show the data from the one amp run and the red data points show the data from the five amp run. The shaded regions show the one sigma confidence region of the fits.

These fits give $A_\mu = 3.41 \pm 0.15$ and $A_\mu = 3.46 \pm 0.15$ for the one amp and the five amp datasets respectively. The reduced χ^2 values for the fits are 0.9 and 1.3 for the one amp and five amp datasets respectively[3].

On looking at Figure 4.20 we wondered whether the result depended strongly on the point at 235μK in the five amp dataset. If we remove this point and reapply the fit we get almost the same result. This is not really surprising because the fit is more constrained by the requirement that the damping drop to zero at zero temperature than by this point. We conclude that our data are well described by the Landau damping theory, taking the coefficient A_μ to be approximately 3.4.

4.4.2 Comparison with the Utrecht Experiment

A similar experiment was carried out in Utrecht [28], but with a fully 3D condensate, compared with our much more tightly confined cloud that approaches the 1D condition. The Utrecht experiment was consistent with a different value of A_μ roughly twice the value we measure.

In the Utrecht experiment a sodium BEC was created in an optical trap of radial frequency 95Hz and axial frequency 1 - 8Hz. The Utrecht experiment then used a weak blue detuned laser to displace the condensate relative to the thermal cloud. This was achieved by setting the potential barrier due to the blue detuned laser to just above the chemical potential - this confined the condensate, but

[3]We give the χ^2 values here to show that the fit is in qualitative agreement, not as a hypothesis test of the model.

allowed the thermal cloud to pass over it. The condensate was then released and allowed to oscillate through the thermal cloud. After a variable time the atoms were released and imaged to measure temperature and position.

The Utrecht experiment is similar to ours in that it measures the relative damping between the COM motion of the condensate and the thermal cloud. The Utrecht experiment use sodium which has a different mass and scattering length, however the important difference is its dimensionality. In the Utrecht experiment the atom number is much higher, in the range 100-500 million and their trapping frequencies are lower. This results in a chemical potential 28-56 times larger than the radial trap frequency, placing the Utrecht experiment firmly in the three-dimensional regime. In our experiment the ratio of chemical potential to radial trap frequency is in the range 1 - 1.8. A comparison of the Utrecht data with ours allows us to see how the dimensionality affects damping.

There are four main energy scales in our system: the two trap frequencies, the chemical potential and the thermal energy. These are summarised in Table 4.1.

	Energy Scale (Hz)		
	1A	5A	Utrecht
Axial Frequency	10	10	1 - 8
Radial Frequency	1400	900	95
Chemical Potential	1600 - 2600	900 - 1800	2700 - 5200
Thermal Energy	3000 - 6000	3000 - 4000	6000 - 10000

Table 4.1: A comparison of the energy scales in each experiment.

In the Utrecht experiment the axial trapping frequency was varied along with temperature and atom number. This allowed the Utrecht experiment to investigate how the damping varied with hydrodynamicity, h, defined as

$$h = \frac{n_{th}\bar{v}_{rel}\sigma}{\omega_z}\,, \tag{4.11}$$

where n_{th} is the average density of the thermal cloud, $\bar{v}_{rel}$ is the average relative velocity of atoms in the thermal cloud and σ is the scattering cross section. Hydrodynamicity is then the average number of collisions per oscillation. The Utrecht experiment found that at low hydrodynamicities, of the order unity, Landau damping theory described the damping well with $A_\mu = 7$. As the hydrodynamicity was increased the rate of damping increased above this level because Landau damping doesn't take into account those collisions.

In our experiment the hydrodynamicity lies in the range 2-5 for the one amp dataset and 3-5 for the five amp dataset. We find that over this range our damping agrees well with Landau theory, but with a smaller A_μ. In Figure 4.21 we plot hydrodynamicity against the ratio of the measured damping rate to the Landau damping rate from equation 4.10 taking $A_\eta = 7$. On this figure we show both the datapoints from our runs and the data taken from the Utrecht experiment [88].

From Figure 4.21 we see that unlike the Utrecht data, our damping agrees well Landau theory over the whole range of hydrodynamicity. The Utrecht data shows a slight increase in damping with hydrodynamicity above what would be predicted by Landau theory.

We believe that the difference lies in our dimensionality. As we lie in the quasi-one-dimensional regime, the effect of the discrete nature of the harmonic trap would be expected to play a role. From 4.1, the chemical potential and the thermal energy in our experiments are of similar size to the quanta of radial excitation energy, whereas for the Utrecht experiment they are much higher.

