UNIVERSITY OF NOTTINGHAM

PHD THESIS

An Apparatus for the Production of Bose-Einstein Condensates in Tunable Geometries on a Chip

Author:
Thomas J. BARRETT

Supervisors:
Professor Peter KRÜGER
Dr. Thomas FERNHOLZ

Examiners:
Dr. Donatella CASSETTARI
Dr. Christopher MELLOR

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

in the

Ultracold Atoms Research Group
School of Physics and Astronomy

1st February 2017

Declaration of Authorship

I, Thomas J. BARRETT, declare that this thesis titled, "An Apparatus for the Production of Bose-Einstein Condensates in Tunable Geometries on a Chip" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.

Signed:

Date:

Abstract

Atom chips are an excellent tool for studying ultracold degenerate quantum gases, due to the high degree of controllability afforded by the precise potentials generated from the current-carrying microfabricated wires on the chip surface. The geometries of the trapping potentials are inherently capable of realising extreme aspect ratios, and therefore creating model systems with effectively reduced dimensionality, particularly the theoretically-tractable one-dimensional Bose gas. In addition, the temporal tunability makes it possible to impart non-adiabatic changes on the trapping potentials, allowing experimental investigation of samples which have been brought out of equilibrium - a situation which is not fully theoretically understood.

This thesis describes the implementation, development and characterisation of an experimental system for producing the first Bose-Einstein condensates of atomic rubidium 87 gas trapped on the surface of an atom chip in Nottingham. Such an apparatus is very complex and requires careful characterisation in order to run in a stable and reliable way. Details of the experimental setup are thoroughly outlined, including the vacuum system, lasers, electronics, computer control and timing, and the optical imaging system.

A newly installed compact two-dimensional magneto-optical trap provides an atom loading rate of $5 \times 10^7 \mathrm{s}^{-1}$ for loading a three-dimensional mirror-magneto-optical trap with 1.5×10^8 atoms, at a temperature of $300\,\mu\mathrm{K}$ within $10\,\mathrm{s}$. The cloud is then sub-Doppler cooled to $50\,\mu\mathrm{K}$, and spin-polarised with 96% purity into the $|F = 2, m_F = +2\rangle$ ground state within $5\,\mathrm{ms}$, in preparation for loading a purely magnetic trap. A millimeter sized copper Z-shaped conductor located beneath the atom chip surface creates a Ioffe-Pritchard magnetic trap, into which the laser cooled cloud is loaded with $\sim 70\%$ efficiency, and can be held with a vacuum-limited lifetime of $40\,\mathrm{s}$. Evaporative cooling then pre-cools the sample to below $20\,\mu\mathrm{K}$ within $10\,\mathrm{s}$, to allow the subsequent loading into potentials created by the atom chip with 100% efficiency. A final evaporation stage then cools the cloud below the phase transition temperature of $\sim 800\,\mathrm{nK}$, resulting finally in pure BECs with 10^5 atoms confined using the atom chip.

Key measurements of various properties of the trapped condensates are presented, which are important in order to characterise the system fully, and to compare with theoretical expectations. In particular, included are the variation of condensate fraction with temperature, the BEC expansion dynamics, and the condensate lifetime in the trap, for example. Finally, it is demonstrated how BECs can be produced on the atom chip without the use of external macroscopic coils, achieved by using novel, integrated sheet structures located beneath the chip surface - unique to this experimental system - to create the necessary bias fields.

"They are all your hats, Mr. Angier."

DAVID BOWIE, as **NIKOLA TESLA**

Acknowledgements

I would like to take the opportunity to thank just a few of the many people who have contributed both directly and indirectly to helping me see this PhD completed, and to making my time and experience in Nottingham such a positive one.

Firstly, sincere thanks go to my supervisor, Peter, for initially taking me on as a summer student having little experience, and then giving me the opportunity to go on and complete a PhD in your group. Your way of thinking has influenced me a great deal, and taught me how to tackle all kinds of problems, distilling them down into their most basic elements. I am continually impressed by your intuitive approach, and ability to explain things to me with useful analogies (using pints of beer as visual aids on more than one occasion!). Most importantly, thank you for making me feel confident to ask basic, silly questions, and for always making time whenever I needed support.

To my post-doc and good friend, Fedja, who has been with me throughout my entire PhD, I really can't thank you enough. Right from my early days in the lab, you have always given up your time and been happy to spend hours on discussions, and as a result I have learned an awful lot from you about many different aspects of experimental physics (not to mention how to drink whisky, and to take my coffee with neither sugar nor milk). Your endless patience is thoroughly appreciated, even when I ask you the same questions every six months or so! – don't worry, one of these days I'll really get to grips with the chemical potential [1].

My appreciation goes to Thomas Fernholz, for being the sage to whom we turned whenever a technical problem eventually got the better of us. Thank you for making me realise that even some things I thought that I knew, I in fact didn't understand well enough! Your aptitude for science and engineering never ceases to impress, and whenever possible I try to emulate your thorough, no-stone-unturned way of dealing with problems. Apologies for disturbing your data-taking next door with our stray magnetic fields!

From a practical standpoint, the experimental setup could not have been constructed without the assistance of the many efficient technical supporting staff in the School of Physics & Astronomy. Thanks to the skilled members of the mechanical workshop - namely Dave Laird, Tommy, and Terry in particular, for manufacturing the intricate structures that comprise the heart of our experiment, and to Pete for always finding a way to squeeze in the many jobs I inundated you with. Big thanks to Ian, Bob, and Andy of the electronics support group, for always having an open door to brainstorm new projects, putting up with the many repairs sent your way, and for coming up with creative, outside-the-box solutions.

The ultracold atoms labs have been made such a welcoming and friendly place by all members of the group, past and present. It was a pleasure to work within such an internationally diverse culture, and one in which everyone is always more than happy

to get stuck in and help each other out with whatever issues are faced. Thank you in particular to Anton and Gal for the large amount of build work and progress made during your time before I arrived, and for showing me the ropes in the early months of my PhD.

There are also a host of people outside the lab who been a supportive presence throughout the last four years. Thanks to my friend George for being such an easy-going flatmate to live with, for sharing in the Bond-a-thon, and for reliably never saying no to ordering takeaway pizza after a late evening in the lab - I have thoroughly enjoyed our PhD years together. Notably, relief from work was often provided by having the company of Emma and Phil – thanks to you two for doing the cooking, and I will miss all those trips to the pictures. A big part of my life in Nottingham when not in the lab has also been in the pool and snooker teams, where we have had some good times with big wins, and I have made great friends.

Thank you to my lovely Rosie for your friendship and amazing support over the years, always being on the end of the phone, and somehow managing to put up with me when I was either talking about physics, or more often just being moody and stressed. You are a real source of inspiration, always helping to give me perspective. I am constantly learning new things from you, and feel so lucky to have you in my life. You are a star.

Finally, although a simple thank you could never do justice, I will say it anyway. The only reason I am where I am today is because of my parents, Jan and Dean. You have both been there to support me constantly throughout my whole life whenever I have needed it, and have always put myself and Oli first in everything you do. You two are the best role models anyone could ask for - thank you for all the love and encouragement, and for helping me along this far.

Contents

Chapter 1

Introduction

Research with ultracold atoms has grown into one of most active, rich and exciting fields in physics over the last several decades. The advent of laser cooling [2, 3] and the magneto-optical trap [4] in the 1980s, resulting from the availability of narrow-band tunable lasers, provided convenient and widespread access to samples of sub-millikelvin gases, and were rewarded with the 1997 Nobel Prize in Physics [5–7]. A long-standing goal then became the achievement of Bose-Einstein condensation (BEC) in a dilute atomic vapour, as predicted 70 years earlier together by Bose and Einstein [8, 9], whereby a macroscopic number of particles are able to occupy the same lowest quantum state. This occurs when the system is cooled below a critical temperature such that the de Broglie wavelengths of particles becomes on the order of their separation, allowing the gas to be described by a single macroscopic wavefunction, and is purely a consequence of quantum statistics. In order to maintain the metastable gaseous state, and prevent the vapour from solidifying, extremely low densities are required (around a million times less dense than air), which necessitates reaching very low temperatures on the order of hundreds of nanokelvin. The breakthrough came in 1995 with the technique of evaporative cooling in magnetic traps [10], and BECs of rubidium, sodium and lithium gases were reported later the same year at JILA, MIT and RICE [11–13]. The achievements were recognised with the Nobel Prize in Physics in 2001 [14, 15].

Many early experiments with BECs focused on macroscopic collective phenomena, perhaps most dramatically with the initial observation of matter-wave interference of two overlapping condensates [16], followed by the so-called "atom laser" [17], and later by direct measurement of long-range spatial phase coherence [18]. The inherent superfluid properties of BEC were demonstrated, most notably by the appearance of vortices [19, 20] and vortex lattices [21], along with dynamical effects of the macroscopic wavefunction such as the generation of solitons [22, 23] and collective excitations [24–26]. On the back of the initial landmark experiments, a multitude of species have since been Bose-condensed, including all the alkali metals,[1] hydrogen, chromium, strontium, calcium, ytterbium and metastable helium. In addition, BECs with molecules of fermionic atoms have been created [28–30], along with more recently a gas of photons [31]. In

[1] With the exception of francium [27].

spite of this, due to their favourable elastic scattering properties and repulsive interactions, combined with a level structure well-suited to laser cooling, rubidium and sodium remain widespread and popular species for use in BEC experiments.

As well as the original magnetically-trapped incarnations, the realisation of optically-confined BECs [32–34] opened the possibility of widely tuning the strength of interactions between atoms using an external magnetic field via Feshbach resonances [35, 36], if the atomic species is chosen appropriately. This added yet another degree of controllability for experimentalists, and eventually paved the way for studies of the BEC-BCS crossover using fermionic molecules [37–39]. With increasingly novel trapping geometries being explored, the potentials created by the standing waves in optical lattices [40, 41] allow production of periodic arrays of multiple degenerate quantum gases. Optical potentials have since become popular for use as quantum simulators for condensed matter and solid state systems, where they can be used to mimic the behaviour of electrons in the lattice potential of a solid, for example. Work along this vein has resulted in the experimental realisation of systems described by the Bose-Hubbard Hamiltonian, accompanied by observation of the transition between the superfluid and Mott insulating phases [42, 43]. Many condensed matter systems have been imitated using cold atoms, including studies of Anderson localisation [44–46] and Josephson junctions [47, 48], to name a few. As well as condensed matter and solid state systems, the high coherence, low speed of sound, and ultralow temperatures of BECs make them candidates for modelling more exotic phenomena, such as early-universe physics [49], acoustic black holes and Hawking radiation [50–52], as well as the dynamical Casimir effect [53].

Ultracold atom systems provide a remarkable degree of control over the internal and external degrees of freedoms of particles using optical, magnetic and electric fields, as well as adjustment of the interaction strengths through either Feshbach resonances or the particle density. In contrast to condensed matter systems, the interactions between atoms can be described in a simple way due to the diluteness of the gas, which allows direct comparison with calculations and predictions. The versatile and precise nature of ultracold atom systems allows engineering of specific Hamiltonians in a controlled, isolated and defect-controlled environment, and provides an excellent testing ground for theoretical models over a wide parameter space. In addition, the relatively large length scales involved means that quantum objects can be investigated directly with optical imaging. The low temperatures and weak interactions also give rise to often relatively slow dynamics, which can be probed on experimentally accessible timescales using rapid imaging methods.

The research reported in this thesis centres around the magnetic trapping of neutral atoms using the so-called *atom chip*, about which there have been written many extensive reviews [54–60]. These devices consist of microfabricated current-carrying conductors on the surface of a substrate (in our case gold on a silicon wafer), and were first employed at the turn of the millenium [61, 62] inspired by the manipulation of

atoms using simple free-standing wire structures [63–66]. The current flowing through the chip wires creates the necessary magnetic field configuration for atom trapping, and by marrying with techniques from the well-established field of microfabrication, the conductor layout can be be tailored into almost arbirary shapes for creating intricate and flexible potential landscapes. As a result of the magnetic field scaling with distance as $|\mathrm{B}| \propto 1/r$, large field gradients and curvatures are offered at close proximity to the trapping wires, and along with the low inductance this allowed a new level of accurate spatial and temporal control - in contrast to the initial experiments with large external coils located outside the vacuum chamber. After the first achievement of on-chip BECs [67, 68], various novel schemes for complex transport with sub-micron precision using magnetic conveyor belts and splitting were developed, and the atom chip began to be seen as a toolbox for manipulating matter-waves near surfaces for atom optics [69, 70]. The coherence of these matter-waves on an atom chip was successfully demonstrated by interference of BECs, which were split coherently and recombined by using radio-frequency-induced adiabatic potentials to create a double well [71], providing a powerful method for phase analysis of trapped quantum gases. Importantly, one of the major benefits of the high field gradients for only modest currents ($< 1\,\mathrm{A}$) with atom chips is the ability to perform rapid evaporative cooling sweeps, allowing fast, robust and efficient BEC production with a state-of-the-art repetition rate of currently up to $1\,\mathrm{Hz}$ [72, 73].

Clouds trapped using atom chips are typically located at distances of $\sim 0.1\,\mu\mathrm{m} - 100\,\mu\mathrm{m}$, which permits the study of atom-surface effects such as, for example, the attractive Casimir-Polder force that is present when a cloud is brought close to the surface [74–76]. However, an important effect that occurs for small distances to the trapping conductor is an undesirable fragmentation of the atomic density [77–79] which has been shown to be caused by the additional field components produced by irregular current flow from the inherent roughness of the wire [80–82]. With higher quality fabrication techniques, the severity of fragmentation is greatly reduced [83], and distances on the order of and below $10\,\mu\mathrm{m}$ are now easily achievable. These roughness effects have actually inspired the use of cold atoms as a magnetic field microscope, allowing the current fluctuations to be measured with both high sensitivity and high resolution [84, 85], and even providing insight into the transport properties of electrons in metals [86]. Another fundamental limitation of atom chips arises due to the randomly-fluctuating magnetic fields (Johnson noise) caused by thermal motion of electrons in the conducting surface [87–91]. This results in increased heating rates and spin-flip losses from the trap for noise in the kHz and MHz ranges, respectively, when approaching distances of less than tens of micrometers.

The strong trapping gradients afforded by atom chips mean that they are capable of producing cylindrically-symmetric traps with extreme aspect ratios (of at least several thousand), and so are naturally well-suited to more fundamental studies of one-dimensional gases and quantum many-body physics [92–94]. The behaviour of

cold gases when moving to lower dimensions is fundamentally different, and several regimes can be realised. In uniform systems, BEC can occur in two dimensions (2D) only at zero temperature, and in one dimension (1D) is not possible even at $T = 0$. Nevertheless, a uniform two-dimensional gas can become superfluid at finite temperature via the so-called Berezinskii-Kosterlitz-Thouless transition [95, 96], due to the appearance of order from the binding of vortex-antivortex pairs, although this effect is driven by interactions and not purely quantum statistics. In practice however, gases are usually confined to harmonic traps, and BECs are then possible in one, two and three dimensions at finite temperature [97], and were initially experimentally demonstrated in [98], attracting more attention to the study of these systems. A one-dimensional gas can be realised in a cylindrically-symmetric trap by providing tight confinement in the radial direction. Quantitatively, if the chemical potential and thermal energy are reduced to below the radial harmonic oscillator level spacing, i.e. $\mu, k_{\mathrm{B}}T \ll \hbar\omega_{\perp}$, then no excitation to higher levels is possible in this dimension, which "freezes out" the motion and makes the gas effectively one-dimensional.

In one-dimensional geometries, the behaviour of the gas depends on its density and the strength of the interactions [99, 100]. For lower densities, the system takes the form of the strongly-correlated Tonks-Girardeau gas, whereby the repulsive interactions localise particles spatially, mimicking the Pauli exclusion principle, and leading to "*fermionisation*" of the bosons. This regime has been reached in optical lattices [101, 102], and its presence can be detected, for example, by a reduction in three-body loss rates [103]. A Tonks gas has not as yet been realised on an atom chip and, in contrast to optical lattices, would allow the production of a single realisation per shot whose momentum distribution dynamics could be studied, for example.

In the opposite high density limit, the gas is weakly-interacting, and it is possible for a BEC to form that is described by the Gross-Pitaevskii equation. However, even at the ultralow temperatures encountered in these systems, thermal fluctuations will reduce the phase coherence length l_{ϕ} [104] and phase fluctuations are observed to occur in a BEC along its axis [105, 106], which are translated into density modulations in time-of-flight through self-interference [107, 108]. However, as the gas is cooled below a characteristic temperature T_{ϕ} (which is below T_c, the degeneracy temperature for BEC), then l_{ϕ} becomes longer than the finite length of the gas, and phase coherence is recovered. For temperatures $T_{\phi} < T < T_c$, the equilibrium state is a so-called *quasi-condensate* that is described by a single macroscopic wavefunction with a fluctuating phase. This quasi-condensate regime was recently reached in our experiment when exploring elongated trapping geometries with the atom chip, and was observed by the signature density ripples in time-of-flight absorption images, described above. For temperatures below T_{ϕ}, phase coherence is established and a *true condensate* is recovered.

In addition to investigation of the equilibrium state of gases in the various regimes as described above, the dynamics of non-equilibrium quantum many-body systems

is currently an active field of study, and is in its relative infancy. A general question is the relaxation of isolated quantum systems to equilibrium, and various experiments have been conducted over the last decade, usually monitoring the evolution of a system after inducing a perturbation or performing a quench [109–113]. Ultracold atoms are ideal for these types of studies, partly because their isolation from the environment effectively renders the system closed, and also because the dynamics occurs on experimentally accessible timescales, essentially dictated by collision times. In fact, in [114] an absence of any sign of thermalisation was observed after perturbing a one-dimensional Tonks gas even after many thousands of collisions, due to the system's near-integrability. Experiments with one-dimensional Bose gases on atom chips have been used to study the phenomenon of *prethermalisation*, whereby an out-of-equilibrium system is seen to rapidly approach an intermediate quasi-stationary state that can appear thermal, but on a timescale much faster than the true expected thermalisation time [115–117]. The ability of the atom chip to create individual realisations of one-dimensional Bose gases in a highly controllable way, as well as directly probing of the phase properties through matter-wave interference, makes them a versatile tool for future studies on the dynamics of many-body quantum systems.

1.1 Thesis Overview

This thesis describes the implementation and characterisation of an experiment to produce the first BEC on an atom chip in Nottingham. The result is a system which is capable of generating on-chip BECs in a stable and repeatable way, which will enable studies of out-of-equilibrium quantum gases in the future by exploiting the flexible potentials generated by the atom chip. The thesis takes over from the work of the first generation of PhD students in this laboratory, Dr. Anton Piccardo-Selg and Dr. Gal Aviv, and further experimental details regarding the earlier stages of the setup and a portable version of the atom trapping assembly can be found in their respective theses [118, 119]. The thesis is organised over the chapters as follows:

Chapter 2 presents an overview of the relevant theoretical concepts which underpin and support the various steps in the experiment, as well as those necessary for interpreting its results. Primary topics include laser cooling and magnetic trapping geometries, a description of BEC in an harmonic trap, and a model for evaporative cooling. The aim is to emphasise key aspects which are directly relevant for later chapters, so the treatments are not fully self-contained and additional references are given where appropriate.

Chapter 3 describes the technical implementation of the experimental apparatus, including vacuum and laser systems, structures for creating various magnetic trapping fields, computer control, the imaging system, and various additional specific hardware. A detailed characterisation of the relevant components was carried out.

Chapter 4 provides an account of the creation of an experimental sequence for trapping, cooling, and producing BECs on the atom chip. Details of technical optimisation processes are given for the various stages, which are comprised essentially of the 2D and 3D magneto-optical traps, magnetic trapping with with macroscopic copper conductors, radio-frequency forced evaporative cooling, and loading to the atom chip.

Chapter 5 presents measurements and characterisation of the salient properties of the BEC, including variation of condensate fraction with temperature, expansion dynamics, lifetime measurements, and the production of a BEC using integrated trapping structures, without external coils.

Chapter 6 offers a brief summary of the work carried out, the current status of the experiment, and gives an outlook for the future.

Chapter 2

Theoretical Background

This chapter provides an overview of the basic theoretical concepts necessary to understand the performance of the experiment. The emphasis is on deriving results and expressions which are directly applicable to interpreting various measurements in later chapters, whilst highlighting the most relevant aspects of the underlying physics. Topics include the fundamental mechanism of the interaction of atoms with both laser light and with external magnetic fields, which underpin the experimental techniques of laser cooling, magnetic trapping, and evaporative cooling. In addition, the important properties of BECs in both the ideal gas and interacting cases are described. The treatment of these topics emphasises key aspects and relates them to the experiments presented later. It therefore does not provide a full self-contained account, and additional references, in which more details can be found, are provided throughout.

2.1 Atomic Structure of Rubidium

The majority of experiments with ultracold neutral atoms are performed using the alkali metals. These elements all have a single valence electron in a zero-angular momentum state outside a number of closed shells. For example, rubidium has four filled shells with the outer electron in the $5s$ state. As a result, the energy levels of these hydrogen-like atoms are particularly simple to analyse, and are determined mostly by the outer electron along with its coupling to the atomic nucleus. It is the magnetic moment of the valence electron which provides sufficiently stong coupling with external magnetic fields to allow the atoms to be magnetically trapped. All the alkali metals have an odd number of protons, and so the number of neutrons determines whether a particular isotope is bosonic or fermionic, altering the quantum statistics at low temperatures. This experiment employs rubidium-87, which is a composite boson, allowing it to undergo Bose-Einstein condensation.

The alkali atoms can be crudely modelled in the same way as the hydrogen atom, consisting of a single electron moving in the central potential due to the confining Coulomb force from the nucleus. When applying the non-relativistic Schrödinger equation, and ignoring all spin coupling interactions, the eigenstates are denoted in the standard way by $|n, L, m_L\rangle$, where n, L and m_L are the principal, azimuthal, and magnetic quantum numbers, respectively. The quantum numbers L and m_L are related to the

eigenvalues of the orbital angular momentum operator $\hat{\boldsymbol{L}}$ through the usual relations [120]

$$\hat{\boldsymbol{L}}^2 |n, L, m_L\rangle = \hbar^2 L(L+1) |n, L, m_L\rangle \tag{2.1}$$

$$\hat{L}_z |n, L, m_L\rangle = \hbar m_L |n, L, m_L\rangle , \tag{2.2}$$

where $\hat{L}_z$ is the projection of $\hat{\boldsymbol{L}}$ onto the z-axis. In this framework, the energy eigenstates E_n are proportional to $1/n^2$, and do not depend on the angular momentum quantum numbers, leading to a degeneracy of levels with different values of L.

However, this level of approximation is far from accurate enough to describe the transitions exploited in cold atom experiments, as there are various effects which conspire to split the energy levels into several manifolds. The fact that the electron also possesses a spin angular momentum, described by $\hat{\boldsymbol{S}}$, leads to an interaction between its magnetic moment and its orbital motion around the nucleus, giving rise to an additional spin-orbit coupling term in the Hamiltonian of the form $\hat{H}_{\text{so}} \propto \hat{\boldsymbol{L}} \cdot \hat{\boldsymbol{S}}$ [120]. This leads to the introduction of the quantum numbers S and m_S, and so a possible basis to use would be $|n, L, S, m_L, m_S\rangle$. However, $\hat{H}_{\text{so}}$ is not diagonal in this basis, because m_L and m_S are not good quantum numbers. This is a reflection of the fact that even though the magnitudes of the total spin and angular momentum are fixed, their relative orientation (and therefore projection onto the z-axis) is not a constant of motion due to their interaction [121]. However, the total angular momentum, $\hat{\boldsymbol{J}} = \hat{\boldsymbol{L}} + \hat{\boldsymbol{S}}$, is conserved - both the spin and oribital angular momenta precess around the total $\hat{\boldsymbol{J}}$, whilst $\hat{J}_z$ remains fixed. A more appropriate basis to use is $|n, L, S, J, m_J\rangle$, as the interaction $\hat{H}_{\text{so}}$ is then diagonal, and m_J is a good quantum number. By writing the interaction in the form

$$\hat{\boldsymbol{L}} \cdot \hat{\boldsymbol{S}} = \frac{1}{2}\left(\hat{\boldsymbol{J}}^2 - \hat{\boldsymbol{L}}^2 - \hat{\boldsymbol{S}}^2\right), \tag{2.3}$$

we can immediately calculate the expectation value of the energy shift due to this coupling term

$$\Delta E_{\text{so}} \propto [J(J+1) - L(L+1) - S(S+1)] , \tag{2.4}$$

from which it can be seen that the energy of the state depends now on J, partly lifting the degeneracy of levels with the same non-zero value of L. This splitting of energy levels is referred to as the *fine structure*.

The energy level structure of the valence electron in ^{87}Rb is shown in Fig. 2.1, depicting transitions from the $L = 0 \rightarrow L = 1$ states - known as the D-line [122]. The spin-orbit interaction splits the 5^2P state into two levels with $J = 1/2$ and $J = 3/2$ - the transitions between which and the single ground state are denoted the D_1 and D_2 lines, respectively. The D_2 line is the one of interest in this work, as it contains a cycling transition which facilitates laser cooling. Since the two lines are separated by $15\,\text{nm}$, and a typical laser linewidth is $\sim 1\,\text{MHz}$, the D_1 line can effectively be ignored for practical purposes.

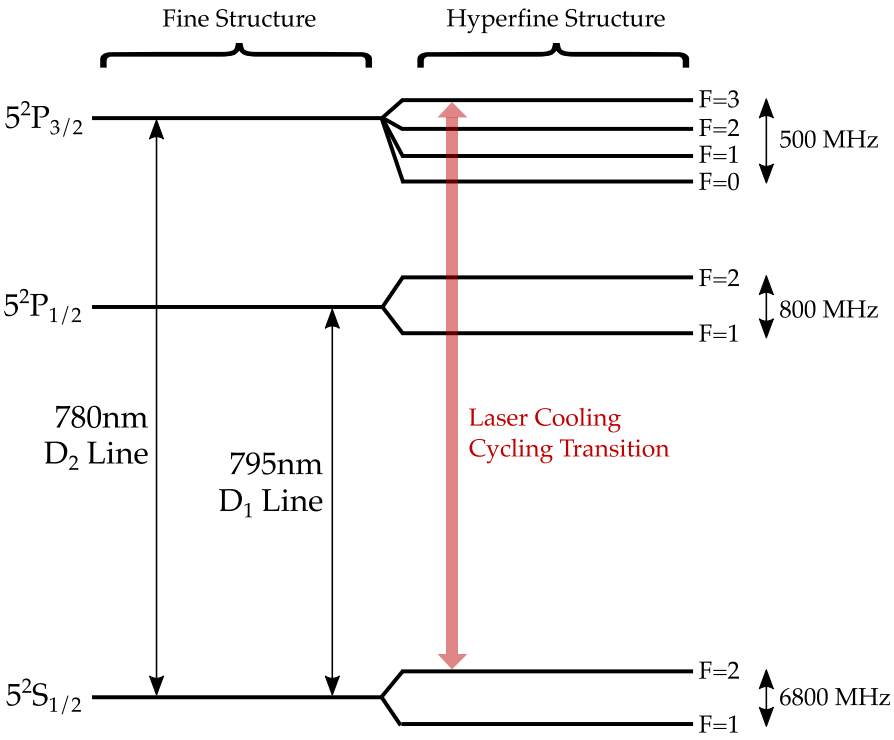

FIGURE 2.1: Rubidium-87 D-line energy levels: The upper 5^2P state is split according to the value of the total electron angular momentum J, resulting in the fine structure comprised of the D_1 and D_2 lines. States are further split into the hyperfine structure according to the total atomic angular momentum F. The transition $|5^2S_{1/2}, F = 2\rangle \rightleftharpoons |5^2P_{3/2}, F' = 3\rangle$ is the cycling transition used for laser cooling, and the large separation between the D_1 and D_2 lines means that the former can be effectively ignored. Values taken from [122].

Figure 2.1 shows that this fine structure doublet is also further split into additional *hyperfine* sublevels, which arise when the intrinsic spin of the nucleus $\hat{\boldsymbol{I}}$ is taken into account. The nuclear spin gives rise to a nuclear magnetic moment, which then couples to both the spin and orbital angular momentum of the valence electron. This magnetic dipole interaction adds a term to the total Hamiltonian of the form $\hat{H}_{\text{hf}} = A_{\text{hf}}\hat{\boldsymbol{I}} \cdot \hat{\boldsymbol{J}}/\hbar^2$, where A_{hf} is the magnetic dipole constant for the particular J level being considered, and has units of energy [120]. The hyperfine interaction is several orders of magnitude smaller than the fine structure - since the nuclear magnetic moment is much smaller than that of the electron, owing to the large mass differences - but still cannot be neglected for cold atom experiments.

The hyperfine interaction causes $\hat{\boldsymbol{J}}$ and $\hat{\boldsymbol{I}}$ to precess around their resultant total angular momentum, $\hat{\boldsymbol{F}} = \hat{\boldsymbol{I}} + \hat{\boldsymbol{J}}$, and so it is now the projection $\hat{F}_z$ which is the conserved quantity. This means that m_F is a good quantum number, instead of m_I and m_J, and the most convenient basis to use when dealing with the hyperfine interaction is $|n, L, S, J, F, m_F\rangle$. In an analagous way to the fine structure, the Hamiltonian $A_{\text{hf}}\hat{\boldsymbol{I}} \cdot \hat{\boldsymbol{J}}$

is diagonal in this basis, and the resultant energy shifts are easily calculated to be

$$\Delta E_{\text{hf}} = \frac{1}{2} A_{\text{hf}} \left[F(F+1) - I(I+1) - J(J+1) \right]. \tag{2.5}$$

Again, the degeneracy is partly removed, and the energy levels for a given J are now split according to their value of F. For ^{87}Rb the nuclear spin is $I = 3/2$, and so the $5^2\mathrm{S}_{1/2}$ state is split into two levels separated by $h \times 6.8\,\mathrm{GHz}$, whilst the $5^2\mathrm{P}_{3/2}$ is split into four hyperfine levels, as shown in Fig. 2.1. The transition $|5^2\mathrm{S}_{1/2}, F = 2\rangle \rightarrow |5^2\mathrm{P}_{3/2}, F' = 3\rangle$ is used almost exclusively in this work for driving a cycling transition for laser cooling, and for imaging the atomic distribution, and forms an almost closed two-level system.

2.2 Interaction with Static Magnetic Fields

External magnetic fields are used almost continuously throughout the experimental cycle to manipulate cold gases, and so a brief description of the physical effect of a field on the atomic states is given here.

An external magnetic field $\boldsymbol{B}$ will couple in general to all of the electron spin and orbital angular momenta, as well as the nuclear magnetic moments, adding an additional interaction term to the total Hamiltonian given by [120]

$$\begin{aligned} \hat{H}_{\text{ext}} &= -\hat{\boldsymbol{\mu}} \cdot \boldsymbol{B} && (2.6)\\ &= -\left(\hat{\boldsymbol{\mu}}_L + \hat{\boldsymbol{\mu}}_S + \hat{\boldsymbol{\mu}}_I\right) \cdot \boldsymbol{B} && (2.7)\\ &= \frac{\mu_B}{\hbar}(g_L \hat{L}_z + g_S \hat{S}_z + g_I \hat{I}_z) B_z, && (2.8) \end{aligned}$$

where the external magnetic field has been taken to be along the z-direction, such that $\boldsymbol{B} = B_z \hat{\boldsymbol{e}}_z$. The factors g_L, g_S and g_I are the g-factors associated with each of the magnetic dipole moments. For the ^{87}Rb D-line, it is sufficient to take $g_S \approx 2$, $g_L \approx 1$, and to neglect the contribution of g_I since it is approximately three orders of magnitude smaller due the large nucelar mass. If the energy shift associated with the applied external field is small compared with the fine structure splitting, then J remains a good quantum number. In this case, the interaction Hamiltonian can be written along with the hyperfine contribution as

$$\hat{H}_{\text{int}} = \hat{H}_{\text{hf}} + \hat{H}_{\text{ext}} = \frac{A_{\text{hf}}}{\hbar^2} \hat{\boldsymbol{I}} \cdot \hat{\boldsymbol{J}} + \frac{\mu_B}{\hbar} g_J \hat{J}_z B_z, \tag{2.9}$$

where

$$g_J = 1 + \frac{J(J+1) + S(S+1) - L(L+1)}{2J(J+1)} \tag{2.10}$$

is the Landé g-factor for the electron [121].

If the energy associated with the magnetic field is small compared with the hyperfine splitting, then $\hat{\boldsymbol{F}}$ precesses around the external field and so F and m_F remain good

quantum numbers. This allows the term $\hat{H}_{\text{ext}}$ to be treated as a perturbation to the zero-field eigenstates of $\hat{H}_{\text{hf}}$, given by $|F, m_F\rangle$, and the energy shift with respect to the field-free case is easily evaluated to be

$$\begin{aligned}\Delta E_{\text{ext}} &= \langle F, m_F| \hat{H}_{\text{ext}} |F, m_F\rangle = \langle F, m_F| \frac{\mu_B}{\hbar} g_F \hat{F}_z B_z |F, m_F\rangle \\ &= g_F \, m_F \, \mu_B \, B_z \end{aligned} \tag{2.11}$$

where [121]

$$g_F = g_J \, \frac{F(F+1) + J(J+1) - I(I+1)}{2F(F+1)}. \tag{2.12}$$

The external applied magnetic field lifts the final layer of degeneracy between the $(2F+1)$ magnetic sublevels within a given F manifold, and for weak magnetic fields the energy of the state is proportional to the quantum number m_F. This linear splitting of the energy levels is called the *Zeeman Effect*. Note that even though we neglected the effect of the nuclear magnetic moment with the external field, the energy splitting still depends on I because the nuclear spin itself is not negligible.

For large magnetic fields which dominate over the hyperfine interaction (but still smaller than the fine structure), $\hat{\boldsymbol{I}}$ and $\hat{\boldsymbol{J}}$ individually couple strongly to $\boldsymbol{B}$, and so m_I and m_J are the appropriate good quantum numbers, rather than m_F. In this case, it is the hyperfine Hamiltonian which can be treated as a perturbation to the eigenstates $|J, I, m_J, m_I\rangle$ of that due to the external field [123]. This is called the *Paschen-Back regime*. For intermediate field strengths, the Hamiltonian Eq. 2.9 should be numerically diagonalised, and the eigenstates will be in general superpositions of the $|F, m_F\rangle$ and $|J, I, m_J, m_I\rangle$ basis states. For the case where $L = 0$ and $J = 1/2$, as for the ^{87}Rb ground state, the energy levels have been solved analytically and are given by the Breit-Rabi formula [124]. For this state, with $I = 3/2$ and $g_I \approx 0$, $g_J \approx 2$, the formula reduces to

$$E_{|I,m_J,m_I\rangle} = \frac{3}{4} A_{\text{hf}} \pm \mu_B B_z \qquad \text{for the states } m = \pm 2 \tag{2.13}$$

$$E_{|I,m_J,m_I\rangle} = A_{\text{hf}} \left(-\frac{1}{4} \pm \sqrt{1 + mx + x^2} \right) \qquad \text{for the states } m = \pm 1, 0 \tag{2.14}$$

with $A_{\text{hf}} = h \cdot 3.417\,\text{GHz}$ and $m = m_I + m_J$. where $x = \mu_B B_z / A_{\text{hf}}$ describes the ratio of the strength of the interaction with the magnetic field to the hyperfine interaction energy. The ^{87}Rb ground state splitting is shown in Fig. 2.2. For magnetic fields below $\sim 300\,\text{G}$, the energy splitting is approximately linear, corresponding to m_F being a good quantum number. This is valid for all values of external fields used in this work, and therefore the relevant description is the low-field Zeeman energy given by Eq. 2.11.

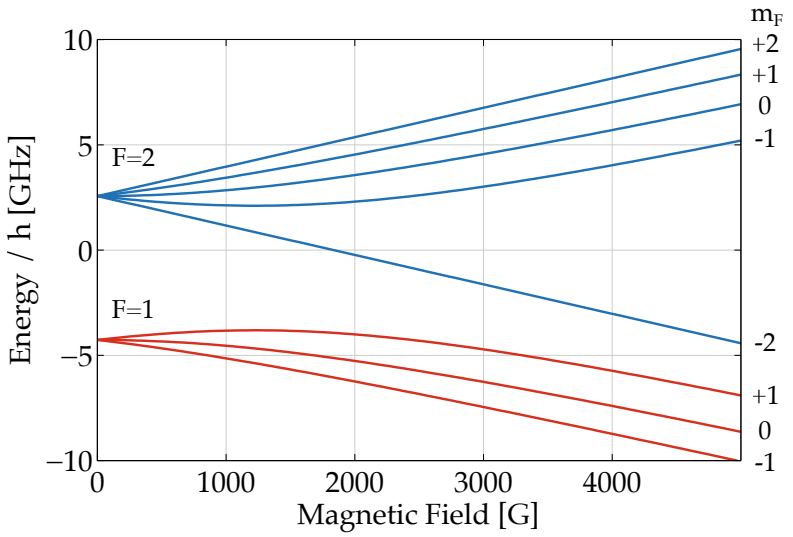

FIGURE 2.2: Breit-Rabi diagram for the ^{87}Rb ground state. The states are labelled according to their weak-field eigenvalues m_F. For fields $\lesssim 300$ G the splitting is approximately linear, given by the Zeeman effect.

2.3 Interaction with Laser Light

As early as 1619 Kepler suggested that the pressure due to sunlight may cause the tails of comets to be deflected away from the sun. Maxwell also showed that a pressure should result when light is absorbed or reflected, as a consequence of his electromagnetic wave equations [125]. Several experiments at the turn of the 20$^{\text{th}}$ century managed to demonstrate the effects of radiation pressure [126, 127], but thermal forces would obscure many of the desirable effects. In 1933, Frisch observed that light truly can have a mechanical effect on matter by deflecting of a beam of sodium atoms using *resonant* light [128]. There was a resurgence of interest in the field in the 1970s with the widespread availability of narrow-band tunable lasers. Ashkin succeeded in trapping micron-sized particles using a laser beam [129], and in addition suggested using the technique to manipulate individual atoms [130]. Much work followed, leading to the development of laser cooling and optical dipole trapping, and to an explosion of the field of ultracold atomic gases. Today, lasers are employed in many experiments around the world for manipulating and studying these atomic systems, and in particular for our experiment they are used for cooling and imaging purposes.

The interaction process can be described well semi-classically, by considering an atom with two quantised energy levels coupled to a classical electromagnetic field [131]. The external electric field causes a displacement of the valence electron, from which arises an induced electric dipole moment that then interacts with the inducing field. Real atoms have many more than two energy levels, but since adjacent levels are much further apart than the transition linewidths the two-level approximation works very well when working near resonance, and in addition narrow laser linewidths allow addressing of individual transitions. The ground and excited states are denoted by $|g\rangle$

and $|e\rangle$, respectively, separated by an energy difference of $\hbar\omega_0$, and the application of monochromatic laser light with angular frequency ω_L introduces the oscillating plane wave electric field $\boldsymbol{E}(t) = \hat{\boldsymbol{\epsilon}} E_0 \cos(\omega_\mathrm{L} t)$, where $\hat{\boldsymbol{\epsilon}}$ is the polarization unit vector and E_0 is the amplitude of the electric field. The spatial dependence has been neglected because the wavelength of the radiation is much larger than the extent of the atom, and so the electric field does not vary significantly on this scale - this is known as the *electric dipole approximation*. The induced dipole $\boldsymbol{d}$ couples to the electric field through an interaction term in the Hamiltonian of the form $H_\mathrm{int} = -\boldsymbol{d} \cdot \boldsymbol{E}$, and the time-dependent Schrödinger equation can be used to calculate the probabilities of finding the atom in each of the two states. This consists of rewriting the Hamiltonian in a frame rotating at the frequency of the applied radiation, and also making the rotating wave approximation (neglecting all fast oscillating terms since they average to zero on the relevant timescales, which is valid when operating near-resonance) [131]. The solution displays Rabi oscillations between the two states as energy is transferred coherently between the atom and external field due to absorption and stimulated emission.

In the above description, an atom prepared in the excited state will remain there indefinitely if the driving field is suddenly removed. Of course, in reality the atom will decay to the ground state via spontaneous emission after some time, with energy being lost to the external environment. The most convenient way to account for spontaneous emission, and the resulting departure from pure quantum states to (incoherent) mixtures, is to move from the wavefunction picture into the formalism of the density matrix. For the two level system, the density matrix is given by [131]

$$\rho = \begin{pmatrix} \rho_{gg} & \rho_{ge} \\ \rho_{eg} & \rho_{ee} \end{pmatrix}, \tag{2.15}$$

where the diagonal elements give directly the probabilities for the atom to be found in a particular state, and the off-diagonal terms are responsible for the coherences. The time-evolution of the density matrix is given by the Liouville-von Neumann equation, to which can be added "by hand" a term describing the decay from the excited state at rate of $\Gamma = 1/\tau$, where τ is the lifetime of the excited state, which for the ^{87}Rb D_2 transition is $\tau = 26\,\mathrm{ns}$ ($\Gamma = 2\pi \times 6.1\,\mathrm{MHz}$). This leads to the set of coupled differential equations know as the *optical Bloch equations*. Taking the steady-state solutions of these equations ($t \to \infty$) leads to the probability of finding the atom in the excited state

$$\rho_\mathrm{ee}(t \to \infty) = \frac{\Omega^2/\Gamma^2}{1 + (2\delta/\Gamma)^2 + 2\Omega/\Gamma^2} = \left(\frac{1}{2}\right) \frac{I/I_\mathrm{sat}}{1 + (2\delta/\Gamma)^2 + I/I_\mathrm{sat}}, \tag{2.16}$$

where $\Omega = E_0/\hbar \langle e| \boldsymbol{d}\cdot\hat{\boldsymbol{\epsilon}} |g\rangle$ is the Rabi frequency describing the strength of the coupling, $\delta = (\omega_\mathrm{L} - \omega_0)$ is the detuning of the laser from the atomic resonance, and the saturation intensity I_sat is defined through the relation $I/I_\mathrm{sat} = 2(\Omega/\Gamma)^2$ [132].

In contrast to the fully coherent case, the inclusion of spontaneous emission causes the Rabi oscillations to be damped out on timescales $t \gg \tau$, after which the system

settles into a steady state population as a result of the balance between absorption and emission. For asymptotically large intensities, the population is distributed equally between the two levels with $\rho_{ee} = 1/2$, and is known as *saturation*.

The rate of decay of the population from the excited state is given by $\Gamma\rho_{ee}$, and since in equilibrium this decay rate equals the rate of absorption, we can write the *scattering rate* from the incident laser beam as

$$R_{sc} = \left(\frac{\Gamma}{2}\right) \frac{I/I_{sat}}{1 + (2\delta/\Gamma)^2 + I/I_{sat}}, \tag{2.17}$$

which is a key quantity experimentally in the context of atom-light interactions. The scattering rate has a Lorenzian profile as a function of detuning (characteristic of a resonance behaviour), and a linewidth given by $\Gamma' = \Gamma\sqrt{1 + I/I_{sat}}$ which exhibits power broadening for increasing intensity.

2.3.1 Doppler Cooling

The use of laser light to cool atoms is based on the idea of the radiation pressure force, which is a result of the conservation of momentum during the emission and absorption of photons. When an atom absorbs a photon of frequency ν, the photon energy $h\nu$ goes into the internal energy of the atom as it is transferred to the excited state. The atom will also receive a linear momentum kick of $\hbar\boldsymbol{k}$ in the direction of propagation of the incoming photon, with an additional momentum kick occuring as the atom later decays to the ground state via spontaneous emission. Since the direction of the emitted photons from spontaneous emission is random, the resulting momentum transfer is isotropic and will average to zero over many absorption-emission cycles. The contribution from the absorption process does *not* average to zero, however, and results in a net change in momentum - and hence velocity - in the propagation direction of the incoming photons. This leads to the *scattering force* $\boldsymbol{F}_{sc} = \hbar\boldsymbol{k}R_{sc}$, given by the product of a single photon momentum and the rate at which photons are absorbed, R_{sc} (Eq. 2.17). Since the maximum scattering rate is $\Gamma/2$, an atom of mass m can experience an impressive maximum acceleration of $a_{max} = \hbar|\boldsymbol{k}|\Gamma/2m \approx 1.1 \times 10^5\ \mathrm{m\,s^{-2}}$ ($\sim 10^4\ g$), for ^{87}Rb atoms on the D_2 line.

In 1975, Hänsch and Schawlow described how atoms could be *cooled* using laser light by exploiting the Doppler effect to make the scattering force *velocity-dependent*, thereby allowing the average translational kinetic energy to be reduced [133]. The idea can be understood simply by considering two counter-propagating laser beams in one dimension, with a frequency red-detuned from the atomic resonance. Atoms moving towards a given beam will be Doppler-shifted closer into resonance, and will scatter preferentially from this beam. The laser beam which is co-propagating with the atomic velocity is shifted even further from resonance, and absorption is reduced. Overall, atoms absorb more photons from whichever beam opposes their motion, leading to a reduction in velocity, known as *Doppler cooling*. The spontaneously emitted photons

are blue-shifted with respect to those which are absorbed, and this is essentially the dissipative mechanism by which the atoms' overall energy is reduced.

The Doppler-shifted frequency seen by an atom moving with velocity v is accounted for by replacing $\delta \rightarrow \delta - \boldsymbol{k} \cdot \boldsymbol{v}$ in Eq. 2.17, and so the total force experienced by the atom is given by the sum of the scattering forces from each beam,

$$\begin{aligned} \boldsymbol{F}_{\text{tot}}(\delta, v) &= \boldsymbol{F}_{\text{sc}}^{(+)} + \boldsymbol{F}_{\text{sc}}^{(-)} \\ &= \frac{\hbar k \Gamma}{2} \left[\frac{I/I_{\text{sat}}}{1 + 4(\delta - kv)^2/\Gamma^2 + I/I_{\text{sat}}} - \frac{I/I_{\text{sat}}}{1 + 4(\delta + kv)^2/\Gamma^2 + I/I_{\text{sat}}} \right] \\ &\simeq \frac{8\hbar k^2}{\Gamma} \frac{I/I_{\text{sat}}}{(1 + 4\delta^2/\Gamma^2 + I/I_{\text{sat}})^2} \delta v, \end{aligned} \tag{2.18}$$

where the final term arises by Taylor expanding in the limit of small velocites. The resulting force is proportional to the velocity, and if the detuning is made negative then the effect is a viscous damping force. By using three pairs of counter-propagating laser beams, the technique is easily extended to obtain cooling in three dimensions, known as *optical molasses*. Laser cooling was initially demonstrated for trapped ions [134, 135], and was then later used to cool an atomic beam [2] and finally to confine neutral atoms in three-dimensions [3]. The developments of laser cooling techniques were recognised with award of the 1997 Nobel Prize in Physics [5–7].

2.3.2 Magneto-Optical Trap

Although optical molasses is a powerful way to cool atoms, eventually they will diffuse out of the region in which the laser beams overlap and become lost. To truly trap atoms in a region of space a *position-dependent* force is needed, in contrast to cooling alone which requries a velocity-dependent force. The magneto-optical trap is the solution to this, first suggested by Dalibard [5], and then demonstrated experimentally using sodium atoms by Raab *et al.* in 1987 [4].

The magneto-optical trap (MOT) is formed by combining the three dimensional arrangement of lasers for optical molasses with a spherical quadrupole magnetic field. The Zeeman shift of the magnetic sublevels, described in Sec. 2.2, together with the correct choice of polarisation of the light beams leads to an imbalance of the forces from opposing beams as a result of optical selection rules between different transitions. The arrangement essentially makes the laser detuning a function of position, with the position-dependent scattering force now providing full three dimensional trapping.

The principle of the MOT is depicted in Fig. 2.3, showing a simplified atom with a J=0 ground state and a J'=1 excited state in one dimension. In the presence of a magnetic field, the degeneracy of the three magnetic sublevels in the excited state manifold is lifted. Since the quadrupole field varies linearly with position, there is a linear Zeeman shift of the energies of the magnetic states. In the schematic, the quantisation axis is defined by the direction of the local magnetic field.

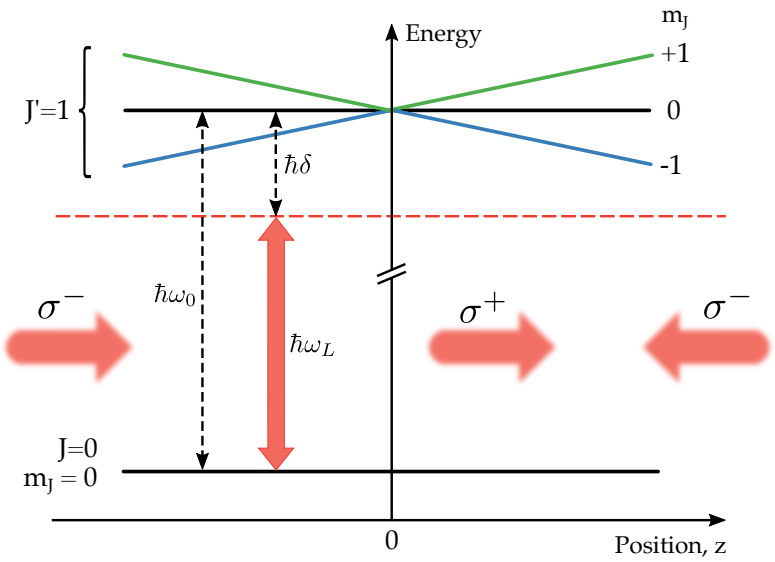

FIGURE 2.3: Principle of operation of the magneto-optical trap in one dimension: A combination of correctly polarised red-detuned laser beams together with the Zeeman shift of magnetic sublevels arising from an externally-applied quadrupole field leads to simultaneous trapping and cooling, as described in the text.

The atom is then illuminated with red-detuned counter-propagating laser beams along the eigenaxes of the quadrupole, and so the direction of propagation of the beams are aligned with the direction of the magnetic field. The beams are chosen to be circularly polarised such that the light can only drive $\Delta m_J = \pm 1$ transitions - and so can be labelled as $\sigma^{\pm}$. Note that the circular polarisation does *not* depend on the magnetic field at all, rather it refers to the rotation of the electric field vector with respect to the wavevector of the beam - known as handedness. Whether this polarisation can then drive σ^+, σ^- or π transitions depends on the rotation of the light's electric field vector with respect to the quantisation axis of the *atom*, i.e. the local magnetic field. The handedness of the circular polarisation is chosen such that both incoming beams can only drive σ^- transitions, and since the field direction flips around at $z = 0$, these beams become σ^+ after passing through the centre, as shown in Fig. 2.3 (even though the handedness remains the same). Thus, an atom will experience both σ^+ and σ^- light at all times, but since the $m_J = -1$ and $m_J = +1$ states are shifted closer and farther from resonance with the red-detuned beams, respectively, there is always preferential scattering from the σ^- beams which provides a net force tending to push the atoms back to the centre of zero field. The magnetic field is not there to provide trapping but rather to break the symmetry, providing an imbalance of absorption from each beam, and adding a position-dependence to the scattering force from the lasers.

The total effective force on an atom in a MOT in one dimension can be expressed in the same way as Eq. 2.18, by replacing the detuning to include the contribution from the Zeeman shift $\delta_{\pm} \rightarrow \delta \mp kv \mp \mu B' z/\hbar$, where μ is the magnetic moment for the excited state and B' is the magnetic field gradient of the quadrupole. The total force is given

by the sum of that from each beam [131]

$$\begin{aligned}\boldsymbol{F}_{\text{tot}}(\delta, v, z) &= \boldsymbol{F}_{\text{sc}}^{(+)} + \boldsymbol{F}_{\text{sc}}^{(-)} \\ &= \frac{\hbar k\Gamma}{2}\left[\frac{I/I_{\text{sat}}}{1+4\delta_+^2/\Gamma^2 + I/I_{\text{sat}}} - \frac{I/I_{\text{sat}}}{1+4\delta_-^2/\Gamma^2 + I/I_{\text{sat}}}\right] \\ &\simeq \frac{8\hbar k}{\Gamma}\frac{I/I_{\text{sat}}}{(1+4\delta^2/\Gamma^2 + I/I_{\text{sat}})^2}\delta\left(kv + \mu B' z/\hbar\right), \end{aligned} \tag{2.19}$$

where the last term applies the Taylor expansion in the limit that the Doppler and Zeeman shifts are small compared with the bare detuning. With red-detuning $\delta < 0$, the total force can be expressed in the form $\boldsymbol{F}_{\text{tot}}(\delta, v, z) = -\alpha v - \beta z$, which illustrates both the velocity and position dependencies, providing cooling and trapping, respectively. The form of the force shows that the motion of an atom displaced from the centre of the MOT is that of a damped simple harmonic oscillator. In fact, for typical detunings, intensities and field gradients, the motion is strongly over-damped, and so a displaced atom is pushed quickly back to the centre without oscillating [121].

The MOT is a very robust trap, not requiring particularly clean polarisations or precise balancing of intensities, and operating best with modest field gradients on the order of ~ 10 - $20\,\text{G}\,\text{cm}^{-1}$ along the strong quadrupole axis, that are easy to generate in the lab with coils or planar structures. It was soon shown [136] that there are enough atoms below the trap's capture velocity in the tail of the Maxwell-Boltzmann distribution to achieve a significant loading rate from a room-temperature vapor, as opposed to using a slowed beam. In addition, the ability to cool large numbers of atoms to temperatures of several hundred microkelvin in a fraction of a second has made the MOT the first step for almost all ultracold atom experiments.

2.4 Magnetic Trapping

Since the energy of a particular magnetic sublevel is a function of the external magnetic field, an atom subjected to an *inhomogeneous* field will experience a spatially-dependent potential energy. This is the fundamental mechanism underlying magnetic trapping of neutral atoms.

When placed in a magnetic field, the magnetic moment of an atom will precess around the field direction, with the frequency given by the Larmor frequency, $\omega_L = \mu_B g_F m_F B_z/\hbar$ [123]. If the motion of the atom is slow compared with ω_L, then the atom sees a slowly varying field and the atomic spin precesses around the *local* direction of the magnetic field. In this case, the magnetic moment can adiabatically follow the direction of the local magnetic field, and the projection of the total spin onto the quantisation axis is a constant of the motion, corresponding to the atom remaining in the same m_F state. With the quantisation axis being defined by the local direction of $\boldsymbol{B}$, the energy shift then depends only on the modulus of the field, $|\boldsymbol{B}|$, and so the energy shift can be written as $\Delta E_{\text{ext}} = g_F\, m_F\, \mu_B\, |\boldsymbol{B}|$ from Eq. 2.11. In this way, the potential

energy seen by the atom is determined proportionally by the design of the magnetic field landscape (with the addition of a small effect due to gravity).

It can be seen that states with $g_F m_F > 0$ will have decreased energy with decreasing magnetic fields, and so are called *low-field seeking* states. Conversely, for states with $g_F m_F < 0$ it is more energetically favourable for the atom to reside in regions of high field, and these are termed *high-field seeking*. A maximum of static magnetic field modulus is forbidden in a source-free region of space as a consequence of Maxwell's equations - a result commonly known as *Wing's theroem* [137] (which is a variant of the Earnshaw theorem for electric fields [138]). However, a magnetic minimum is possible and so low-field seeking states can be magnetically trapped. In the ^{87}Rb ground state, the g-factors are $g_F = -\frac{1}{2}$ and $g_F = +\frac{1}{2}$ for the $F = 1$ and $F = 2$ hyperfine states, respectively, and so it is possible to trap the states $|1, -1\rangle, |2, +1\rangle$ and $|2, +2\rangle$.

Adiabaticity

An atom will be lost from the trap if it makes a transition either to an anti-trapped state ($g_F m_F < 0$), or to an un-trapped state ($m_F = 0$). To remain trapped, the projection of its spin onto the quantisation axis (given by the local field direction) must remain constant, keeping it in the same m_F state. The atom can adiabatically follow the field direction provided that the variation of the magnitude and direction of the field is slow when compared with the Larmor frequency. An upper bound on the field variation is given by the trap frequency ω_T, leading to the adiabatic condition [54]

$$\omega_T \ll \omega_L = \frac{g_F m_F \mu_B}{\hbar} |\boldsymbol{B}|. \tag{2.20}$$

This condition is violated in regions of very small magnetic field, where the quantisation axis for the atom is not well-defined, leading to a projection into other Zeeman levels and to a loss rate from the trap known as *Majorana transitions* [139]. The loss rate depends on the temperature of the cloud, because colder atoms spend more time around the region of zero-field [140]. For this reason, magnetic traps are designed such that there is a residual static field at the trap minimum, which "plugs the hole" at the bottom, ensuring a well-defined quantisation axis is always provided. An offset field of $\sim 1\,\mathrm{G}$ is typically used for this bottom field, which is more than sufficient to suppress this type of spin-flip losses [118]. A measurement of spin-flips with a BEC in our magnetic trap due to a region of small field is presented in Sec. 5.5.

2.5 Trapping Geometries

In order to construct a magnetic trap for neutral atoms, the task is to design a field configuration that contains a local minimum in the modulus $|\boldsymbol{B}|$, towards which atoms in low field-seeking states will be attracted. In general, there are two widely used

classes of traps - those which have a zero field at the minimum, and those with a finite offset field.

2.5.1 Quadrupole Trap

The simplest configuration exhibiting a local minimum for trapping is the 3D quadrupole. The field around the minimum is of the form $\boldsymbol{B}(x, y, z) = B'_x\, x\, \hat{\boldsymbol{e}}_x + B'_y\, y\, \hat{\boldsymbol{e}}_y + B'_z\, z\, \hat{\boldsymbol{e}}_z$, where B'_i characterises the gradient in the i^{th} direction, and Maxwell's equations require $B'_x + B'_y + B'_z = 0$ due to $\nabla \cdot \boldsymbol{B} = 0$ [141]. Since the potential energy, $U \propto |\boldsymbol{B}| = \sqrt{(B'_x x)^2 + (B'_y y)^2 + (B'_z z)^2}$, the potential varies *linearly* with distance from the minimum, and this type of linear trap provides the strongest possible confinement for the atoms. A quadrupole field is easily produced using a pair of coils in the anti-Helmholtz configuration (currents flowing in opposite senses), as shown in Fig. 2.4 (a), which leads to cylindrical symmetry with $B'_x = B'_y = -\frac{1}{2}B'_z$. Indeed, this arrangement was used in the first demonstration of magnetic trapping of neutral sodium atoms [142].

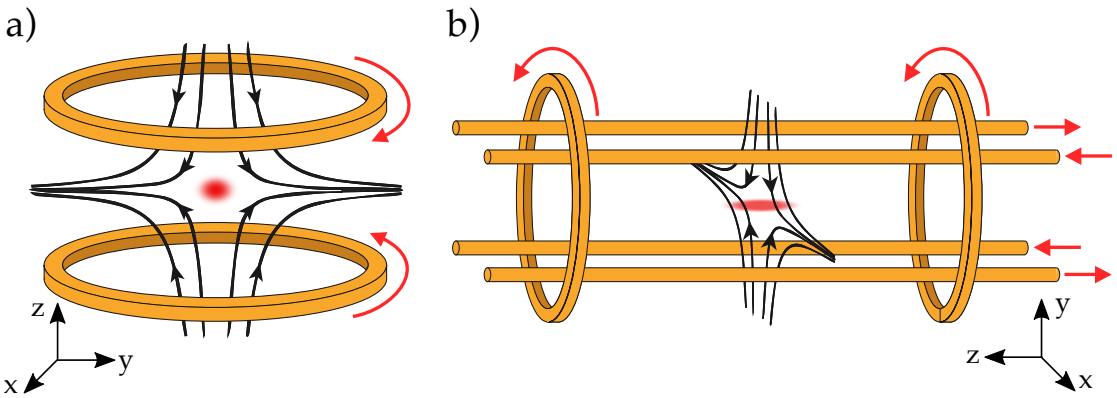

FIGURE 2.4: Magnetic Trapping Geometries: a) A quadrupole trap produced by anti-Helmholtz coils, with currents propagating in opposite senses (shown by the red arrows), provides a field that is rotationally symmetric around the z-axis, having a zero at the centre, and with the strength of the field gradient along the z-axis being twice that in the x-y plane due to Maxwell's equations. b) The original incarnation of the Ioffe-Pritchard trap, exhibiting three-dimensional confinement with a non-zero field at the minimum. Four so-called *Ioffe bars* with electrical currents flowing as indicated produce a two-dimensional quadrupole field in the x-y plane, which is translationally-invariant along the z-axis for distances close to the centre, providing radial confinement. The pair of coils with currents flowing in the same sense have a separation greater than their radius, which caps the ends of the potential by providing weak harmonic axial confinement along the z-axis, and leads to a cigar-shaped trapping geometry.

The major drawback with the quadrupole trap is the field zero at the minimum, leading to spin-flip losses due to a lack of sufficient quantisation axis, as discussed in Sec. 2.4. This effect becomes more significant for colder atoms, with the loss rate scaling as T^{-2} [143], since colder atoms will sample the zero-field region of the trap more frequently. Various methods have been devised to "plug the hole" in the bottom of the

trap. For example, the Time-Orbiting Potential (TOP) trap uses a rotating bias field to shift the effective location of the zero faster than the atoms can respond [144]. Alternatively, a blue-detuned laser beam can be used to repel atoms from the centre away from the field zero [12]. The most common way to avoid spin-flip losses is to design a field configuration which has a finite value at the minimum, instead of a zero. This is the method used throughout this thesis, and will be discussed in the next section.

The quadrupole trap is useful for handling relatively hot atoms, when the spin-flip loss is negligible. For example, it can be used as an intermediate step, for transporting or transferring atoms into another trap. It is also widely used for magneto-optical trapping, where it is required that the direction of the field reverses either side of the zero. This is in contrast to purely magentic traps, where only the modulus of the field is important.

2.5.2 Ioffe-Pritchard Trap

Spin-flip losses in static magnetic traps can be avoided by using a field with a non-zero value at the minimum. It is not enough to just add an additional offset field to a 3D quadrupole, as this will simply shift the position of the zero point. The most popular design is the Ioffe-Pritchard trap, which consists of a 2D quadrupole in the radial plane superposed with a magnetic bottle field to plug the axial confinement [145], and the original geometry used in experiments is shown schematically in Fig. 2.4 (b). The form of the field is axially symmetric, and is given by [55]

$$\boldsymbol{B}(x,y,z) = \underbrace{B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2+y^2) \end{pmatrix}}_{\text{"Magnetic Bottle"}} + \underbrace{B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix}}_{\text{2D quadrupole}}, \tag{2.21}$$

where B'' describes the curvature of the bottle field and B_0 is the offset field remaining at the minimum. The modulus of the field is then given by

$$\begin{aligned} |\boldsymbol{B}|(x,y,z) &= \sqrt{\left(B'x - \frac{B''}{2}xz\right)^2 + \left(B'y + \frac{B''}{2}yz\right)^2 + \left(B_0 + \frac{B''}{2}\left[z^2 - \frac{1}{2}(x^2+y^2)\right]\right)^2} \\ &\simeq B_0 + \frac{1}{2}\left[\frac{B'^2}{B_0} - \frac{B''}{2}\right]r^2 + \frac{B''}{2}z^2, \end{aligned} \tag{2.22}$$

where r is the radial coordinate, and the second expression describes the field form only close to minimum, arising after retaining only terms which are quadratic and below in the spatial dimensions. The potential energy is proportional to the magnetic field, and so this magnetic trap provides harmonic confinement

$$U(r,z) = g_F\, m_F\, \mu_B\, |\boldsymbol{B}| = U_0 + \frac{1}{2}m\Big(\omega_\perp^2 r^2 + \omega_{/\!/}^2\, z^2\Big), \tag{2.23}$$

where $U_0 = g_F m_F \mu_B B_0$, m is the mass of an atom, and the axial and radial trap frequencies are given respectively by

$$\omega_{/\!/} = \sqrt{\frac{g_F\, m_F\, \mu_B}{m} B''} \qquad \omega_\perp = \sqrt{\frac{g_F\, m_F\, \mu_B}{m}\left[\frac{B'^2}{B_0} - \frac{B''}{2}\right]}. \tag{2.24}$$

Consequently, the radial trap frequency (and therefore the aspect ratio $\omega_\perp/\omega_{/\!/}$) can be tuned by adjusting B_0, for example with an additional external "anti-Ioffe" field. Reducing the value of the trap bottom has the effect of increasing the confinement, but care must be taken to ensure always sufficient field for preventing spin-flip losses. The harmonic approximation arising from the expansion in Eq. 2.22 is valid only for small distances from the trap centre. This description therefore applies to cold clouds with a small spatial extension, whilst for hotter clouds the potential looks essentially linear. Both regimes are explored in our experiment - depending on the exact trap, the potential is typically harmonic for clouds below the order of $\sim 10\,\mu\mathrm{K}$ and for all condensates, whilst it is linear at the initial stages when we deal with temperatures of $\sim 300\,\mu\mathrm{K}$.

2.5.3 Single Wire Guide

Initially, cold atom and BEC experiments used large macroscopic coil structures located outside the vacuum system, requiring tens or hundreds of amps of current due to the large distance between the current source and the trapping region. In 1995, it was pointed out that microstructured traps on substrates would be suited to producing much tighter traps, because the magnetic field and its gradient scale as I/s and I/s^2, respectively, where s is the characteristic size of the system [63]. Much work then followed, trapping neutral atoms on free-standing wire structures [64, 65], and then eventually with microstructures on substrates, or *atom chips* [61, 62]. Achievement of BECs on atom chips soon followed in 2001 [67, 68], and today many groups around the world employ atom chips to create degenerate quantum gases. Tighter magnetic traps allow evaporative cooling to be performed faster, which relaxes vacuum requirements for an experiment. Additionally using atom chips, higher trap gradients than those produced by macroscopic structures can be achieved for a much lower power consumption.

Many different types of trap are used on atom chips, varying mainly in the layout of the conductors which carry the current. The simplest arrangement that can be used to produce a minimum in field is the single wire guide, shown in Fig. 2.5, which consists of a long, thin, current-carrying conductor superposed with an external homogeneous bias field.

The magnetic field, its gradient, and curvature produced by the infinitely thin and long wire alone carrying a current I are given respectively by

$$B = \frac{\mu_0 I}{2\pi r}, \qquad B' = -\frac{\mu_0 I}{2\pi r^2}, \qquad \text{and} \qquad B'' = \frac{\mu_0 I}{\pi r^3}, \tag{2.25}$$

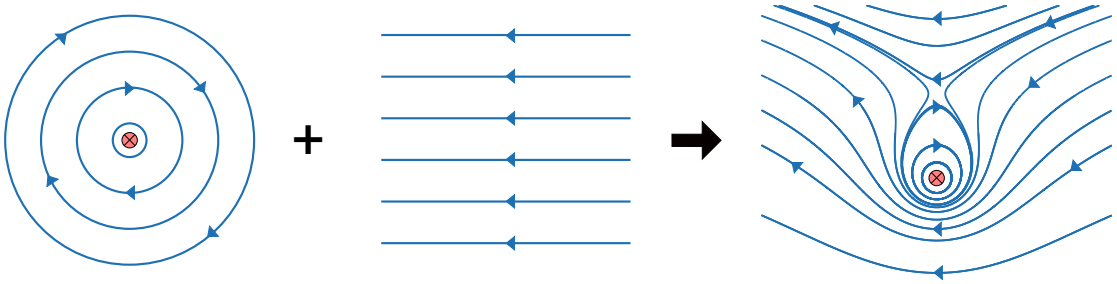

FIGURE 2.5: Single wire guide: The magnetic field produced by a thin wire (left) combined with a cancelling external homogeneous field (centre) gives rise to a minimum at a distance above the wire given by the ratio of the wire current and the bias field.

where r is the distance from the wire. An external homogeneous bias field B_{bias} can be used to cancel the field of the wire at some position, producing a line of zero field at a height above the wire given by

$$r_0 = \frac{\mu_0}{2\pi}\frac{I}{B_{\text{bias}}}. \tag{2.26}$$

In the vicinity of the zero point the field configuration is well-described by a two-dimensional quadrupole with gradient B', although it is clear from Fig. 2.5 that the field deviates with increasing distance from the minimum.

2.5.4 Z-Wire Trap

The single wire guide described above does not constitute an atom trap, since there is no confinement along the length of the wire. To provide full three-dimensional trapping, there must be some spatial depenence of the field along the axial direction. One of the most popular ways to achieve this is to bend the ends of the single wire into a Z-shape [61], leading to the so-called *Z-wire trap*, shown in Fig. 2.6. The central part of the Z-wire behaves in the same way as the single wire guide: the transverse field component being cancelled by an external field in the x-direction, B_{bias}, resulting in a 2D quadrupole in the x-z plane. Now however, the co-propagating currents in each of the end-cap wires leads to a maximum of the y-component of the field above each one with a minimum inbetween, giving rise to axial confinement. The overall minimum of field is geometrically centrered at $x = y = 0$, and $z = z_0$, where z_0 is determined by Eq. 2.26.

The vertical z-component of the fields from the end-cap wires cancel at the centre, whilst the y-components add together, leading to a residual *Ioffe field* at the trap minimum which provides a quatisation axis for the atoms and prevents spin-flip losses. The field at the miniumum can then be adjusted by applying an external *anti-Ioffe field*, B_{AI}, without changing the spatial position of the minimum.

The final trapping potential follows the form of an Ioffe-Pritchard trap for distances close to the minimum, as described in Sec.2.5.2, with full three-dimensional harmonic confinement that is cylindrically-symmetric. The axial frequency is set predominantly

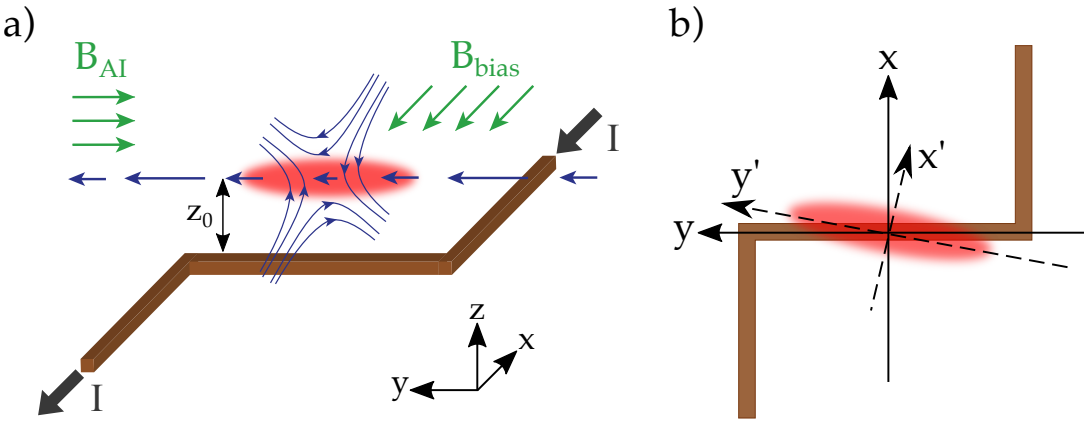

FIGURE 2.6: Z-Wire trap configuration: a) A current I passing through the conductor combined with an external bias field B_{bias} creates a 2D quadrupole in the radial (xz) plane with its minimum located at z_0 above the wire. The leg wires provide harmonic confinement along the axial (y) direction with a non-zero offset at the minimum, the value of which can be tuned using the second external anti-Ioffe field B_{AI}. b) Due to a vertical field component produced by the leg wires that varies along y, the isopotential is rotated in the xy plane around the z-axis, with the new eigenaxes labelled x' and y'.

by the distance between the end-cap wires, and is typically much weaker than the radial frequencies, leading to elongated cigar-shaped traps.

The vertical fields produced by each of the end-cap wires cancel only at the trap minimum, and so the vertical component increases with axial distance from the centre. This leads to a twisting in the isopotential around the z-axis, as shown in Fig.2.6 (b). In general, the new eigenaxes of the trap, $\{x', y', z'\}$, are found by diagonalising the Hessian matrix at the position of the trap minimum $\boldsymbol{r} = \boldsymbol{r}_0$ [146]

$$\boldsymbol{\mathcal{H}} = \begin{pmatrix} \frac{\partial^2|\boldsymbol{B}|}{\partial x^2} & \frac{\partial^2|\boldsymbol{B}|}{\partial x \partial y} & \frac{\partial^2|\boldsymbol{B}|}{\partial x \partial z} \\ \frac{\partial^2|\boldsymbol{B}|}{\partial y \partial x} & \frac{\partial^2|\boldsymbol{B}|}{\partial y^2} & \frac{\partial^2|\boldsymbol{B}|}{\partial y \partial z} \\ \frac{\partial^2|\boldsymbol{B}|}{\partial z \partial x} & \frac{\partial^2|\boldsymbol{B}|}{\partial z \partial y} & \frac{\partial^2|\boldsymbol{B}|}{\partial z^2} \end{pmatrix} \Bigg|_{\boldsymbol{r}=\boldsymbol{r}_0} \xrightarrow{\text{diagonalise}} \begin{pmatrix} \frac{\partial^2|\boldsymbol{B}|}{\partial x'^2} & 0 & 0 \\ 0 & \frac{\partial^2|\boldsymbol{B}|}{\partial y'^2} & 0 \\ 0 & 0 & \frac{\partial^2|\boldsymbol{B}|}{\partial z'^2} \end{pmatrix}. \quad (2.27)$$

Due to symmetry in the Z-wire geometry, there is always an eigenaxis perpendicular to the plane of the wire which is exactly aligned with the canonical laboratory axis, so that $z = z'$. The x'-y' plane is then rotated around the z axis by some angle, typically a few degrees. The amount of rotation increases with distance of the minimum from the wire surface, when the effect of the end-cap wires becomes more significant. For close distances the atoms essentially only see the central wire of the conductor, and the twisting angle will be very small.

The eigenvalues of the Hessian matrix in Eq. 2.27 correspond exactly to the curvatures in each direction at the minimum, which are then directly related to the trap

frequencies along the eigenaxes, x_i', through[1]

$$\frac{\omega_i}{2\pi} = \sqrt{\frac{\mu_B g_F m_F}{m} \frac{\partial^2 |\boldsymbol{B}|}{\partial x_i'^2}}. \tag{2.28}$$

Throughout this thesis, simulations of the trapping potentials are carried out using numerical diagonalisation of the Hessian, and so the quoted trap frequencies always correspond to those along the principal axes of the trap, accounting for the twisting of the potential.

2.6 Evaporative Cooling

The most relevant parameter during the various cooling stages of the experimental cycle is the *phase-space density*, defined by PSD $= n(0)\lambda_{\text{dB}}^3$, where $n(0)$ is the peak atomic density of the cloud and $\lambda_{\text{dB}} = \sqrt{2\pi\hbar^2/mk_BT}$ is the thermal de Broglie wavelength. The phase-space density describes the number of particles contained within a volume equal to the cube of the de Broglie wavelength, providing a quantitative measure of the significance of quantum statistics in the gas, and as shown later in Sec. 2.7 must be increased to on the order of unity for the transition to BEC to occur. After laser cooling the cloud has PSD $\sim 10^{-6}$, and the final increase is obtained using the powerful technique of evaporative cooling. This process was first suggested as a means of achieving BEC in 1986 [147, 148], and was initially observed in magnetically trapped hydrogen [149]. Evaporative cooling of alkali atoms was demonstrated in 1995 [10, 144], and became a key step leading to the first observations of BEC later that same year [11, 12].

The fundamental principle of the cooling process is the selective removal of atoms from the high-energy tail of the energy distribution function. Since these atoms carry away higher than average energy, the total energy per particle in the distribution is now lower. Crucially, with the aid of elastic collisions between particles, the distribution rethermalises to a new lower temperature. The process is illustrated schematically in Fig. 2.7.

If the distribution is truncated at an energy ϵ_t, one of the key variables in the process is the *truncation parameter* $\eta = \epsilon_t/k_BT$, quantifying how deep into the distribution the cut is made. If ϵ_t is held constant, then η will increase as the gas cools. The rate of evaporation will become more and more suppressed as fewer atoms can acquire enough energy to leave the trap, and the process will eventually stagnate. Therefore, to keep the evaporation going, one should maintain η as the temperature drops by simultaneously reducing the trap depth - so-called *forced evaporative cooling*.

[1] This relation can be obtained by equating the magnetic potential energy with the general harmonic form $\mu_B g_F m_F |\boldsymbol{B}| = \frac{1}{2} m\omega_i^2 x_i^2$.

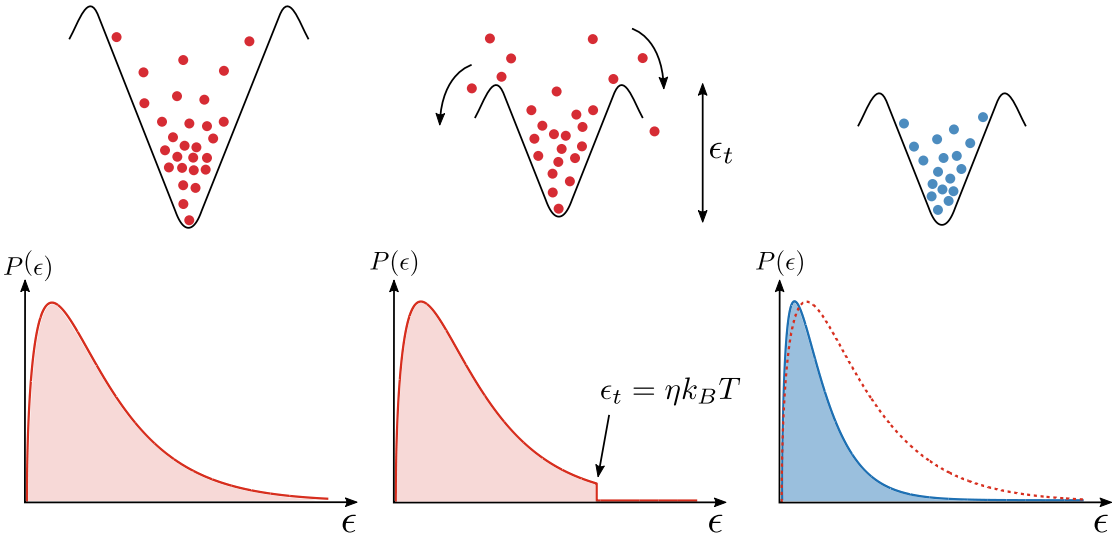

FIGURE 2.7: Principle of evaporative cooling: An initially "deep" trap contains a gas of atoms which have an (almost) thermal distribution (left). Truncation of the distribution at an energy $\epsilon_t = \eta k_B T$ selectively removes particles with higher than the average energy (centre), where the truncation parameter η controls the severity of the cut. The total energy of the system is now reduced, and after many rethermalising collisions reaches a new state which can be described by an altogether lower temperature (right).

It is the task of the experimentalist to choose the trajectory of the truncation energy $\epsilon_t(t)$ such that the most efficient cooling results, i.e with the smallest loss in atom number. As an extreme example, η could be set arbitrarily high, and one could wait for the rare event in which a single atom acquires all the energy of the entire system. This atom would then be evaporated, and the system would be very efficiently cooled to absolute zero at the expense of one atom. In practice of course, all the atoms would have been lost due to collisions with the background gas in the vacuum chamber long before this event occurred, and so due to other loss processes there is a trade-off between the cooling speed and efficiency.

The standard figure of merit for evaporative cooling is the phase-space density efficiency parameter [150]

$$\gamma_{\text{eff}} \equiv -\frac{d\ln(\text{PSD})}{d\ln(N)}, \tag{2.29}$$

which measures the relative fractional change in PSD for a given change in atom number N. Since PSD $\propto n(0)T^{-3/2}$, it is possible to obtain an increase in phase-space density if the increasing density (due to the shrinking volume during cooling) and the reduction in temperature dominate over the necessary loss in N.

2.6.1 A Quantitative Model

Since its inception, there have been a variety of approaches to the modelling of the evaporative cooling process, and several comprehensive reviews are available [150,

151]. One of the first, and simplest, was studied by Ketterle's group at MIT, and considered the evaporation as a series of discrete truncations of the distribution, with each step followed by rethermalisation in an infinitely deep trap [152]. Although this allows derivation of many analytical expressions, it does not reflect the situation in real experiments, where the cooling process is typically performed using a continuous sweep. A more accurate description was developed in the group of Walraven in Amsterdam, which directly accounts for the non-equilibrium dynamics by numerically intergrating the kinetic Boltzmann transport equation for the phase-space distribution function [153]. This is the most thorough model available, and the important results which can be applied to understanding our experimental system will be briefly outlined here.

Several underlying assumptions are required, which are well-approximated in realistic experimental situations:

- The model assumes the system exhibits "sufficient ergodicity", meaning the distribution function depends only on the particles' energies, and so an atom with $\epsilon > \epsilon_t$ will sample an untrapped state with unit probability.
- Only classical statistics are used, i.e. the distribution is Maxwell-Boltzmann, since the effects of quantum statistics only become significant when close to the BEC transition, at the very end of the cooling sweep.
- "Full power evaporation" is assumed, i.e. a particle with energy $\epsilon > \epsilon_t$ will be lost immediately with unit probability, without any more collisions with other particles.

With the assumption of ergodicity, the distribution function is replaced $f(\boldsymbol{r}, \boldsymbol{p}) \rightarrow f(\epsilon)$, which greatly simplifies the Boltzmann equation. Although this equation is solved numerically, the results show that the distribution function is almost indistinguishable from a Boltzmann which has been truncated at the trap depth, i.e. the evaporation preserves the thermal nature of the distribution [153]. Adopting this "truncated Boltzmann approximation", many analytical expressions are obtained in terms of incomplete gamma functions, which arise from the truncation. The distribution function is then simply assumed to be of the form

$$f(\epsilon) = \exp\left(\frac{\mu - \epsilon}{k_B T}\right) \Theta(\epsilon_t - \epsilon), \tag{2.30}$$

where $\Theta(\epsilon_t - \epsilon)$ is the Heaviside function, which sets the occupation probability to zero for $\epsilon > \epsilon_t$, and the chemical potential serves to normalise to the total number of atoms

$$N = \int_0^{\epsilon_t} \exp\left(\frac{\mu - \epsilon}{k_B T}\right) g(\epsilon) d\epsilon, \tag{2.31}$$

where $g(\epsilon)$ is the density of states. Here, T is no longer the true temperature, since a finite depth trap will be empty in equilibrium - there will always be a small evaporation rate out of the potential. Strictly speaking then, T is a "quasi-temperature", which

coincides with the thermodynamic equilibrium temperature only in the limit of a deep trap ($\eta \to \infty$).

The time evolution of the total atom number can be written as a combination of several terms

$$\dot{N} = \dot{N}_{\text{ev}} + \dot{N}_{\text{bg}} = -\frac{N}{\tau_{\text{ev}}} - \frac{N}{\tau_{\text{bg}}}, \tag{2.32}$$

where $\dot{N}_{\text{ev}}$ describes the loss of atoms due to the evaporation process, and $\dot{N}_{\text{bg}}$ is the rate of loss due to collisions with the background gas in the vacuum chamber. Any higher order losses are neglected, since the majority of the cooling takes place when the trapped gas is at sufficiently low density.

It is convenient to introduce the parameter α_{ev}, which describes how the temperature changes as atoms are lost from the system. Assuming that only evaporated atoms change the temperature of the sample (since losses due to to the background gas are indiscriminant, and independent of density) α_{ev} is defined as

$$\alpha_{\text{ev}} \equiv \frac{d\ln(T)}{d\ln(N)} = \frac{\dot{T}/T}{\dot{N}_{\text{ev}}/N}, \tag{2.33}$$

The evaporation rate is actually composed of two contributions. Firstly, there is the loss due to atoms being promoted to energies above the trap depth via elastic collisions - known as *plain evaporation*. Secondly, during forced evaporative cooling there is a loss rate due to the fact that the truncation point is being moved deeper into the distribution over time. These atoms are not lost due to elastic collisions, and this effect is known as *spilling*. We can then write

$$\dot{N}_{\text{ev}} = \dot{N}_{\text{plain}} + \dot{N}_{\text{spill}} = -\frac{N}{\tau_{\text{plain}}} - \frac{N}{\tau_{\text{spill}}}. \tag{2.34}$$

Under the truncated Boltzmann approximation, Luiten *et.al.* [153] calculated the plain evaporation rate to be

$$\dot{N}_{\text{plain}} = -\left[n_0 \overline{v} \sigma \overline{V}(\eta) e^{-\eta}\right] N, \tag{2.35}$$

where n_0 is the density at the trap centre, $\overline{v} = \sqrt{8k_B T/\pi m}$ is the average thermal velocity of an atom, $\sigma = 8\pi a^2$ is the temperature-independent s-wave scattering cross-section, and $\overline{V}$ is an effective volume ratio which arises due to the truncation of the distribution. Since the elastic collision rate is given by [150]

$$\Gamma_{\text{el}} \equiv (\tau_{\text{el}})^{-1} = n_0 \sigma \sqrt{2}\, \overline{v}, \tag{2.36}$$

this model provides a direct way to link the time constant for plain evaporation with that due to elastic collisions through

$$\tau_{\text{plain}} = \frac{\sqrt{2} e^{\eta}}{\overline{V}} \tau_{\text{el}}. \tag{2.37}$$

The Amsterdam treatment [151] also accounts for spilling by introducing the parameter ξ, which makes the connection between the atom loss rate due to spilling and the ramp trajectory of the truncation energy during forced evaporative cooling through

$$\frac{\dot{N}_{\text{spill}}}{N} = \xi \frac{\dot{\epsilon}_t}{\epsilon_t}. \tag{2.38}$$

In the case where η is held constant (a situation which is usually fullfilled in the experiment), the energy distribution is only rescaled in each time step, and the effect of spilling is simply to modify the time constant for plain evaporation according to

$$\tau_{\text{ev}} = (1 - \alpha_{\text{ev}}\xi)\, \tau_{\text{plain}}, \tag{2.39}$$

and so spilling is entirely accounted for by the parameter ξ.

2.6.2 Runaway Regime

The efficiency of the cooling can only be maintained if the time required for the distribution to rethermalise is shorter than the timescale in which atoms are lost due to background collisions, otherwise at some point the background losses would dominate and the efficiency drops to zero. Therefore, an important quantity is $\mathcal{R} = \tau_{\text{bg}}/\tau_{\text{el}}$, often called the "ratio of good to bad collisions". Since the elastic collision rate scales as $\sim n_0 T^{1/2}$ from Eq. 2.36, we can see that in order to maintain the cooling efficiency it is not a stringent enough requirement that the density remain constant - it must instead *increase* at least stronger than $T^{1/2}$, due to the temperature dependence in the thermal velocity.