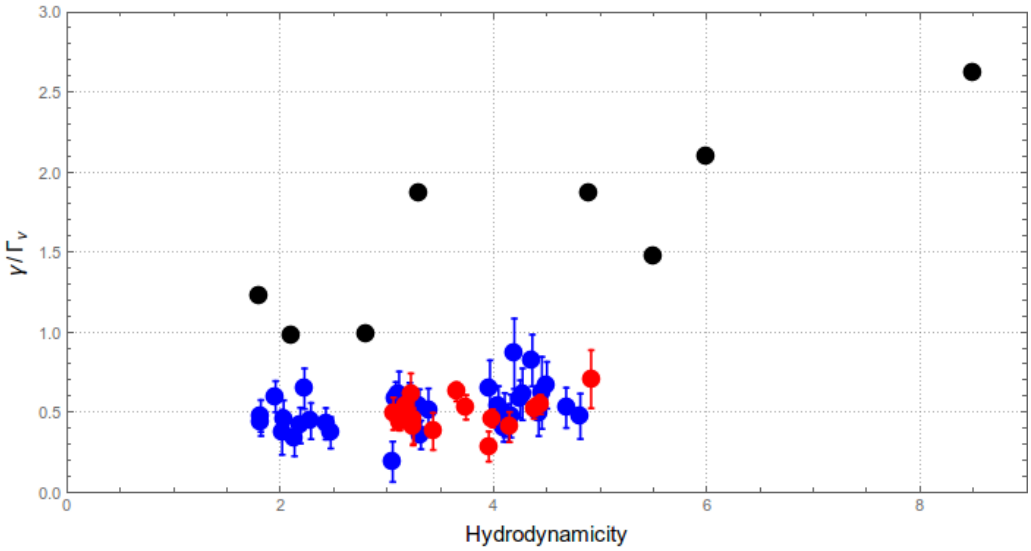

Figure 4.21: Ratio of γ, the measured damping rates and Γ_η the rates predicted by equation 4.10 if we take $A_\eta = 7$. The black points show the data from the Utrecht experiment (see text). The blue and red points show our data, taken from the one amp and five amp datasets respectively.

Using the assumption of Landau theory, the damping is completely described by equation 4.9. To get from this equation to the form in equation 4.10, an integral over all modes was used. In the quasi-one-dimensional case, this would have to be replaced with a sum, and we would expect the damping to be reduced relative to the three-dimensional case.

CHAPTER 5

CONCLUSION

5.1 Summary

In this work we have presented new data on the damping of the COM motion of a quasi-one-dimensional condensate moving relative to a thermal cloud.

We created a Bose Einstein condensate of rubidium 87 though a sequence of laser cooling followed by magnetic trapping and forced evaporative cooling. The result was several tens of thousands of atoms at a few hundred nanokelvin.

The trapping potentials that hold our BEC are generated using wires etched on an atom chip. Combining the magnetic field from these wires with an external bias field creates a highly elongated trap about 100μm from the surface of our atom chip. The use of an atom chip allows precise control over the aspect ratios of our trap, allowing us to vary it between 140 and 90, while still allowing us to vary the temperature of the system.

The potential experienced by the atoms is anharmonic when close to the atom chip, due to imperfections in the current carrying wires. The result is a system in which the condensed atoms and the thermal atoms experience different trapping frequencies. This allowed us to resonantly excite the condensed cloud, while keeping the thermal cloud stationary.

The relative motion between the condensate and the thermal cloud results in damping of the system. By setting the system in motion, and then probing the system after variable times of evolution we build up a picture of the damping rate of the system. For condensates in three-dimensions the damping has been shown to be described by Landau damping, and in one-dimension there is expected to be no damping. Our condensate lies somewhere between the two regimes.

The chemical potential and the thermal energy of our system are comparable with the radial harmonic energy, the regime where the discrete nature of the excitation would be expected to have an effect, but are not small enough that the radial dimension is completely frozen out. In this regime we have measured how COM damping of our condensate varies as a function of temperature for two different aspect ratios. The results are well described by Landau theory, up to a constant, A_μ, which we found to be about 3.5.

A similar experiment carried out at Utrecht, but with a condensate in the three-dimensional regime, found that damping in their case is also well described by Landau damping, but with A_μ of twice what we find. A second finding from that experiment was that they measured a deviation from Landau damping as the hydrodynamicity, a measure of the thermal gas collision rate, was increased. We did not see a similar change in our experiment.

5.2 Discussion

The main result of our experiment is that the damping of the COM motion of a condensate appears to be reduced in highly elongated systems. Previous work has looked at damping in similar systems and found that in three-dimensional systems the damping was well described by Landau damping [22][23][28]. In purely one-dimensional system, examples of systems which do not damp way have been measured [36]. To the best of our knowledge, our experiment is the first to attempt to map out the parameter space between the two.

It is worth mentioning that in describing the damping mechanism we have used relationships developed to describe Landau damping. In doing so we have assumed that Landau damping is the correct mechanism. Given we are operating outside the regime that it was developed for, it is surprising that this would be the case. However, the scaling relationships between damping, temperature and number that we find are in agreement with those described by Landau damping, suggesting that Landau damping captures some aspect of the system.