For power-law traps of the form

$$U(x,y,z) \sim |x|^{1/\delta_1} + |y|^{1/\delta_2} + |z|^{1/\delta_3} \qquad \text{with} \quad \delta = \sum_i \delta_i, \tag{2.40}$$

the volume occupied scales as $V \propto T^{\delta}$ [151], and so the elastic collision rate scales as $\Gamma_{\text{el}} \propto N T^{1/2-\delta}$. This allows an expression to be obtained for the fractional time-dependence of Γ_{el}, given by [150]

$$\frac{1}{\Gamma_{\text{el}}}\frac{d\Gamma_{\text{el}}}{dt} = \frac{1}{\tau_{\text{el}}}\left(\frac{\alpha_{\text{ev}}\,(\delta - 1/2) - 1}{\lambda} - \frac{1}{\mathcal{R}}\right), \tag{2.41}$$

where the parameter λ expresses the ratio of the evaporation time to the elastic collision time

$$\lambda \equiv \frac{\tau_{\text{ev}}}{\tau_{\text{el}}} = (1 - \alpha_{\text{ev}}\xi)\,\frac{\sqrt{2}e^{\eta}}{\overline{V}}. \tag{2.42}$$

In order to maintain or increase the elastic collision rate, we must have the right hand side of the expression Eq. 2.41 positive, meaning that the value of $\mathcal{R}$ must be at least

greater than some minimum value

$$\mathcal{R} \geq \mathcal{R}_{\text{min}}(\eta) = \frac{\lambda}{\alpha_{\text{ev}}\left(\delta - 1/2\right) - 1}. \tag{2.43}$$

When this condition is met, the elastic collision can increase *despite* the loss in atom number, and the evaporation rate is able to grow rapidly, entering the so-called *runaway regime*.

After obtaining a high enough initial ratio of good-to-bad collisions for entering the runaway regime and thus ensuring that the evaporation can be sustained, the overall goal of the process is to increase the phase-space density with minimal atom loss. The PSD scales as $\sim NT^{-(\delta+3/2)}$, which allows the efficiency parameter of Eq. 2.29 to be expressed as [150]

$$\gamma_{\text{eff}}(\eta) = \frac{\alpha_{\text{ev}}\left(\delta + 3/2\right)}{1 + \lambda/\mathcal{R}} - 1. \tag{2.44}$$

By adopting the truncated Bolzmann approximation, Walraven [151] was able to calculate closed expressions for the parameters $\alpha_{\text{ev}}, \overline{V}$ and ξ for the case of forced evaporative cooling at constant η in a power-law trap. These expressions are given in terms of incomplete gamma functions $P\left\{a, \eta\right\}$ [2] by

$$\begin{cases} \alpha_{\text{ev}} = \dfrac{\eta + \kappa + \xi - \delta - 3/2}{3/2 + \delta + \xi\left(\kappa - 1\right)} & (2.45) \\ \overline{V} = \eta - \left(5/2 + \delta\right)\dfrac{P\left\{5/2 + \delta, \eta\right\}}{P\left\{3/2 + \delta, \eta\right\}} & (2.46) \\ \xi = \left(3/2 + \delta\right)\left[1 - \dfrac{P\left\{5/2 + \delta, \eta\right\}}{P\left\{3/2 + \delta, \eta\right\}}\right], & (2.47) \end{cases}$$

where

$$\kappa = 1 - \frac{P\left\{7/2 + \delta, \eta\right\}}{P\left\{3/2 + \delta, \eta\right\}}\overline{V}. \tag{2.48}$$

These expressions can then be used to calculate the experimentally relevant quantities $\mathcal{R}_{\text{min}}$ and γ_{eff} as a function of the truncation parameter, as shown in Fig. 2.8. The function $\mathcal{R}_{\text{min}}$ is shown for three trapping geometries $\delta = \frac{3}{2},\ \delta = 3$ and $\delta = \frac{5}{2}$, corresponding to 3D harmonic, 3D linear, and Ioffe-Pritchard (2D linear + 1D harmonic) traps, respectively. It can be seen that in order to reach the runaway regime, a more favourable ratio of good-to-bad collisions is required when using harmonic traps as opposed to linear ones. For the Ioffe-Pritchard trap used in the experiment, the curve shows that a ratio of at least $\mathcal{R} \approx 140$ is required, and that in addition for truncation

[2] The lower incomplete gamma fucntion is defined as $P\left\{a, \eta\right\} = \int_0^\eta x^{a-1}e^{-x}dx$, which can be written as a series expansion $P\left\{a, \eta\right\} = e^{-\eta}\eta^a \sum_{m=0}^{\infty} \frac{\eta^m}{\Gamma\left\{m + a + 1\right\}}$, where $\Gamma\left\{b\right\} = \int_0^\infty t^{b-1}e^{-t}dt$ is the Gamma function.

parameters of less than $\eta \approx 3.5$ it is impossible to reach the runaway regime, regardless of the value of $\mathcal{R}$.

Once in the runaway regime, the optimum efficiency is obtained by maximising γ_{eff}, as shown in Fig. 2.8 (lower). The plot shows three different values of $\mathcal{R}$, all for the same Ioffe-Pritchard type trap. It can be seen that with larger values of $\mathcal{R}$ the sweep can afford to be performed slower (i.e. using larger values of η), which results in overall higher efficiencies. A typical value of $\mathcal{R}$ achieved in our experiment at the very beginning of the evaporative cooling stage is ~ 500, and so the maximum efficiency should be obtained by selecting a truncation parameter of $\eta \sim 6$.

In general, the elastic collision rate can always be increased by applying a compression of the trapping potential. If the compression is adiabatic, then PSD will remain constant despite the accompanying increase in temperature. To reach the highest efficiencies the trap should be compressed as much as possible, up until the point where density-dependent three-body collisions begin to lead to additional losses. This typically can only arise at the very end of the cooling sweep, and in tight trapping geometries.

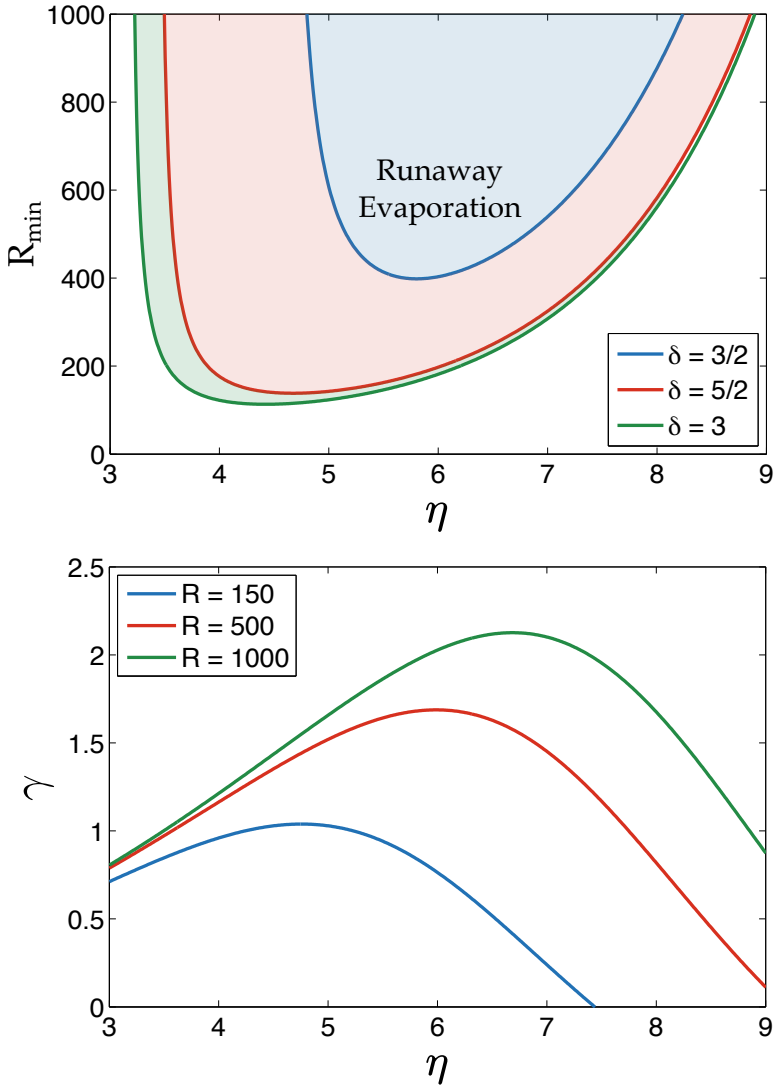

FIGURE 2.8: Optimum parameters for evaporative cooling using the Amsterdam model [151, 153]: The upper plot shows the minimum collision ratio $\mathcal{R}_{\text{min}}$ required to reach runaway evaporation as a function of the truncation parameter η in three types of trap with $\delta = 3/2$ (3D harmonic), $\delta = 5/2$ (Ioffe-Pritchard), and $\delta = 3$ (linear). The shaded regions indicate the regime in which it is possible to achieve runaway in each case. The lower plot shows the efficiency parameter γ_{eff} for values of collision ratios of $\mathcal{R} = 150$, 500 and 1000, indicating that greater ratios permit slower and more efficient evaporation sweeps.

2.7 Bose-Einstein Condensation

This section aims to give an overview of the important theoretical concepts and physics underlying the mechanism and of BEC, and to obtain the relevant parameters necessary for describing the experimental observations and measurements in later chapters. More detailed information can be found in the additional references where given, and there are several excellent reviews and textbooks available [154–156].

2.7.1 Ideal Bose Gas in a Harmonic Trap

The first and simplest situation is to consider an ideal gas of bosons, neglecting atomic interactions, trapped in the three-dimensional anisotropic harmonic oscillator potential

$$V_{\text{ext}}(\boldsymbol{r}) = \frac{1}{2}m\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right), \tag{2.49}$$

with single-particle energy eigenvalues given by

$$\epsilon_i = \left(n_x + \frac{1}{2}\right)\hbar\omega_x + \left(n_y + \frac{1}{2}\right)\hbar\omega_y + \left(n_z + \frac{1}{2}\right)\hbar\omega_z, \tag{2.50}$$

where the i^{th} state is characterised by the quantum numbers n_x, n_y and n_z, which can takes the values 0, 1, 2... . The mean occupation number of the state i with energy ϵ_i for a gas of bosons in thermodynamic equilibrium at temperature T is given through the framework of the grand canonical ensemble by the Bose-Einstein distribution function [157]

$$f(\epsilon_i) = \frac{1}{e^{(\epsilon_i - \mu)/k_B T} - 1}, \tag{2.51}$$

where μ is the chemical potential [1] acting to conserve the total particle number through the relation

$$N_{\text{tot}} = \sum_i f(\epsilon_i). \tag{2.52}$$

At high temperatures the mean occupation number of any given state is $\ll 1$, and so the chemical potential is very large and negative. As the temperature is reduced the chemical potential increases, but is constrained to be less than the value of the ground state energy $\mu < \epsilon_0 = \frac{1}{2}\hbar\left(\omega_x + \omega_y + \omega_z\right)$, since this would result in unphysical negative occupation numbers in Eq. 2.51. As $\mu \to \epsilon_0$ from below the occupation of the ground state, given by

$$N_0 = f(\epsilon_0) = \frac{1}{e^{(\epsilon_0 - \mu)/k_B T} - 1}, \tag{2.53}$$

becomes increasingly large. This is the mechanism of Bose-Einstein condensation - the macroscopic occupation of a single-particle energy state.

Critical Temperature

The total number of atoms can be written as a sum of the number in the ground state plus those in all other excited states

$$N_{\text{tot}} = N_0 + N_{\text{ex}} = N_0 + \sum_{i \neq 0} \frac{1}{e^{(\epsilon_i - \mu)/k_B T} - 1}. \tag{2.54}$$

The value of N_{ex} is bounded from above because of the constraint on the chemical potential $\mu < \epsilon_0$, and the maximum number of particles in excited states, $N_{\text{ex}}^{\text{max}}$, occurs when $\mu = \epsilon_0$ (which safely results in only positive occupation numbers because the ground state has been explicity separated out). If particles are then added to the system such that $N_{\text{tot}} > N_{\text{ex}}^{\text{max}}$, then the additional particles must go into the ground state condenstate, on which there is no restriction on the allowed number of particles, since the particles are bosons. In some sense then, the phenomenon of Bose-Einstein condensation can be viewed as a saturation of the available excited states in the system.

We can therefore determine the condition for the onset of Bose-Einstein condensation by setting the total number of atoms equal to the maximum number allowed in excited states

$$N_{\text{tot}} = N_{\text{ex}}^{\text{max}} = \sum_{i \neq 0} \frac{1}{e^{(\epsilon_i - \epsilon_0)/k_B T_c} - 1}, \tag{2.55}$$

which is reached when the system is cooled to the *critical temperature* T_c. When the number of particles is sufficiently large ($N_{\text{tot}} \to \infty$) and the temperature is much larger than the harmonic oscillator energy level spacing ($k_B T \gg \hbar\omega_i$), the discrete sum in Eq. 2.55 can be replaced by an integral across the continuum, and in addition the ground state energy can be neglected $\epsilon_0 = \mu = 0$ (inclusion of the ground state energy results in a finite size correction to the critical temperature on the order of $\sim 2\%$, as described in Sec. 5.2). Making use of the density of states, $g(\epsilon)$, the integration can be performed over energy, so that

$$N_{\text{tot}} = \int_0^\infty \frac{g(\epsilon)}{e^{\epsilon/k_B T_c} - 1} d\epsilon. \tag{2.56}$$

For the three-dimensional anisotropic harmonic oscillator, the density of states is given by [156]

$$g(\epsilon) = \frac{\epsilon^2}{2\hbar^3 \omega_x \omega_y \omega_z}, \tag{2.57}$$

and so after evaluating the integral Eq. 2.56 the critical temperature of the ideal gas in a harmonic trap is given by[3]

$$T_c = 0.94 \frac{\hbar}{k_B} \omega_{\text{ho}} N_{\text{tot}}^{1/3}, \tag{2.58}$$

where $\omega_{\text{ho}} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean of the three trapping frequencies. We

[3] Using the general result $\int_0^\infty x^{\alpha-1} / (e^x - 1)\, dx = \Gamma(\alpha)\zeta(\alpha)$.

can see that having either more atoms or a tighter trap raises T_c, and so less cooling has to be performed in order to reach the BEC transition.

The dependence of the condensate fraction N_0/N_{tot} on temperature can be found by combining Eq. 2.54 and Eq. 2.58, noting that below T_c the total number of atoms remains fixed but the number in the ground state increases with decreasing temperature. This dependence for the harmonic trap is given by

$$\frac{N_0}{N_{\text{tot}}} = 1 - \left(\frac{T}{T_c}\right)^3, \tag{2.59}$$

and is in contrast to the result for a uniform system which follows a $(T/T_c)^{3/2}$ form as opposed to $(T/T_c)^3$.

This approach describes the critical temperature remarkably well [158], considering all interactions have been neglected. In a typical cylindrically-symmetric trap in our experiment, with $\{\omega_\perp, \omega_{/\!/}\} \approx 2\pi \times \{1\,\text{kHz}, 20\,\text{Hz}\}$ and 10^5 atoms, critical temperatures for the transition to BEC are on the order of $\sim 500\,\text{nK}$, given by Eq. 2.58.

Another important quantity of interest is the spatial density distribution of the atoms which are not in the condensate - the *thermal component*, since information on the atomic cloud (particularly temperature) is gained from this density distribution via absorption imaging. The density distribution for the gas of non-interacting bosons is given in general by $n(\boldsymbol{r}) = \sum f(\epsilon_i)|\phi_i(\boldsymbol{r})|^2$, where $\phi_i(\boldsymbol{r})$ is the wavefunction for the i^{th} state [156]. This expression requires knowledge of the individual wavefunctions for the potential, and so usually a simpler approach is to employ a *semi-classical approximation*. Then it is assumed that the state of each particle can be described by its position $\boldsymbol{r}$ and momentum $\boldsymbol{p}$ using a distribution function in phase space $f_{\boldsymbol{p}}(\boldsymbol{r}) = \left[\exp\left([\epsilon_{\boldsymbol{p}}(\boldsymbol{r}) - \mu]/k_B T\right) - 1\right]^{-1}$, where $\epsilon_{\boldsymbol{p}}(\boldsymbol{r}) = p^2/2m + V_{\text{ext}}(\boldsymbol{r})$ is the particle energy. The spatial density of particles in excited states is then obtained by integrating the distribution function over momentum space $n_{\text{ex}}(\boldsymbol{r}) = (2\pi\hbar)^{-3} \int f_{\boldsymbol{p}}(\boldsymbol{r}) d\boldsymbol{p}$, which has been evaluated to be [97]

$$n_{\text{ex}}(\boldsymbol{r}) = \frac{1}{\lambda_{\text{dB}}^3} g_{3/2}\left\{\exp\left[\frac{\mu - V_{\text{ext}}(\boldsymbol{r})}{k_B T}\right]\right\}, \tag{2.60}$$

where the thermal de Broglie wavelength λ_{dB} and polylogarithm function $g_\alpha(x)$ are defined by

$$\lambda_{\text{dB}} = \sqrt{\frac{2\pi\hbar^2}{m k_B T}} \qquad \text{and} \qquad g_\alpha(x) = \sum_{k=1}^{\infty} \frac{x^k}{k^\alpha}. \tag{2.61}$$

At high temperatures, the chemical potential is large and negative, and only the first term in the polylogarithm sum in Eq. 2.60 contributes, in which case the density distribution reduces to a Gaussian. As the temperature is reduced and the higher order

terms in the sum become increasingly significant, the density at the centre of the gas increases due to *Bose enhancement* - a consequence of the quantum statistics. This effect has been observed in the experiment just before the onset of condensation, shown in Sec. 5.1, and must be accounted for when analysing density distribution images of clouds.

The condition for the transition to BEC can now be expressed in terms of the phase-space density of the cloud, PSD $= n(0)\lambda_{\mathrm{dB}}^3$, by setting $\mu, \boldsymbol{r} = 0$ in Eq. 2.60, leading to PSD $= g_{3/2}(1) \approx 2.612$. This condition reflects the fact that BEC occurs when the thermal de Broglie wavelength becomes on the order of the distance between particles, at which point the atomic wavepackets overlap and can be described by a single wavefunction.

The function in Eq. 2.60 describes the thermal component of the gas well - even for temperatures below T_c - but is not appropriate for atoms in the condensed state, which must be treated separately when integrating over momentum states. The density of the ground state can be obtained by putting all the atoms in the lowest state $\phi_0(\boldsymbol{r})$ with energy ϵ_0, and so the distribution, given by $n_0(\boldsymbol{r}) = N_0|\phi_0(\boldsymbol{r})|^2$, reflects the single-particle ground state wavefunction, and grows with the number of particles. The normalised wavefunction is determined in the non-interacting case by the Schrödinger equation, and is given for the harmonic trap by the Gaussian [155]

$$\phi_0(\boldsymbol{r}) = \Big(\frac{m\omega_{\mathrm{ho}}}{\pi\hbar}\Big)^{3/4} \exp\left[-\frac{x^2}{2a_x^2} - \frac{y^2}{2a_y^2} - \frac{z^2}{2a_z^2}\right], \tag{2.62}$$

where $a_i = \sqrt{\hbar/m\omega_i}$ are the harmonic oscillator lengths. These harmonic oscillator lengths, which determine the in-situ size of the consensate, are *independent* of the number of particles - a result which does not hold when the effect of atom-atom interactions is considered.

2.7.2 Effect of Interactions

Although the description of the ideal bose gas provides a way to understand the mechanism of BEC, it is not sufficient to describe the observations in experiments accurately. Even though the cold gases are extremely dilute, in reality there are always interactions between atoms which play an important part in describing the various properties, including the equilibrium density distribution and the dynamics of the condensate.

A gas of N interacting bosons in an external potential is described in the framework of second quantisation, which provides a natural way to account for interactions, with the Hamiltonian for the system given by [154]

$$\begin{aligned}\hat{H} = &\int d\boldsymbol{r}\hat{\Psi}^\dagger(\boldsymbol{r})\left[-\frac{\hbar^2}{2m}\nabla^2 + V_{\mathrm{ext}}(\boldsymbol{r})\right]\hat{\Psi}(\boldsymbol{r}) \\ &+ \frac{1}{2}\int d\boldsymbol{r}d\boldsymbol{r}'\hat{\Psi}^\dagger(\boldsymbol{r})\hat{\Psi}^\dagger(\boldsymbol{r}')U_{\mathrm{int}}(\boldsymbol{r}-\boldsymbol{r}')\hat{\Psi}(\boldsymbol{r}')\hat{\Psi}(\boldsymbol{r}),\end{aligned} \tag{2.63}$$

where $\hat{\Psi}^{\dagger}(\boldsymbol{r})$ and $\hat{\Psi}(\boldsymbol{r})$ are the bosonic field operators which create and annihilate a particle at position $\boldsymbol{r}$, respectively. Since the range of the atomic interaction is smaller than the interparticle distance, processes involving three or more particles can be neglected, and so the third term describes a two-body interaction potential U_{int}. For very low temperatures such as those below T_c, the collisions involve low momenta and take place in the s-wave scattering regime. In this case, the interactions are well-described by an effective elastic "hard sphere" contact potential

$$U_{\text{int}}(\boldsymbol{r}-\boldsymbol{r}') = g\delta(\boldsymbol{r}-\boldsymbol{r}'), \tag{2.64}$$

where the coupling constant g is given by

$$g = \frac{4\pi\hbar^2 a}{m}. \tag{2.65}$$

The interactions do not depend on the specifics of the interaction potential, but are characterised simply by a single parameter through the s-wave scattering length a. For ^{87}Rb, the interactions are repulsive, with $a = 5.2\,\text{nm}$ [159].

The Hamiltonian in Eq. 2.63 can be solved numerically using, for example Monte Carlo techniques, but becomes inefficient for large numbers of particles. The standard approach to overcome this is to use a mean-field, in which the field operator is separated into two parts

$$\hat{\Psi}(\boldsymbol{r}) = \hat{\Psi}_0(\boldsymbol{r}) + \delta\hat{\Psi}(\boldsymbol{r}), \tag{2.66}$$

where $\hat{\Psi}_0(\boldsymbol{r})$ represents the macroscopic condensed part and $\delta\hat{\Psi}(\boldsymbol{r})$ accounts for the small perturbation due to non-condensed atoms.

If there are a large number of particles in the same quantum state in the condensed part ($N_0 \gg 1$), then $\hat{\Psi}_0(\boldsymbol{r})$ can be treated as a classical field $\Phi_0(\boldsymbol{r})$. This is known as the *Bogoliubov approximation* [160], and $\Phi_0(\boldsymbol{r})$ is called the "wavefunction of the condensate". For very low or zero temperature there is no depletion of the condensate and the term $\delta\hat{\Psi}(\boldsymbol{r})$ can be neglected, replacing $\hat{\Psi}(\boldsymbol{r}) \to \Phi_0(\boldsymbol{r})$ directly. This means that the density of the condensate is given by

$$n_0(\boldsymbol{r}) = |\Phi_0(\boldsymbol{r})|^2, \tag{2.67}$$

with the condensate wavefunction being normalised to the total number of atoms $N_0 = \int |\Phi_0(\boldsymbol{r})|^2 d\boldsymbol{r}$. This is analagous to quantum electrodynamics, in which the microscopic description of photons using operators can be replaced by the classical electric and magnetic fields, if there are a large number of photons in the same quantum state.

The equation governing the time-evolution of the condensate wavefunction can be determined by writing the Heisenberg equation for the field operator $\hat{\Psi}(\boldsymbol{r})$ using the Hamiltonian in Eq. 2.63 together with the interaction of Eq. 2.64, giving rise to the

time-dependent Gross-Pitaevskii equation (GPE) [155]

$$i\hbar\frac{\partial}{\partial t}\Phi_0(\boldsymbol{r},t) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(\boldsymbol{r}) + g|\Phi_0(\boldsymbol{r},t)|^2\right)\Phi_0(\boldsymbol{r},t). \tag{2.68}$$

By writing the condensate wavefunction in terms of its spatial and time-dependent parts as $\Phi_0(\boldsymbol{r},t) = \Phi(\boldsymbol{r})\exp(-i\mu t/\hbar)$, a time-independent version is obtained

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{ext}}(\boldsymbol{r}) + g|\Phi_0(\boldsymbol{r})|^2\right)\Phi_0(\boldsymbol{r}) = \mu\,\Phi_0(\boldsymbol{r}), \tag{2.69}$$

where the chemical potential μ is fixed by normalisation on the total number of atoms. The equation governing the condensate wavefunction at zero temperature is therefore a Schrödinger equation with a non-linear term, proportional to the atomic density, that arises due to the effect of atom-atom interactions. The in-situ density distribution of the condensate differs greatly from the Gaussian of the non-interacting case given in Eq. 2.62. The effect of the interactions in the repulsive case is to reduce the central density of the cloud and to expand the physical size of the trapped gas, and agrees very well with the experimentally measured profiles [161].

2.7.3 Thomas-Fermi Regime

The interaction energy term in the GPE Eq. 2.68 will increase for larger numbers of atoms in the condensate, since this term is proportional to the density $|\Phi_0(\boldsymbol{r})|^2$. For repulsive interactions ($g > 0$), increasing interactions leads to an expansion of the in-situ size of the gas, and therefore to a more slowly varying density profile. The kinetic energy term in the GPE, which is proportional to the gradient of the wavefunction $\nabla^2\Phi_0(\boldsymbol{r})$, will therefore become negligible at some point with respect to the interaction energy. Dropping the kinetic energy term completely is known as the *Thomas-Fermi approximation*, and in fact describes the properties of the gas very well in most BEC experiments. In this case the equilibrium density distribution is determined by the balance between the repulsive interactions and the external confining potential. The Thomas-Fermi approximation describes the bulk part of a trapped condensate very well, but breaks down however near the surface of the cloud where the density becomes so low that the kinetic energy term is no longer respectively negligible.

The proper condition for applying the Thomas-Fermi approximation can be seen by writing the GPE in terms of scaled variables [155], which leads to the requirement $N_0\, a/a_{\text{ho}} \gg 1$, where $a_{\text{ho}} = \sqrt{\hbar/m\omega_{\text{ho}}}$ and a is the scattering length. For a typical BEC in our experiment with 10^4 atoms in a trap of strength $\{\omega_\perp, \omega_{/\!/}\} \approx 2\pi \times \{1\,\text{kHz}, 20\,\text{Hz}\}$, the parameter $N_0\, a/a_{\text{ho}} \sim 80$, and so the condensate is in the Thomas-Fermi regime.

When the kinetic energy term is dropped, the GPE gives the simple analytic solution for the condensate density distribution

$$n_0(\boldsymbol{r}) = \begin{cases} \frac{1}{g}\left(\mu - V_{\text{ext}}(\boldsymbol{r})\right), & \text{for } \mu > V_{\text{ext}}(\boldsymbol{r}) \\ 0, & \text{otherwise,} \end{cases} \tag{2.70}$$

and so the atoms will effectively fill the trap up to the chemical potential, with the density directly reflecting the shape of the external trapping potential. For the anisotropic harmomic oscillator with the potential given by Eq. 2.49, the density takes the form of the inverted parabola for $\mu > V_{\text{ext}}(\boldsymbol{r})$

$$n_0(\boldsymbol{r}) = \frac{\mu}{g}\left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2}\right), \tag{2.71}$$

where the point at which the density vanishes is given by the *Thomas-Fermi radii* along the i^{th} direction

$$R_i = \sqrt{\frac{2\mu}{m\omega_i^2}}. \tag{2.72}$$

The chemical potential is fixed by normalising the density profile in Eq. 2.71 to the total number of condensed atoms

$$\mu = \frac{1}{2}\hbar\omega_{\text{ho}}\left(\frac{15a}{a_{\text{ho}}}N_0\right)^{2/5}, \tag{2.73}$$

and the density distribution can be written in terms of the total BEC atom number

$$n_0(\boldsymbol{r}) = \frac{15}{8\pi}\frac{N_0}{R_x R_y R_z}\left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2}\right). \tag{2.74}$$

An estimate of the effect of the repulsive interactions can be seen by using the above equations to compare the size of the in-situ cloud, R_i, with the non-interacting case, a_i in the same trap. For a typical trap with 10^4 atoms and $\{\omega_\perp, \omega_{/\!/}\} \approx 2\pi \times \{1\,\text{kHz}, 20\,\text{Hz}\}$, we have $R_\perp \approx 2a_\perp$ and $R_{/\!/} \approx 15a_{/\!/}$, and so the interactions increase the size of the cloud, with the largest effect occuring along the direction of weakest confinement. Additionally, the central density is reduced by a factor of ~ 20. It can also be seen from Eqs. 2.72 and 2.73 that the size of the trapped gas grows with the number atoms $R_i \propto N_0^{1/5}$, in contrast to the non-interacting case in which the ground state size is independent of N_0.

Chapter 3

Experimental System

The production of BECs requires a complex apparatus, the implementation and characterisation of which is described in this chapter. The system consists of an ultra-high vacuum environment in which all the trapping and cooling takes place. A frequency stabilised laser system is used for interaction with the atomic clouds, primarily for laser cooling and probing the samples through a custom optical microscope. In addition the means of producing and regulating all magnetic trapping fields, both using external coils and integrated trapping structures, is described. All these elements must also be synchronised together using a computer-controlled timing system to achieve microsecond resolution.

3.1 Vacuum System

In order to achieve efficient evaporative cooling for BEC production, it is essential to minimise the loss rate due to collisions with residual hot background atoms, when compared with the rate of rethermalising elastic collisions between trapped atoms, as described in Sec. 2.6. Typically, this means a magnetic trap lifetime on the order of at least $10\,\mathrm{s}$, which requires an ultra-high vacuum (UHV) system with a pressure below 10^{-10} mbar. In contrast however, for initially capturing a large number of laser cooled atoms it is desirable to have a high partial pressure of trappable rubidium. In our system, these two competing criteria are addressed through the use of a two-dimensional magneto-optical trap (2D MOT) [162–164], which is connected to the main chamber via a low conductance tube and allows a pressure difference to be establised between the two regions. In this way, UHV is maintained in the science chamber whilst simultaneously dispensing rubidium into the 2D MOT cell.

The vacuum setup is depicted in Fig. 3.1. The main science chamber is a spherical octagon[1] made from the non-magnetic 316L stainless steel to minimise distortion of the applied trapping fields. There are two large DN100CF openings and eight DN40CF openings around the outside - five of which house anti-reflection coated viewports to enable optical access to the cold atomic clouds. The bottom of the chamber is fitted with a large viewport, providing entry for the MOT laser beams, whilst the top opening

[1] Kimball Physics MCF600.

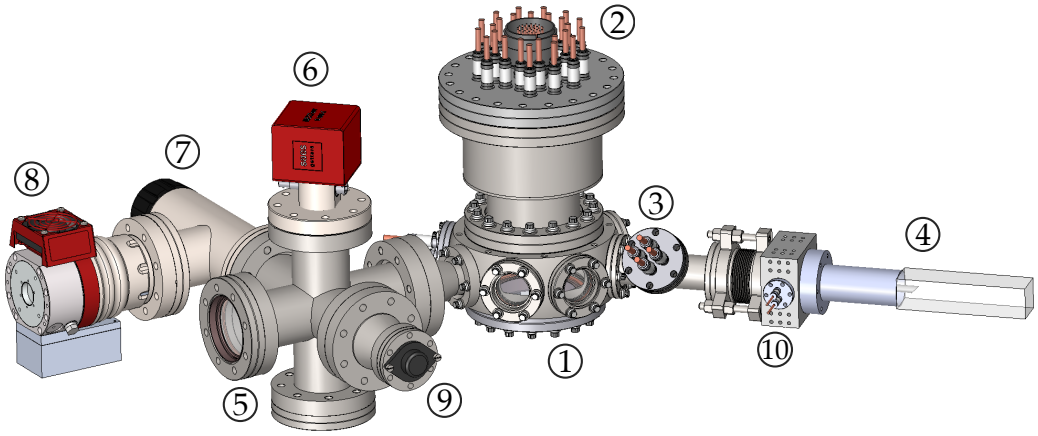

FIGURE 3.1: Vacuum System Overview: ① Main chamber, ② Large feedthrough, ③ 3D MOT rubidium dispensers, ④ 2D MOT connected through a differential pumping tube, ⑤ Six-way cross, ⑥ SAES NEXTorr pump, ⑦ All-metal valve, ⑧ Turbo pump, ⑨ Ion gauge, ⑩ 2D MOT rubidium dispensers.

houses a feedthrough for sending electrical currents to the trapping structures inside the vacuum chamber, as described in Sec. 3.3.

A six-way cross is fitted to the science chamber through a conical reducer, which allows connection of the pumping equipment. An all-metal valve[2] is the only opening to atmosphere, and allows attachment of a turbomolecular drag pump[3] during the initial pump-down stages. The pressure can be estimated using an ion gauge, but this becomes unreliable at increasingly lower values (below $\sim 10^{-11}$ mbar) and ultimately the vacuum quality is assessed by the lifetime of atoms in the magnetic traps.

Originally, pumping was accomplished through a combination of both two $55\,\mathrm{l\,s^{-1}}$ ion getter pumps and a titanium sublimation pump. However in mid 2015, all three of these sizeable pumps were removed and replaced by a single compact NEXTorr[4] device [165], which combines a large $500\,\mathrm{l\,s^{-1}}$ non-evaporable getter (NEG) element with a small $6\,\mathrm{l\,s^{-1}}$ ion pump. The NEG element is primarily responsible for removing active gases such as hydrogen, whilst the ion pump removes noble gases and methane that are not pumped by the NEG. Since the getter element is a passive device, the system is able to maintain UHV even in the event of a power failure. By replacing the old large ion pumps and titanium sublimation pump with the NEXTorr, the total volume of the vacuum chamber and the complexity was considerably reduced.

The source of rubidium-87 for the main 3D MOT is an intense beam of pre-cooled atoms, produced from a very compact and modular 2D MOT system. The trapping region is a rectangular cell manufactured by Japan Cell Co., and is made from borosilicate glass which has a high degree of flatness and a resistance to thermal shock making it ideal for the bakeout process. The glass cell has dimensions $3.5\,\mathrm{cm} \times 3.5\,\mathrm{cm} \times 14\,\mathrm{cm}$,

[2]Kurt Lesker VZCR60R, bakeable up to $450\,^{\circ}\mathrm{C}$.
[3]Pfeiffer TMU 071.
[4]SAES Getters D 500-5.

and is anti-reflection coated for $780\,\mathrm{nm}$ on the outside, whilst the inside is left uncoated. To connect the glass cell to the main chamber, we use a custom made flange with an integrated differential pumping tube, produced by LewVac. There is no additional pump for the glass cell - vacuum is maintained only by the main pump through a small $800\,\mu\mathrm{m}$ aperture in the differential pumping tube, which allows a pressure difference to be established between the glass cell and the science chamber. The glass cell is mounted onto one of the DN40CF openings on this custom flange, which is in turn connected to the main science chamber via the use of a port aligner[5], as shown in Fig. 3.1. Crucially, the port aligner gives the ability to the fine-tune the angle of the entire 2D MOT module for optimising the position of the atomic beam and therefore loading of the 3D MOT. The custom flange has three smaller DN16CF openings, two of which are used to mount alkali metal dispensers to introduce rubidium into the glass cell. In addition there are also rubidium dispensers located directly in the main chamber, in order to perform optimisation of the 3D MOT without the use of the 2D MOT.

The geometry of the differential pumping tube is illustrated in Fig. 3.2. Its purpose is to maintain a difference in pressure between the main chamber and the 2D MOT cell, made possible due to the large mean free path of particles at UHV pressures which is larger than the dimensions of the system.

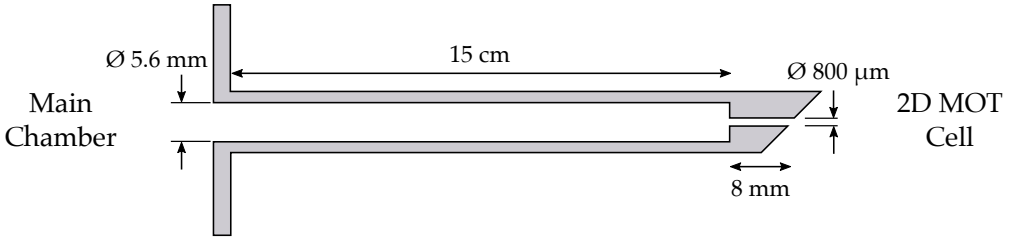

FIGURE 3.2: Geometry of the Differential Pumping Tube: By modelling the tube as two successive pipes, it is estimated that this geometry is able to maintain a pressure difference of five orders of magnitude between the main chamber and the 2D MOT cell.

The ratio of the pressures in each region can be estimated by [166]

$$\frac{P_{\text{main}}}{P_{\text{cell}}} = \frac{C_{\text{tot}}}{S}, \tag{3.1}$$

where C_{tot} is the total conductance of the tube, and S is the effective pumping speed on the main chamber side. In the molecular flow regime, the conductance of a pipe of length L and diameter d is given by [166]

$$C = 12.1\frac{d^3}{L}[\mathrm{l\,s^{-1}}], \tag{3.2}$$

[5]Lesker PA35-H

where d and L are measured in centimetres. As can be seen in Fig. 3.2, the tube can be modelled as two successive pipes with a total conductance given by

$$\frac{1}{C_{\text{tot}}} = \frac{1}{C_1} + \frac{1}{C_2}. \tag{3.3}$$

Due to the location of the pump in the setup, the effective pumping speed is reduced from the nominal $500\,\mathrm{l\,s^{-1}}$ down to $S \simeq 21\,\mathrm{l\,s^{-1}}$ at the entrance to the differential pumping tube, as a result of the combined conductances of the connecting pipes. Using Eq. 3.1 and 3.2 the maximum pressure ratio is then calculated to be $P_{\text{cell}}/P_{\text{main}} \sim 2500$, which means that we can afford to increase the 2D MOT pressure up to $\sim 10^{-7}$ mbar whilst ensuring that the science chamber remains below 10^{-10} mbar to enable long magnetic trap lifetimes. However, at pressures higher than $\sim 10^{-7}$ mbar the mean free path of particles becomes on the order of the length of the glass cell [163], and collisional losses will begin to inhibit the performance of the 2D MOT.

3.1.1 Bakeout Procedures

To achieve UHV pressures, there are many careful steps which have to be taken. Importantly, all components must be thoroughly cleaned to remove oils and dust particles, which would otherwise lead to outgassing and limit the final pressure obtainable. For smaller parts, the cleaning begins by placing them in an ultrasonic cleaner with a lab soap, mainly to remove residual oils. Next, after rinsing in distilled water, they are then cleaned again using methanol, which evaporates without leaving a residue. All parts are handled using latex gloves, and are kept in oil-free aluminium foil until they are ready to be used.

After assembling the entire vacuum system, a helium test is performed to check for any leaks around viewports and other sensitive connections. Using the turbo pump, backed by an oil-free roughing pump, we are able to pump down the system down to around 10^{-6} mbar. In order to go further and achieve UHV, the system must then be baked at around $200\,^{\circ}\mathrm{C}$, in order to accelerate the outgassing process and remove water and other trapped gases from the inside walls of the chamber.

By measuring the pressure in the chamber using the ion gauge, and by monitoring all critical parts of the setup using thermocouples, a typical characterisic time-evolution of the system is shown in Fig. 3.3. The temperature is controlled in such a way as to ensure there are no large gradients, which is especially important around glass-to-metal transitions, for example, that could lead to failure if the rates of expansion were too high. Everything is heated up in a slow, controlled way up to $200\,^{\circ}\mathrm{C}$ over the course of a day, and is left at this temperature for another five days until outgassing subsides and the pressure stabilises. Mid-way through the bake, the dispensers are activated by running a current of $5\,\mathrm{A}$ for several minutes, and the ion pump within the NEXTorr is degassed. The system is then cooled back down to room temperature slowly over the course of a full day. Once below $100\,^{\circ}\mathrm{C}$, the magnets can be mounted back on the

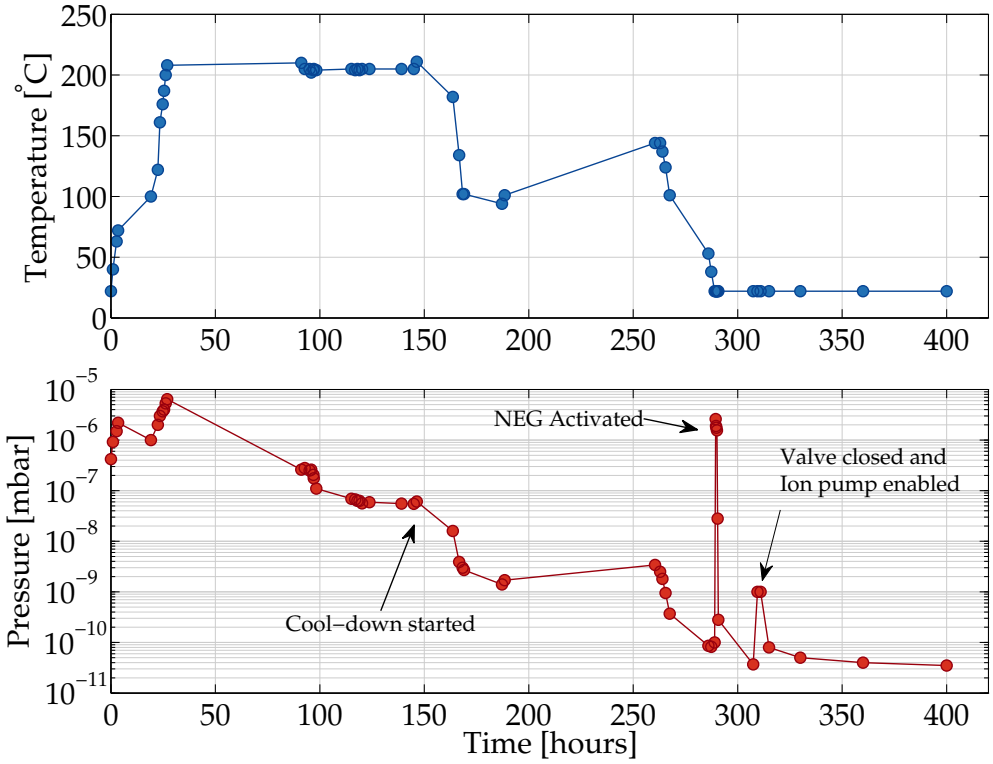

FIGURE 3.3: System evolution during bakeout: After baking at 200 °C, UHV is achieved with an ultimate pressure of 3×10^{-11} mbar.

NEXTorr pump and the NEG can be activated. Finally, the ion pump is switched on, the all-metal valve sealed with a torque wrench, and the turbo pump is switched off to minimise the effect of acoustic vibrations on the chamber. The ultimate pressure obtained with this proceduce is 3×10^{-11} mbar, as measured with the ion gauge, and leads to a measured magnetic trap lifetime in excess of 80 seconds.

3.2 2D MOT

For a long time, our system was run as a single chamber setup, by using only rubidium dispensers in the main science chamber. With this approach, a lot of time was spent characterising pulsing of these dispensers [167], in order to obtain a sufficient number of atoms in the initial MOT whilst allowing the vacuum to recover enough to perform efficient evaporative cooling. Nevertheless, we struggled to achieve this in a reliable way - for example, changing the duty cycle of the pulsing would lead to a complete change in the system characteristics. Eventually, to overcome these issues, a 2D MOT was installed in 2015, consistently facilitating a large number of atoms with sufficient magnetic trap lifetime. A 2D MOT works by trapping and cooling along two dimensions, but leaving the third dimension unconfined. This allows a stream of pre-cooled atoms to exit the trapping region to be recaptured by a main 3D MOT.

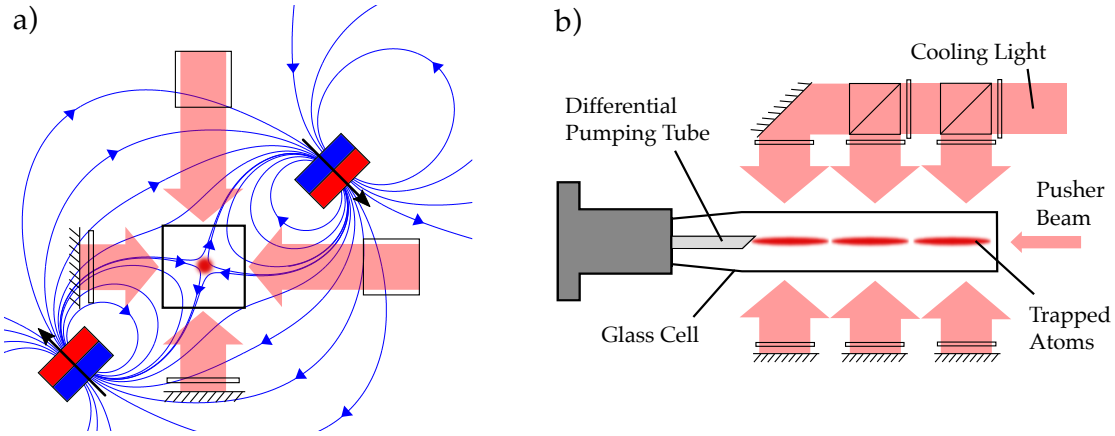

FIGURE 3.4: 2D MOT system: a) Front view. Two rows of permanent magnets generate the 2D quadrupole field with eigenaxes aligned with the propagation of the laser light. b) Side view. Cooling light is brought to the cell and then distributed and retroreflected to form three trapping regions. A pusher beam along the axis increases the flux which passes through the differential pumping tube.

Our 2D MOT is based on a design from the University of Heidelberg [168, 169], and was implemented at Nottingham by PhD student Jorge Ferreras [170]. This particular design has many advantages - it is a very compact, unimposing device, with total dimensions of only $20 \times 20 \times 40\,\mathrm{cm}^3$, including all optics and magnetic fields, and is relatively straightforward to implement. The setup is shown schematically in Fig. 3.4. A row of permanent magnets with their magnetisations anti-aligned generates a 2D quadrupole field that is translationally invariant along the length of the glass cell, as shown in Fig. 3.4 (a). The generation of the magnetic field is described in detail in Sec. 3.2.1. By placing the magnets along a $45°$ plane, the propagation direction of the cooling beams is collinear with the eigenaxes of the magnetic quadrupole at the centre of the glass cell, giving the correct polarisations for the operation of the MOT.

The available distance for incident cooling light between the end of the glass cell and the tip of the differential pumping tube is approximately $12\,\mathrm{cm}$, and so we installed three separate cooling regions of 1" diameter each, depicted in Fig. 3.4 (b). They are placed as close together as possible, to minimise the areas in between where there is no cooling light. The use of three circular cooling regions removes the need for any cylindrical lenses, which simplifies the setup. All optics are mounted on a detachable cage, which is supported directly by the custom flange, allowing it to be removed easily for alignment purposes or during a bakeout. Cooling and repumping light is overlapped into two optical fibres - one for vertical beams and one for horizontal - and brought to outcouplers on the 2D MOT cage. From there, it is distributed equally between the three cooling regions using polarizing beam splitting cubes and half-wave plates. Light is circularly polarized before entering the cell with quarter-wave plates, and is retroreflected by mirrors after exiting the cell, with additional quarter-wave plates to give the correct handedness on the return pass. In total, we use around $100\,\mathrm{mW}$ of cooling

power and $5\,\mathrm{mW}$ of repumping power, leading to approximately $17\,\mathrm{mW}$ and $1\,\mathrm{mW}$ of cooler and repumper in each of the three trapping regions, respectively. A photograph of the 2D MOT structure is shown in Fig. 3.5.

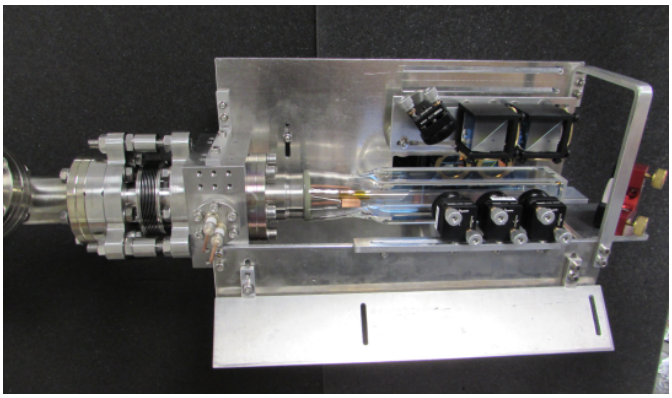

FIGURE 3.5: Photograph of the 2D MOT: Cooling light is distributed into the glass cell using polarising beam splitter cubes and waveplates mounted onto the detachable aluminium cage structure.

In addition, a so-called pusher beam with $500\,\mu\mathrm{W}$ of power is mounted on the front of the glass cell and directed along the long axis of the 2D MOT [163, 171]. Its purpose is to improve the flux of atoms passing through the differential pumping tube, by turning around atoms which would otherwise exit the wrong end of the axis in which there is no confinement. In principle then, the pusher beam should be red-detuned with respect to the atomic resonance so as to only address atoms which are moving towards it. However, after optimisation we find a large improvement with the pushing beam in fact blue-detuned, as presented in Fig. 4.12 in Sec. 4.2.

3.2.1 Magnetic Field Design

The two-dimensional magnetic quadrupole field required is produced entirely by permenant magnets, as shown in Fig. 3.4 (a), which has several advantages. Firstly, it can achieve the necessary field gradients for magneto-optical trapping very easily, which for 2D MOTs are usually produced by racetrack-type coils needing additional current supplies, increasing both power consumption and cost [162]. Here, the entire field is generated by a total of 18 magnets, making it space-effective, simple to arrange and construct, and requiring little alignment.

The neodymium magnets[6] used in our setup are rectangular with dimensions $25 \times 25 \times 5\,\mathrm{mm}^3$. To be able to simulate the fields, the magnetization of a single magnet must first be determined, which is not given in the manufacturer's specifications. Then, the total quadrupole field is calculated by simply summing field components due to individual magnets. The magnetic field components arising from a rectangular magnet

[6] First4Magnets F425.

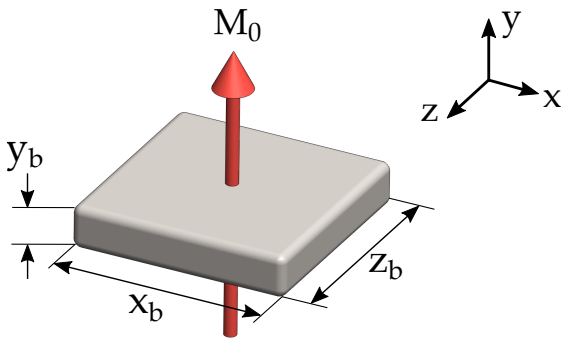

FIGURE 3.6: Schematic of the geometry of a single rectangular permenant magnet, providing definitions for Eq. 3.4. The side lengths are denoted by x_b, y_b and z_b, where the magnetization M_0 (shown by the red arrow) is defined to be along the y-axis, and the origin of the coordinate system is at the geometric centre of the magnet. For the magnets used in our 2D MOT setup, the dimensions are $x_b = z_b = 25\,\mathrm{mm}$ and $y_b = 5\,\mathrm{mm}$.

of side-lengths $\{x_b, y_b, z_b\}$ and with a constant magnetization M_0 aligned with the y-direction, as depicted in Fig. 3.6 is given by the analytical function [172]

$$
\begin{aligned}
B_x(x,y,z) &= \frac{M_0\mu_0}{4\pi}\sum_{k,l,m=1}^{2}(-1)^{k+l+m}\ln\Big\{z+(-1)^m z_b+g(x,y,z)\Big\},\\
B_z(x,y,z) &= \frac{M_0\mu_0}{4\pi}\sum_{k,l,m=1}^{2}(-1)^{k+l+m}\ln\Big\{x+(-1)^k x_b+g(x,y,z)\Big\},\\
B_y(x,y,z) &= \frac{M_0\mu_0}{4\pi}\sum_{k,l,m=1}^{2}(-1)^{k+l+m}\frac{[y+(-1)^l y_b][x+(-1)^k x_b]}{|y+(-1)^l y_b||x+(-1)^k x_b|}\\
&\quad\times\arctan\left\{\frac{|x+(-1)^k x_b|[z+(-1)^m z_b]}{|y+(-1)^l y_b|\,g(x,y,z)}\right\},
\end{aligned}
\tag{3.4}
$$

where

$$
g(x,y,z)=\sqrt{[x+(-1)^k x_b]^2+[y+(-1)^l y_b]^2+[z+(-1)^m z_b]^2}.
$$

To determine M_0 for a single magnet, the y-component of the magnetic field was measured as a function of the spatial y-coordinate using a computer-controlled translation stage and an automated Gaussmeter. Using Eq. 3.4, the parameter M_0 was varied to best fit the measured data, as shown in Fig. 3.7. Along with a measurement along the axis of the magnet a second was done along an axis displaced by $24\,\mathrm{mm}$, allowing the centre position of the magnet to be calibrated more precisely, shown in Fig. 3.7 (b). The resulting magnetization was determined to be $(9.95 \pm 0.05) \times 10^5 \mathrm{A/m}$.

Now that M_0 is known, a suitable way to arrange the magnets can be established. There is a very strong attractive force between individual magnets, and so they are mounted inside two custom-made non-magnetic aluminium holders in order to rigidly fix the arrangement. There are two requirements on the magnetic field - firstly, the gradient along each radial direction must be 10 - $20\,\mathrm{G\,cm^{-1}}$ for magneto-optical trapping; secondly, the variation of this radial gradient should be small along the axial direction,

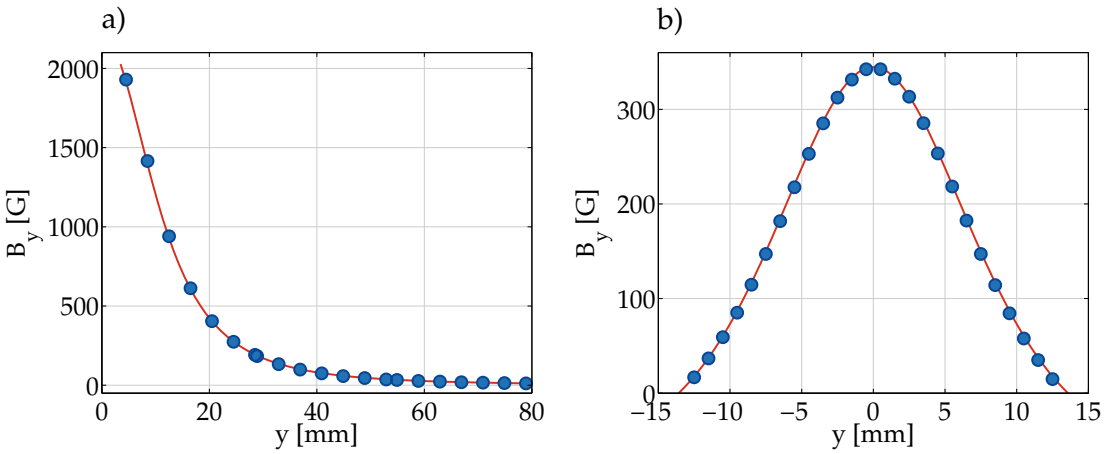

FIGURE 3.7: Characterisation of a single magnet: a) Axial field component B_y measured as a function of distance from the centre of the magnet y. b) B_y vs. y along an axis displaced by $24\,\mathrm{mm}$, allowing the magnet centre to be precisely determined. Data points correspond to measured magnetic field, whilst the solid lines are a calculation according to Eq. 3.4, with the magnetisation as the single fit parameter. The result of the measurement gives $M_0 = (9.95 \pm 0.05) \times 10^5 \mathrm{A\,m^{-1}}$.

i.e. it should not vary too much over the length of the glass cell. After optimisation, the best results are achieved with 9 magnets arranged in each holder, shown in Fig. 3.8. The magnets that are situated towards the outside ends are positioned progressively closer to the glass cell, in order to make the gradient more constant along the axial direction. In addition, it also helps the homogeneity to include three magnets stacked on top of each other at the two extreme outside positions.

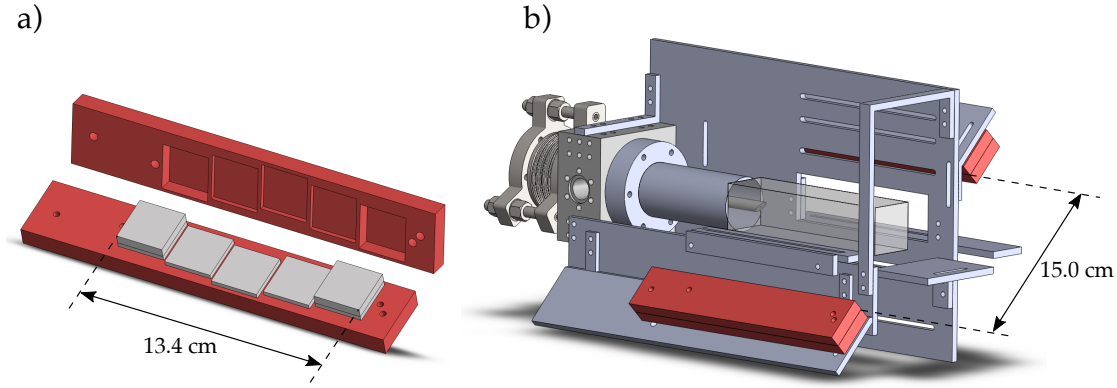

FIGURE 3.8: 2D MOT magnet arrangement: a) Two aluminium holders are used to fix the positioning of a total of 18 permanent magnets. b) The holders are attached on a $45°$ plane onto the mounting structure, producing the 2D quadrupole along the axis of the glass cell.

The two magnet holders are mounted along a 45° axis with respect to the optics structure, shown in Fig. 3.8 (b), with a series of slots to allow the separation between them to be varied for adjusting the absolute value of the field gradient. The separation

can range from $13\,\mathrm{cm}$ to $15.5\,\mathrm{cm}$, giving rise to gradients of $13\,\mathrm{G\,cm^{-1}}$ to $21\,\mathrm{G\,cm^{-1}}$, respectively. At $15\,\mathrm{cm}$, the radial field gradient is $15\,\mathrm{G\,cm^{-1}}$, which is ideal for producing a MOT. The calculated magnetic field in a radial plane at the centre of the glass cell is shown in Fig.3.9 (a) for this geometry.

After mounting the holders onto the 2D MOT cage, the 2D quadrupole field was measured in a radial plane at the centre before attaching the glass cell. Each of the vertical and horizontal field components was measured independently by scanning a 2D grid using a computer-controlled translation stage, and recording the field using a Gaussmeter. The resultant magnetic field is shown in Fig. 3.9 (b), and agrees reasonably well with the calculation. It should be noted that there is a slight tilt of $11°$ of the measured quadrupole eigenaxes with respect to the laboratory frame. This arises because the holders are not mounted exactly on the same plane but have a small offset, and is reproduced well with a more detailed simulation.

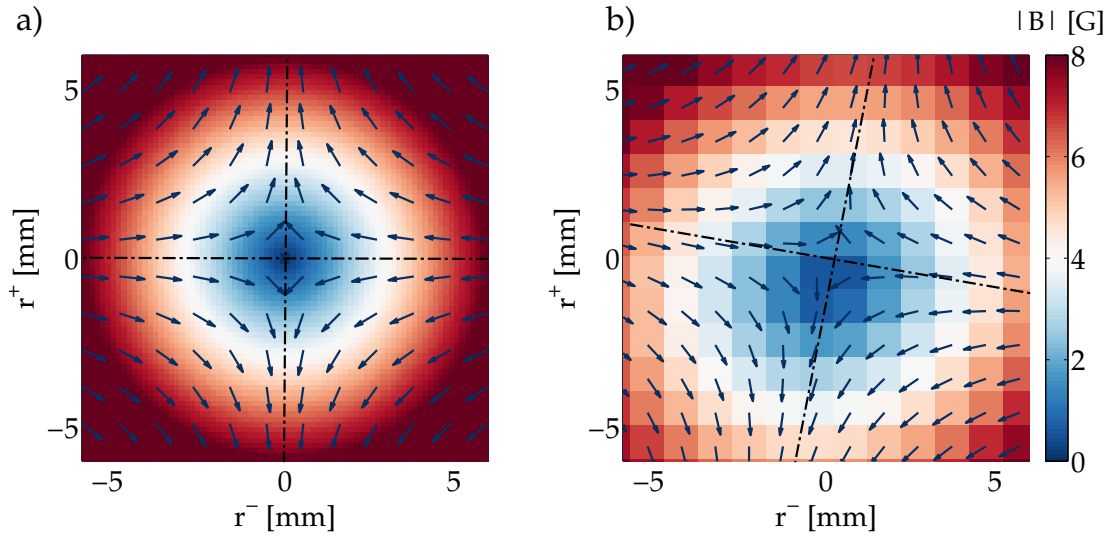

FIGURE 3.9: Magnetic field for the 2D MOT in the radial plane: a) Calculation of the field at the centre of the glass cell. The two orthogonal coordinates are labelled r^+ and r^-, and the eigenaxes of the quadrupole are shown by the dash-dotted lines. b) Measurement of the field at the centre of the glass cell is consistent with the calculated field gradient of $15\,\mathrm{G\,cm^{-1}}$. The measured eigenaxes are tilted by $\sim 11°$ as described in the text.

In addition to the radial plane, the field gradient was checked along the axis of the glass cell. The calculation shown in Fig. 3.10 confirms that the magnet arrangement described above leads to a satisfactory homogeneity along the axial direction.

Because the alignment of the holders is fairly coarse, it is impossible to fix the magnets by hand whilst maintaining the zero line of magnetic field collinear with the axis of the differential pumping tube aperture. To provide fine-tuning of the position of this zero line, we added two small sets of rectangular offset coils in approximately Helmholtz configuration, allowing displacement of the 2D atom cloud by several millimeters in each of the horizontal and vertical directions. This turned out to be critical

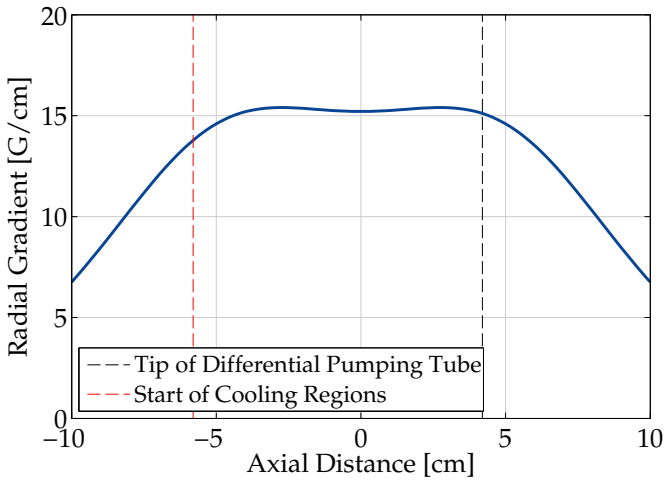

FIGURE 3.10: 2D MOT radial magnetic field gradient along the axial direction: Due to the geometry of the magnets inside the holders, the field gradient remains relatively constant at $15\,\mathrm{G\,cm^{-1}}$ over the range in which cooling light is present.

for obtaining a flux of atoms into the 3D MOT, as shown in Fig. 4.10 in Sec. 4.2.

3.3 Main Flange Assembly

The central part of the experiment is the atom chip and its custom designed mounting. The main purpose of the mount is to align the atom chip at the centre of the vacuum chamber, providing mechanical stability and heat sinking, as well as to facilitate electrical connections between the chip wires and the outside of the vacuum system. In addition, the mounting also incorporates the "understructure", which is a variety of copper structures located below the chip surface. These are used in the earlier parts of the experimental cycle for pre-cooling atoms, to enable efficient transfer to chip traps.

A photograph of the mount, before attaching the atom chip itself, is shown in Fig. 3.11 (a). It was designed primarily by previous PhD student Anton Piccardo-Selg, and more details can be found in his PhD thesis [118]. The entire assembly is attached to a single DN160CF custom flange, which has 25 large 1/4" copper feedthroughs rated at $180\,\mathrm{A}$ each. There is also a smaller 35-pin feedthrough connector in the centre of the flange, providing connections for the atom chip wires. The macroscopic copper trapping structures, shown in Fig. 3.11 (b), are attached to the large feedthroughs via 12 copper rods, and all fit into the chip holder together, using sheets of Kapton film to ensure electrical isolation. The chip holder is made from Macor, which is a machineable ceramic with good thermal conductivity and is UHV compatible. The entire assembly is $20\,\mathrm{cm}$ long, and is mounted to the top opening of the science chamber with the help of a DN160CF to DN100CF adapter, as can be seen in Fig. 3.1 of Sec. 3.1. In this way, the atom chip is mounted upside down so that atoms fall away from the surface

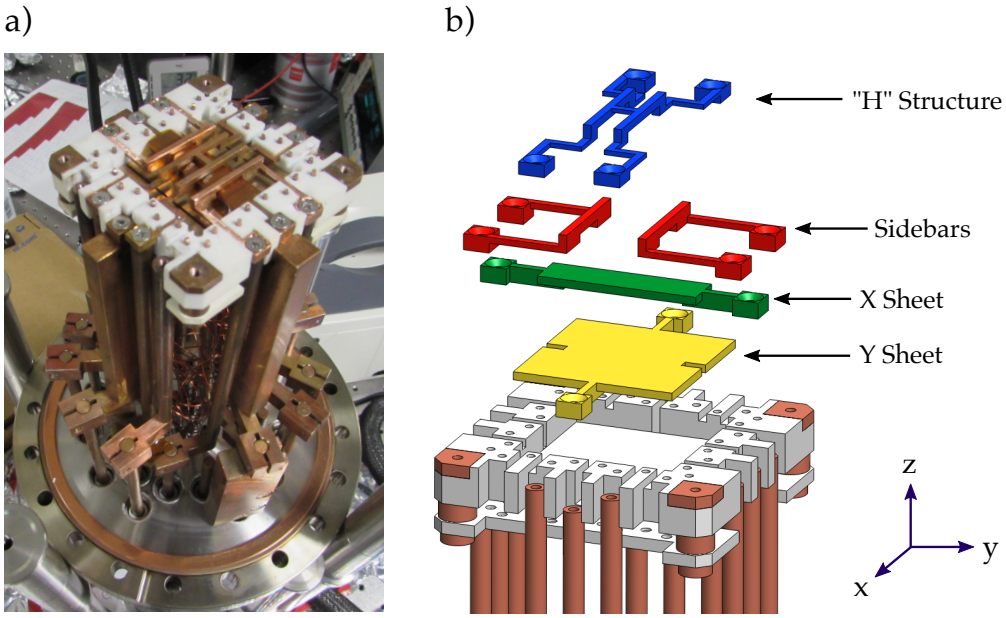

FIGURE 3.11: Main flange and understructure: a) Photograph of the main flange assembly. The mounting structure for the atom chip consists of a series of copper structures for producing magnetic fields, all mounted directly onto a large DN160CF flange via copper rods. b) Exploded view showing the various trapping structures. There are 3 types of geometry: the H-shaped conductor, the sidebars, and two orthogonal current sheets, all of which are stacked inside the Macor ceramic chip holder.

after the trapping fields have been removed, enabling time-of-flight measurements to be performed.

3.3.1 Copper Trapping Structures

Several different shaped copper wires are available to use in the understructure, all of which can be seen in Fig. 3.11 (b). Firstly, closest to the atom chip surface we have an H-shaped conductor. In normal operation, opposite diagonal terminals are connnected allowing current to flow in a Z configuration, which creates an Ioffe-Pritchard trap of the type described in Sec. 2.5.4. The cross-section of the leg wires is $2 \times 2\,\mathrm{mm}^2$, and they have a centre-to-centre distance of $7.7\,\mathrm{mm}$. The central bar is narrower than the two legs, being $1.5\,\mathrm{mm}$ wide by $1\,\mathrm{mm}$ high, which leads to higher current densities and therefore higher magnetic field gradients. Due to low resistance and large cross-section of these understructure wires, substantial currents on the order of $100\,\mathrm{A}$ are able to be pushed through them. This allows the Z conductor to produce a magnetic trap with a depth of $\sim 1\,\mathrm{mK}$ at a distance of several millimetres from the atom chip surface, which is ideal for capturing sub-Doppler cooled clouds directly from a MOT with high efficiency.

In addition to the H-structure, there are two orthogonal sheet conductors, denoted X-Sheet and Y-sheet, as shown in Fig. 3.11 (b). A broad conductor produces a magnetic

field component with very low gradient perpendicular to the current flow, and in fact exhibits vanishing curvature at a distance of $0.21W$, where W is the conductor width [118]. The field tends to homogeneous in the limit of an infinitely wide conductor, and sheet structures can therefore be used to produce homogeneous bias fields if they are made wide enough. The Y-Sheet has a width of $40\,\mathrm{mm}$, a height of $2\,\mathrm{mm}$, with its surface located $7.6\,\mathrm{mm}$ below the atom chip surface, and produces a magnetic field in the y-direction. Since this field is parallel to the central bar of the H-structure, the Y-Sheet can be used to adjust the Ioffe field at the trap bottom when using the Z-trap. Fig. 3.12 (a) shows that the Y-Sheet produces a magnetic field at the chip surface of $86\,\mathrm{mG\,A^{-1}}$, and so a modest current of $12\,\mathrm{A}$ is sufficient to provide the typical $1\,\mathrm{G}$ Ioffe offset field. Using planar sheets to create homogeneous bias fields has the advantage of fast switching times due to their low inductance, in contrast to the use of external Helmholtz coils.

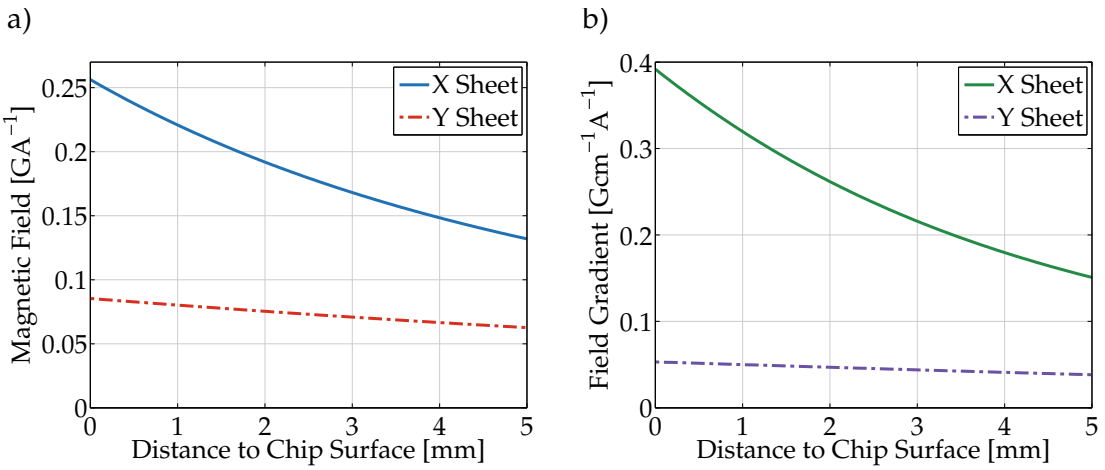

FIGURE 3.12: Magnetic fields produced by copper sheet structures: a) The component of the magnetic field perpendicular to the flow of current per unit current, as a function of distance to the atom chip surface. The X-Sheet and Y-Sheet produce fields of $0.26\,\mathrm{G\,A^{-1}}$ and $0.086\,\mathrm{G\,A^{-1}}$ at the plane of the chip, respectively. b) Gradients of the fields shown in a). The X-Sheet and Y-Sheet produce gradients of $0.4\,\mathrm{G\,cm^{-1}\,A^{-1}}$ and $0.05\,\mathrm{G\,cm^{-1}\,A^{-1}}$ at the chip surface, respectively. The stronger gradient of the X-Sheet is due to its narrower width.

The X-Sheet is narrower than the Y-Sheet with a width of $10\,\mathrm{mm}$, and its surface is located $4.0\,\mathrm{mm}$ below the chip surface. Due to the smaller width, this sheet is able to produce more field per unit current with $0.26\,\mathrm{G\,A^{-1}}$, but exhibits 8 times the gradient of the wider Y-Sheet. This sheet was made narrower in order to provide a stronger radial gradient for producing the magneto-optical trap, but since the X-Sheet generates a field perpendicular to the central bar of the H-structure, it can be used as the main bias field for the Z-trap. This allows creation of a magnetic trap - and even BEC - without the use of large external coils (demonstrated in Sec. 5.6).

The final conductors located beneath the atom chip are the sidebars, shown in Fig. 3.11. By running currents either parallel or anti-parallel to each other, both horizontal

(y-direction) or vertical (z-direction) fields can be generated, respectively. The sidebars are used to provide the third axis of confinement for the magneto-optical trap, by running the currents in the anti-parallel configuration.

Due to the large length ($26\,\mathrm{cm}$) of the connection leads for the copper structures, it is important to consider the relative change in length due to thermal expansion when passing high currents. We must ensure that enough gap is allowed between all structures and the underside of the chip so that they do not come into contact in routine operation of the experiment. This is especially important for the H-structure, which sits closest to the bottom of the atom chip, and must not be allowed to expand and touch the chip. However, we also wish to have the structures as close as possible to provide the largest magnetic fields for a given current at the position of the atoms. To determine the gap which must be left, the amount of expansion of the H-structure was measured for various currents. This was done by attaching the feedthrough to the science chamber, without the atom chip mounted, and observing the position of the H-structure on the CCD camera. A moderate vacuum was obtained of $\sim$1 mbar, in order to remove any effect of convection cooling, which would not be present in UHV. Various currents from $80\,\mathrm{A}$ to $160\,\mathrm{A}$ were passed through the wire, allowing the temperature equilibriate for ~ 30 minutes, and the position of the conductor was compared to that when cold. The result can be seen in Fig. 3.13, showing that the maximal expansion for $160\,\mathrm{A}$ is $300\,\mu\mathrm{m}$. In normal operation of the experiment, we apply only $80\,\mathrm{A}$ with a duty cycle of 50%, and a conservative gap of $350\,\mu\mathrm{m}$ was adjusted between the top of the H-structure and the bottom of the chip surface.

For calculation of the magnetic fields produced by the understructure wires, there are several approaches that can be taken. For a crude approximation, a stick model can be used (see, for example, [173, 174]), but since the distance to the atoms is on the order of the width of the conductor itself, this model is not good enough. It can be greatly improved by approximating the wires as finite-sized rectangular conductors with a uniform current density, analytical expressions for which are listed in Appendix A. However, to accurately calculate the magnetic fields produced by these millimetre-sized structures finite element methods should be used. An example of such a calculation is shown for the H-structure connected as a Z configuration in Fig. 3.14, performed using a finite element solver software package.[7] Whenever there is an abrupt change in the geometry, for example at the 90° bends, there will be a non-uniform current density which produces additional magnetic field components. In addition, it can be seen that hot-spots arise at the sharp corners, where the most power will be dissipated. Once the current density has been obtained, the magnetic field is easily calculated by integration using the Biot-Savart law. All calculations of magnetic fields in this thesis produced by the macroscopic copper understructure are performed using this method.

For verification of the finite element simulations, the magnetic field produced by the various structures was measured with a Gaussmeter in the vicinity of the chip surface.

[7] COMSOL Multiphysics.

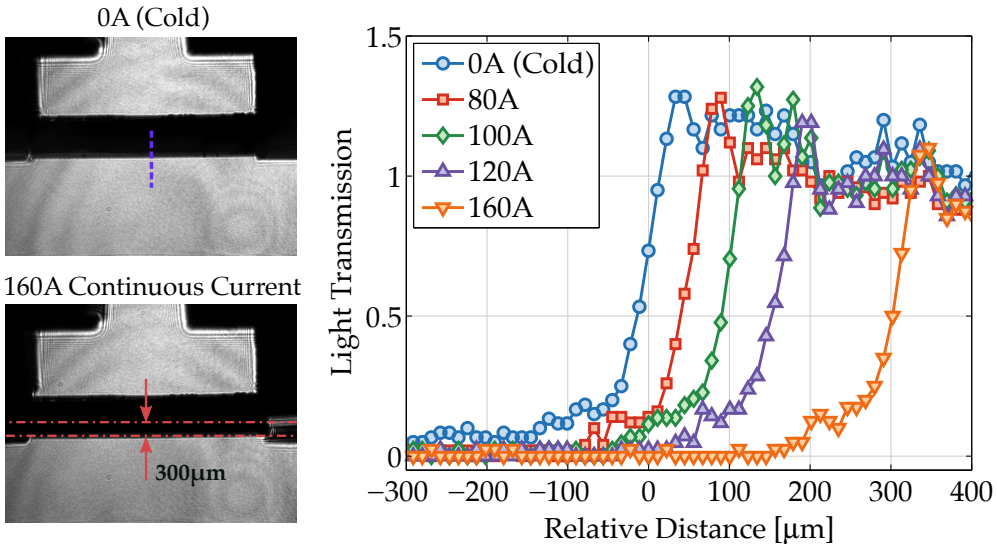

FIGURE 3.13: Thermal expansion of the copper H-structure: Currents of 80 A (□), 100 A (◇), 120 A (△) and 160 A (▽) were passed through the structure continuously. A weak laser beam was directed perpendicular to the central part of the H-structure along the x-direction, and the position of the shadow cast (black region in the images) was observed on a CCD camera relative to the position when cold (○), providing an indicator of the thermal expansion. Examples of the images when cold (left, top) and with 160 A (left, bottom) are shown. The plot (right) shows the normalised light transmission along the blue dashed line indicated in the left, top image. For 160 A continuous current, the maximal relative expansion is 300 µm.

The results are displayed in Fig. 3.15, and the agreement between the calculation and measurement is excellent.

3.3.2 Atom Chip Mounting

The atom chip is mounted directly above the copper understructure onto a custom-made carrier piece, shown in Fig. 3.16. The chip carrier is made from boron nitride, which is a UHV-compatible ceramic with exceptional thermal conductivity,[8] can be baked to high temperatures, and is easily machinable. The atom chip is aligned and glued to the chip carrier by using a small amount of thermally-conductive epoxy[9] in each of the four corners. The boron nitride chip carrier is finally attached to the main Macor understructure housing using four screws, along with a copper rim clamp which distributes the force of the screws to prevent the boron nitride from cracking.

Electrical connections to the atom chip wires are made with the help of a custom printed circuit board (PCB), designed by Dr. Fedja Oručević, which is glued onto the chip carrier using the same epoxy. The PCB is 500 µm thick and is double layered with 70 µm copper tracks on each side, connected by vias. The atom chip wires are connected to the PCB pads by wire bonding, using the ball bonding machine in the

[8] 27 W K^{-1} m^{-1}.
[9] EPO-TEK 920.

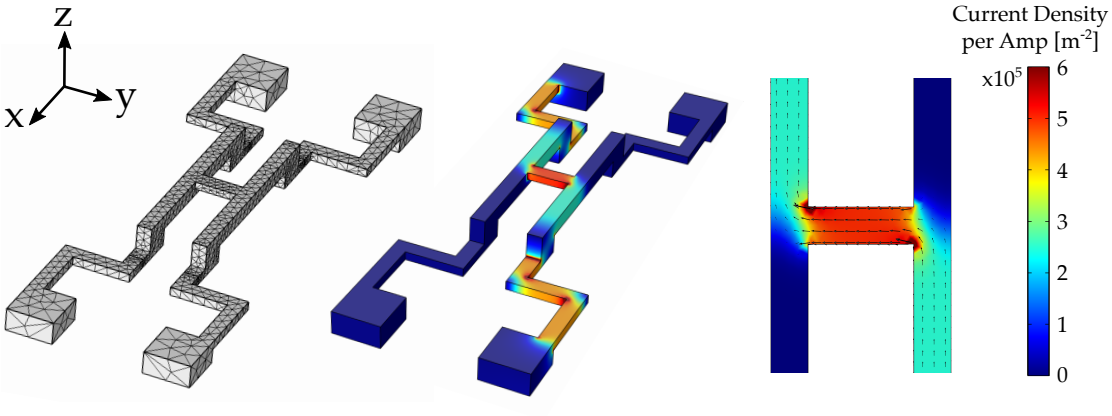

FIGURE 3.14: Current density distribution in the H-structure: For complicated geometries, the current density must be calculated using finite element analysis. This consists of first creating a mesh of the geometry (left), and then solving Maxwell's equations in each element for the various electromagnetic quantities, subject to supplied boundary conditions. The resultant current density distribution is shown in the centre image, with an enlargement of the relevant central region on the right, illustrating the hot spots which arise around the sharp geometrical features.

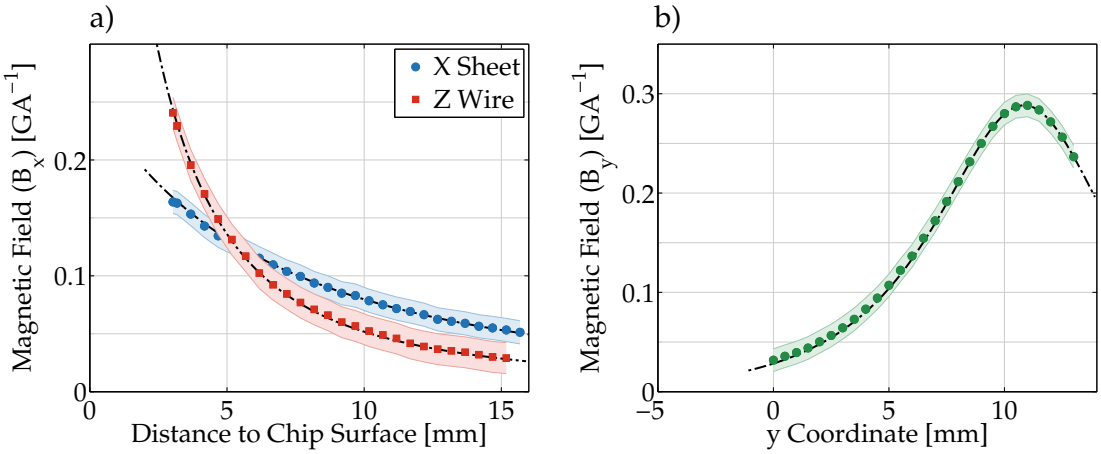

FIGURE 3.15: Measurements of understructure fields: a) The measured magnetic field component perpendicular to the current flow (B_x) is shown as a function of distance from the atom chip surface for both the X-Sheet (●) and the H-structure connected in a Z configuration (■). b) The measured parallel field component (B_y) as a function of the parallel y-coordinate for a single sidebar structure (●) , measured at a distance of $z = 3.0\,\mathrm{mm}$ above the chip surface. The origin of the y coordinate is located at the geometric centre of the H-structure shown in Fig. 3.11, and is halfway between the two sidebars which have a centre-to-centre separation of $20\,\mathrm{mm}$. In both plots, the shaded regions correspond to the uncertainty of the Gaussmeter, and the dashed lines are the result of the finite element calculation.

school cleanroom. This technique works by utilising heat, pressure and ultrasonic energy to weld micrometer-sized wires between the two surfaces. Since it gives a cleaner, stronger connection, the ball side of the bond is attached to the atom chip pad, whilst

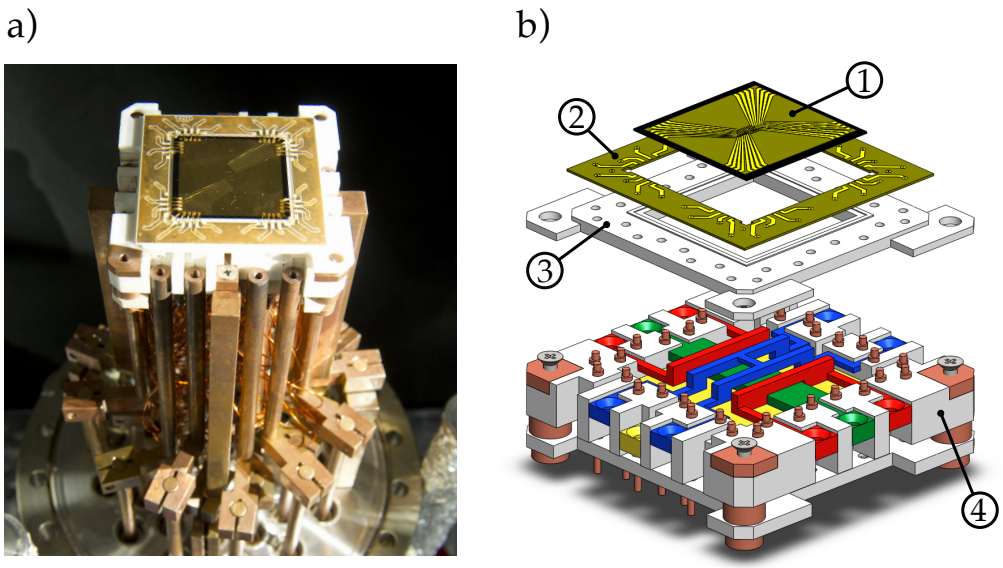

FIGURE 3.16: Atom chip mounting: a) Photograph of the main flange structure after fully assembling the atom chip, just before inserting into the vacuum chamber. b) Exploded view of the components for connecting the atom chip. Electric connections to the atom chip ① are made by wire bonding 32 pads to the PCB ②. The PCB and chip are then directly glued to the boron nitride chip carrier ③, which is finally attached to the Macor chip holder ④ by screws in each of the four corners.

the fishtail side is positioned on the PCB [175]. In our case, 25 µm diameter gold wire is used to create typically 25 bond wires per pad, and thus allowing the connection to withstand currents of several amps. Once the optimum parameters for the particular surfaces have been found - namely bonding pulse time, force, power and temperature - the task of bonding the total of around 800 wires can be completed in less than two days. A scanning electron microscopy (SEM) image of the finished bonded connections is shown in Fig. 3.17.

After the atom chip has been wire bonded to the PCB, we have to finally bring connections from outside the vacuum chamber. This is done using 32 spring-loaded copper pins which are held in the Macor understructure housing, and can be seen in Fig. 3.16. When the boron nitride chip carrier is attached, these pins make contact with small individual copper pieces, which are designed to divert the force produced by the springs away from the PCB and onto the chip carrier instead. A small drop of electrically-conductive glue fixes the copper pieces to the underside of the PCB, ensuring a homogenous, low-resistance contact. Finally, Kapton-isolated copper wires are attached between the 35-pin feedthrough and the underside of the spring-loaded pins using barrel connectors.

3.3.3 Chip Characteristics

The atom chip currently being used in the expeiment was produced at Ben Gurion University, Israel. In the future, more atom chips can be, and indeed already have

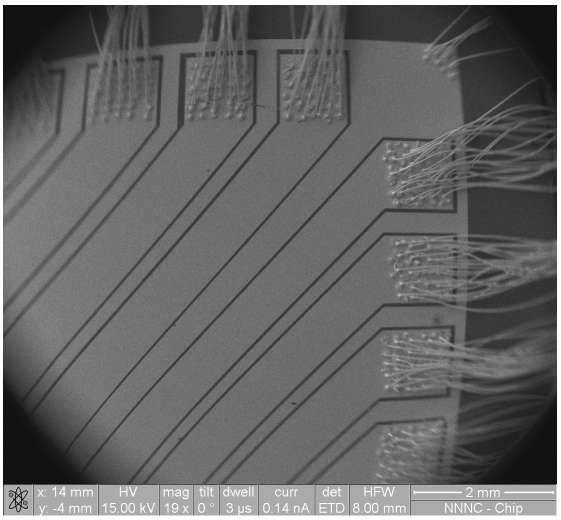

FIGURE 3.17: Completed wire bonds: An SEM image of the corner of the atom chip, after bonding is finished. It is possible to fit approximately 25 bonds of $25\,\mu\mathrm{m}$ wire per pad, allowing currents of several amps to be supported.

been, produced within the school of physics at Nottingham, by fabrication postdoctoral research associate Dr. Jessica Maclean. The current atom chip, designated BGU3, was produced with evaporated gold via photolithography, whereby a custom design is transferred to a photoresist through the use of a mask [176]. The substrate is a $35\,\mathrm{mm} \times 35\,\mathrm{mm} \times 500\,\mu\mathrm{m}$ silicon wafer, which is boron-doped to improve the heat conductivity, and has a $100\,\mathrm{nm}$ layer of silicon dioxide (SiO_2) to provide electrical isolation between neighbouring wires. A thin $30\,\mathrm{nm}$ layer of titanium is applied on top of the SiO_2 to provide a sticky surface, which helps the gold to adhere to the substrate during evaporation. Photoresist is then spun onto the titanium, and then exposed to ultraviolet light through a lithography mask to develop the resist. Finally, $2\,\mu\mathrm{m}$ of gold is evaporated onto the surface, and then the remaining resist pattern is removed during lift-off, leaving behind only the desired wire pattern. Gold is used because of its low electrical resistance, allowing higher currents to be passed through the wires. It also acts as a very good mirror - the optical reflectivity of the current chip was measured to be $>98\%$ for a $780\,\mathrm{nm}$ laser beam incident at $45°$ - allowing the implementation of the Mirror-MOT.