It is believed that the differences measured between our experiment and the one carried out in Utrecht are due to us operating in the quasi-one-dimensional regime. In this regime the discrete nature of radial excitations would be expected to reduce the damping of the condensate, as it reduces the number of states into which scattering events can occur due. It is this scattering that causes the damping. However we know of no such theory that covers this regime.

5.3 Outlook

We hope that these experiments will motivate future work, both experimental and theoretical, to study damping mechanisms in BEC systems as the dimensionality changes. While we have mapped out some of the parameter space, there are still large areas left to explore. In this search it will also be interesting to investigate how other modes of the system damp.

APPENDIX A

RATES OF EVAPORATIVE COOLING

Following on from section 2.3.3, evaporative cooling has been essential in the achievement of BEC in atomic vapours. In this appendix we derive some rate equations to describe the process. What we hope to find is a connection for how readily measurable experimental quantities N, the total number of atoms in the cloud, and T the temperature of the cloud are expected to change.

In order to find a rate equation for how the relevant quantities change we are going look at how the ratio $\rho = \frac{N}{T}$ varies in time. By partial differentiation we get

$$\frac{\dot{\rho}}{\rho} = \frac{\dot{N}}{N} - \frac{\dot{T}}{T}\,. \tag{A.1}$$

The rate of change of number in the trap is given by

$$\dot{N} = -N(\gamma_e + \gamma_b)\,, \tag{A.2}$$

where γ_b is the background loss rate, and will just be a constant, γ_e is the rate at which atoms are 'evaporated'. In order for us to find an express for γ_e, let us assume that there is an cut off energy in the trap such that all atoms that have an energy above this will be removed immediately. This cut off energy is denoted as $\eta k_B T$, where η is known as the truncation parameter[1].

The evaporation rate is going to vary in time as both the atom number and temperature of the gas changes. To calculate evaporation rate we use the idea of detailed balance: the rate at which particles are scattered into high energy states is equal to the rate at which high energy particles would be scattered[2]. This rate is approximated as

$$\gamma_e = n_0 v(\eta)\sigma f(\eta)\,, \tag{A.3}$$

where $n_0 \propto \frac{N}{T^{\frac{3}{2}}}$ is the peak density of the trap, $v(\eta) \propto \sqrt{\eta T}$ is the velocity of an atom with energy $\eta k_B T$, σ is the scattering rate and $f(\eta)$ is the fraction of particles with energy above the cut off in a full thermal distribution. Putting this together we find that $\gamma_e \propto \rho$, allowing us to write $\gamma_e = \gamma_{e0}\rho$, where γ_{e0} is the initial scattering rate. The final rate of change of number is then

$$\frac{\dot{N}}{N} = -\gamma_{e0}\rho - \gamma_b\,. \tag{A.4}$$

To find how the temperature changes with time, we use the fact that in a power law potential the

[1] It's worth noting that this definition of the cut off energy includes T, which changes in time. To maintain a fixed value of η during evaporation we need to be lowing the lab frame cut off as a function of time.

[2] The assumption here is that when a collision involves a particle of high energy, the result is almost always results in two particles of lower energy.

energy and temperature are related by $E = \alpha N k_B T$, where α is a dimensionless constant depending on the type of trap, with $\alpha = 3$ for a three dimensional harmonic potential. This allows us to relate temperature change to energy change by

$$\frac{\dot{E}}{E} = \frac{\dot{N}}{N} + \frac{\dot{T}}{T}\,. \tag{A.5}$$

We can then relate the rate of change in energy to the rate of change in number using

$$\frac{\partial E}{\partial t} = \frac{\partial E}{\partial N}\frac{\partial N}{\partial t}\,, \tag{A.6}$$

and the fact that

$$\frac{\partial E}{\partial N} = \frac{\gamma_e \epsilon_e + \gamma_b \epsilon_b}{\gamma_e + \gamma_b}\,, \tag{A.7}$$

where ϵ_e is the average energy carried away by an evaporating atom, and ϵ_b is the energy carried away by an average atom.

These equations express the fact that we assume the only way to change the energy of the system is to remove atoms of specific energies from the system - either by evaporation or by standard loss. Combining these with the rate of change of number, the relative rate of change of energy is given by

$$\frac{\dot{E}}{E} = -\gamma_e \frac{\epsilon_e}{\epsilon_b} - \gamma_b\,. \tag{A.8}$$

Equation A.4 gives us the rate of change in number, equation A.8 gives us the rate of change of energy and A.5 relates these quantities to temperature change. Putting them all into A.1 we get the rate equation

$$\begin{aligned}\frac{\dot{\rho}}{\rho} &= 2\frac{\dot{N}}{N} - \frac{\dot{E}}{E} \\ &= -\rho\gamma_{e0}(2 - \frac{\epsilon_e}{\epsilon_b}) - \gamma_b \\ &= A\rho - \gamma_b\,,\end{aligned}$$

where $A = \gamma_{e0}(\frac{\epsilon_e}{\epsilon_b} - 2) = 2\gamma_{e0}(\eta - 1)$. To calculate the initial scattering rate and the ratio of evaporation energy to average energy integrals over the thermal distribution can be preformed. Both values will depend on the choice of truncation parameter.