The layout of the current version of the atom chip is shown in Fig. 3.18, illustrating the different types of wires available for manipulating the atomic clouds. The central part of the chip contains three $10\,\mu\mathrm{m}$ wide and two $20\,\mu\mathrm{m}$ wide Z-shaped wires, with a leg-to-leg separation of $1.8\,\mathrm{mm}$, which are used to generate Ioffe-Pritchard fields for magnetic trapping. Secondly, there are two $50\,\mu\mathrm{m}$ wide ⊢-shaped conductors, which depending on the connection points can be run in either a Z or U configuration, or also as straight sidebars. There are also two independent large $200\,\mu\mathrm{m}$ wide sidebars, which can be used to tune the axial confinement of the trapped gas. Finally, the chip design includes a $50\,\mu\mathrm{m}$ wide comb-shaped structure. The original intention of this

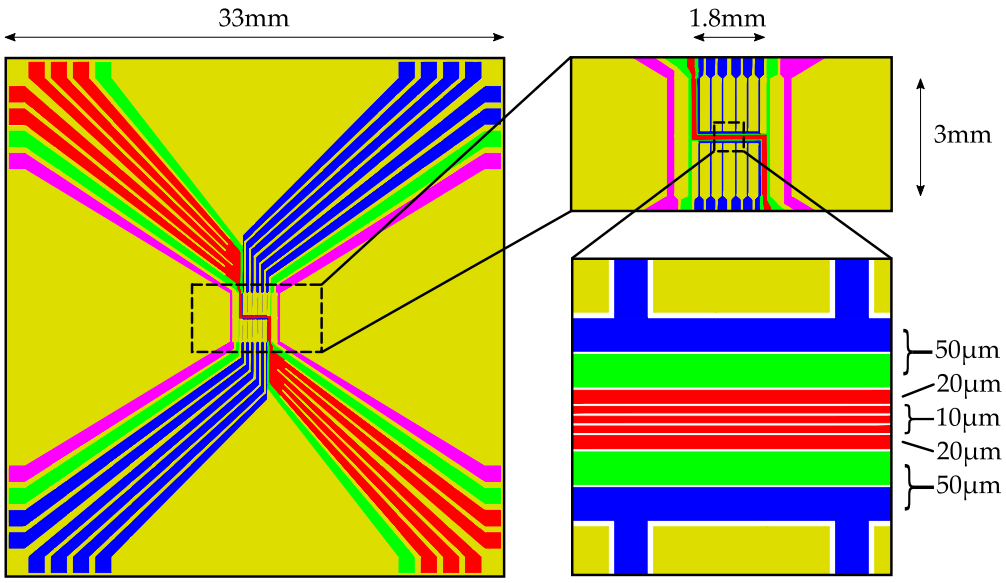

FIGURE 3.18: Layout of the atom chip wires: Several classes of structure are available for manipulating cold clouds. Centrally, there are five Z-shaped wires ( ) for producing Ioffe-Pritchard traps. The outer two have widths of 20 µm, whilst the inner three have widths of 10 µm. Two ⊢- shaped structures ( ) are available, with widths of 50 µm. In addition, there are two independent large 200 µm sidebars ( ) which can be used to adjust the axial confinement of a trapped cloud. Finally, the comb-structure ( ) provides the ability to create flexible RF potentials.

comb layout was to implement a dynamical splitting of the elongated trapped cloud through the use of radio-frequency dressed potentials [71, 119]. However, this has not been tested in our setup to date.

As part of the characterisation process, the resistances of all chip wires were measured before the chip was mounted in the vacuum chamber. This was done by a four-wire measurement technique using a high-precision multimeter and a dedicated probe station to ensure a consistent contact for each wire. These data were taken after the atom chip had already been bonded to the PCB, allowing the probes to be attached to the PCB pads and to prevent damage and scratching of the soft gold chip surface. Resistances were also measured after each additional stage in the mounting - including after connecting the small copper pegs, the barrel connectors, after the feedtrhough, and after a custom feedthrough-to-BNC breakout box. The values are listed in Table 3.1.

Finally, it is important to know how much current can safely be passed through the chip wires. Larger current densities through the Z wires is always beneficial, since it increases the depth of the trap, making it easier to load hotter clouds. However, too high a current will melt the gold, rendering the wire (along possibly with neighbouring wires) inoperative. A useful quantity for monitoring is the resistance of the wire, which increases when power is dissipated due to the temperature-dependent resistivity, and is therefore an indicator of the amount of heating. Before the chip was placed in the

TABLE 3.1: Ohmic resistances of various chip wires.

Wire Type	Bare Resistance (Ω)	Including All Connections (Ω)
10 µm Z	8.4 ± 0.2	12 ± 2
20 µm Z	4.53 ± 0.01	17.8 ± 0.4
50 µm Z	2.00 ± 0.01	3.6 ± 0.1
200 µm Sidebar	0.80 ± 0.01	2.3 ± 0.1

vacuum chamber, various currents were passed through the narrowest ($10\,\mu m$) and widest ($200\,\mu m$) structures and the voltage drop was recorded after it had been allowed to equilibriate, providing a measure of the steady-state resistance. The relative increase of the resistance with respect to the room-temperature value is shown in Fig. 3.19. This situation represents a worst-case scenario whereby current is passed continuously - in practice, the current is pulsed for around $5\,s$ with a cool-down time between pulses of about $20\,s$, i.e. a 20% duty cycle. A typical maximum allowable resistance increase for similar atom chip experiments is 50% [177, 178]. In practice, we limit ourselves to 20%, to avoid a risk of burning the wire, which occurs for currents of $500\,mA$ and $1.6\,A$ for the $10\,\mu m$ and $200\,\mu m$ wires, respectively.

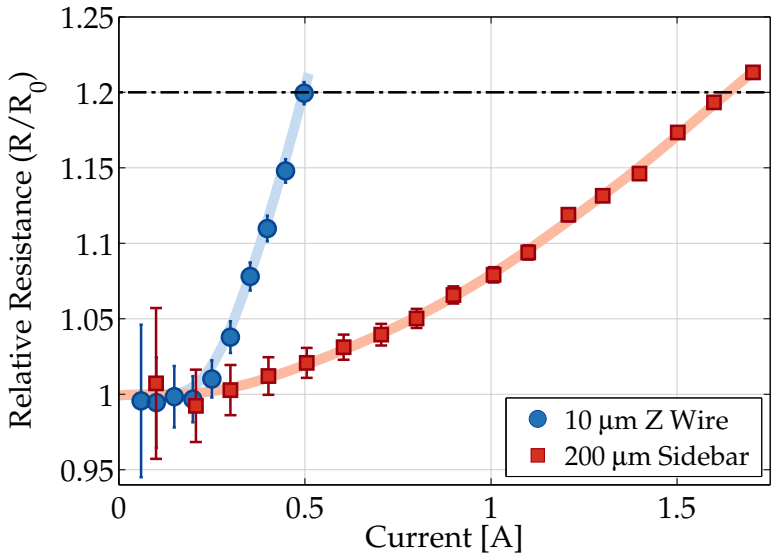

FIGURE 3.19: Heating of the chip wires: The relative change in resistance R/R_0 is measured as a function of continuously applied current, for the narrowest $10\,\mu m$ (●) and widest $200\,\mu m$ (■). The "cold" resistances are $R_0 = 8.67\,\Omega$ and $0.75\,\Omega$, with the 20% increase level reached with $500\,mA$ and $1.6\,A$, respectively.

3.4 External Coils

In addition to the inhomogeneous magnetic fields produced by the understructure and atom chip wires, a variety of homogeneous fields are required throughout the experimental cycle. Although low-gradient fields can indeed be produced by the sheet-like conductors in the understructure, as described in Sec. 3.3.1, we also have a total of 4 sets of external coils in Helmholtz configuration, and their arrangement is shown in Fig. 3.20.

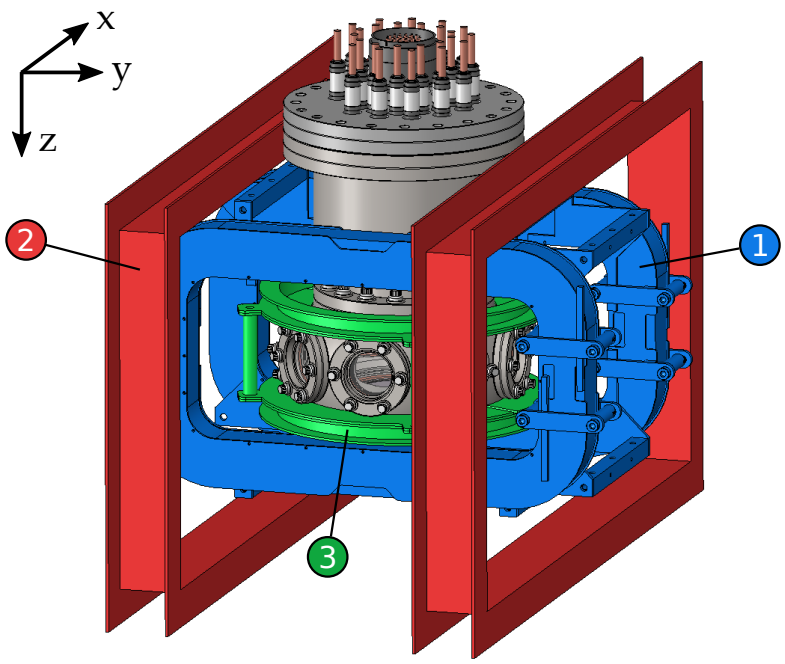

FIGURE 3.20: External coils in the experiment: Three coil holders were designed, producing fields in the three orthogonal directions. These coils are labelled X Bias ①, Y Bias ②, and Z Bias ③. A fourth separate coil, designated X Bias MOT, is wound on top of the main X Bias and creates a bias field for the MOT.

The largest homogeneous field, acting as the bias for the copper Z magnetic trap, must produce fields up to $\sim 70\,\mathrm{G}$, and requires the most thought from a design standpoint. This coil is designated "X Bias", and can be seen in Fig. 3.20. The coil should produce as large a field as possible at the position of the atoms for a given amount of current, since higher current increases the power dissipation. More field for the same current can be obtained by adding additional turns to the coil, but this increases the inductance and therefore limits the achievable switching times. In addition, a larger number of windings will increase the total resistance of the coil, again raising the power dissipation. There is an interplay therefore, between the inductance of the coil and the power dissipation, which must be considered carefully.

The final design uses kapton-insulated copper wire[10] with a square cross-section of

[10]S&W Wire SQ197, resistance = $0.85\,\mathrm{m}\Omega$ per metre.

$5 \times 5\ \mathrm{mm}^2$, and a circular hollow core of diameter $2.7\ \mathrm{mm}$ to allow the option of water-cooling, depicted in Fig. 3.21 (left). Since the wire has a relatively large cross-section however, we have not found water-cooling to be necessary, as the equilibrium temperature of the coils reaches less than $40\ ^\circ\mathrm{C}$ during routine operation of the experiment. In total there are 36 windings per coil, resulting in an inductance of $1.0\ \mathrm{mH}$ and a magnetic field of $1.65\ \mathrm{G\ A}^{-1}$ at the centre of the chamber, shown in Fig. 3.21 (right), agreeing exactly with the calculated value for this geometry. The holder has a series of crossbars to ensure consistent separation of the coils each time they are removed from the setup, for example during a bakeout. It is mounted directly onto the optical table and not on the chamber, in an effort to minimise vibrations from the magnetic kick produced during switch-off. It also turned out to be critical to cut slots in each side of the holder to prevent Eddy currents from circulating whenever the field switched.

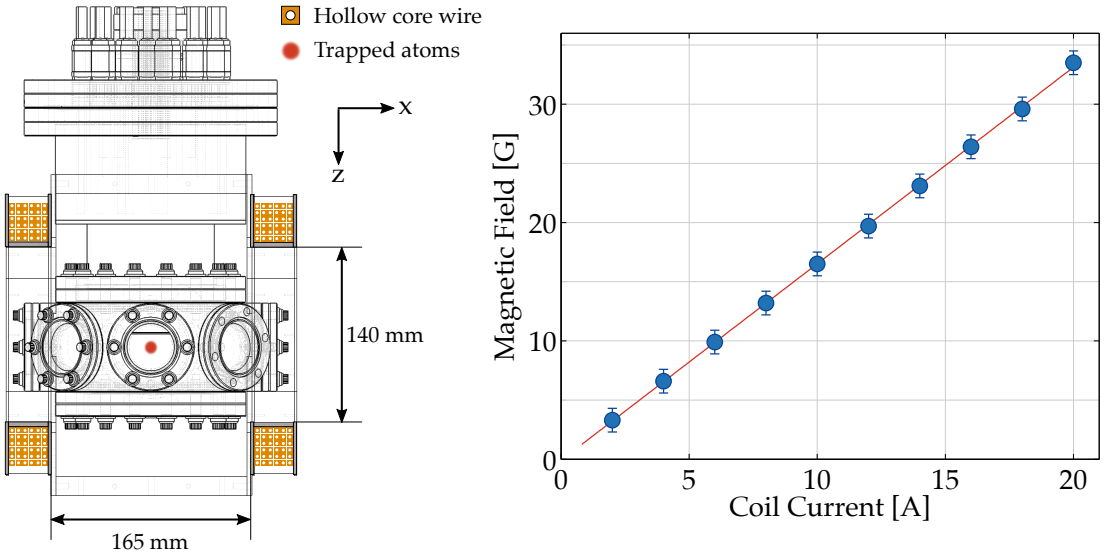

FIGURE 3.21: External "X Bias" coil: (Left) The coil holder is placed as close as possible to the chamber. The coil itself is wound from 36 windings hollow-core copper wire per side, to facilitate water cooling if necessary. (Right) The axial magnetic field component is measured at the centre of the Helmholtz coil pair as a function of applied current. A linear fit (solid line) to the data shows that the coil geometry produces a magnetic field of $(1.65 \pm 0.02)\mathrm{G\ A}^{-1}$.

The design of the three remaining coils is less critical, since they are not required to produce such large fields. There are three coils producing fields in all three orthogonal directions, used throughout the experimental cycle, shown in Fig. 3.20, denoted X Bias, Y Bias and Z Bias. There is also a fourth coil which is wound directly on top of the X Bias, named X Bias MOT, which is oriented in the opposite sense to the X Bias and is used to provide the bias field for the MOT stage. The measured magnetic fields produced by all four coils are given in Table 3.2.

TABLE 3.2: Magnetic fields produced by each coil in the experiment.

Coil	Magnetic Field ($\mathrm{G\,A^{-1}}$)
X Bias	1.65
Y Bias	1.80
Z Bias	3.61
X Bias MOT	1.72

3.5 Laser Setup

A cold atom experiment requires precise control over the frequencies, powers, and timing of several lasers simultaneously. Laser cooling is done on the D_2 transition of rubidium, which contains a cycling transition, at $780\,\mathrm{nm}$. Light with this near-infrared wavelength is conveniently obtainable with semiconductor diode lasers [179], which are readily available and cost-effective since they are routinely used, for example, in CD players. The natural linewidth of the rubidium D_2 line is $6.1\,\mathrm{MHz}$, and so a laser linewidth on the order of $\sim 1\,\mathrm{MHz}$ is required to be able to reliably address the individual hyperfine transitions. Since the linewidth of a free-running laser diode is much larger than this (at least several tens of megaHertz) [180], an external cavity is used for narrowing it down below $1\,\mathrm{MHz}$ [181, 182].

The employed transitions on the ^{87}Rb D_2 line are shown in Fig. 3.22. There are two lasers used in the experiment, both of which are bought commercially.[11] The largest amount of power required is that which is used for driving the cycling transition for laser cooling on the line $|5^2\mathrm{S}_{1/2}, F = 2\rangle \rightarrow |5^2\mathrm{P}_{3/2}, F' = 3\rangle$. This is provided by a TA Pro, which is a grating-stabilised master diode laser with an integrated tapered amplifier and maximum output power of $3\,\mathrm{W}$ at $780\,\mathrm{nm}$.

With such large intensities used for cooling in the MOT, there is also a significant probability of off-resonant excitation on the transition $|F = 2\rangle \rightarrow |F' = 2\rangle$. From here, atoms can decay by spontaneous emission into the F=1 ground state. Since the two ground states are separated by $6.8\,\mathrm{GHz}$, these atoms are no longer addressed by the cooling laser and would be rapidly lost from the trap. Therefore, a second laser is required to pump atoms back into the cooling cycle. This "repumping" light is provided by a DL Pro, which provides a maximum power of $80\,\mathrm{mW}$ tuned to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. With the inclusion of an optical isolator (to prevent back reflections into the laser cavity), the actual useable output of this laser is only $28\,\mathrm{mW}$ and after fibre-coupling this is not sufficient power for both the 3D and 2D MOTs, as well as for the frequency locking. Therefore, the repumper laser power is further increased by using the standalone commercial tapered amplifier BoosTA, which is able to amplify a small $20\,\mathrm{mW}$ seed beam up to a maximum of $1.5\,\mathrm{W}$, whilst preserving the spectral

[11] Toptica Photonics

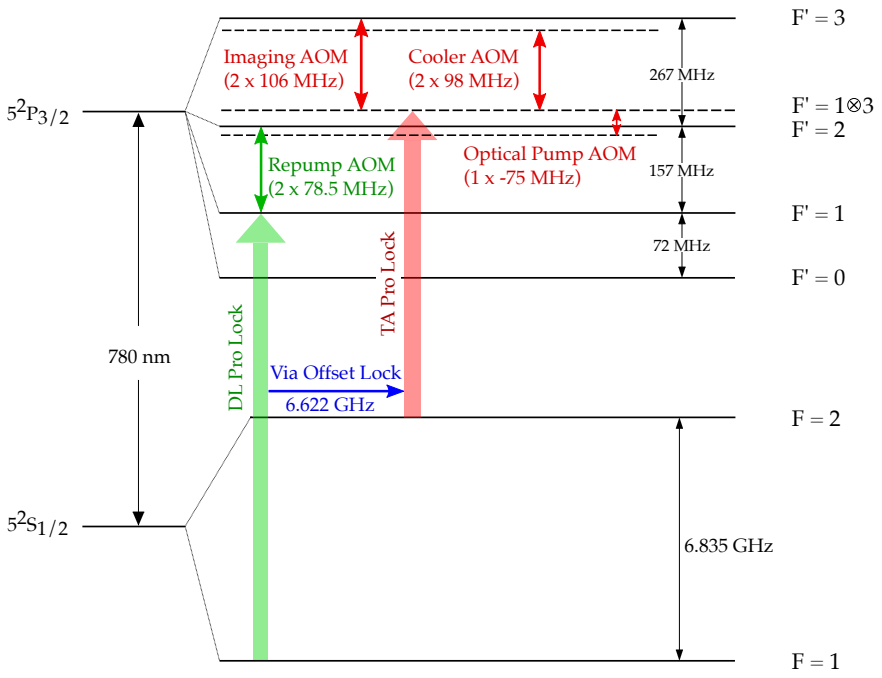

FIGURE 3.22: Transitions in the D_2 line of ^{87}Rb: The DL Pro laser provides repumping light for the experiment, and is stabilised to the $|F = 1\rangle \rightarrow |F' = 1\rangle$ transition. The TA Pro is offset-locked with respect to the DL Pro using a digital beat locking technique. For convenience, this offset is chosen to be $6.622\,\mathrm{GHz}$, to bring the TA Pro on resonance with the $|F = 2\rangle \rightarrow |F' = 1 \otimes 3\rangle$ crossover peak. All of cooling, optical pumping and imaging light are derived from this TA Pro laser, and are further fine-tuned using their individual AOMs. For clarity, the 2D MOT cooling and repumping lines have been omitted.

properties of the master laser. This amplifier is simple to setup, and requires no extra control electronics.

In addition to the cooling and repumping light for the MOTs, two other frequencies are required. The first is for state preparation before loading the purely magnetic trap, known as optical pumping. This requires a beam of $400\,\mu\mathrm{W}$ detuned by $20\,\mathrm{MHz}$ from the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition as will be described in Sec. 4.4. Secondly, for absorption imaging of the atomic cloud a beam of $1\,\mathrm{mW}$ is needed, which should be resonant with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition in order to maximize the absorption signal to noise ratio. Both optical pumping and imaging beams are derived from the main TA Pro laser, and are then independently detuned using their respective AOMs, as described in Sec. 3.5.1.

3.5.1 Frequency Stabilisation

It is important to have stable laser frequencies over an extended period of time when performing measurements. To reduce noise on the lasers, everything is mounted on a floating optical table, and enclosed in a matte aluminium box to minimise air currents

and reflective etalon effects. This helps to reduce detrimental acoustic vibrations and temperature variations in the laboratory. However, to prevent long term drifts, as well as to combat electrical noise, we apply active stabilisation of the frequencies ("locking") using a feedback loop with a proportional-integral-derivative (PID) controller. A schematic of the layout for the frequency locking is depicted in Fig. 3.23. The DL Pro laser is locked to a rubdium vapour cell using standard frequency modulation spectroscopy [183], whilst the TA Pro laser is then offset locked with respect to the DL Pro. Both methods in the scheme are described in more detail below.

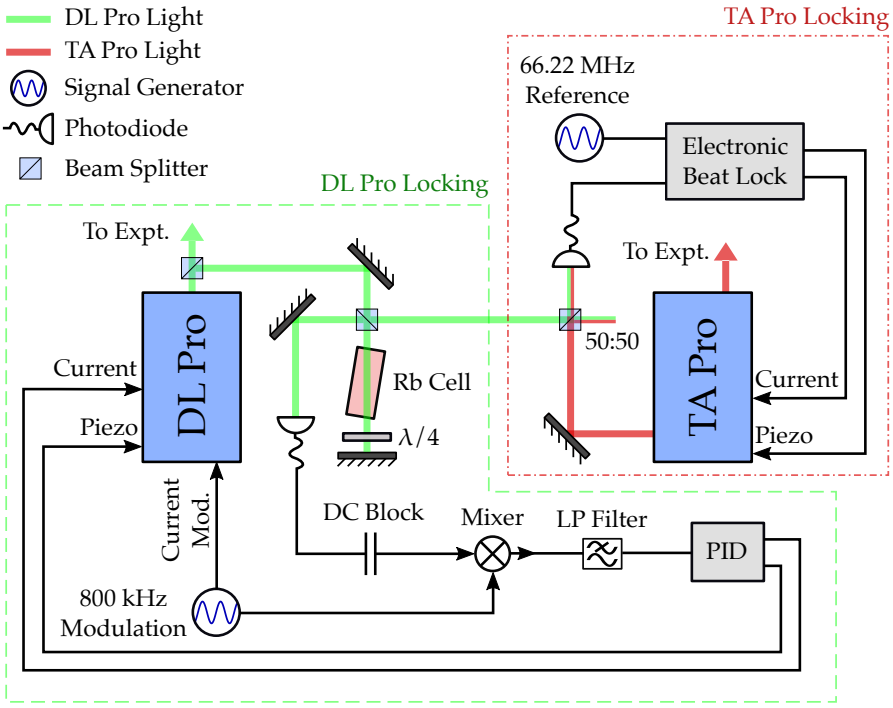

FIGURE 3.23: Overview of the frequency locking scheme: The DL Pro laser is stabilised by dithering of the diode current, which leads to a modulation of the transmission of intensity through a Rb vapour cell. The signal is demodulated again by mixing it with the original dithering frequency. This creates the derivative of the absorption profile, which is then used in a PID feedback loop to lock to an atomic transition. The TA Pro laser is then offset-locked with respect to the DL Pro by overlapping a portion of light from each laser to generate a beat signal. This provides a direct way to measure the frequency difference between the lasers, which is then compared to a reference and allows an error signal to be produced using a digital beat lock technique, closing the feedback loop.

The atomic reference frequency for locking the DL Pro laser is produced by performing Doppler-free saturation spectroscopy [184] on the rubidium vapour cell. This technique uses counter-propagating "pump" and "probe" beams to resolve the individual hyperfine transitions, which would otherwise be smeared out due to the Doppler shifts of the atomic velocities in the cell. The Doppler-free spectra obtained for rubidum are shown in Fig. 3.24.

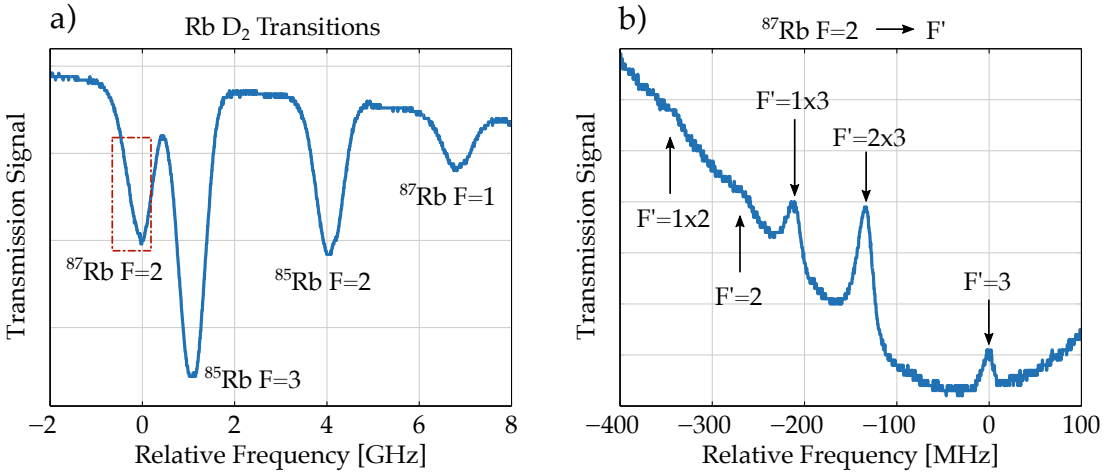

FIGURE 3.24: Spectroscopy signals in rubidium: a) The Doppler-broadened D_2 transitions for both rubidium 87 and 85 isotopes at $780\,\mathrm{nm}$ are measured by monitoring the transmission through a vapour cell whilst scanning the laser piezo over several gigahertz. b) A zoom of the boxed region indicated in (a), showing resolved hyperfine transitions using Doppler-free saturation spectroscopy in the ^{87}Rb $|F = 2\rangle \rightarrow |F'\rangle$ manifold. The TA Pro master laser is tuned to the $F' = 1 \otimes 3$ peak by using a digital offset lock with respect to the DL Pro.

In order to generate an error signal for the PID controller, the frequency of the laser is modulated by "dithering" the diode current with a low-amplitude $800\,\mathrm{kHz}$ sine wave. This dithering translates into a modulation of the intensity of light transmitted through the vapour cell, $I_\mathrm{T}(\nu)$, which is measured using a photodiode, and is directly proportional to the derivative of the absorption profile to first order. This can be seen by Taylor expanding the signal on the photodiode $S_\mathrm{pd}(t)$ after modulation at frequency Ω_m with amplitude A_m,

$$S_\mathrm{pd}(t) \propto I_\mathrm{T}(\nu + A_\mathrm{m} \sin \Omega_\mathrm{m} t) \simeq I_\mathrm{T}(\nu) + (A_\mathrm{m} \sin \Omega_\mathrm{m} t) \frac{dI_\mathrm{T}}{d\nu} + ... \tag{3.5}$$

The derivative of the modulated absorption profile is extracted by demodulating the signal using a lock-in amplifier [185]. The first term in Eq. 3.5 is removed using a DC block, and the photodiode signal is then mixed with the original dithering frequency Ω_m. The output of the mixer is then given by

$$S_\mathrm{O}(t) = A_\mathrm{m} \sin(\Omega_\mathrm{m} t + \phi_2) \times \left[A_\mathrm{m} \sin(\Omega_\mathrm{m} t + \phi_1) \frac{dI_\mathrm{T}}{d\nu} \right] \tag{3.6}$$

$$= B \frac{dI_\mathrm{T}}{d\nu} \left[\cos(\phi_1 - \phi_2) - \cos(2\Omega_\mathrm{m} t + \phi_1 + \phi_2) \right], \tag{3.7}$$

where arbirary phases have been allowed between the two signals. If $\phi_1 = \phi_2$ and the high frequency second term is removed using a low-pass filter, the result is a DC signal which is proportional to the derivative of the absorption profile, and this is shown in Fig. 3.25. In practice, since the lengths of the cables used for the locking electronics are

short when compared with the wavelength of the signal ($\sim 400\,\mathrm{m}$), there is no appreciable phase difference accumulated between the two waves sent into the mixer (i.e. $\phi_1 \approx \phi_2$), and therefore no explicit phase shifting is required in the setup. The derivative signal exhibits zeroes at the peaks of the absorption spectrum, and changes sign at either side. It can therefore be directly used as an error signal for the input of a PID controller, allowing locking of the frequency to a zero-crossing by closing the feedback loop to correct both the current of the laser diode and the piezo of the external cavity. This is a powerful locking method which is less susceptible to ambient laboratory $1/f$ noise, since the derivative signal is essentially shifted to the higher $800\,\mathrm{kHz}$ modulation frequency, and the laser is able to remain locked to an atomic resonance for many days.

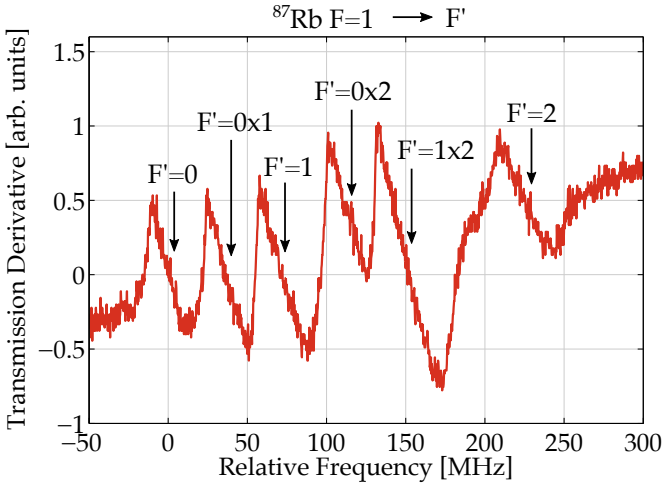

FIGURE 3.25: Derivative spectroscopy: The derivative of the atomic absorption profile for the ^{87}Rb $|F=1\rangle \rightarrow |F'\rangle$ manifold is obtained by dithering the DL Pro laser diode current and then subsequently demodulating the signal using a lock-in amplifier. This derivative curve is then used as the error signal for a PID controller, and the DL Pro laser is locked to the $|F=1\rangle \rightarrow |F'=1\rangle$ zero-crossing.

After stabilising the frequency of the DL Pro, the TA Pro laser is fixed by using an offset lock to maintain a frequency difference of $6.622\,\mathrm{GHz}$ between the two lasers. This is accomplished through a homebuilt digital "beat lock" box, the design of which is taken from [186]. A small $3\,\mathrm{mW}$ portion of each laser is picked off and overlapped on a 50:50 non-polarising beam splitter cube, and the resulting beat signal contains components at both the sum and difference of the absolute frequencies of the two lasers. The sum term is too fast for the electronics to detect, but the difference can be measured on a fast photodiode, and is then sent to the beat lock box. The device then electronically divides the signal down by a nominated factor (usually 100), and compares the result with a reference wave. A digital error signal is then produced which is used to close the feedback loop with the TA Pro piezo controller. A large tuning range can be achieved with this method - from several megahertz up to $7\,\mathrm{GHz}$. Additionally, the digital division process means that the reference need only be of the relatively low frequency

of $66.22\,\mathrm{MHz}$, and so no expensive microwave generator is required as in conventional offset locking schemes [187]. The disadvantage is that the box offers no proportional feedback, and so is limited in its ability to remove noise from the laser. However, we do not find this to be a problem so long as the reference laser is ensured to be spectrally stable.

Ultimately, the DL Pro is locked to the $|F = 1\rangle \rightarrow |F' = 1\rangle$ transition using the frequency modulation spectroscopy method, and then an offset of $6.622\,\mathrm{GHz}$ is applied to the TA Pro laser. This brings it into resonance with the $|F = 2\rangle \rightarrow |F' = 1 \otimes 3\rangle$ crossover peak, which allows all the desired frequencies of light in the experiment to be comfortably reached using AOMs.

Acousto-Optic Modulators

To allow fine control over the powers, frequencies and timings of the various beams, each one has its own dedicated acousto-optic modulator (AOM). These devices use a piezoelectric transducer to generate a radio-frequency sound wave inside a crystal through which the laser beam is passing. The sound wave creates a diffraction grating for the beam due to periodic modulation of the index of refraction, and the diffracted beam is deflected by a certain angle, with the frequency of the beam being shifted by exactly the frequency of the incident RF wave, due to conservation of energy and momentum between the scattering of phonons and photons. Having precise control of the RF wave therefore provides precise control of the frequency of the laser beams, with adjustments on the scale of below $1\,\mathrm{MHz}$ being possible. Our AOMs are setup in a double-pass configuration [188], whereby the change in deflection angle for a given frequency can be compensated on a second passage through the crystal. This allows alignment into the optical fibres to be maintained when applying different detunings at various stages of the experiment.

In addition, the power in the beams can be varied from 0% to 80%, by changing the amplitude of the incident RF wave. Controllable, sinusoidal RF waves are generated by a series of homemade AOM drivers, shown schematically in Fig. 3.26. Each one consists of a voltage-controlled oscillator (VCO), a voltage-controlled attenuator (VCA), an RF switch, and a power amplifier, allowing creation of $70\,\mathrm{MHz}$ to $150\,\mathrm{MHz}$ waves with up to $3\,\mathrm{W}$ of power. In total, there are 7 AOMs to control the frequencies of both the cooling and repumping beams for the 2D and 3D MOTs, as well as for optical pumping and absorption imaging. The frequency shifts applied to each corresponding AOM are given in Table 3.3.

The AOMs allow fast switching of the beams on a timescale of $\sim 10\,\mu\mathrm{s}$, but do not completely extinguish 100% of the light (contrast ratios are typically $> 60\,\mathrm{dB}$). Therefore, in order to completely shield the experiment from stray resonant light, a series of mechanical shutters are also used in series. These have a slower response time of

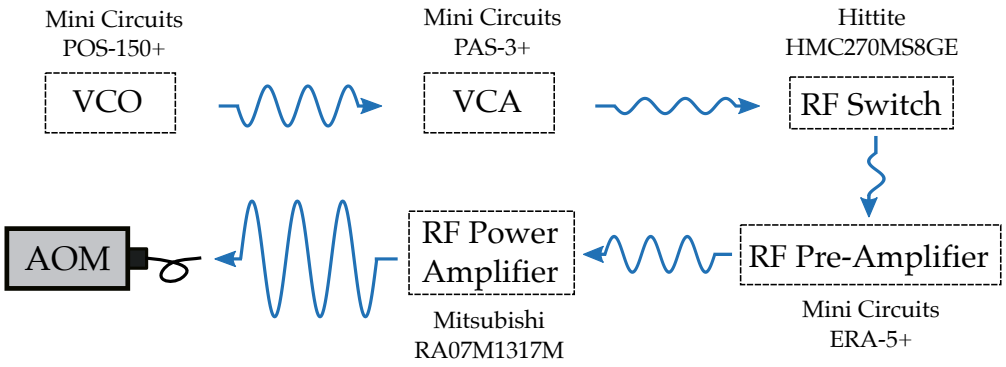

FIGURE 3.26: Schematic of the homemade AOM driver unit: A voltage-controlled oscillator (VCO) followed by a voltage-controlled attenuator produce a RF wave with adjustable frequency and amplitude, which can then be switched in $\sim 50\,\mathrm{ns}$ using an RF switch. A pre-amplifier follwed by a power amplifier then increase the power in the wave up to a maximum of 3 W, to then be delivered to the AOM crystal.

TABLE 3.3: Frequency shifts applied by AOMs in the experiment. The integer in parentheses indicates whether an AOM is configured in double or single passage.

AOM	$\Delta\nu$ [MHz]
3D MOT Cooler	$+(2\times)$ 98.0
3D MOT Repumper	$+(2\times)$ 78.5
2D MOT Cooler	$+(2\times)$ 100.5
2D MOT Repumper	$+(2\times)$ 78.5
2D MOT Pusher	$+(2\times)$ 112.5
Imaging	$+(2\times)$ 106.0
Optical Pumping	$-(1\times)$ 75.0

$\sim 5\,\mathrm{ms}$ but completely block all residual light. The shutters are mounted to the optical table on a layer of shock absorbing material,[12] which serves to damp vibrations produced when opening and closing.

3.5.2 Distribution and Optics

After the frequency of all light has been prepared using the respective AOMs, the beams are divided and distributed before being coupled into polarisation maintaining optical fibres to be transferred across to the vacuum chamber, as shown in Fig. 3.27. The use of optical fibres allows decoupling of the laser layout from the vacuum system, providing more flexibilty in beam positioning, whilst also serving to clean the spatial mode of the laser beams. It is important to precisely align the input linear polarisation with the fast axis of the optical fibre in order to minimise variation of the output polarisation. If this is not the case, then slight changes in ambient temperature in the

[12] Sorbothane

lab or any stress on the fibre can introduce a delay between the two orthogonal polarisation components, leading to some elliptical polarisation at the fibre output. We have observed this to cause slow oscillations in the MOT atom number, and so now ensure that changes in polarisation are kept below 5% which completely eliminates any detectable variations in atom number arising from polarisation. The output of the cooling laser tapered amplifier has an irregular elliptical mode shape, which leads to a coupling efficiency into the optical fibres of 50%. Therefore, to improve this a cylindical telescope is placed before the fibres to better match the two modes, increasing the efficiency to better than 70%.

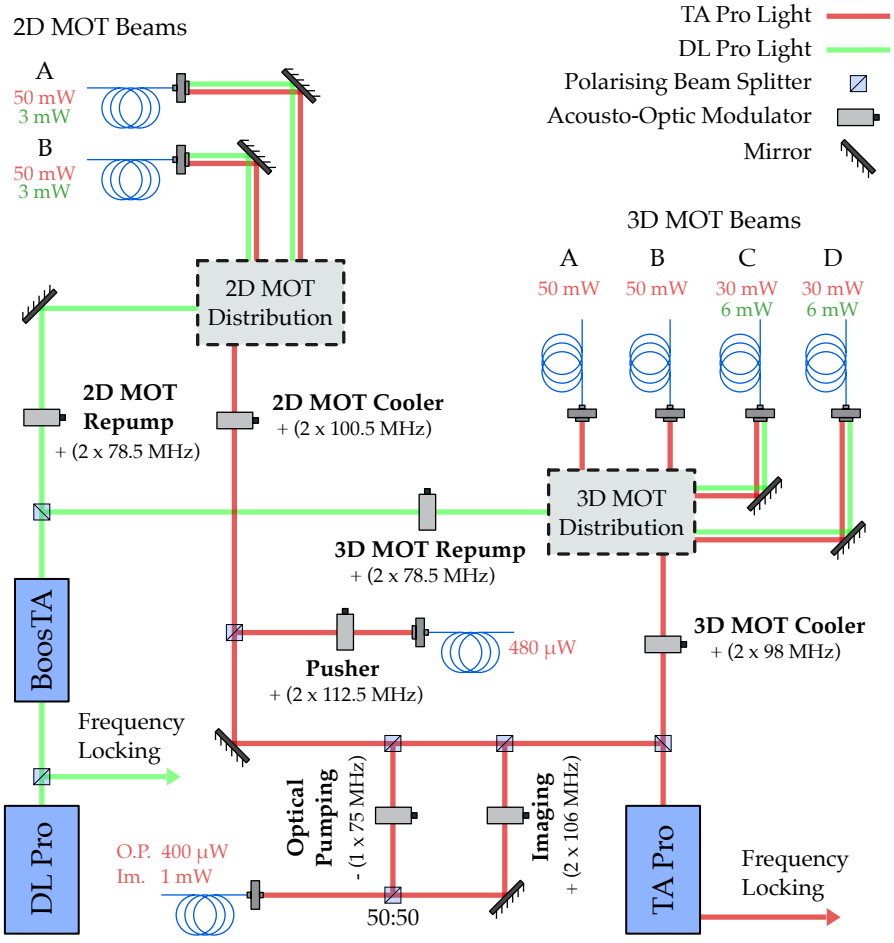

FIGURE 3.27: Schematic of preparation of the laser light: Light for both 2D and 3D MOT cooling, along with the 2D MOT pusher, optical pumping and imaging, are all derived from the TA Pro laser. Repumping light for the MOTs is provided by the DL Pro in series with a BoosTA amplifier. After each line has been frequency-prepared using the respective AOMs, the light is distributed using a series of polarising beam splitter cubes and half-wave plates. Finally, the light is coupled into a collection of optical fibres to be transferred to the vacuum system.

For the 3D MOT a total of 160 mW of cooling light is divided between four beams labelled A, B, C, and D, with a total of 12 mW of repumper shared equally between beams C and D. After exiting the fibres, the MOT beams are expanded and collimated with

a waist ($1/e^2$ radius) of $10\,\mathrm{mm}$, before being circularly polarized using quarter-wave plates to provide the correct arrangement for the MOT. The beams A and B enter the vacuum chamber through the large bottom viewport at an angle of $32°$ from the vertical and are reflected from the atom chip surface, as shown in Fig. 3.28. This is narrower than the optimal $45°$ as required for the mirror-MOT configuration due to geometrical constraints of the vacuum chamber, but does not seem to cause significant degradation of the MOT performance. The second pair of beams C and D are counter-propagating and travel parallel to the surface of the atom chip. Finally, the optical pumping and imaging light is also parallel to the chip surface but propagates orthogonally to the MOT beams.

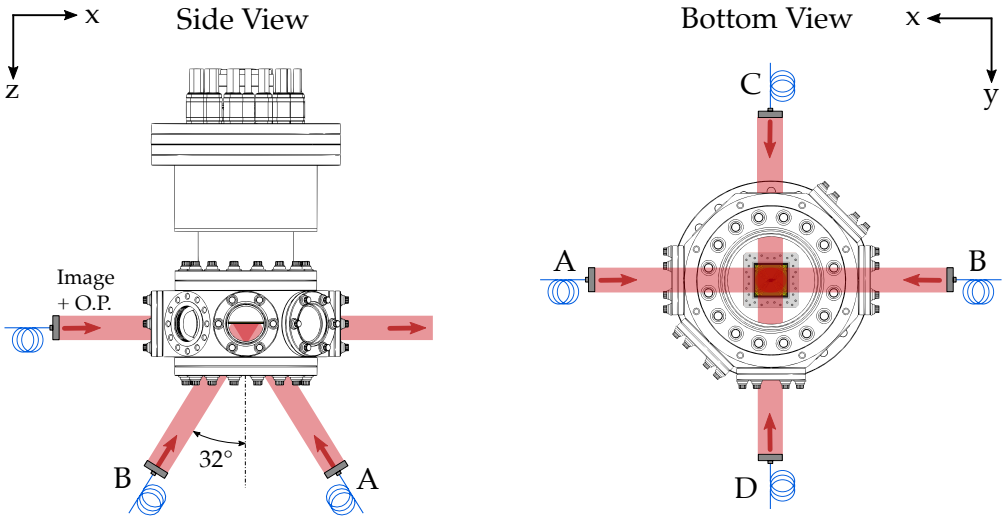

FIGURE 3.28: Beam arrangement: Four beams are required for laser cooling in the mirror-MOT configuration. The two beams A and B enter the bottom large DN100CF viewport at an angle of $32°$ to the vertical, and are reflected from the surface of the atom chip positioned at the centre of the vacuum chamber. Cooling in the third dimension is provided by beams C and D propagating parallel to the atom chip surface, which are also orthogonal to the optical pumping and imaging light.