By putting the equation in this form it is possible to solve it exactly to get

$$\rho(t) = \frac{e^{-\gamma_b t}\rho_0}{1 - \gamma_b^{-1} A\rho_0\,(1 - e^{-\gamma_b t})}\,, \tag{A.9}$$

where $\rho_0 = \rho(0)$.

Having found a solution for ρ, it is possible to use this to find expressions for N, T and the phase space density by solving the rate equations given above. The results are detailed and have been repeated elsewhere [64]. Instead we can gain some insight by looking at equation A.9.

If we set A to zero (the limit of no evaporation), we recover an exponential decay of number as expected. From its definition the quantity $0 \leq \gamma_b^{-1} A\rho_0 \leq 1$ for $\eta \geq 1$ (situations where you are removing hotter than average atoms). For non zero values, ρ decreased slower than for the zero case.

As we are necessarily removing more atoms in the case we must be cooling the system as well.

APPENDIX B

EXTRACTING AN ERROR WHEN FITTING

In several places in this thesis we extract information from a set of data by fitting a model to it. Whenever extracting information from such a process it is worth assigning an error to that information. In this appendix we outline how this is done. The method for presented here for calculating errors is taken from [89].

The fitting problem can be defined by assuming we have a set of pairs of data points of the form (x_i, y_i) that may be described by a function f such that $y_i = f(x_i, \alpha)$, where α can be one or more parameters.

B.1 Fitting with Errors in One Variable

For the first case, I will assume that x is error free. This roughly means that we set x and measure y to get our data[1]

We start by assuming that there are Gaussian errors on each point, so that what we really measured was $y_i = f(x_i, \alpha^*) + n(0, \sigma_i)$, where α^* are the real fitting parameters and $n(0, \sigma_i)$ is Gaussian noise with mean 0 and standard deviation σ_i. The maximum likelihood estimator for the parameters α are those that minimise the value of χ^2, defined by

$$\chi^2(\alpha) = \sum_i \frac{(y_i - f(x_i, \alpha))^2}{\sigma_i^2} \,. \tag{B.1}$$

In the special case where f is linear in the parameters α it is possible to minimise χ^2 analytically, however in some cases we are interested in it must be minimised numerically. We do this using using Mathematica.

Having obtained the values α_m that minimise χ^2, we can, in the case of normal errors on our initial data, use $\chi^2(\alpha)$ to calculate confidence intervals on our fitted parameters. To do this we note that the probability of obtaining some set of values $\mathbf{x}$, given a true set of parameters is

$$P(\mathbf{x}|\alpha^*) \propto \exp(-\chi^2(\mathbf{x}, \alpha^*)) \,. \tag{B.2}$$

Using Bayes' theorem, the probability of obtaining α_m, assuming a uniform prior, is given by

[1] In reality there is noise in our value of x however for the situations some situations it is negligible. The two main areas we make this assumption are in fitting to the COM oscillations - here x is the time and we are measure effects over 10s of millisecond, whereas the timing error in our equipment is on hundreds of nanoseconds level. The other case is fitting to density profiles of the cloud. Here x is position of each CCD pixel on our camera where differences are also negligible.

$$P(\alpha_m|\mathbf{x}) \propto \exp(-\chi^2(x,\alpha^*)\,. \tag{B.3}$$

We now make the approximation that around the minimum value we can expand χ^2 as

$$\chi^2(\alpha^*+\delta\alpha) \approx \chi^2(\alpha^*) + \frac{1}{2}\delta\alpha\mathbf{A}\delta\alpha^T\,, \tag{B.4}$$

where $A_{ij} = \frac{\partial^2\chi^2}{\partial\alpha_i\partial\alpha_j}$ and we have expanded using a Taylor series to second order.

Under this approximation the distribution of the outcome of our estimator, α_m, is given a multinomial distribution about α^*. This approach is the equivalent of assuming that we can linearise the function $f(\mathbf{x},\alpha)$ from B.1 in α.

The confidence intervals can be found by computing regions of constant χ^2. Under the above approximation the covariance matrix of α is given by A^{-1}, and the one sigma confidence interval for parameter α_i is given by $\sqrt{A^{-1}_{ii}}$. We refer to this approach as the parametric approach.