3.6 Current Control

Throughout the experimental cycle, a variety of magnetic fields used for manipulating the atoms must be switched and ramped precisely. Magnetic fields are produced by passing current through both low inductance loads such as the atom chip wires and the understructure, as well as through the higher inductance external coils.

All currents except for the copper Z wire and the atom chip wires are controlled by an external analog input voltage to their respective power supplies. The power supplies are operated in constant voltage mode, since constant current mode is known to be inherently noisier for these types of PSUs. The problem with running power supplies in this way is that if the resistance of the load changes - for example, because

the contact resistance of the connecting wires varies, or if the temperature increases - then the current will change accordingly. For this reason, active feedback circuits are implemented for both the copper Z trapping wire and the atom chip wires, which ensures that the current flowing is always maintained at the desired value, regardless of changes in resistance. For the copper Z, the circuit is based around a precise $200\,\mathrm{A}$ Hall effect current transducer, which is used to measure the current and adjusts the gate voltage of a bank of MOSFETs via a feedback loop. For the atom chip wires, a similar circuit is implemented but using sense resistors instead since Ohmic heating of the sensor is not an issue at these relatively lower currents ($< 2\,\mathrm{A}$). The feedback for the atom chip wires is able to switch on and off a typical current of $1\,\mathrm{A}$ within $50\,\mu\mathrm{s}$ (90%-10%).

Fast Switch-Off for Inductive Loads

Imaging of the atomic cloud is performed after releasing atoms from the traps, and so magnetic fields must be removed quickly (non-adiabatically) in order to obtain the true in-situ momentum distribution. In practice, this means switching of fields should be achieved in less than $1\,\mathrm{ms}$, as we wish to take time-of-flight measurements from around $1\,\mathrm{ms}$ to $30\,\mathrm{ms}$.

Depending on the values of the currents to be switched, which can be from $\sim 100\,\mathrm{mA}$ up to $100\,\mathrm{A}$, we use home-made current shutters containing banks of up to 16 insulated-gate bipolar transistors (IGBTs) in parallel, controlled by an optoisolated TTL signal, a schematic of which is shown in Fig. 3.29 (a). External coils have a large inductance, on the order of $1\,\mathrm{mH}$, and so there is a significant amount of energy stored in the magnetic field which resists changes in current due to Lenz's law. When the switch is opened, there is a large negative voltage induced across the coil (known as *back EMF*, or *flyback voltage*), according to $V = -LdI/dt$. This voltage continues to rise until eventually arcing occurs across the IGBTs. The flyback voltage can be several thousands of volts for even a few amps through the coil, and would destroy the IGBTs (which have a nominal breakdown of $600\,\mathrm{V}$).

One solution to prevent the large back EMF is to add a simple flyback diode in parallel with the load, which provides an alternative route for the power to dissipate [189]. However, using this approach the induced voltage can only reach approximately one diode drop, and so the rate of change of current is small, leading to slow switch-off times of many milliseconds. To achieve faster switching the induced voltage must be allowed to build up as much as possible without causing damage to the IGBTs. For this reason, a small snubber circuit, shown in Fig. 3.29 (a), is placed around the coil using a zener diode. The purpose of this circuit is to allow the voltage spike to build up but capping it at $-200\,\mathrm{V}$, which is still below the damage limit for the IGBTs. An example of typical switching characteristics for the main X Bias coil are shown in Fig. 3.29 (b), illustrating a current of $50\,\mathrm{A}$ being reduced to zero in $260\,\mu\mathrm{s}$. In general, all

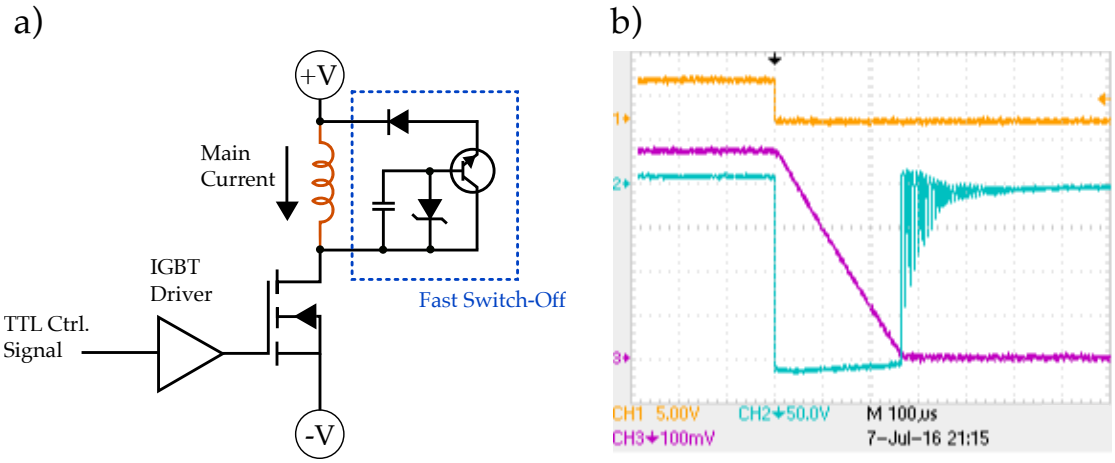

FIGURE 3.29: Fast switching-off of inductive loads: a) Schematic of the switching circuit. The main current through the coil is switched using a bank of IGBTs controlled by an optocoupled IGBT driver. Fast switch-off is achieved by a small snubber circuit around the load to provide dissipation of the energy stored in the magnetic field. The zener diode allows the induced flyback voltage to build up to $-200\,\mathrm{V}$, where it is capped. b) Performance of the switch for a $1\,\mathrm{mH}$ load. The current through the coil (magenta trace) falls from $50\,\mathrm{A}$ to zero in $260\,\mu\mathrm{s}$ after the trigger (yellow) is received. The induced voltage across the coil (cyan) rises and is capped at $-200\,\mathrm{V}$ by the snubber circuit. Since this voltage is constant, the current through the coil reduces linearly. As the coil's magnetic energy is gradually dissipated through the resistive elements in the circuit, the induced voltage drops until eventually below the zener breakdown voltage, and ultimately down to zero. The remnants of an RLC oscillation can also be seen at the very end.

currents fall to zero within seveal hundred microseconds, depending on their inductance. However, due to unavoidable metallic parts around the vacuum chamber, the fast changes in magnetic flux induce Eddy currents which lead to ringing fields lasting several milliseconds. A measurement of these Eddy currents is given in Fig. 3.38 in Sec. 3.9.1, and therefore, time-of-flight measurements are usually performed from $3\,\mathrm{ms}$ or longer release times.

3.7 Computer Control System

This experiment requires precise control over the operation of many different pieces of equipment in parallel. Primarily, we must control a wide range of currents passing through various wires and coils, along with manipulating the frequencies, powers, and timing of many different laser beams. In addition, other equipment such as radio-frequency sources, CCD cameras, rubidium dispensers, and mechanical shutters must be controlled. The timings required from these devies can range from several microseconds to several tens of seconds. Each experimental run must be automatically repeated hundreds of times, in order to build statistics.

This level of control is achieved with a client-server model system, in which the client is used to send details of a particular experimental sequence to a dedicated server, which then carries out the required tasks. The advantage of this setup is that the server can run a real-time operating system (RTOS), which provides consistency and reliability in terms of the timings of all processes, minimizing jitter between different events. Our sever is comprised of a National Instruments PXIe-1065 chassis, containing a PXIe-8130 controller and several analog and digital output cards. In total, we have four PXI-6733 analog output cards, providing a total of 24 analog channels which can output signals from -10 V to +10 V at a rate of 1 MHz, with 16-bit resolution. Should more channels be required, further cards can be added to this chassis with this modular system. Additionally, there is a single PXI-7813R digital output card with a built in FPGA, providing the ability to create pulses with a resolution of up to 1 µs across 80 channels. The FPGA also provides a clock signal along the backplane of the chassis, which is used to synchronize the outputs of all digital and analog channels to better than 100 ns. The PXI controller is running the LabVIEW Real-Time operating system, which is loaded with software written in-house in the C programming language by James Clewett and German Sinuco. The job of this software is to listen for commands sent by the client, and then interpret them to generate outputs on the various cards.

The client is a lightweight, standalone application named *PyPlayer*, and was written jointly by PhD students Matthew Jones and Asaf Paris-Mandoki in the Python programming language. The client application is loaded onto a standard desktop PC, and communicates with the PXI server through TCP/IP. PyPlayer provides a very nice graphical user interface for prototyping experimental sequences quickly and easily. It features the abiliy to construct modular groups of experimental events, which can be attached together depending what is being investigated. For example, one can create groups for producing a MOT, or for absorption imaging. PyPlayer also enables very useful visualization of the waveforms to be outputted on every channel, via the built-in *WaveViewer*. On analog channels, several different waveforms can be output, including linear ramps, sine waves, and exponential curves. There is also the functionality of defining variables, which can be scanned in a nested manner when optimizing over a particular parameter space. The PyPlayer also provides an interface for managing

which digital and analog channels are included in a particular sequence, allowing new pieces of equipment to be added and controlled very easily. More detailed information regarding both the client and the server can be found in the thesis of Matt Jones [190].

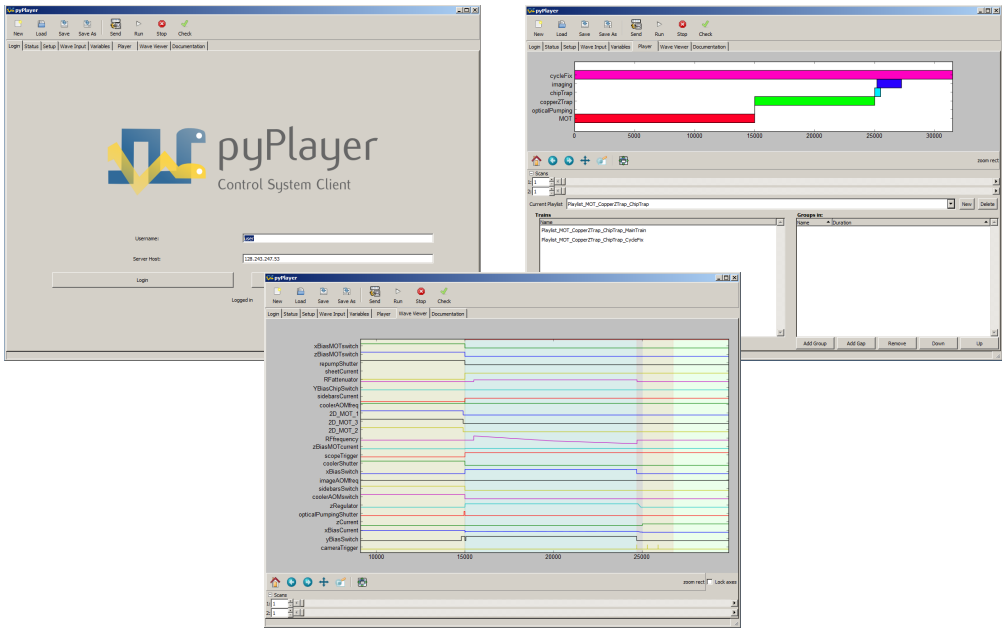

FIGURE 3.30: Control system client: The client, named *PyPlayer*, is a lightweight application developed in the programming language Python by PhD students at Nottingham to control the cold atom experiments. Left: The PyPlayer log-in window for connection to the server over TCP/IP. Centre: The *WaveViewer* allows visualisation of the timings and signals for all output channels in a particular sequence. Right: Each sequence can be contructed in a modular way, enabling elements to be created for producing a MOT or for absorption imaging, for example.

Output Isolation

Many different types of equipment are connected directly to the outputs of the PXI system, including both high voltage and high current power supplies and radio-frequency sources. It is extremely important to electrically isolate different devices, both from each other and from the PXI itself. This keeps various types of noise and ground loops to a minimum, and also protects the expensive PXI system from potential current or voltage spikes. Digital isolation is achieved through a custom box, designed and built by Bob Chettle of the School's electronics support group, which routes all digital channels thorough individual optocouplers before passing onto the various devices. The box can also provide additional current, for directly driving $50\,\Omega$ loads. Another feature is the ability to latch any channel to a desired state, which is useful for example to protect against spurious noise, or to keep a device switched on between cycles, such as keeping AOMs warm, for example. Decoupling of analog signals is typically more

difficult to achieve, but is implemented easily using a small circuit based on an isolation amplifier chip,[13] providing galvanic isolation between the input and output sides through digital modulation.

Image Analysis Application

After each experimental sequence, it is necessary to analyse and interpret the resulting absorption images on-the-fly. This is achieved through an image analysis tool, written in MATLAB by the author, which is running on the same PC as the control system client. The tool will import any new images whenever they are retrieved from the camera, and process them to obtain the optical density distribution. It provides the ability to select regions of interest, and will then apply various types of fitting routines - including both one and two dimensional Gaussian and bimodal fits. Results of the analysis, such as atom number, cloud size and position, condensate fraction, temperature, integrated profiles, and one-dimensional slices, are displayed to the user, along with a history of the previous experimental runs for comparison. A screenshot of the application is shown in Fig. 3.31.

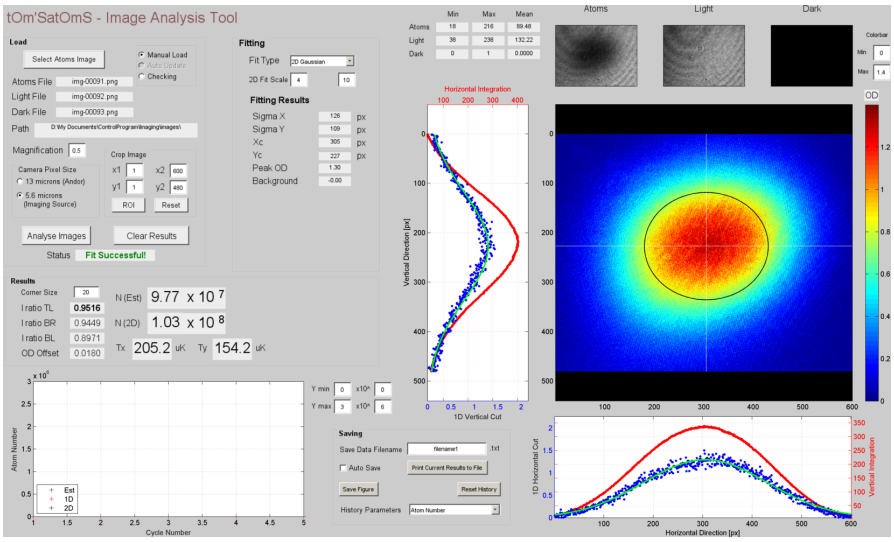

FIGURE 3.31: Screenshot of the image analysis tool: The tool automatically analyses any new absorption images in real time, and calculates the optical density. Various fitting routines can be applied to determine the physical properties of the atomic cloud.

[13] Texas Instruments ISO124

3.8 Forced Radio Frequency Evaporation

After laser cooling the atoms to below $1\,\mathrm{mK}$, the temperature is then further reduced by using evaporative cooling, a technique which is used in almost all experiments with degenerate quantum gases.[14] Forced evaporation, as described in Sec. 2.6, requires a way of regulating the depth of the trapping potential, allowing the most energetic atoms to escape. One way to achieve this is to simply lower the height of the external trapping field, as is usually implemented in optical dipole trap experiments [33, 192, 193], and as well with the early efforts at evaporation in magnetic traps [194]. However, this is ordinarily accompanied by a relaxation of the trapping frequencies, which in turn reduces the elastic collision rate and therefore the evaporation efficiency. Several techniques have been developed to overcome this inherent decompression and reach the runaway regime in optical traps [195, 196]. For magnetic traps, although it is possible to lower the depth whilst maintaining the trapping frequencies [197], almost all experiments employ a form of radiative evaporation, whereby an external radiation field is used to induce transitions of atoms into anti-trapped states [10].

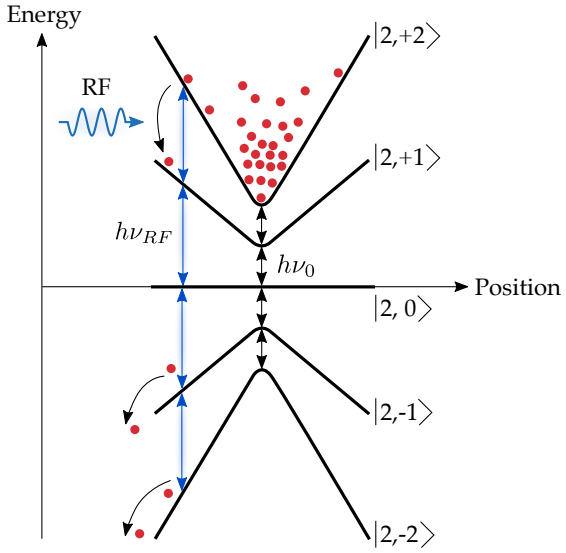

FIGURE 3.32: Working principle of radio frequency evaporative cooling: The magnetic substates of the $|F = 2, m_F\rangle$ ground state manifold are shown. Atoms initially trapped in the state $|2, +2\rangle$ can be outcoupled from the trap by driving transitions to anti-trapped levels through the use of an external radio frequency field with frequency ν_{RF} and polarisation such that it is able to drive $\Delta m_F = \pm 1$ transitions. Only atoms with energy above a certain threshold are able to climb the potential and sample those magnetic fields which Zeeman shift the atomic levels into resonance with the RF, and therefore the method is energy-selective.

In this experiment, magnetic trapping takes place in the $F = 2$ ground state, and the energy levels of the five magnetic substates are shown in Fig. 3.32, where initially all atoms are prepared in the $|F = 2, m_{\mathrm{F}} = +2\rangle$ low field-seeking state. Application

[14] Although a BEC has been created using laser cooling as the only cooling mechanism [191].

of a correctly polarized oscillating magnetic field is able to transfer atoms from the $|+2\rangle$ state and cascade them down to the lower states. Atoms in the lower states are then repelled from the field minimum and are lost from the trap. Crucially, only the atoms are addressed which sample a position in space where the magnetic field is able to Zeeman shift the levels into resonance with the external field. Since only the most energetic atoms are able to sample the spatial regions of highest magentic field, this technique is energy-selective and allows the "hottest" atoms to be outcoupled from the trap. For the temperature ranges and magnetic fields used in the experiment, the frequency of the external radiation required to match the separation of neighbouring Zeeman levels is in the radio frequency range. Specifically, atoms are brought into resonance when the magnetic field, $|\boldsymbol{B}|$, which they experience fulfills the condition

$$g_F \mu_B |\boldsymbol{B}| = h\nu_{\mathrm{RF}}, \tag{3.8}$$

where ν_{RF} is the frequency of the RF radiation. The truncation energy ϵ_t (trap depth), as described in Sec. 2.6, of the $|+2\rangle$ state is then given by

$$\epsilon_t = 2h(\nu_{RF} - \nu_0), \tag{3.9}$$

where ν_0 is the RF frequency which corresponds to the energy separation of the states due to the magnetic field at the trap minimum, and is given by

$$\nu_0 = \frac{g_F \mu_B |\boldsymbol{B}_0|}{h} = 0.70\,\mathrm{MHz\,G^{-1}}. \tag{3.10}$$

Since this process removes particles with energy greater than ϵ_t, it is often referred to as an *RF knife*. The truncation parameter, which determines how far into the energy distribution to cut and determines the efficiency of the cooling process, can therefore be written as

$$\eta = \frac{\epsilon_t}{k_B T} = \frac{2h(\nu_{RF} - \nu_0)}{k_B T}, \tag{3.11}$$

where T is the temperature of the gas. The dynamics of evaporation is then governed by the trajectory of the frequency and amplitude of the RF radiation, $\nu_{\mathrm{RF}}(t)$, which is experimentally simple to control. For an RF frequency of $20\,\mathrm{MHz}$ and a trap bottom of $|\boldsymbol{B}_0| = 1\,\mathrm{G}$ the corresponding truncation energy is $\epsilon_t/k_B = 1.8\,\mathrm{mK}$, which corresponds to $\eta = 6$ for a $300\,\mu\mathrm{K}$ cloud.

Forced RF evaporative cooling is an extremely powerful method, since it decouples the reduction of the trap depth from the trap frequencies, allowing optimum confinement to be maintained throughout the process. In addition, for the Ioffe-Pritchard traps used, the resonance condition of Eq. 3.8 is fulfilled at an ellipsoidal shell of equipotential energy around the minimum, and so the evaporation occurs efficiently in three-dimensions.

RF Setup

After loading the magnetic trap, the temperature of the gas is several hundred µK and must be reduced to below 1 µK. This requires an oscillating RF field that can be ramped from $\sim 20\,\mathrm{MHz}$ down to below 1 MHz, according to Eq. 3.11.

The RF setup is shown schematically in Fig. 3.33. The initial signal is produced by a function generator,[15] which has a modulation input to allow control of the frequency using an external analog voltage. In this way, the frequency ramp of the RF can be arbitrarily shaped from 20 MHz down to D.C. Precise timing of switching the signal is achieved by then passing through an RF switch,[16] which has a typical switching time of 10 ns and an isolation of -100 dB up to 1000 MHz.

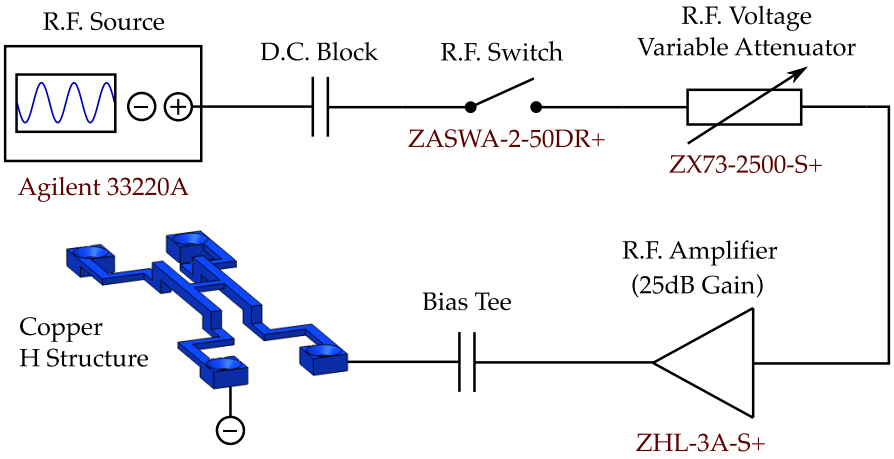

FIGURE 3.33: Schematic for producing the RF field for evaporative cooling: A signal generator creates the initial wave, with a frequency which can be modulated from 20 MHz to D.C. via an external analog voltage input from -5 V to +5 V. Fast timing and amplitide control are provided by an RF switch and voltage variable attenuator in series. The power is then increased with an RF amplifier with a fixed 25 dB gain. The wave is then coupled through a bias tee into the copper H structure, which sits below the chip surface and acts as an antenna.

As well as controlling the frequency, the amplitude must be reduced as the cloud becomes colder to ensure that the linewidth of the resonance is not power-broadened, which would reduce the resolution of the RF knife. Arbitrary control of the amplitude in parallel with the frequency is obtained with the use of a voltage variable attenuator,[17] providing attenuation from around -80 dB to -3 dB (corresponding to insersion loss of the device) according to an external analog voltage from 0 V to +10 V.

Sufficient power is then obtained by passing the signal through an RF amplifier[18] with a fixed gain of +25 dB and a maximum output power of +30 dBm. A home-made bias tee is then used to couple the amplified wave into the copper H structure in a U-configuration, and ensures that the high D.C. currents of $\sim 100\,\mathrm{A}$ used for magnetic

[15] Agilent 33220A
[16] Mini-Circuits ZASWA-2-50DR+
[17] Mini-Circuits ZX73-2500-S+
[18] Mini-Circuits ZHL-3A-S+

trapping cannot back-couple into the RF line. The copper U conductor acts as an antenna, located at $\sim 1-2\,\mathrm{mm}$ from the trapped cloud, and produces a linearly-polarized RF magnetic field oriented only in the directions perpendicular to the static Ioffe field at the trap minimum (the quantisation axis for the atoms). This leads to $\sigma^{\pm}$ - polarized radiation capable of driving $\Delta m_F = \pm 1$ transitions, and therefore the energy-selective spin-flip trap loss necessary for efficient evaporative cooling.

Although everything in the RF line - from the signal generator up to the output of the amplifier - is impedance-matched to $50\,\Omega$ for maximum power transfer, no special care was taken to match from the outside of the vacuum feedthrough onwards. This may lead to reflections and resonances in transmitted RF power as a function of frequency. We attemped to identify any features in the power radiated from the U-wire antenna by using a $5\,\mathrm{cm}$ pick-up coil (which was measured to have an almost flat response over the frequency range of interest) close to the bottom window of the science chamber, together with a network analyser. The signal picked up by the coil indicated that indeed the transmission depends on the frequency, with less efficent coupling at the lower frequency values. It is difficult in practice to correct for this, for example by using a tailored impedance-matching circuit, and so observing the behaviour of the cold atomic clouds and the efficiency of evaporative cooling is ultimately the best way to ascertain if sufficient RF power is being radiated.

3.9 Absorption Imaging

Quantitative information about the trapped gases is almost exclusively gained through absorption imaging [198], a technique which relies on the interaction with an incident probe beam. Imaging is typically performed using a *time-of-flight* measurement, whereby the trapping potential is switched off non-adiabatically and the cloud is allowed to expand ballistically for some variable time before the imaging process is executed. For short expansion times the measured density provides details of the in-situ spatial distribution, whilst for longer times-of-flight the density reflects the momentum distribution in-trap.

3.9.1 Imaging System

The technique of absorption imaging requires an incident resonant probe laser beam to be directed at the cloud, creating a shadow as atoms absorb some of the light. This shadow must then collected by an imaging system, and recorded onto a charge-coupled device (CCD) camera. A custom optical microscope is used to implement this scheme, consisting of two lenses[19] L_1 and L_2 with focal lengths f_1 and f_2, with the arrangement shown schematically in Fig. 3.34.

[19] All lenses are 2" diameter spherical achromatic doublets, anti-reflection coated for near-infrared wavelengths (Thorlabs AC508-xxx-B-ML, where xxx denotes focal length in mm).

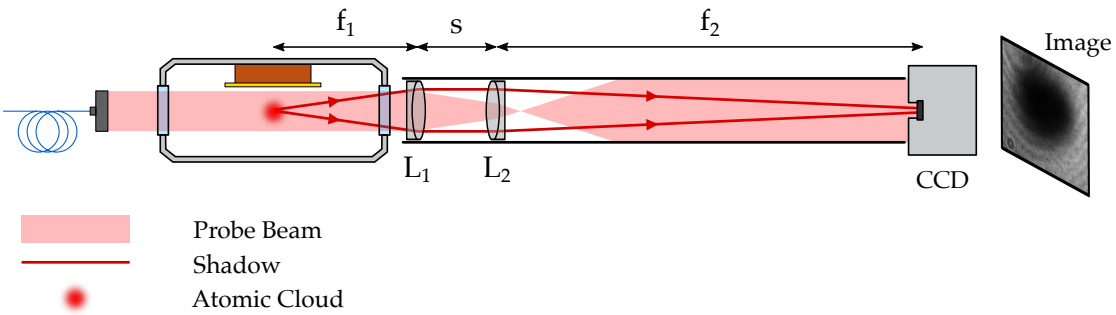

FIGURE 3.34: Schematic of the absorption imaging system : The collimated probe beam propagates through the vacuum chamber parallel to the surface of the atom chip and is directed at the atomic cloud, which then partially absorbs the light and creates a shadow image. A two-lens optical telescope then focusses the shadow image onto the CCD of the camera. The first lens L_1 is positioned close to the vacuum viewport to maximise the numerical aperture and resolution of the optical system. With the first lens fixed, the focal length of the second lens L_2 then determines the resultant magnification.

The spatial extent of the cloud in our experiment ranges from $\sim 5\,\mathrm{mm}$ for atoms initially captured in the MOT, down to $\sim 10\,\mu\mathrm{m}$ for expanded BECs. To be able to image clouds on such varying scales we use two different cameras, each with its own eyepiece lens, providing the ability to quickly switch between two different magnifications depending on the stage of the experimental cycle. Generally, the first magnification system is used for imaging larger clouds of atoms in the MOT and macroscopic copper Z-trap, whilst the second is used for the smaller atom chip traps. In both cases, imaging is performed perpendicular to the central part of the Z-wire traps, yielding the integrated density along the transverse direction of the clouds. We plan to implement additional imaging along the longitudinal axis of the traps in the near future.

To maximise the numerical aperture of the system, thereby increasing the obtainable resolution, the first lens L_1 should be as close to the atomic cloud as possible. In our case this is limited by the distance between the centre of the vacuum chamber and the viewport, which is $10\,\mathrm{cm}$, so the focal length of the first lens is chosen to have $f_1 = 150\,\mathrm{mm}$, and is shared by both imaging systems. An estimate of the maximum achievable resolution r_{max} of the system is given by the Rayleigh criterion $r_{\mathrm{max}} = 1.22\lambda f/D = 0.61\lambda/NA$, where λ is the radiation wavelength, and f, D and NA are the focal length, diameter and numerical aperture of the first lens, respectively. For our system, this gives a diffraction-limited resolution of $2.8\,\mu\mathrm{m}$. However, in practice there are other geometrical limitations such as the diameter of the vacuum viewport and clipping due to the atom chip surface, in addition to imperfect optics and probe light. These factors lead to a resolution of $\sim 10\,\mu\mathrm{m}$, measured using a resolution target. The focal length of the second lens L_2 then determines the resultant magnification. A large translation stage is used to mount the entire imaging system, allowing the camera to be moved together with both lenses in order to accurately focus the image of the cloud, as shown in Fig. 3.35. We simply take the in-focus position to be the point

at which the image is sharpest (other more sophisticated techniques are available, that allow the focal plane to be determined at the $\approx 2\,\mu\mathrm{m}$ level [199]). It can be seen that the appearance of a larger ring-like structure is an indication that the imaging system is out-of-focus.

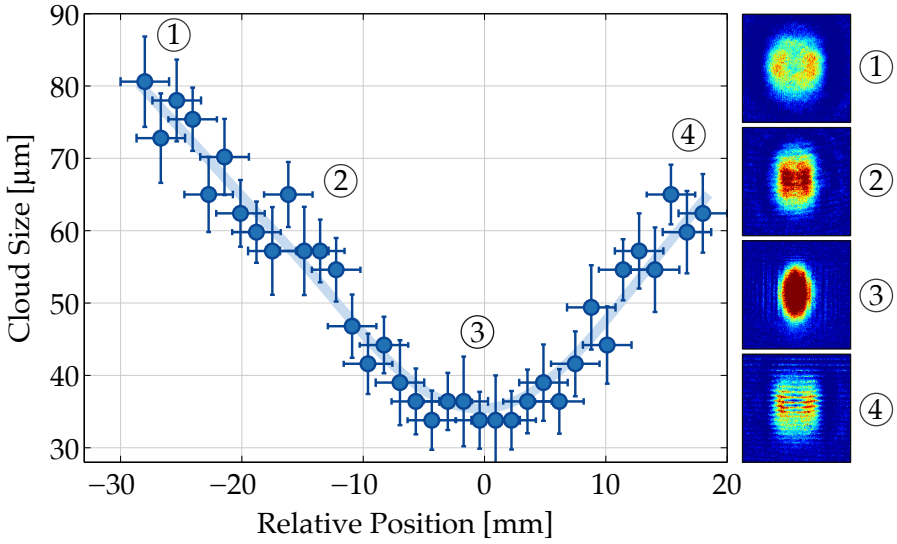

FIGURE 3.35: Imaging System Focussing: By translating the entire telescope-camera system over a distance of $\sim 5\,\mathrm{cm}$, the image of the cloud can be accurately focussed onto the CCD chip. The plot shows the RMS width of images of a BEC after $25\,\mathrm{ms}$ time-of-flight as a function of the position relative to the optimally focussed case, which is taken to be when the extension is minimised [200]. Examples of optical density images are shown in the right-most panels at various indicated relative focus positions, with the sharpest image ("in-focus") shown in ③.

Imaging Source System

Images of the larger clouds are captured with an Imaging Source DMK-21BF04 8-bit CCD camera. The sensor of this camera is 640×480 pixels, with each pixel being $5.6\,\mu\mathrm{m}$ square. Since this camera is used primarily for imaging larger clouds, the second lens is chosen to have $f_2 = 75\,\mathrm{mm}$, leading to a demagnification of $m \sim 0.5$ in order to match the size of the image to that of the CCD.

This system is arranged in the so-called *afocal* scheme [201], and the centre-to-centre distance between the lenses is set to be $238\,\mathrm{mm}$, calculated using a ray-tracing optical design software package.[20] This ensures that the resultant magnification is constant even if the cloud moves out of the focal plane of the first lens - indeed, the MOT can be offset by up to several millimeters from the geometric centre.

Andor System

Cold clouds and BECs are imaged using the more sophisticated Andor iXon3 888 16-bit camera. This camera has an electron-multiplying (EM) CCD, which multiplies the

[20] Optics Software for Layout and Optimization (OSLO)

signal before reaching the amplifier thereby reducing read-out noise. The sensor has an array of 1024 $\times$ 1024 pixels of size $13\,\mu\mathrm{m}$ square, and a quoted quantum efficiency of 82% at $780\,\mathrm{nm}$. Peltier cooling is used to reduce the sensor temperature down to $-80\,^{\circ}\mathrm{C}$, which reduces the dark current noise produced by thermal electrons.

For this system the second lens has $\mathrm{f}_2 = 750\,\mathrm{mm}$, giving a larger zoom of $m \sim 5$, and so employment of the afocal arrangement in this case would require a total length of the imaging system of $\sim 1.65\,\mathrm{m}$, which is impractical to implement. However, since the cold trapped clouds are always localised and geometrically centred on the atom chip, the image is always collimated between the lenses, as shown in Fig. 3.34, and we can allow a more compact separation of $s = 10\,\mathrm{cm}$.

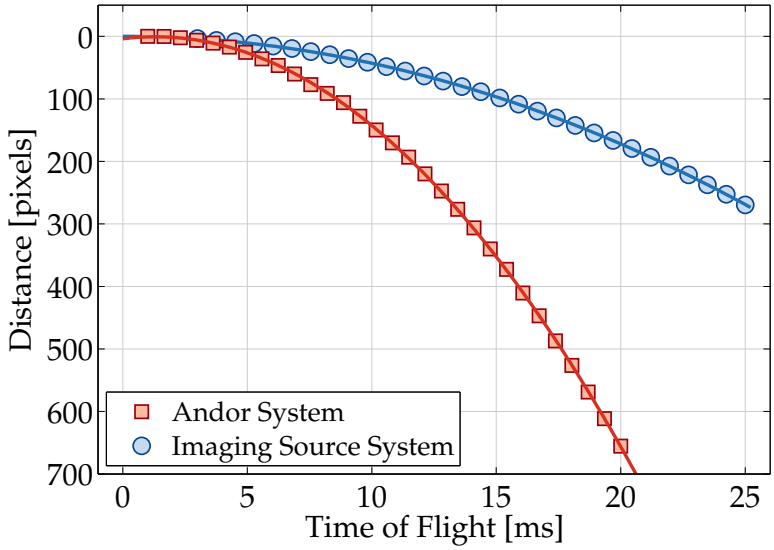

FIGURE 3.36: Magnifications of each of the imaging systems: A cold ($\sim 1\,\mu\mathrm{K}$) cloud is prepared using evaporative cooling in the magnetic trap, and then the position of the centre-of-mass is measured in pixels on the CCD image as a function of drop time. By fitting the result according to the influence of gravitational acceleration (solid lines), the magnifications are measured for the Andor (□) and Imaging Source (○) systems to be 4.97 ± 0.04 and 0.489 ± 0.004, respectively.

The magnification of each imaging system is precisely measured by switching off the trapping potential and observing the centre-of-mass position of a cold cloud as it falls during time-of-flight. The data are fit according to free-fall gravitational acceleration, with allowance for a non-zero initial velocity in the event that the cloud receives a kick during the switch-off process (although this effect appears to be negligible). The result is shown in Fig. 3.36 indicating magnifications of 0.489 ± 0.004 and 4.97 ± 0.04 - and therefore object plane pixel sizes of $11.5\,\mu\mathrm{m}$ and $2.6\,\mu\mathrm{m}$ - for the Imaging Source and Andor systems, respectively.

Probe Beam

A single-mode optical fibre is used to spatially filter the imaging probe beam, which is then expanded at the output to a waist ($1/e^2$ radius) of $w = 10.2\,\mathrm{mm}$. The collimated beam is then directed at the atomic cloud parallel to the atom chip surface in the UHV chamber through a viewport, as can be seen in Fig. 3.34. The total power in the beam is $P = 1\,\mathrm{mW}$, leading to a peak intensity of $I_0 = 2P/\pi w^2 = 0.6\,\mathrm{mW\,cm^{-2}} \approx 0.3\ I_{\mathrm{sat}}$. The effect of the imaging beam intensity on the signal-to-noise ratio for a given atomic density is characterised in detail for our system in [119]. The atomic cloud is illuminated with a flash of the probe light, whilst the shadow is exposed onto the camera's CCD. Typically, pulse durations from 30 - 100 µs are used, chosen as a compromise between exploiting a large dynamic range of the camera versus the unwanted recoil blurring of the image that results from long imaging pulses.

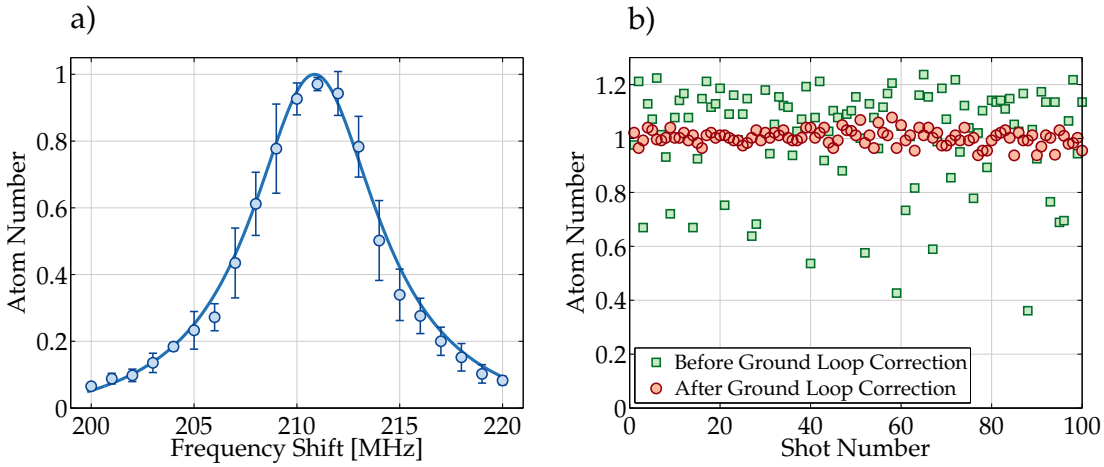

FIGURE 3.37: Control of the imaging probe laser frequency: a) The response of the normalised atom number as a function of the frequency shift applied to the probe beam using the AOM. The data are fit with a Lorentzian profile, with the resulting on-resonance frequency shift being (210.8 ± 0.6)MHz, which agrees with the value expected from the location of the laser lock point. Uncertainty bars represent the standard error after several scans. b) The presence of electrical ground loops causes the noise on the frequency of the probe laser, which leads to shot-to-shot fluctuations in the normalised atom number on the order of $\pm 22\%$ around the mean value (□). Taking precautions to minimise these ground loops reduces the fluctuation to $\pm 2.8\%$ (○).

In order to obtain the maximum signal-to-noise ratio, the probe light must be resonant with the atomic transition frequency, as shown earlier. Using an AOM in double-passage configuration the probe frequency can be precisely detuned, and an example of the response of the detected signal is shown in Fig. 3.37 (a). The profile is Lorentzian with a measured linewidth (FWHM) of $6.8\,\mathrm{MHz}$, which is the result of a combination of the natural atomic linewidth ($6.1\,\mathrm{MHz}$) and the linewidth of the probe laser ($< 1\,\mathrm{MHz}$). The larger error bars on the slopes of the profile due to shot-to-shot variations are indicative of either ambient magnetic field noise or laser frequency noise. Indeed, it was

found that the presence of electrical ground loops around the laser controller electronics and locking equipment gives rise to a large amount of $50\,\mathrm{Hz}$ mains noise on the probe laser, as measured by using a scanning Fabry-Perot interferometer to transform the frequency variations into intensity variations. By taking care to eliminate these ground loops, the shot-to-shot variation in detected atom number was drastically improved from $\pm 22\%$ down to $\pm 2.8\%$, as illustrated in Fig. 3.37 (b).

Finally, it is important to note that the resonant frequency for imaging as described in Fig. 3.37 (a) will in general depend on the magnetic field at the position of the atomic cloud. Most frequently, measurements are performed after some time-of-flight, with the magnetic trapping fields removed. However, a fast changing magnetic flux leads to Eddy currents circulating in the metallic parts of the vacuum structure, which then Zeeman shifts the position of the atomic resonance for short times-of-flight. In order to obtain reliable atom numbers the cloud should typically be imaged after at least $2\,\mathrm{ms}$ time-of-flight, after which time the resonance stabilises as shown by the measurement in Fig. 3.38.

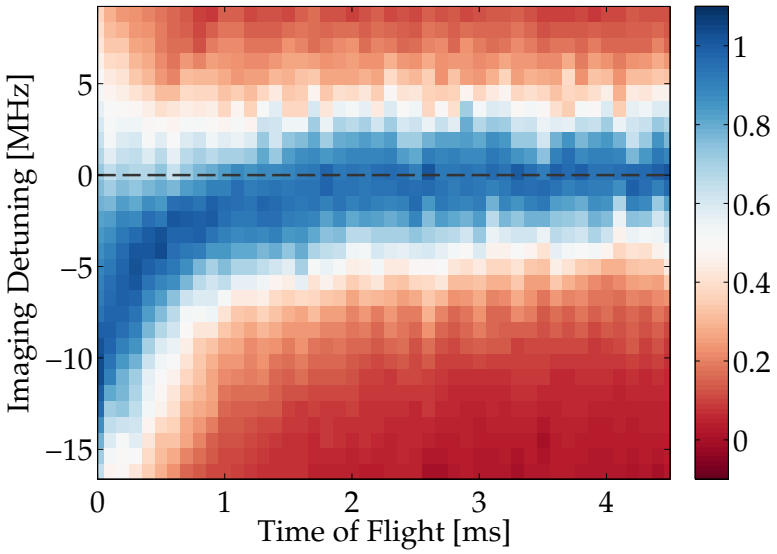

FIGURE 3.38: Measurement of Eddy currents after trap switch-off. A cloud in the atom chip trap is pre-cooled to $1\,\mu\mathrm{K}$, and the normalised detected absorption signal is recorded as a function of both imaging beam detuning and time-of-flight. After $\sim 2\,\mathrm{ms}$ the magnetic fields subside and the resonance returns to its zero-field value.