B.2 Non Parametric Confidence Intervals

The above procedure is valid only in the case of Gaussian errors. It also requires a knowledge of the size of the error to be known (we need to know the σ_i's in B.1). A non parametric approach offers an alternative. Here by non parametric we mean we make no assumptions about the distribution of the noise - we still use the same fitting function.

A standard method for finding non parametric confidence intervals is the bootstrap[90]. It involves resampling the original data with replacement to get a set of new datasets and then applying an estimator to each of these datasets to estimate the sampling distribution. From this approximate sampling distribution we estimate the confidence intervals.

For our data, the order of the data is important (the ordering in time of the oscillation data, and the ordering in space of the density profiles). This means that directly applying the bootstrap will not work, as resampling the data points could result in the new data not covering the same range as the old data, drastically changing the fit. Instead the approach is to resample the residuals of the fit, and apply the bootstrap to them. This assumes the the noise on each point is identical.

To carry out this approach we take our data d, still arranged in pairs (x_i, y_i) and from them estimate our model parameters, α_m using some estimator $T(d)$ (in our case we numerically minimise $\sum_i(y_i - f(x_i,\alpha))^2$ with respect to α). Next we calculate the residuals of this fit, $e_i = y_i - f(x_i,\alpha_m)$. For every values x_i, we can now produce a new value $y_{bi} = y_i + e_{bi}$, where each e_{bi} is drawn with replacement from the residuals of our original fit. This process results in a new dataset, d_b, which when fed into our estimator $T(d_b)$ produces a new set of fitter parameters α_b. Repeating this process many times over gives us the sampling distribution of T from which we get the confidence interval by finding the one sigma region of this distribution.

The bootstrap technique has been applied to the various situations where we fit and the result is identical to the errors found using the parametric technique.

REFERENCES

[1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of bose-einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, 1995.

[2] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969–3973, Nov 1995.

[3] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. Evidence of bose-einstein condensation in an atomic gas with attractive interactions. *Phys. Rev. Lett.*, 75:1687–1690, Aug 1995.

[4] S. N. Bose. Planks law and the light quantum hypothesis. *Z. Phys*, 1924.

[5] A. Einstein. quantentheorie des einatomigen idealen gases. *Preuss. Akad. Wiss.*, 1924.

[6] F. London. The λ-Phenomenon of Liquid Helium and the Bose-Einstein Degeneracy. *Nature*, 141:643–644, April 1938.

[7] Fritz London. *Superfluids*. Dover Publications, 1938.

[8] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of Superconductivity. *Physical Review*, 108:1175–1204, December 1957.

[9] J. Klaers, J. Schmitt, F. Vewinger, and M. Weitz. Bose-Einstein condensation of photons in an optical microcavity. *Nature*, 468:545–548, November 2010.

[10] J. Kasprzak, M. Richard, S. Kundermann, A. Baas, P. Jeambrun, J. M. J. Keeling, F. M. Marchetti, M. H. Szymańska, R. André, J. L. Staehli, V. Savona, P. B. Littlewood, B. Deveaud, and L. S. Dang. Bose-Einstein condensation of exciton polaritons. *Nature*, 443:409–414, September 2006.

[11] S. O. Demokritov, V. E. Demidov, O. Dzyapko, G. A. Melkov, A. A. Serga, B. Hillebrands, and A. N. Slavin. Bose-Einstein condensation of quasi-equilibrium magnons at room temperature under pumping. *Nature*, 443:430–433, September 2006.

[12] M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. Collective excitations of a bose-einstein condensate in a magnetic trap. *Phys. Rev. Lett.*, 77:988–991, Aug 1996.

[13] D. S. Jin, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Collective excitations of a bose-einstein condensate in a dilute gas. *Phys. Rev. Lett.*, 77:420–423, Jul 1996.

[14] C.J. Pethick and H. Smith. *Bose-Einstein Condensates in Dilute Gases*. Cambridge University Press, 2001.

[15] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. Oxford University Press, 2003.

[16] Franco Dalfovo, Stefano Giorgini, Lev P. Pitaevskii, and Sandro Stringari. Theory of bose-einstein condensation in trapped gases. *Rev. Mod. Phys.*, 71:463–512, Apr 1999.

[17] S. Stringari. Collective excitations of a trapped bose-condensed gas. *Phys. Rev. Lett.*, 77:2360–2363, Sep 1996.

[18] M. R. Andrews, D. M. Kurn, H.-J. Miesner, D. S. Durfee, C. G. Townsend, S. Inouye, and W. Ketterle. Propagation of sound in a bose-einstein condensate. *Phys. Rev. Lett.*, 79:553–556, Jul 1997.

[19] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. Observation of interference between Bose-Einstein condensates. In *APS April Meeting Abstracts*, page J1901, April 1997.

[20] D. M. Stamper-Kurn, H.-J. Miesner, S. Inouye, M. R. Andrews, and W. Ketterle. Collisionless and hydrodynamic excitations of a bose-einstein condensate. *Phys. Rev. Lett.*, 81:500–503, Jul 1998.

[21] D. S. Jin, M. R. Matthews, J. R. Ensher, C. E. Wieman, and E. A. Cornell. Temperature-dependent damping and frequency shifts in collective excitations of a dilute bose-einstein condensate. *Phys. Rev. Lett.*, 78:764–767, Feb 1997.

[22] L. P. Pitaevskii and S. Stringari. Landau damping in dilute Bose gases. *Physics Letters A*, 235:398–402, February 1997.

[23] P. O. Fedichev, G. V. Shlyapnikov, and J. T. M. Walraven. Damping of low-energy excitations of a trapped bose-einstein condensate at finite temperatures. *Phys. Rev. Lett.*, 80:2269–2272, Mar 1998.

[24] P. Öhberg, E. L. Surkov, I. Tittonen, S. Stenholm, M. Wilkens, and G. V. Shlyapnikov. Low-energy elementary excitations of a trapped bose-condensed gas. *Phys. Rev. A*, 56:R3346–R3349, Nov 1997.

[25] S. Giorgini. Damping in dilute Bose gases: A mean-field approach. *Phys. Rev. A*, 57:2949–2957, April 1998.

[26] P. O. Fedichev and G. V. Shlyapnikov. Finite-temperature perturbation theory for a spatially inhomogeneous bose-condensed gas. *Phys. Rev. A*, 58:3146–3158, Oct 1998.

[27] Onofrio Maragò, Gerald Hechenblaikner, Eleanor Hodby, and Christopher Foot. Temperature dependence of damping and frequency shifts of the scissors mode of a trapped bose-einstein condensate. *Phys. Rev. Lett.*, 86:3938–3941, Apr 2001.

[28] R. Meppelink, S. B. Koller, J. M. Vogels, H. T. C. Stoof, and P. van der Straten. Damping of superfluid flow by a thermal cloud. *Phys. Rev. Lett.*, 103:265301, Dec 2009.

[29] A. Görlitz, J. M. Vogels, A. E. Leanhardt, C. Raman, T. L. Gustavson, J. R. Abo-Shaeer, A. P. Chikkatur, S. Gupta, S. Inouye, T. Rosenband, and W. Ketterle. Realization of bose-einstein condensates in lower dimensions. *Phys. Rev. Lett.*, 87:130402, Sep 2001.

[30] Z. Hadzibabic, P. Krüger, M. Cheneau, B. Battelier, and J. Dalibard. Berezinskii-Kosterlitz-Thouless crossover in a trapped atomic gas. *Nature*, 441:1118–1121, June 2006.

[31] M. A. Cazalilla, R. Citro, T. Giamarchi, E. Orignac, and M. Rigol. One dimensional bosons: From condensed matter systems to ultracold gases. *Rev. Mod. Phys.*, 83:1405–1466, Dec 2011.

[32] S. Trotzky, Y.-A. Chen, A. Flesch, I. P. McCulloch, U. Schollwöck, J. Eisert, and I. Bloch. Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional Bose gas. *Nature Physics*, 8:325–330, April 2012.

[33] M. Cheneau, P. Barmettler, D. Poletti, M. Endres, P. Schauß, T. Fukuhara, C. Gross, I. Bloch, C. Kollath, and S. Kuhr. Light-cone-like spreading of correlations in a quantum many-body system. *Nature*, 481:484–487, January 2012.

[34] T. Langen, R. Geiger, M. Kuhnert, B. Rauer, and J. Schmiedmayer. Local emergence of thermal correlations in an isolated quantum many-body system. *Nature Physics*, 9:640–643, October 2013.

[35] M. Gring, M. Kuhnert, T. Langen, T. Kitagawa, B. Rauer, M. Schreitl, I. Mazets, D. A. Smith, E. Demler, and J. Schmiedmayer. Relaxation and Prethermalization in an Isolated Quantum System. *Science*, 337:1318–, September 2012.

[36] T. Kinoshita, T. Wenger, and D. S. Weiss. A quantum Newton's cradle. *Nature*, 440:900–903, April 2006.

[37] Bess Fang, Giuseppe Carleo, Aisling Johnson, and Isabelle Bouchoule. Quench-induced breathing mode of one-dimensional bose gases. *Phys. Rev. Lett.*, 113:035301, Jul 2014.

[38] Part rb/nf/6.4/17/ft.

[39] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005.

[40] Daniel A Steck. Rubidium 87 d line data. 2001.

[41] A. J. Leggett. BEC: The alkali gases from the perspective of research on liquid helium. 477:154–169, June 1999.

[42] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999.

[43] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6(11):2023–2045, Nov 1989.

[44] S. Chu. Nobel Lecture: The manipulation of neutral particles. *Reviews of Modern Physics*, 70:685–706, July 1998.