3.9.2 Quantitative Analysis of Absorption Images

The effect on the probe beam when passing through the gas of atoms can be seen by modelling the cloud as a dielectric medium with a complex refractive index $\tilde{n} = n_\mathrm{r} + i n_\mathrm{i}$. The complex electric field of a probe beam propagating along the z-direction through

the gas can then be written as the modified plane wave

$$\boldsymbol{E}(z,t) = \boldsymbol{E}_0 \exp\Big(i\left[\tilde{n}k_0 z - \omega_{\mathrm{L}} t\right]\Big) = \boldsymbol{E}_0 \underbrace{\exp\Big(-n_{\mathrm{i}}k_0 z\Big)}_{\text{attenuation}} \underbrace{\exp\Big(i[n_{\mathrm{r}}k_0 z - \omega_{\mathrm{L}} t]\Big)}_{\text{phase shift}}, \quad (3.12)$$

where $k_0 = \omega_{\mathrm{L}}/c$ is the vacuum wave number of the laser. This illustrates that there are two effects on the probe beam. The real part of the refractive index leads to dispersion, imparting a phase shift onto the beam, whilst the imaginary part attenuates the amplitude. The phase-shift can be used to perform other types of non-destructive imaging [202], whereas absorption imaging exploits the attenuation effect.

Since the laser intensity is given by $I = \frac{1}{2}\epsilon_0 c|\boldsymbol{E}|^2$, it can be seen from Eq. 3.12 that the intensity of probe beam decays with distance z through the medium according to $I_{\text{out}} = I_{\text{in}} e^{-az}$, where the absorption coefficient is given by $a = 2n_{\mathrm{i}}k_0$, and the fractional intensity lost is governed by the differential equation

$$\frac{dI}{dz} = -aI. \quad (3.13)$$

To interpret the measured change in intensity and relate it to the physical properties of the cloud, an expression for the absorption coefficient must be derived, which can be done by connecting the microscopic properties of the gas with $\tilde{n}$. The refractive index is given by $\tilde{n} = \sqrt{\epsilon_{\mathrm{r}}\mu_{\mathrm{r}}} \approx 1 + \frac{\chi_e}{2}$, where ϵ_{r} and μ_{r} are the relative permittivity and permeability of the gas, respectively, and χ_e is the frequency-dependent electric susceptibility [132]. Here, an approximation assumes that the magnetic response is negligible for the optical frequencies considered ($\mu_{\mathrm{r}} \approx 1$), and in addition that $|\chi_e| \ll 1$ for a dilute vapour. For an homogeneous and isotropic medium, the polarisation is aligned with the applied electric field, and assuming a linear response is given by $\boldsymbol{P} = \epsilon_0 \chi_e \boldsymbol{E}$. Treating the cloud as an ensemble of two-level atoms, the polarisation is defined as the number density of electric dipoles, and so can also be expressed in the form $\boldsymbol{P} = n_{3\mathrm{D}}\boldsymbol{d} = n_{3\mathrm{D}}\alpha\boldsymbol{E}$, where $n_{3\mathrm{D}}$ is the density of the gas, and $\boldsymbol{d}$ is the induced dipole moment of a single atom. For the two-level atom, the semi-classical scalar polarisability is given by [34, 203]

$$\alpha = \frac{3\epsilon_0\lambda^3}{4\pi^2} \frac{i - (2\delta/\Gamma)}{1 + (2\delta/\Gamma)^2 + I/I_{\text{sat}}}, \quad (3.14)$$

where λ is the transition wavelength, Γ is it's natural linewidth, δ is the probe beam detuning, and I_{sat} is the polarisation-dependent saturation intensity. In this way, the connection is made between the microscopic and macroscopic quantities through $\tilde{n} = 1 + n_{3\mathrm{D}}\,\alpha/2\epsilon_0$. The real and imaginary parts of the refractive index are shown in Fig. 3.39, illustrating that on-resonance the probe beam is most strongly attenuated, whilst the dispersive effect dominates for larger detunings. For this reason the probe beam is set to resonance in the experiment, since this maximises the signal-to-noise ratio and simultaneously reduces the phase-shift imparted onto the beam.

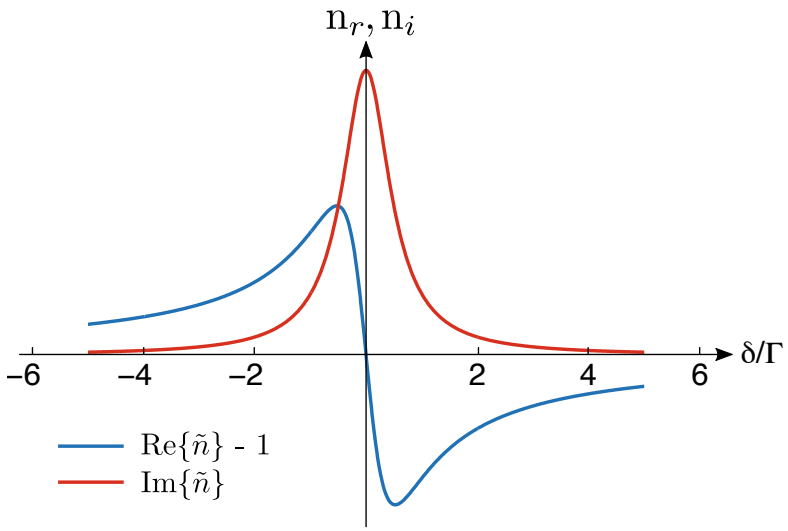

FIGURE 3.39: Refractive index of a gas of two-level atoms: The normalised real and imaginary parts of the refractive index n_r and n_i are illustrated for a probe beam with $I/I_\mathrm{sat} = 0.1$, as a function of reduced detuning δ/Γ. The imaginary part is responsible for attenuation of the beam during absorption imaging, and has the Lorentzian profile characteristic of resonance behaviour. The dispersive real part causes a phase shift of the beam, which is minimised when imaging is performed with zero detuning.

The absorption coefficient describing the probe beam's attenuation is then directly related to the atomic polarisability through

$$a = 2k_0 \,\mathrm{Im}\,\{\tilde{n}\} = \frac{k_0}{\epsilon_0} n_\mathrm{3D} \,\mathrm{Im}\,\{\alpha\} = n_\mathrm{3D} \cdot \sigma_\mathrm{sc}(\delta, I), \tag{3.15}$$

where the scattering cross-section $\sigma_\mathrm{sc}(\delta, I)$ describes the probability of absorption of an incoming photon

$$\sigma_\mathrm{sc}(\delta, I) = \frac{\sigma_0}{1 + (2\delta/\Gamma)^2 + I/I_\mathrm{sat}} \tag{3.16}$$

with $\sigma_0 = 3\lambda^2/2\pi = 2.90 \times 10^{-13}\,\mathrm{m}^2$ being the on-resonance, weak intensity absorption cross-section. In general, there are many magnetic sublevels and the scattering rate depends on the polarisation of the probe beam. However, if σ^+ light is used on the transition $|F = 2\rangle \rightarrow |F' = 3\rangle$, i.e. circular polarisation and propagation direction parallel to the quantisation axis of the atoms, then it is only possible to drive $m_F = +2 \rightarrow m'_F = +3$, and the simplicity of the two-level atom is recovered. In the case in which unpolarized light is used (or equivalently without the use of a magnetic field providing an explicit quantisation axis), then all possible transitions must be considered, which involves taking an average of the relevant Clebsch-Gordan coefficients and ultimately a modification of the scattering cross section by a factor of 7/15 [204].

The experimentally accessible quantity is the resultant intensity after passing the entire extension of the cloud. Inserting Eq. 3.15 into Eq. 3.13 and integrating over the

length of the gas in the z-direction gives an expression for the column density

$$n_{2D}(x,y) = \frac{1}{\sigma_0}\left\{\ln\left(\frac{I_{in}(x,y)}{I_{out}(x,y)}\right)\left[1+\left(\frac{2\delta}{\Gamma}\right)^2\right] + \frac{I_{in}(x,y)-I_{out}(x,y)}{I_{sat}}\right\}, \tag{3.17}$$

where

$$n_{2D}(x,y) = \int\limits_{\substack{\text{cloud}\\ \text{size}}} n_{3D}(x,y,z)\,dz. \tag{3.18}$$

If the intensity of the probe beam is weak such that $I_{in} \ll I_{sat}$, and is tuned to the atomic resonance, with $\delta = 0$, then the scattering cross-section is constant at $\sigma_{sc} = \sigma_0$ and the beam intensity simply decays exponentially with distance through the sample. In this case, the intensity loss is described by the Beer-Lambert law [205, 206], and the two-dimensional column density in Eq. 3.17 simplifies to

$$n_{2D}(x,y) = \frac{1}{\sigma_0}\ln\left(\frac{I_{in}(x,y)}{I_{out}(x,y)}\right) = \frac{1}{\sigma_0}\mathrm{OD}(x,y), \tag{3.19}$$

where the *optical density*, $\mathrm{OD}(x,y)$, is defined through the relation

$$I_{out}(x,y) = I_{in}(x,y)\,e^{-\mathrm{OD}(x,y)}. \tag{3.20}$$

Equation 3.19 demonstrates that knowledge of the absolute intensity in the beam is not required in order to extract the physical column density - it is sufficient to measure the relative intensity before and after absorption.

The process of absorption imaging in practice is illustrated in Fig. 3.40. First the cloud of atoms is illuminated with the probe beam, and the image of the shadow is then recorded on a CCD camera, giving $I_{in}(x,y)$. Next, after the cloud has fallen out of the field of view, the second image $I_{out}(x,y)$ is recorded, typically $100\,\mathrm{ms}$ later, containing only the probe light distribution. We also take a third image $I_{bg}(x,y)$ at the end, which contains the contribution from stray background light, and subtract this from both $I_{in}(x,y)$ and $I_{out}(x,y)$. The effect due to the presence of the cloud is then obtained by constructing the optical density image using Eq. 3.19, and any artifacts or interference fringes in the resultant distribution are largely removed using the algorithm developed in Ref. [207], which is outlined in detail in Appendix B.

Extracting Atom Number

All information on the atom cloud is gained from two-dimensional optical density images, measuring the attenuation of a probe beam passing through the sample. Properties of the cloud, including temperature, atom number, spatial extension and density, are inferred from these images, and it is therefore important to interpret them correctly. Different types of fitting procedures are applied based on theoretical predictions, depending on the various regimes in which the cloud can be.

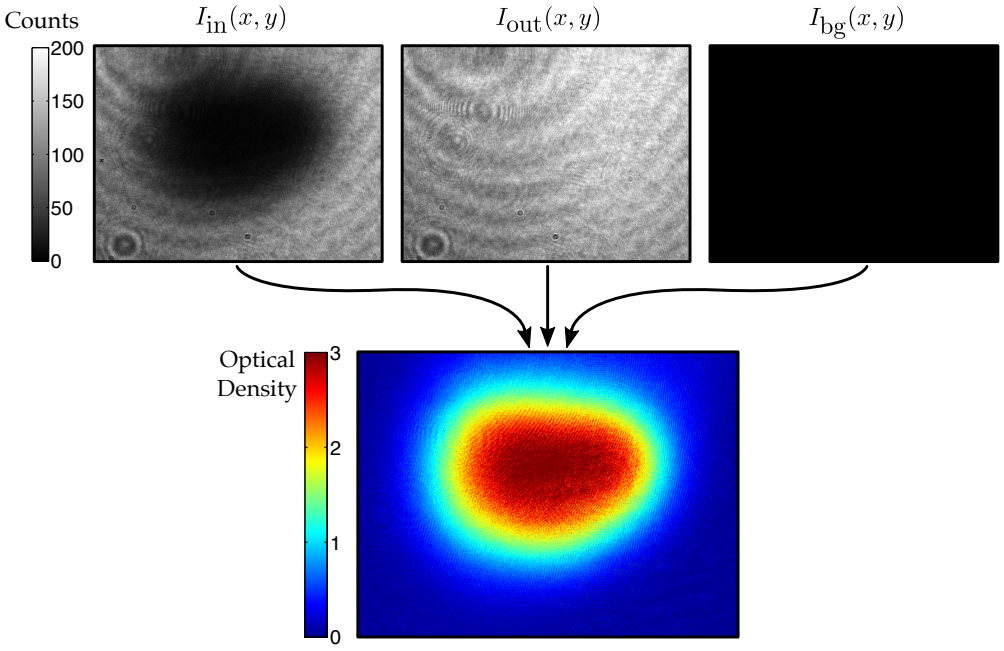

FIGURE 3.40: Construction of the optical density using absorption imaging: Three images are required to obtain the OD - an image with the atoms $I_{\text{in}}(x,y)$, an image with the probe light but without the atoms $I_{\text{out}}(x,y)$, and an image without either the atoms or the probe light $I_{\text{bg}}(x,y)$. The optical density is obtained by first subtracting the background from the first two images to remove the effect of stray light, and then taking the logarithmic quotient of the resulting images according to Eq. 3.19 as described in the text.

Due to the finite pixel size of the CCD camera, the optical density in practice is constructed pixel-wise as a matrix, and so the column density at each pixel obtained using Eq. 3.19 is in fact an average over that pixel. To calculate the actual number of atoms, it is then crucial to account for the sizes of the pixels. The number of atoms which have interacted with the laser beam along the z-direction contained within the area of a pixel centred at position (x,y) is given by

$$\begin{aligned}\bar{N}(x,y) &= \int_{x-l_x/2}^{x+l_x/2}\int_{y-l_y/2}^{y+l_y/2}\int_{-\infty}^{\infty} n_{\text{3D}}(x,y,z)\,dz\,dy\,dx \\ &= \int_{x-l_x/2}^{x+l_x/2}\int_{y-l_y/2}^{y+l_y/2} n_{\text{2D}}(x,y)\,dy\,dx \\ &\simeq l_x\,l_y\,n_{\text{2D}}(x,y) = \frac{A}{M^2}\,n_{\text{2D}}(x,y) = \frac{A}{M^2}\frac{1}{\sigma_0}\text{OD}(x,y),\end{aligned} \tag{3.21}$$

where l_x and l_y are the dimensions of each pixel in the object plane, A is the physical area of a pixel on the CCD chip and M is the magnification of the optical system. Here, the final line arises by making the approximation that the column density is constant over the size of a pixel, and in the limit of infinitesimally small pixel sizes the expression becomes exact. The total number of atoms can then be obtained in a simple way by

summing over all pixels in the resultant image through

$$N = \sum_{\text{pixels}} \bar{N}(x,y) = \sum_{\text{pixels}} \frac{A}{M^2}\frac{1}{\sigma_0}\text{OD}(x,y). \tag{3.22}$$

In addition to merely the number of atoms, it is important to know the actual three-dimensional density distribution, which for a classical, ideal gas with temperature T confined in a trapping potential $U(\boldsymbol{r})$ is given by [208]

$$n_{3\text{D}}(\boldsymbol{r}) = n_{3\text{D}}(0)\,\exp\left(-\frac{U(\boldsymbol{r})}{k_B T}\right). \tag{3.23}$$

For purely thermal clouds in an harmonic potential (which applies to both magnetically-trapped atoms and also describes well the distribution in the MOT [209]), far above the critical temperature ($\text{T} \gg \text{T}_\text{c}$), the density follows the three-dimensional Gaussian

$$n_{3\text{D}}(\boldsymbol{r}) = n_{3\text{D}}(0)\,\exp\left[-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2} - \frac{z^2}{2\sigma_z^2}\right], \tag{3.24}$$

where $n_{3\text{D}}(0)$ is the peak density of the cloud, $\sigma_i = \sqrt{k_B T/m\omega_i^2}$ is the extension in the i^{th} dimension with harmonic trapping ω_i. Normalisation on the total number of atoms N then dictates that[21]

$$\int n_{3\text{D}}(\boldsymbol{r})\; d^3\boldsymbol{r} \equiv N \implies n_{3\text{D}}(0) = \frac{N}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_z}. \tag{3.25}$$

However, the absorption imaging technique in the experiment does not measure full three-dimensional densities, but rather the optical density, as described earlier. Integrating Eq. 3.24 over z results in the two-dimensional column density

$$n_{2\text{D}}(x,y) = \frac{N}{2\pi\sigma_x\sigma_y}\,\exp\left[-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right], \tag{3.26}$$

which is also Gaussian, and so using 3.19 the measured optical density is then given by

$$\text{OD}(x,y) = \text{OD}(0)\,\exp\left[-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}\right]. \tag{3.27}$$

This functional form is then used to apply a two-dimensional fitting routine to the obtained images for determining the peak optical density and cloud dimensions, and the total number of atoms can be calculated from the resulting fitted parameters using

$$N = \left(\frac{7}{15}\right)^{-1}\frac{1}{\sigma_0}\frac{M^2}{A}2\pi\sigma_x\sigma_y\,\text{OD}(0), \tag{3.28}$$

[21] Using the Gaussian integral $\int_{-\infty}^{+\infty}\exp\left(-\frac{x^2}{\beta}\right)dx = \sqrt{\pi\beta}$.

which is valid when a resonant, weak-intensity probe beam is used to image the cloud in a magnetic field-free environment. Here, the factor of M^2/A appears due to finite pixel size introduced by Eq. 3.21.

Finally, in order to calculate the three-dimensional density of Eq. 3.25, information is needed about the extension of the cloud along the direction of the imaging beam. Absorption imaging provides only the number of particles integrated along a given "column", and it is necessary therefore to assume some symmetry in the trapped clouds. For all instances in this experiment, it is assumed that $\sigma_z = \sigma_y$ when calculating the 3D density. This assumption is accurate for the cigar-shaped cylindrically-symmetric Ioffe-Pritchard magnetic traps described in Sec. 2.5.2, but will be less precise when imaging the MOT.

Extracting Temperature

The temperature of classical, thermal clouds can be inferred by switching off the trapping potential and observing the expansion of the gas due to the thermal velocity of constituent particles. The density distribution of the intially Gaussian cloud after an expansion time t can be calculated by assuming that particles fly ballistically from their initial positions according to their in-trap momentum. This leads to the cloud maintaining its Gaussian distribution, but with a RMS width that evolves in the i^{th} direction according to [210, 211]

$$\sigma_i(t) = \sqrt{\sigma_i^2(0) + \frac{k_B T}{m} t^2}. \tag{3.29}$$

At large times-of-flight the initial size $\sigma_i(0)$ can be neglected and the average expansion velocity is just proportional to the square root of the temperature, as expected from the equipartition of energy $\frac{1}{2}k_B T = \frac{1}{2}mv^2$. The temperature of the gas is then calculated by repeatedly releasing it from the trapping potential and measuring the expansion as a function of time, and fitting according to Eq. 3.29. An example of a typical measurement is shown in Fig. 3.41. Each optical density image is fit with the two-dimensional function of Eq. 3.27 to extract the RMS widths of the cloud, and the effect of gravity is accounted for in each image by moving into the frame of reference in which the centre-of-mass position is stationary.

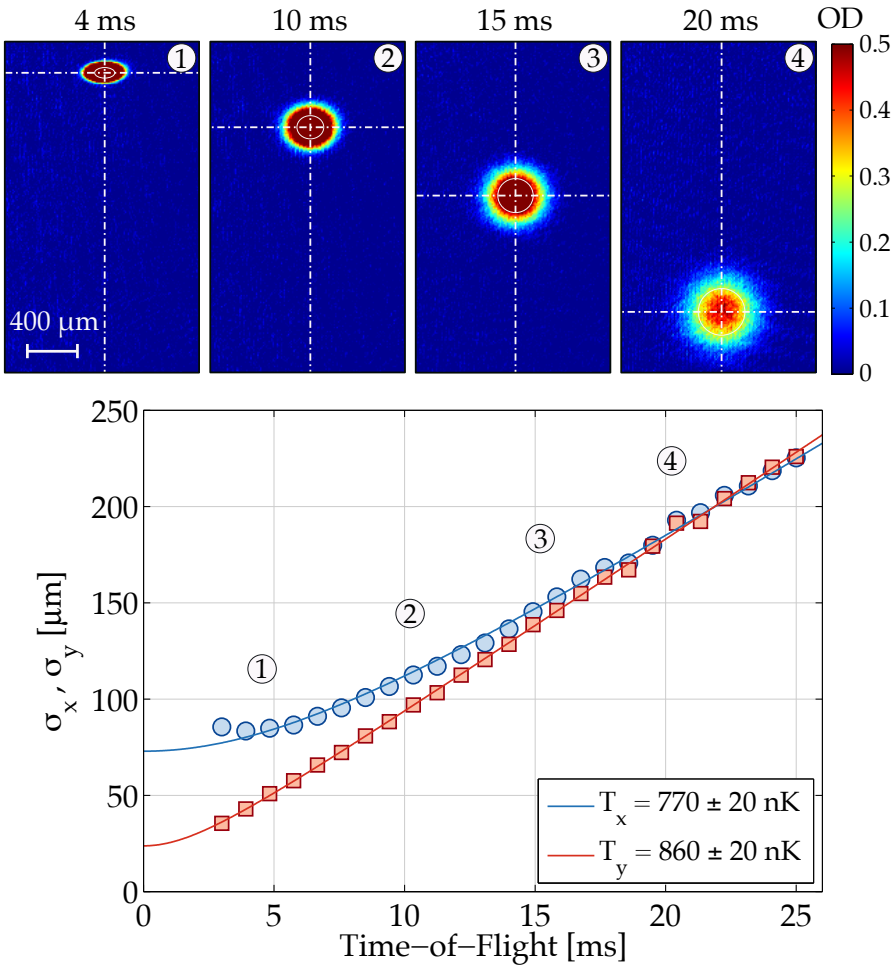

FIGURE 3.41: Example of a temperature measurement: The data correspond to a cloud at the end of an evaporative cooling sweep in the copper Z trap. Temperature can be calculated by repeatedly releasing clouds from the trap and measuring the rate of expansion of their size, which is dependent on the thermal energy. The lower panel shows the measured RMS widths in each of the axial and radial directions σ_x (○) and σ_y (□), respectively. The derived temperatures are $T_x = 770 \pm 20\,\mathrm{nK}$ and $T_y = 860 \pm 20\,\mathrm{nK}$. Selected examples of optical density absorption images are illustrated in the upper row, with the dashed lines marking the fitted centre-of-mass position as the cloud falls under the influence of gravity, and the elliptical contours depicting the RMS width.

Chapter 4

BEC Production on the Atom Chip

This chapter describes in detail the development of an experimental procedure for creating a ^{87}Rb BEC on the atom chip. The production of a BEC requires a complex and carefullly orchestrated series of steps, the characterisation and detailed measurements of which made up a large part of this PhD work. Each stage has been carefully optimised, with the ultimate goal being to increase the phase-space density as efficiently as possible throughout, and leading eventually to the highest number of atoms in the condensed state.

Section 4.1 decribes the main 3D mirror-MOT, and its loading from a newly-installed 2D MOT then follows in Sec. 4.2. Several steps are then performed, such as sub-Doppler cooling (Sec. 4.3) and optical pumping (Sec. 4.4), which improve greatly the efficiency of the transfer into the purely magnetic copper-Z trap (Sec. 4.5). Much time was spent on achieving efficient evaporative cooling in the copper-Z trap, which is described in Sec. 4.6, and the first observation of BEC in our lab is documented in Sec. 4.7. Finally, a scheme was implemented for loading clouds into the atom chip (Sec. 4.8.1), and the process of producing an on-chip BEC is outlined in Sec. 4.8.2.

Overview of An Experimental Cycle

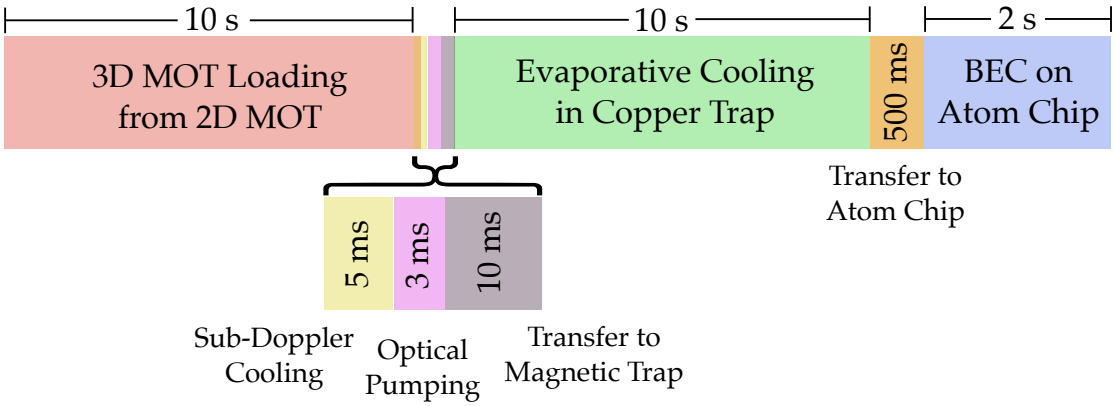

FIGURE 4.1: Temporal Overview of the Experimental Sequence.

The key stages in the experimental sequence are as follows:

- **2D MOT** (Sec. 4.2): In the first step, a cold atomic beam produced by the 2D MOT is directed from the high-pressure glass cell region, where it is used to load the

3D MOT situated in the UHV science chamber. The goal here is to produce a high flux of atoms with velocities below the capture velocity of the 3D MOT.

- **3D Mirror-MOT** (Sec. 4.1): A large cloud of $\sim 2 \times 10^8$ atoms with a temperature of $300\,\mu\mathrm{K}$ is captured from the 2D MOT within $10\,\mathrm{s}$ located several millimeters above the atom chip surface. Note that this stage has been described first in this chapter, since this is the way initial optimisation is carried out in practice.
- **Sub-Doppler Cooling** (Sec. 4.3): The cloud is then subjected to several milliseconds of further-detuned, low intensity light, reducing the temperature by an order of magnitude in preparation for transfer to the magnetic trap.
- **Spin-Polarisation** (Sec. 4.4): A pulse of several milliseconds of optical pumping light prepares the sample in the maximally-trappable low field-seeking magnetic state $|F = 2, m_F = +2\rangle$ with a purity of $> 96\%$.
- **Transfer to Copper Z-trap** (Sec. 4.5): The laser cooled cloud is transferred into a purely magnetic trap with $\sim 70\%$ efficiency, in preparation for further cooling below the photon recoil limit. This transfer process is crucial, and the goal is to minimise any loss in phase-space density occuring due to non-adiabatic effects.
- **Evaporative Pre-Cooling** (Sec. 4.6): A $10\,\mathrm{s}$ forced RF evaporative cooling sweep pre-cools the cloud to $\sim 20\,\mu\mathrm{K}$.
- **BEC on the Atom Chip** (Sec. 4.8): Typically, 6×10^6 atoms are transferred with 100% efficiency over $500\,\mathrm{ms}$ to the micropotentials created by the atom chip, where a final $2\,\mathrm{s}$ cooling phase results in BECs of 10^5 atoms.

4.1 Mirror MOT

The first stage of trapping and cooling in the experimental cycle is the magneto-optical trap (MOT) - a powerful technique which allows atoms from a hot source at $600\,^\circ\mathrm{C}$ to be captured and cooled to $\sim 100\,\mu\mathrm{K}$. This step provides the largest increase in phase space density of the entire experimental procedure. For a gas of room temperature rubidium atoms with a partial pressure of 10^{-9} mbar in the vacuum chamber, the de Broglie wavelength is $\lambda_{\mathrm{dB}} = 0.01\,\mathrm{nm}$ and the phase space density is 10^{-22}. The MOT increases these values to $\lambda_{\mathrm{dB}} = 10\,\mathrm{nm}$ and PSD $= 10^{-8}$ in just several tens of milliseconds.

The MOT stage begins by injecting a source of rubidium atoms, which is done using alkali metal dispensers.[1] Although a 2D MOT is installed in the system, and presently acts as the source of atoms, all optimisation of the 3D MOT parameters was done using a set of dispensers mounted in the main science chamber. The reason for this is that we do not want any alignment of the atomic beam from the 2D MOT to influence the

[1] SAES Getters, rubidium alkali metal dispensers.

optimisation process. Additionally, long trap lifetimes are not required for the MOT, allowing the dispenser to be run continuously.

Our system implements the so-called *mirror-MOT* configuration [61, 62, 212], common to atom chip experiments, whereby two of the required beams for laser cooling are provided by reflection from the chip surface. The mirror-MOT allows integrated trapping close to the chip surface, facilitating good transfer into the purely magnetic traps, without the need for an additional transfer stage from anti-Helmholz coils. Four trapping beams are used in total, each with Gaussian waists ($1/e^2$ radii) of $10\,\mathrm{mm}$ and clipped at 1" diameter to ensure that the beams fit within the area of the chip surface. Two beams enter through the large DN100CF viewport on the bottom of the science chamber at an incident angle of $32°$ to the vertical, which is slightly narrower than the optimum $45°$ due to geometrical constraints of the chamber. These beams are reflected from the surface of the chip (which was measured to have better than 98% reflectivity at $780\,\mathrm{nm}$), obtaining the correct polarisations upon reflection due to a reversal of helicity, and thus forming two orthogonal pairs. The third remaining pair of beams is incident through the side viewports, parallel to the chip surface.

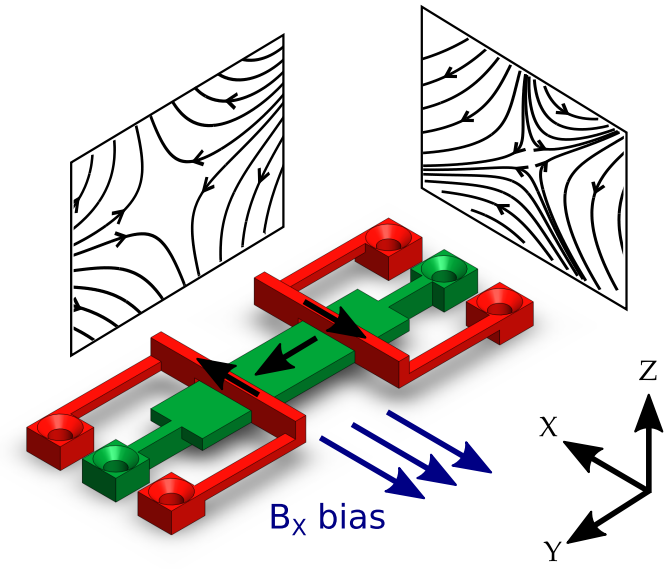

FIGURE 4.2: Schematic of the structures for creating the 3D MOT magnetic fields: The combination of a current through the X-Sheet structure (green) and an external bias field B_x produces a two dimensional $45°$ rotated quadrupole field for trapping in the radial plane. Axial confinement is introduced by passing anti-parallel currents through the sidebars structure (red).

To exert a trapping force in all directions, the MOT requires a three-dimensional quadrupole field with eigenaxes oriented along the wave vector of the laser beams. A $45°$ tilted quadrupole can be produced using a combination of the macroscopic copper understructure wires located below the surface of the atom chip, along with an external bias field [212]. Specifically in our setup, a current flowing in the y-direction through

the X-sheet structure generates a magnetic field gradient which is cancelled by an opposing external bias field B_x in the X-direction at some height above the surface. This produces a tilted two-dimensional quadrupole in the radial plane, illustrated schematically in Fig. 4.2. Axial confinement is introduced by running currents in an anti-parallel configuration through the sidebars structures, leading to a field gradient that changes sign when passing through the minimum, as required for operation of a MOT.

A calculation of the magnetic field in two orthogonal planes is shown in Fig. 4.3. Typically the MOT is operated with $80\,\mathrm{A}$ though the X-sheet, $20\,\mathrm{A}$ through the sidebars, and an external field of $B_x = -5.5\,\mathrm{G}$. The resultant cloud forms $6\,\mathrm{mm}$ below the surface of the atom chip, although we observe that the position of the MOT can be varied over several millimeters using only alignment of the beams, and is not always formed at the magnetic minimum. In general, it is the interplay of alignment of both the magnetic and light fields which determines the final position and atom number.

The purpose of the MOT loading stage is to increase the phase-space density enough to efficiently transfer the gas to a purely magnetic trap for further cooling. This means capturing a large number of atoms at a low enough temperature, in as short a time as possible to increase the repetition rate of cycles. This can always be achieved by increasing the dispenser current, but which comes at the expense of a degraded vacuum. Therefore, considerable time was spent optimising the available MOT parameters, whilst maintaining a moderate continuous dispenser current of $4\,\mathrm{A}$. Firstly, it is ensured that there is enough power in both cooling and repumping beams, as shown in Fig. 4.4 where it can be seen that approximately $170\,\mathrm{mW}$ and $4\,\mathrm{mW}$ of cooling and repumping power, respectively, is required to saturate the captured atom number. The MOT is usually operated with $50\,\mathrm{mW}$ of cooler in each of the reflected beams, and $30\,\mathrm{mW}$ in each of the side beams, ensuring saturation of the MOT number whilst running the tapered amplifier of the trap laser at $\sim$2/3 of its maximum power. Having more power in the 45° beams was determined empirically using balancing half-wave plates, and may be due to the fact that the chip is not perfectly reflecting. In addition, $5\,\mathrm{mW}$ of repumping power is overlapped with each of the two side beams.

For Doppler cooling, the trap beams must be red-detuned with respect to the cycling transition to provide a frictional force, as described in Sec. 2.3.1. To maximise atom number the cooling beam is set with a frequency of $-16\,\mathrm{MHz}$ ($-2.6\,\Gamma$) below the $|F = 2\rangle \rightarrow |F' = 3\rangle$ line, as shown in Fig. 4.5, whilst the repumping laser is placed on resonance with the transition $|F = 1\rangle \rightarrow |F' = 2\rangle$.

It is observed that the maximum atom number in the MOT is obtained when the two $45°$ beams are slightly misaligned from being completely collinear. This may be due again to the fact that the radiation force from the reflected beam can never fully equal the incident beam if they are aligned exactly, because of the imperfect reflectivity. We observe an increase in atom number of a factor of 3 by misaligning the beams slightly. However, when the beams are indeed aligned, the cloud position lies exactly at the calculated magnetic quadrupole centre, indicating that the trapping forces are

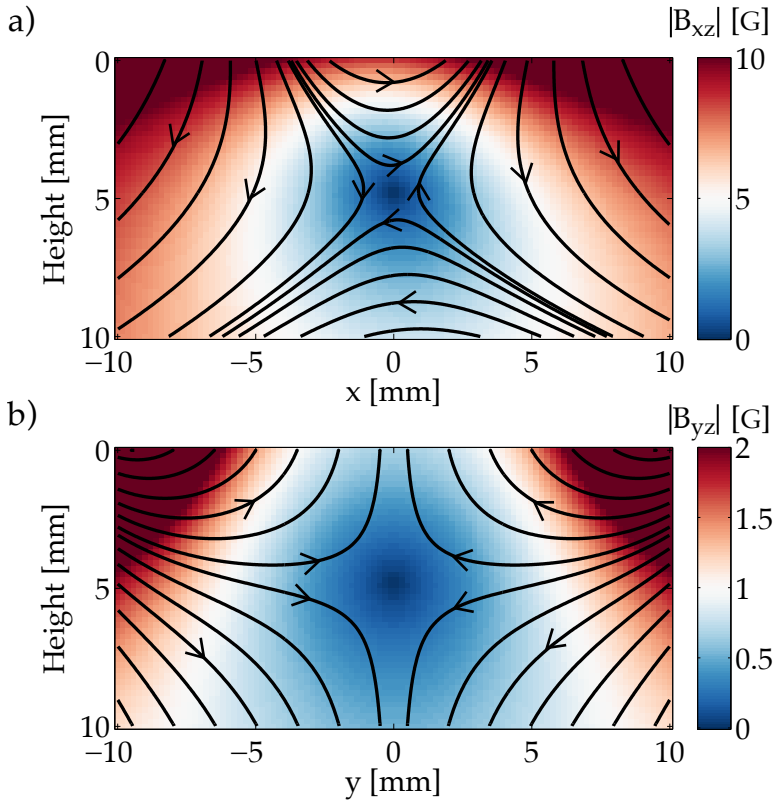

FIGURE 4.3: Calculated magnetic quadrupole field for the MOT: The magnetic field is calculated using finite element methods for $80\,\mathrm{A}$ through the X-sheet structure, $20\,\mathrm{A}$ in the sidebars, and an external bias field of $B_x = -5.5\,\mathrm{G}$. A small $0.8\,\mathrm{G}$ vertical bias field is also used to centre the field zero into the $x = 0$ plane. a) The magnitude of the field components projected on to the $x - z$ plane, where z represents the height above the atom chip surface. In this plane, the radial quadrupole is formed at $z = 5\,\mathrm{mm}$, and is rotated by $45°$ so that the eigenaxes coincide with the wavevectors of the laser beams. b) The field components projected onto the $y - z$ plane, providing the confinement in the axial dimension, with eigenaxes oriented horizontally to coincide with the side laser beams parallel to the chip surface.

balanced. To confirm simulations of the trapping structures, the position of the cloud was recorded for various perpendicular and vertical bias fields, and compared with the calculated height of the magnetic minimum from the chip surface, shown in Fig. 4.6. Increasing the perpendicular bias field moves the cloud closer to the chip surface, simply because it now cancels the X-sheet field at a closer distance. Changing the vertical bias field rotates the magnetic minimum around an axis perpendicular to the imaging direction, resulting in a maximum in the distance of the cloud from the chip surface. The data show very good agreement between the measured and calculated positions.

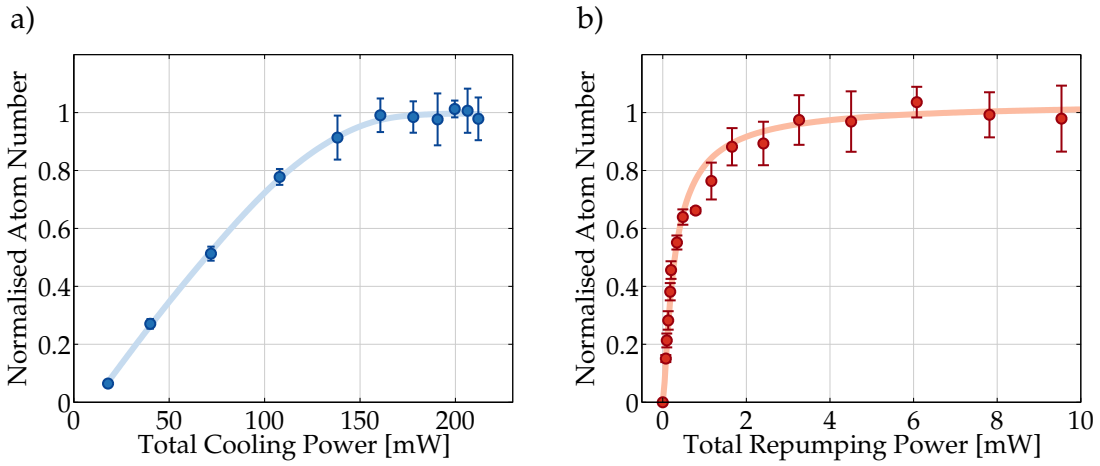

FIGURE 4.4: Laser powers required for the MOT: a) The normalised total number of atoms in a fully loaded MOT is measured as a function of the total amount of cooling laser power sent into the vacuum chamber, which is shared amongst the four trap beams. A total power of $\sim 170\,\mathrm{mW}$ is required to ensure saturation of the curve. b) Normalised atom number as a function of total repumping power, where a minimum power of $\sim 4\,\mathrm{mW}$ is required. In both cases, the laser powers are varied using the amplitude of the RF wave delivered to their respective AOMs. Error bars represent the standard error over several repetitions, whilst solid lines are guides to the eye.

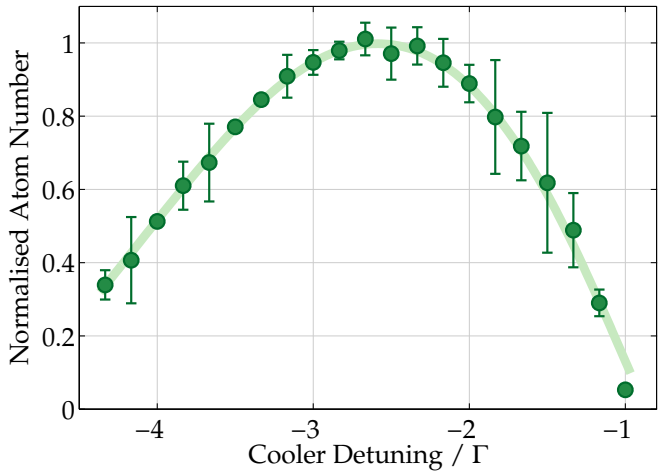

FIGURE 4.5: Optimal 3D MOT cooler detuning: The normalised atom number captured in a fully loaded MOT is measured as a function of the red-detuning of the cooling laser beam from the cycling transition. The frequency is varied whilst maintaining the power approximately constant using an AOM in double-passage configuration. Maximum atom number is obtained with $-2.6\,\Gamma$.

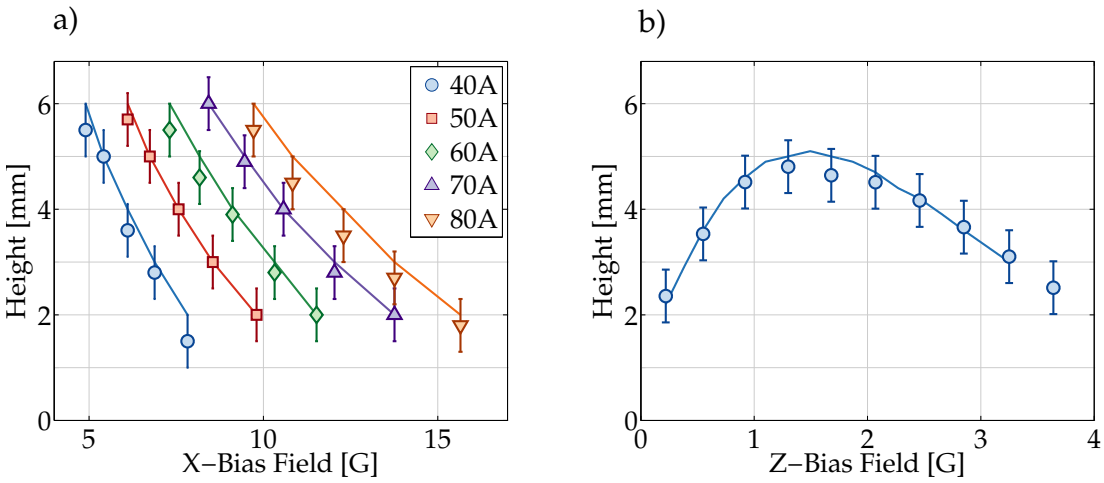

FIGURE 4.6: Comparison between calculated and measured MOT positions: a) The height of the MOT above the atom chip surface is measured as a function of the main X-Bias field for five different currents in the X-Sheet structure. With increasing sheet current, larger values of the bias field are required to produce a minimum at the same height. b) Height of the MOT as a function of an external vertical Z-Bias field. This field rotates the minimum off-centre around the x-direction (see Fig. 4.2 for coordinate definitions), resulting in passing through a maximum in the observed position on the CCD camera. For both (a) and (b) the solid lines are the result of finite element calculations, and show good agreement with the data.

4.1.1 MOT Loading

The evolution of the number of trapped atoms, $N(t)$, in the MOT is governed by a rate equation of the form [213–215]

$$\frac{dN(t)}{dt} = R - \gamma N - \beta \int_{\substack{\text{trap}\\ \text{volume}}} n_{\text{MOT}}^2(\boldsymbol{r}, t)\, d\boldsymbol{r}, \tag{4.1}$$

where R is the loading rate, γ is the loss rate due to collisions with the hotter background gas particles, and β describes losses due to two-body inelastic collisions occuring between trapped atoms, which therefore depends on the MOT particle density $n_{\text{MOT}}(\boldsymbol{r}, t)$. The one-body loss rate is given by $\gamma = n_{\text{bg}}\sigma v_{\text{rms}}$ [136], where σ is the cross-section for collisions between trapped and untrapped atoms, and n_{bg} and v_{rms} are the density and RMS velocity of background vapour particles, respectively.