[45] W. D. Phillips. Nobel Lecture: Laser cooling and trapping of neutral atoms. *Reviews of Modern Physics*, 70:721–741, July 1998.

[46] C. N. Cohen-Tannoudji. Nobel Lecture: Manipulating atoms with photons. *Reviews of Modern Physics*, 70:707–719, July 1998.

[47] Harald F. Hess. Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen. *Phys. Rev. B*, 34:3476–3479, Sep 1986.

[48] W. Ketterle and N.J. van Druten. Evaporative cooling of atoms. *Adv. At. Mol. Opt. Phys.*, 37, 1996.

[49] J.T.M. Walraven. Atomic hydrogen in magnetostatic traps. *Proc. Scott. Univ. Summer. Sch. Phys.*, 44, 1996.

[50] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. *Phys. Rev. A*, 53:381–389, Jan 1996.

[51] J Söding, D Guéry-Odelin, P Desbiolles, F Chevy, H Inamori, and J Dalibard. Three-body decay of a rubidium bose–einstein condensate. *Applied physics B*, 69(4):257–261, 1999.

[52] M. V. Berry. The Levitron^{TM}: an Adiabatic Trap for Spins. *Royal Society of London Proceedings Series A*, 452:1207–1220, May 1996.

[53] O. Zobay and B. M. Garraway. Two-dimensional atom trapping in field-induced adiabatic potentials. *Phys. Rev. Lett.*, 86:1195–1198, Feb 2001.

[54] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger. Matter-wave interferometry in a double well on an atom chip. *Nature Physics*, 1:57–62, October 2005.

[55] B. E. Sherlock, M. Gildemeister, E. Owen, E. Nugent, and C. J. Foot. Time-averaged adiabatic ring potential for ultracold atoms. *Phys. Rev. A*, 83:043408, Apr 2011.

[56] R. Bücker, A. Perrin, S. Manz, T. Betz, C. Koller, T. Plisson, J. Rottmann, T. Schumm, and J. Schmiedmayer. Single-particle-sensitive imaging of freely propagating ultracold atoms. *New Journal of Physics*, 11(10):103039, October 2009.

[57] D. A. Smith, S. Aigner, S. Hofferberth, M. Gring, M. Andersson, S. Wildermuth, P. Krüger, S. Schneider, T. Schumm, and J. Schmiedmayer. Absorption imaging of ultracold atoms on atom chips. *Optics Express*, 19:8471, April 2011.

[58] C. F. Ockeloen, A. F. Tauschinsky, R. J. C. Spreeuw, and S. Whitlock. Detection of small atom numbers through image processing. *pra*, 82(6):061606, December 2010.

[59] P. Öhberg, E. L. Surkov, I. Tittonen, S. Stenholm, M. Wilkens, and G. V. Shlyapnikov. Low-energy elementary excitations of a trapped bose-condensed gas. *Phys. Rev. A*, 56:R3346–R3349, Nov 1997.

[60] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. Phase-fluctuating 3d bose-einstein condensates in elongated traps. *Phys. Rev. Lett.*, 87:050404, Jul 2001.

[61] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. Regimes of quantum degeneracy in trapped 1d gases. *Phys. Rev. Lett.*, 85:3745–3749, Oct 2000.

[62] R. Sewell. Matter wave interference on an atom chip. 2009. Ph.D. Thesis, Imperial College London.

[63] F. Baumgärtner. Measuring the acceleration of free fall with an atom chip interferometer. *Ph.D. Thesis, Imperial College London*, 2010.

[64] B. Yuen. Production and oscillations of a bose einstein condensate on an atom chip. *Ph.D. Thesis, Imperial College London*, 2013.

[65] R. Folman, P. Krüger, J. Schmiedmayer, J. Denschlag, and C. Henkel. Microscopic atom optics: From wires to an atom chip. *Advances in Atomic Molecular and Optical Physics*, 48:263–356, 2002.

[66] Jakob Reichel and Vladan Vuletic. *Atom Chips*. Wiley, 2011.

[67] R. J. Sewell, J. Dingjan, F. Baumgärtner, I. Llorente-García, S. Eriksson, E. A. Hinds, G. Lewis, P. Srinivasan, Z. Moktadir, C. O. Gollasch, and M. Kraft. FAST TRACK COMMUNICATION: Atom chip for BEC interferometry. *Journal of Physics B Atomic Molecular Physics*, 43(5):051003, March 2010.

[68] E. A. Hinds, C. J. Vale, and M. G. Boshier. Two-wire waveguide and interferometer for cold atoms. *Phys. Rev. Lett.*, 86:1462–1465, Feb 2001.

[69] M. P. A. Jones, C. J. Vale, D. Sahagun, B. V. Hall, C. C. Eberlein, B. E. Sauer, K. Furusawa, D. Richardson, and E. A. Hinds. LETTER TO THE EDITOR: Cold atoms probe the magnetic field near a wire. *Journal of Physics B Atomic Molecular Physics*, 37:L15–L20, January 2004.

[70] J. Estève, C. Aussibal, T. Schumm, C. Figl, D. Mailly, I. Bouchoule, C. I. Westbrook, and A. Aspect. Role of wire imperfections in micromagnetic traps for atoms. *Phys. Rev. A*, 70:043629, Oct 2004.

[71] P. Krüger, S. Wildermuth, S. Hofferberth, L. Mauritz Andersson, S. Groth, I. Bar-Joseph, and J. Schmiedmayer. Cold atoms close to surfaces: measuring magnetic field roughness and disorder potentials. *Journal of Physics Conference Series*, 19:56–65, January 2005.

[72] A. E. Leanhardt, Y. Shin, A. P. Chikkatur, D. Kielpinski, W. Ketterle, and D. E. Pritchard. Bose-Einstein Condensates near a Microfabricated Surface. *Physical Review Letters*, 90(10):100404, March 2003.

[73] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A simple extended-cavity diode laser. *Review of Scientific Instruments*, 69(3), 1998.

[74] C. P. Pearman, C. S. Adams, S. G. Cox, P. F. Griffin, Smith D. A., and I. G. Hughes. Polarization spectroscopy of a closed atomic transition: applications to laser frequency locking. *J. Phys. B: At. Mol. Opt. Phys.*, 35:5141, Jun 2002.

[75] G. Ritt, G. Cennini, C. Geckeler, and M. Weitz. Laser frequency offset locking using a side of filter technique. *Applied Physics B*, 79(3):363–365, 2004.

[76] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman. Low-velocity intense source of atoms from a magneto-optical trap. *Phys. Rev. Lett.*, 77:3331–3334, Oct 1996.

[77] C. Sinclair. Bose-einstein condensation in microtraps on videotape. *Ph.D. Thesis, Imperial College London*, 2005.

[78] M. Succo. An integrated optical-waveguide chip for measurement of cold-atom clouds. *Ph.D. Thesis, Imperial College London*, 2011.

[79] Y. Castin and R. Dum. Bose-einstein condensates in time dependent traps. *Phys. Rev. Lett.*, 77:5315–5319, Dec 1996.

[80] M. Schellekens, R. Hoppeler, A. Perrin, J. V. Gomes, D. Boiron, A. Aspect, and C. I. Westbrook. Hanbury Brown Twiss Effect for Ultracold Quantum Gases. *Science*, 310:648–651, October 2005.

[81] S. Dettmer, D. Hellweg, P. Ryytty, J. J. Arlt, W. Ertmer, K. Sengstock, D. S. Petrov, G. V. Shlyapnikov, H. Kreutzmann, L. Santos, and M. Lewenstein. Observation of phase fluctuations in elongated bose-einstein condensates. *Phys. Rev. Lett.*, 87:160406, Oct 2001.

[82] A. Imambekov, I. E. Mazets, D. S. Petrov, V. Gritsev, S. Manz, S. Hofferberth, T. Schumm, E. Demler, and J. Schmiedmayer. Density ripples in expanding low-dimensional gases as a probe of correlations. *PRA*, 80(3):033604, September 2009.

[83] S. Manz, R. Bücker, T. Betz, C. Koller, S. Hofferberth, I. E. Mazets, A. Imambekov, E. Demler, A. Perrin, J. Schmiedmayer, and T. Schumm. Two-point density correlations of quasicondensates in free expansion. *PRA*, 81(3):031610, March 2010.

[84] G. Lamporesi, S. Donadello, S. Serafini, F. Dalfovo, and G. Ferrari. Spontaneous creation of Kibble-Zurek solitons in a Bose-Einstein condensate. *Nature Physics*, 9:656–660, October 2013.

[85] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. *eprint arXiv:cond-mat/9904034*, April 1999.

[86] R. Dubessy, C De Rossi, T. Badr, L. Longchambon, and H. Perrin. Imaging the collective excitations of an ultracold gas using statistical correlations. *Arxiv*, Oct 2014.

[87] Stephen R. Segal, Quentin Diot, Eric A. Cornell, Alex A. Zozulya, and Dana Z. Anderson. Revealing buried information: Statistical processing techniques for ultracold-gas image analysis. *Phys. Rev. A*, 81:053601, May 2010.

[88] R. Meppelink. Hydrodynamic excitations in a bose-einstein condensate. *Ph.D. Thesis, University Utrecht*, 2009.

[89] W. H. Press, B. P. Flannery, and S. A. Teukolsky. *Numerical recipes. The art of scientific computing.* 1986.

[90] Larry Wassermann. *All of Statistics.* Springer, 2003.