Models of the loading process in the MOT are usually described in terms of the *capture velocity* v_c, and it is assumed that any particle with speed $v < v_c$ is cooled and trapped. In general, the capture velocity is determined by the ability of the trap to decelerate particles, which depends on the force imparted onto an atom through the laser powers and detunings, as well as the magnetic field gradient. Typical values for the capture velocity of 3D MOTs are several tens of $\mathrm{m\,s^{-1}}$. With the capture velocity defined, the loading rate can be estimated by calculating the number of particles in a

Maxwell-Boltzmann distribution with $v < v_c$ entering the trapping region (the intersection of the laser beams) in a unit of time, resulting in [136, 215]

$$R \propto \frac{v_c^4}{v_{\text{th}}^3} A\, n, \tag{4.2}$$

where A is the surface area of the trapping region, and n and v_{th} are the density and average thermal velocity of trappable particles in the background gas, respectively. It can be seen from Eq. 4.2 that it is advantageous to have large diameter trapping beams, and a high density of slow particles with respect to the capture velocity (which is improved through the use of the 2D MOT in Sec. 4.2).

In order to describe the dynamical loading characteristics of the MOT, Eq. 4.1 must be solved for $N(t)$. The equation can be rewritten in the form [216]

$$\frac{dN(t)}{dt} = R - \gamma N(t) - \beta \bar{n}(t) N(t), \tag{4.3}$$

where the average density in the MOT is given by

$$\bar{n} = \frac{\int n_{\text{MOT}}^2(\boldsymbol{r})\, d\boldsymbol{r}}{\int n_{\text{MOT}}(\boldsymbol{r})\, d\boldsymbol{r}} = \frac{1}{N} \int n_{\text{MOT}}^2(\boldsymbol{r})\, d\boldsymbol{r}. \tag{4.4}$$

The behaviour of the MOT can be broadly described in two different regimes [217, 218]. For small numbers of atoms ($< 10^6$), the trap volume remains fixed and the average density grows with increasing atom number, such that $\bar{n}(t) = N(t)/V$. However, in most realistic experimental situations with larger atom numbers, there is multiple scattering of spontaneously emitted photons inside the cloud, leading to a repulsive radiation pressure which ultimately limits the density of the cloud. In this regime, the average density remains *constant* and the volume grows proportionally with the number of atoms as the MOT is loaded. The transition between the two regimes was experimentally investigated in [218]. With the assumption of constant average density, Eq. 4.3 can be integrated to find the evolution of the number of trapped atoms with time

$$N(t) = N_{\text{max}} \left(1 - e^{-t/\tau}\right), \tag{4.5}$$

where the characteristic time for the exponential loading process $\tau = 1/(\gamma + \beta\bar{n})$ arises due to both collisions with background particles and two-body inelastic losses within the cloud, and the maximum number of trapped atoms $N_{\text{max}} = R/(\gamma + \beta\bar{n})$ therefore results from the equilibrium when the loading and loss rates are balanced. An example of a MOT loading curve is shown in Fig. 4.7, measured by recording the number of atoms at various times after switching on the trap, along with a fit to Eq. 4.5 which describes the data well. At $t = 0$ all losses are negligible, and it can be seen by differentiating Eq. 4.5 that the intial slope of the loading curve is a direct measure of the

loading rate into the trap

$$\frac{dN(t)}{dt}\Big|_{t=0} = R, \tag{4.6}$$

as illustrated by the dash-dotted line in Fig. 4.7.

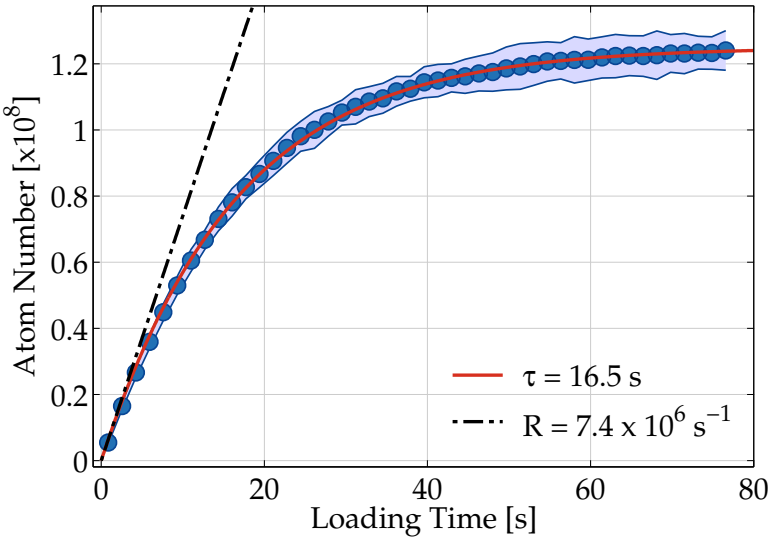

FIGURE 4.7: MOT loading curve: The number of trapped atoms is measured as a function of loading time. A fit to Eq. 4.5 (solid line) results in a loading time of $\tau = 16.5\,\mathrm{s}$, whilst a fit to the initial slope (dash-dotted line) is a direct measure of the loading rate $R = 7.4 \times 10^6 \mathrm{s}^{-1}$. The shaded area represents the statistical deviation between shots.

Typically, the MOT is not loaded to saturation, but for a time of $\sim \tau$, giving the most efficient collection of atoms with the available time. After all the above optimisation of the 3D MOT, $\sim 1.2 \times 10^8$ atoms are collected at $6\,\mathrm{mm}$ from the atom chip surface, shown in Fig. 4.8. The peak density at this point is $1.5 \times 10^{10}\mathrm{cm}^{-3}$ and the temperature is measured to be $320\,\mu\mathrm{K}$ before any additional cooling is applied. Note again that all the optimisation presented in this section is done using dispensers in the main chamber, and that when switching to the use of the 2D MOT the final atom number is increased to $> 3 \times 10^8$.

4.2 2D MOT

A two-dimensional magneto-optical trap (2D MOT) was installed into the experiment in 2015, to improve the atom number and cycle stability. The 2D MOT has to be optimised to provide a high flux of cold atoms for loading the 3D MOT in the main science chamber. There are many parameters which contribute to the efficiency, several of which are coupled with each other.

The setup of the 2D MOT is described in Sec. 3.2. Initial alignment of the optics is done with the cage structure removed from the glass cell. Beams for the three cooling

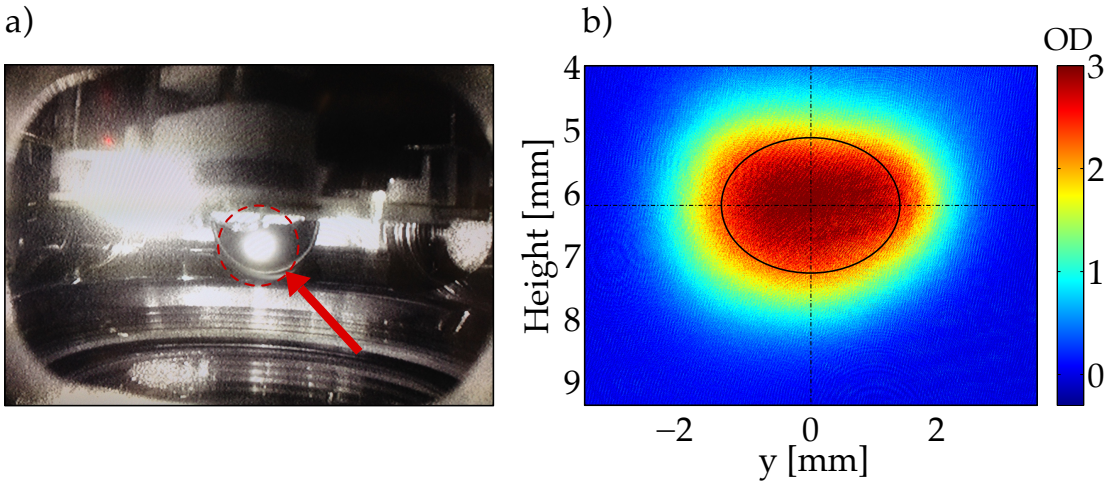

FIGURE 4.8: Images of the 3D MOT: a) The atomic fluorescence in the presence of the cooling beams is viewed in real time for diagnostic purposes, and is captured on an inexpensive CCTV camera (lacking an infrared filter) through a side vacuum viewport. The cloud of atoms can be seen as the white area in the centre of the image, suspended below the atom chip surface, and is indicated by the circled region. b) Absorption image providing quantitative analysis of the same cloud in (a) with $3\,\mathrm{ms}$ time-of-flight after switching off the trapping potential. The measured atom number is 1.3×10^8.

regions are overlapped, and the correct circular polarisations are obtained by adjusting each quarter-wave plate. The cage is then installed along with the two holders of permenant magnets to provide the 2D quadrupole field. A total of $100\,\mathrm{mW}$ of cooling light and $5\,\mathrm{mW}$ of repumping light is used, shared equally amongst the three trapping regions, and frequencies of the lasers are set using individual AOMs. When initially installing the 2D MOT, the integrated density can be viewed on a fingertip camera along the long axis of the cloud (which provides the largest signal), and a photograph of the first observation of the 2D MOT is shown in Fig. 4.9.

It can be difficult to initially get atoms from the 2D MOT into the 3D MOT through the $800\,\mu\mathrm{m}$ aperture of the differential pumping tube, but once there is some signal the optimisation is an iterative process. The figure of merit is the loading rate of the 3D MOT, which is a direct indicator of the number of atoms below the capture velocity being directed at the 3D MOT region. This loading rate is given by the inital slope of the loading curve (Eq. 4.6), and is measured using fluorescence imaging by pulsing the 3D MOT cooling beams with a $1\,\mathrm{Hz}$ rate, giving a real-time indication of the atom flux. All optimisation is done initially with only the most critical cooling region closest to the differential pumping tube, with the remaining two added afterwards, and without the use of the pushing beam. The first step in the process is alignment of the entire structure using the adjustable bellow, to ensure that the cold atomic beam is directed optimally at the 3D MOT trapping region. Then, the atoms are made to pass fully through the differential pumping tube aperture by displacing the 2D cloud in its radial plane, which is done using the horizontal and vertical offset coils set up in Helmholtz configuration.

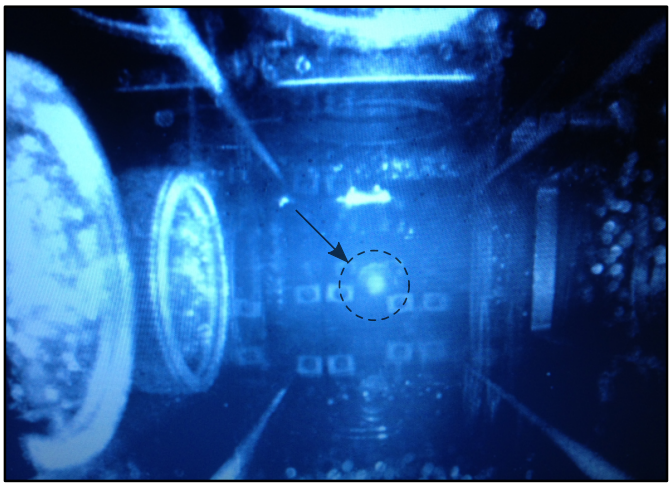

FIGURE 4.9: First observation of the 2D MOT: The integrated fluorescence along the axial direction is viewed on an inexpensive CCTV camera, sensitive to infrared light. The cloud (circled) can be seen at the centre of the glass cell, in front of the differential pumping tube aperture. The square structures on each side of the cloud are small mounting clamps for rubidium dispensers on the feedthroughs.

The result of this process is shown in Fig. 4.10, indicating optimal currents of $2.0\,\mathrm{A}$ and $3.7\,\mathrm{A}$ for the vertical and horizonatal coils, respectively.

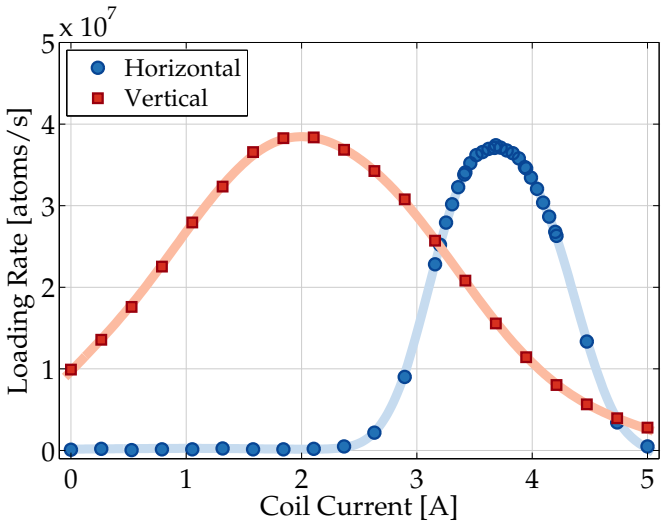

FIGURE 4.10: Optimisation of 2D MOT offset fields: The position of the 2D cloud is adjusted by varying the current through both the horizontal (●) and vertical (■) offset coils, whilst observing the loading rate of the 3D MOT. The flux dropping to zero indicates the cloud completely missing the aperture of the differential pumping tube.

Additionally, the frequencies of both the cooling and repumping beams are varied using individual double-passage AOMs, in order to maximise the 3D MOT loading rate. As seen in Fig. 4.11 (a), the cooling beam is set to $9.7\,\mathrm{MHz}$ ($-1.6\,\Gamma$) red-detuned from the $|F = 2\rangle \rightarrow |F' = 3\rangle$ line, whilst the repumper is optimised on resonance with the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. Additionally, we test the importance of the rubidium

pressure in the 2D MOT glass cell, shown also in Fig. 4.11 (b). The pressure itself is not directly known, as there is no gauge on the glass cell side, and so the dispenser current is used instead. For these data, the measured parameter was the atom number that remained at the end of an evaporative cooling ramp, since this also then takes into account the lifetime of atoms in the science chamber in addition to the loading rate. When operating the 2D MOT dispenser continuously in daily operation, the pressure on the ion gauge in the main chamber remains below 10^{-10} mbar. We can see that the performance initially improves with increasing dispenser current, whereby there are simply more atoms available to contribute to the 2D MOT atom number inside the glass cell. However, after a maximum at $3.9\,\mathrm{A}$, the number begins to decrease again, consistent with other observations that collisions between atoms in the cell become detrimental when the mean free path becomes on the order of the trap region length [163, 164, 219]. To prolong the life of the dispenser, it is usually run with a lower current of $3.4\,\mathrm{A}$, which provides sufficient atoms for our purposes.

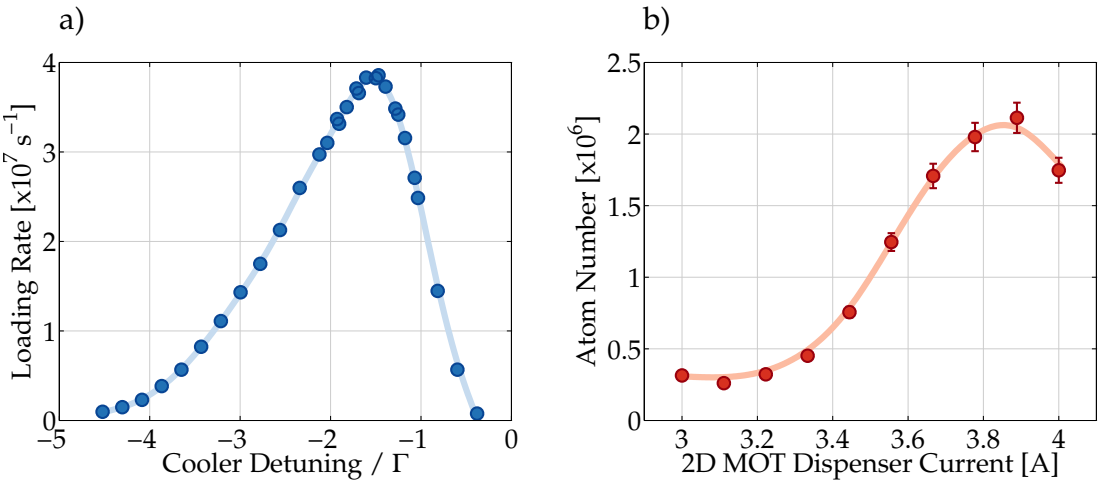

FIGURE 4.11: a) Loading rate of the 3D MOT measured as a function of the 2D MOT cooling beam detuning. Optimal flux is achieved with a red-detuning of $-9.7\,\mathrm{MHz}$ ($-1.6\,\Gamma$). b) Atom number measured at the end of the evaporative cooling sweep as a function of continuous through a single dispenser in the 2D MOT glass cell. The optimum current is $3.9\,\mathrm{A}$, with higher currents leading to a lower flux as described in the text.

Pushing Beam

Most 2D MOT setups use some form of pushing beam directed along the axis towards the differential pumping tube, with a view to cooling atoms which are travelling in the wrong direction, eventually turning them around so they can contribute to the overall flux. These setups require a red-detuned laser beam with enough power to impart an appreciable force on the atoms. Since the power and frequency of the beam, along with its diameter, are coupled, we investigated a large range of parameters, as shown in Fig. 4.12. It can be seen that a red-detuned push beam does in fact improve the loading

rate, but requires higher laser powers - approximately $10\,\mathrm{mW}$ for $-45\,\mathrm{MHz}$. However, the best improvement comes with lower powers and closer to resonance. For many values in this parameter space, the push beam harms the loading rate and in some cases completely extinguishes it - especially for large powers and large blue-detunings.

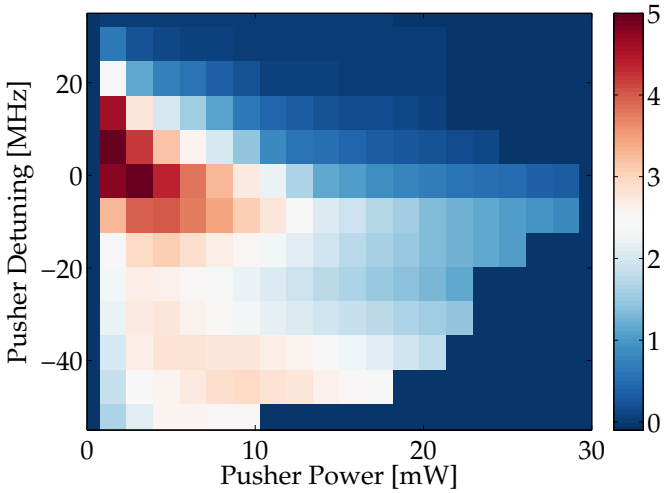

FIGURE 4.12: Optimisation of the 2D MOT pushing beam: The loading rate of the 3D MOT is measured as a function of both the total power and detuning from the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition. The loading rates are normalised to the case where the pusher is switched off, illustrating the improvement arising from the push beam, which is typically a factor of 5 for low powers and small detunings.

This low-power, low-detuning region is explored in more detail in Fig. 4.13. The optimum parameters for the pusher are $500\,\mu\mathrm{W}$ of power at a detuning of $+14\,\mathrm{MHz}$ *blue*-detuning, increasing the loading rate of the 3D MOT by more than a factor of 5 relative to the case without the use of a pushing beam. Most groups operate 2D MOTs with a red-detuned push beam, although blue-detunings have been reported (for example, [220]). We believe that blue-detuning works in our case because it helps to accelerate those atoms through the aperture that would otherwise not reach the 3D MOT, due to the opening of the differential pumping tube being manufactured too small. Further investigation has not been undertaken since the current improvement factor from the pusher is satisfactory for our purposes.

After all the above optimisation, in daily operation the 3D MOT is typically loaded for $10\,\mathrm{s}$, capturing 3×10^8 atoms whilst maintaining better than 10^{-10} mbar vacuum in the main science chamber, leading to a magnetic trap lifetime of $\sim 20\,\mathrm{s}$ that is sufficient for evaporative cooling. One of the most convenient features of the 2D MOT is the ability to switch on the dispenser in the morning and leave it running continuously all day. After a period of 20 - 30 minutes to allow the glass cell to fill with rubidium after starting the dispenser, as shown in Fig. 4.14, the 2D MOT provides a reliable loading rate for the whole day, and is independent of the duty cycle of the main experimental sequence.

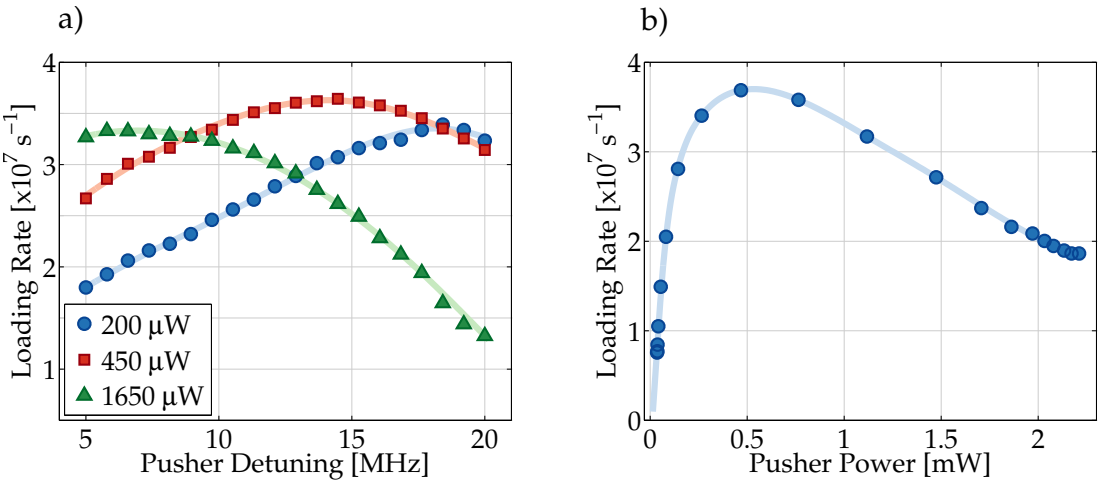

FIGURE 4.13: Fine optimisation of 2D MOT pushing beam: a) The loading rate of atoms into the 3D MOT is measured as a function of the push beam detuning for three different total pusher powers - 200 µW (●), 450 µW (■), and 1.65 mW (▲). b) Loading rate of the 3D MOT as a function of push beam power for a fixed detuning of +14 MHz. The final global optimum push beam parameters are 500 µW of power and a blue-detuning of +14 MHz.

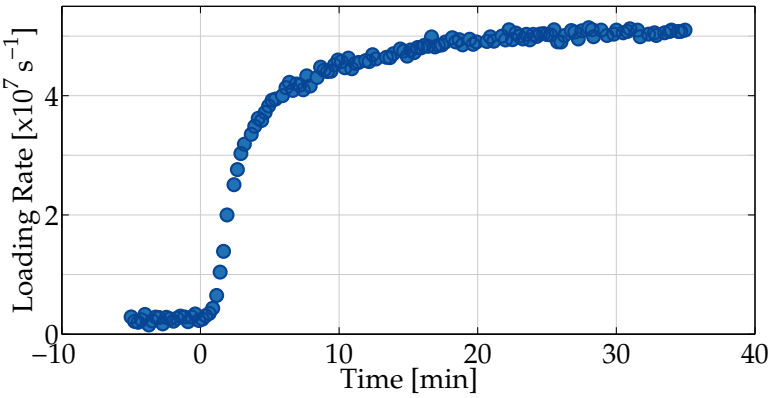

FIGURE 4.14: "Warm up" of the 2D MOT: The loading rate of the 3D MOT is measured as a function of time after switching the 2D MOT dispenser to 3.4 A in the morning. A period of 20 - 30 minutes is required for the glass cell rubidium partial pressure to accumulate.

4.3 Compression and Sub-Doppler Cooling

After loading a sufficient number of atoms into the MOT, the next stage is to transfer to a purely magnetic trap, ready to perform cooling below the recoil limit through evaporative cooling. The distance of the MOT from the atom chip surface at this point is 6 mm, which is where the maximum number of atoms is captured, and the temperature is 250 µK. However, it is not possible to form a sufficiently deep magnetic trap at this distance with the copper Z structure, for any reasonable current. Therefore, we perform two additional steps to improve the transfer. Firstly, the cloud is moved closer to the

chip surface where a deeper trap can be established, and secondly the cloud is sub-Doppler cooled and compressed to allow the maximum transfer efficiency.

Compression Using Magnetic Fields

The cloud height is reduced to $1.5\,\mathrm{mm}$ above the chip surface by increasing the X-bias field from $-5.2\,\mathrm{G}$ to $-15.8\,\mathrm{G}$ at the end of the MOT loading phase. Additonally, the vertical Z-bias field is increased to $3.0\,\mathrm{G}$ to rotate the quadrupole around the longitudinal axis for matching to the position of the magnetic trap minimum. If the beams are reasonably well-balanced, this process can be done rather slowly (over $\sim 1\,\mathrm{s}$) without loss of atoms, but is usually ramped more quickly over $200\,\mathrm{ms}$. This step is critical, as any mismatch between the position of the MOT and the magnetic trap minimum will lead to sloshing and non-adiabatic heating during the transfer, and consequently a loss in phase-space density.

The inhomogeneous fields required for the MOT are produced by a combination of the X-sheet and Sidebars structures, as described in Sec. 4.1 and illustrated in Fig. 4.2. A nice feature with this arrangement is that, because there are separate currents running through each structure, the axial and radial confinements can be changed independently. This is in contrast to the standard U-wire mirror-MOTs, in which a single current is passed in a U-shape, fixing the ratio of the confinements in each direction [212]. In this way, we are able to adjust the shape of the cloud during the ramp phase, by changing the current though the Sidebars, as shown in Fig. 4.15, allowing the aspect ratio to be continuously varied from 2 to 0.5. This may prove to be of use when loading the magnetic trap, as the highest phase-space density is achieved when the effective frequencies of the two traps are matched, although a thorough characterisation beyond the data in Fig. 4.15 has so far not been performed.

Cooling

To further reduce the temperature of the MOT before transfer, we apply a short sub-Doppler cooling stage. The initial experiments with optical molasses observed a temperature six times lower than that predicted by the Doppler cooling theory [221], and several proposals of new cooling mechanisms to explain the observation soon followed [222, 223]. These mechanisms are a result of several effects; firstly that the atoms are not simple two-level systems as assumed for the standard Doppler cooling theory, but that the ground and excited states are comprised of many sublevels amongst which optical pumping can take place. The other important aspect is that the counter-propagating trapping laser beams interfere with each other, giving rise to polarisation gradients. The motion of an atom in these polarisation gradient landscapes gives rise to a stronger frictional force, and therefore a more efficient cooling rate [224]. In our experiment, the trapping beams are oppositely circularly polarised, and so the active cooling mechanism is the so-called "σ^+-σ^-" scheme [222]. This theory predicts a temperature that

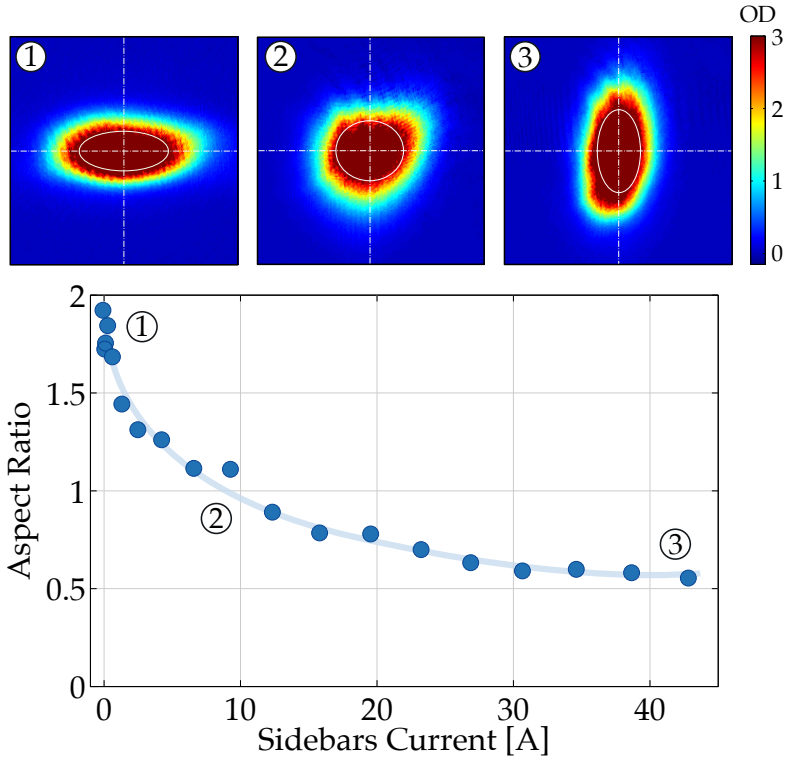

FIGURE 4.15: Variation of MOT aspect ratio: The independence of the currents running through the X-Sheet and Sidebars that produce the MOT quadrupole field in our experiment allows continuous adjustment of the cloud aspect ratio. The lower plot shows the measured aspect ratio - defined as the ratio of the RMS widths from a fitted 2D Gaussian function - as a function of the Sidebars current in the compression stage. The aspect ratio can be tuned smoothly from 2.0 to 0.5. The upper panels show selected examples of corresponding optical density images for aspect ratios of 2.0 ①, 1.0 ②, and 0.5 ③. For all these data, the cloud is imaged in the y-z plane, i.e. the same field of view as in Fig. 4.8 b), and axes definitions can be found in Fig. 4.2.

depends on the laser intensity I and detuning δ as

$$T \propto \frac{I}{\delta}, \tag{4.7}$$

for small detunings and low intensities - a relationship that was soon experimentally verified [225].

In the experiment, the detuning of the trap laser is increased for several milliseconds at the very end of the magnetic field compression ramp. This is automatically accompanied in parallel by a decrease in the beam intensity, as a result of the changing efficiency of the AOM. This reduces the scattering rate of photons inside the cloud, partly alleviating the density limitations from photon pressure, and allowing the cloud to be compressed due to the increased field gradients from the ramp. Many groups use a full optical molasses technique in which the magnetic fields are first extinguished

(see, for example, [73]), since they Zeeman shift the substates and disturb the optical-pumping processes that are involved in sub-Doppler cooling. However, this proves difficult for two reasons: precise cancellation of any stray fields is required; secondly, all beams must be well-balanced, which is difficult to achieve in the mirror-MOT configuration. However, since the atoms in a MOT are confined near the quadrupole field zero, the atoms experience only small magnetic fields and the cooling process can still be very efficient [226, 227]. We see a sufficient drop in temperature with only the detuning technique, which has the advantage that it can be performed whilst not losing magnetic confinement, and therefore maintaining density.

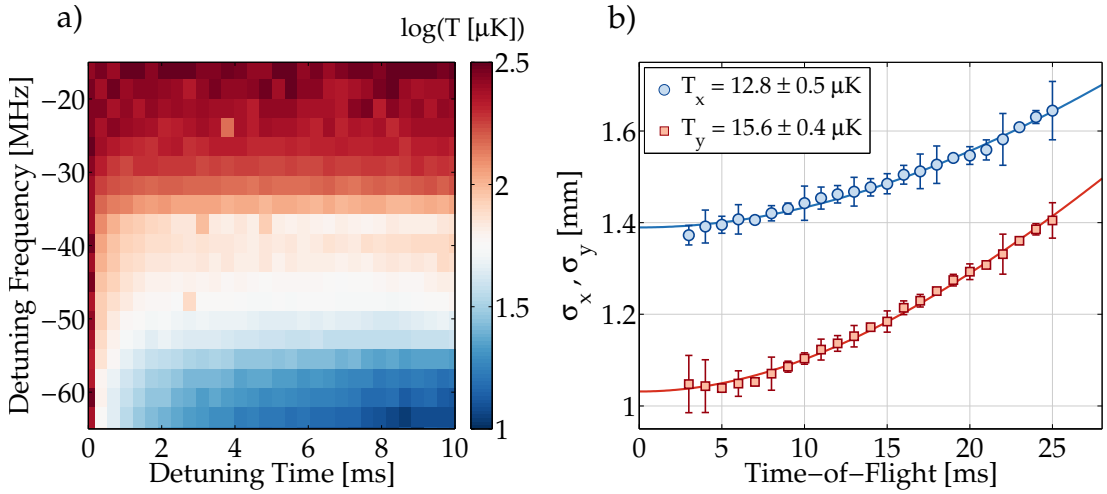

FIGURE 4.16: Sub-Doppler cooling of the MOT: a) The detuning of the cooling beams is ramped linearly from $-16\,\mathrm{MHz}$ to various values up to $-65\,\mathrm{MHz}$, and for ramp durations of up to $10\,\mathrm{ms}$, measuring the 2D parameter space of temperature. The sub-Doppler cooling requires at least several milliseconds to be most effective, whilst larger detunings are always better (for the measured range). b) A time-of-flight measurement for the lowest temperature achieved in (a), with a detuning of $-65\,\mathrm{MHz}$ and a ramp time of $5\,\mathrm{ms}$, showing a minimum temperature of $15\,\mu\mathrm{K}$.

As a result of the optimisation process shown in Fig. 4.16, we typically ramp the cooling beam frequency over $5\,\mathrm{ms}$ at the end of the $200\,\mathrm{ms}$ compression phase from the steady-state value of $-16\,\mathrm{MHz}$ ($-2.6\,\Gamma$) to $-65\,\mathrm{MHz}$ ($-10.7\,\Gamma$), which is the limit of the AOM driver. Since the AOM is optimised for transmission at the initial frequency, the intensity of the beam also greatly reduces automatically in parallel when ramping the frequency. With this additional cooling technique, the MOT is able to be cooled from $250\,\mu\mathrm{K}$ to a value of $15\,\mu\mathrm{K}$ in our setup. The compression of the magnetic fields increases this temperature slightly, and at this point in the cycle, we have $\sim 3 \times 10^8$ atoms at $50\,\mu\mathrm{K}$, positioned at $1.5\,\mathrm{mm}$ from the atom chip surface, which is ideal for transferring to the magnetic trap.

4.4 Optical Pumping

At the end of the MOT loading stage, the atoms are distributed over all five magnetic spin states in the $F = 2$ hyperfine ground state. However, only two of these states (the $|F = 2, m_F = +1\rangle$ and $|F = 2, m_F = +2\rangle$) are low-field seeking and can be recaptured in a purely magnetic trap, as described in Sec. 2.4. Transferring this cloud directly to a magnetic trap would therefore then result in losses on the order of 60% (assuming the trap is capable of capturing both spin states). In addition, atoms in the $|2, +2\rangle$ state will experience twice the strength of confining potential than those in the $|2, +1\rangle$ state, due to its higher magnetic moment. For these reasons, we spin-polarise the cold cloud by transferring the atoms into the $|2, +2\rangle$ state. After loading, this also reduces losses due to collisions between atoms of different spin states.

Polarisation of the sample is achieved by illuminating the laser-cooled cloud with a short pulse of circularly polarised σ^+ light (driving $\Delta m_F = +1$ transitions), whose frequency is tuned to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ hyperfine transition. This results in exchange of angular momentum between the atoms and photons, which tends to align the atomic dipole (the total coupled nuclear and electronic spins, F) with an external magnetic field. This process is known as *optical pumping* [228, 229], and the net result is that atoms are transferred gradually to the state with maximum m_F number (termed the *stretched* state), as shown in Fig. 4.17. Once transferred, they can then no longer absorb the σ^+ optical pumping beam, and so no unnecessary heating occurs.

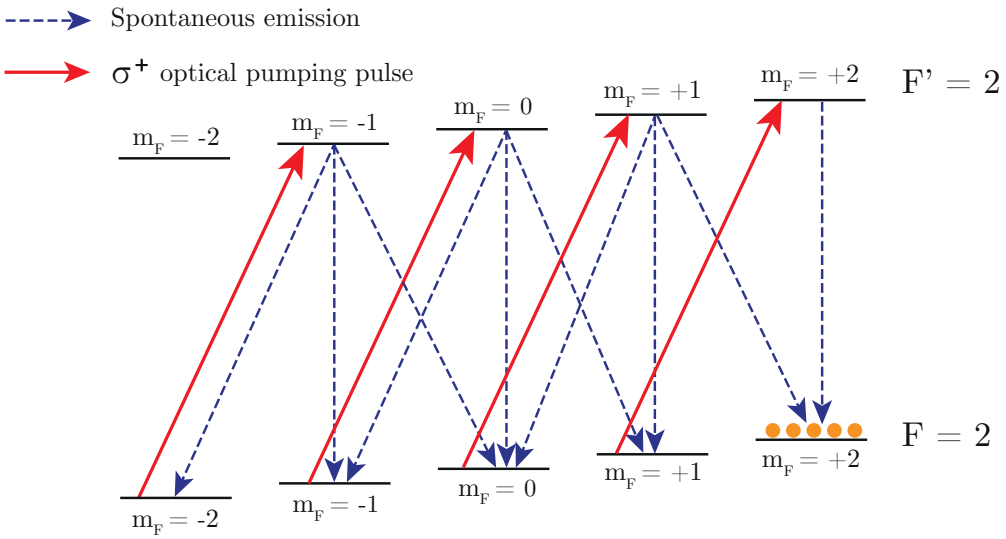

FIGURE 4.17: Optical pumping mechanism: Illumination of the cloud with a σ^+ polarised laser beam allows only absorption on transitions for which $\Delta m_F = +1$. After absorption, each atom can decay by spontaneous emission of a photon of random polarisation, but the net result is for atoms to accumulate in the stretched state $m_F = +2$. Once in this state, it is no longer possible to absorb the σ^+ photons and the sample is polarised with the atomic dipoles aligned with the magnetic field.

The optical pumping beam is derived from a the main TA Pro laser, which is locked to the $|F = 2\rangle \rightarrow |F' = 1 \otimes 3\rangle$ transition, as detailed in Sec. 3.5.1. A single passage AOM then provides a shift of $-75\,\mathrm{MHz}$, leading to a red-detuning of $-18\,\mathrm{MHz}$ from the $|F = 2\rangle \rightarrow |F' = 2\rangle$ line. In order to achieve σ^+ light, all cooling light and magnetic fields are extinguished after the MOT compression and sub-Doppler cooling phase, and then a homogeneous bias field of around $3\,\mathrm{G}$ is applied to provide a quantisation axis for the atoms during illumination with the laser pulse. A polarisation-maintaining fibre is used to bring linearly-polarised light to the vacuum chamber, where it is made circular using a quarter waveplate, and the wavevector of the beam is directed parallel to the external magnetic field, as required to create σ^+ polarisation. Critically, the repumping beam is also turned on simultaneously to pump back atoms which have decayed to the $|F = 1\rangle$ ground state.

4.4.1 Optimisation of Optical Pumping

There are many parameters which have to be optimised in order to achieve efficient optical pumping, thereby transferring as many atoms as possible into the $|2, +2\rangle$ state. These include the duration of the light pulse, the total power in the beam and its intensity, the frequency of the light, and the alignment and polarisation of the beam with respect to the quantisation field.

There are two ways in which these parameters are optimised. The first is to simply look at the improvement in the number of atoms captured in the magnetic trap. Although this method can certainly determine if the optical pumping is working correctly, it cannot distinguish between atoms in the $|2, +1\rangle$ and $|2, +2\rangle$ states, which can both be magnetically trapped. For a more direct measurement of the populations of the substates, we perform a Stern-Gerlach type experiment. This consists of applying a strong magnetic field gradient of $\sim 400\,\mathrm{G\,cm^{-1}}$ to the atoms using the copper Z wire, resulting in a force on each atom that depends on its magnetic substate according to Eq. 2.11. After a short time ($\sim 5\,\mathrm{ms}$) in this magnetic gradient, atoms in different states will separate spatially, allowing measurement of the relative populations, as shown in Fig. 4.18. In order to optimise the optical pumping stage, we wish to have as many atoms as possible in the $|2, +2\rangle$ state. For a quantitative measure, the images are fit with a triple two-dimensional Gaussian funciton which allows the total number of atoms in each sublevel to be extracted. Inefficient pumping results in a small fraction of atoms remaining in the $|2, +1\rangle$ state, and so should be minimised.

To obtain the most pure polarisation of the beam, both the angle of the quarter waveplate and the magnitude of the applied homogenous external field are adjusted, as shown in Fig. 4.19 (a) and (b). If the magnetic field is too small, the quantisation axis is not well defined. Conversely, too large a field results in different resonant frequencies for each transition due to the differential Zeeman splitting, which means that not all states can be addressed for a fixed detuning. The most efficient pumping occurs with a magnetic field of $3\,\mathrm{G}$.

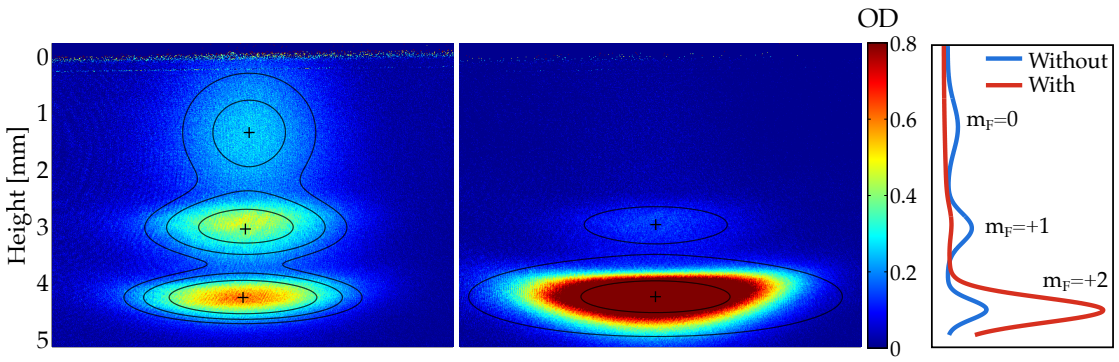

FIGURE 4.18: Stern-Gerlach type experiment: The atoms are positioned at $1\,\mathrm{mm}$ from the atom chip surface and subjected to a magnetic field gradient of $\sim 400\,\mathrm{G\,cm^{-1}}$ for $7\,\mathrm{ms}$. This results in spatial separation of the five magnetic substates due to the different Zeeman energy shifts. The left image shows the optical density distribution without any optical pumping laser pulse. Three of the magnetic states $m_\mathrm{F} = 0, +1, +2$ can be seen, whilst the $m_\mathrm{F} = -1, -2$ have been accelerated into the chip surface. The states $|2, +1\rangle$ and $|2, +2\rangle$ are accelerated downwards, whilst the $|2, 0\rangle$ is unaffected by the field gradient and remains at the original position. Application of a triple two-dimensional Gaussian fit reveals approximately equal atom numbers in each state. The central image shows the case when an optimised optical pumping pulse is applied, resulting in 96% of the atoms being pumped into the stretched state. The rightmost plot shows corresponding line profiles integrated over the horizontal direction for the two cases both with and without applying an optical pumping pulse.

The detuning of the beam with respect to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition is varied in Fig. 4.19 (c), and it is found that an offset of $18\,\mathrm{MHz}$ provides the best transfer of atoms into the $|2, +2\rangle$ state. We have observed that the efficiency does not depend on whether the beam is red or blue detuned, and therefore believe that a frequency detuning is necessary because the cloud is optically thick for resonant light, leading to atoms at the center not being addressed. Indeed, for higher laser powers, we have also seen that resonant light is actually able to blow away atoms from the cloud, along with severely heating it.

Finally, it is important to determine the power and duration of the light pulse required in order to fully optically pump the sample. Both of these parameters should be kept as low as possible, in order to avoid unecessary heating of the sub-Doppler cooled cloud. As shown in Fig. 4.19 (d), a lower power requires a longer pulse, but this leads to a loss in phase space density since the cloud is not confined during this phase of the experimental cycle. We typically illuminate the gas for $3\,\mathrm{ms}$ with a power of 300-$400\,\mu\mathrm{W}$, corresponding to a peak intensity of $\sim 200\,\mu\mathrm{W\,cm^{-2}}$. All the above procedures lead to spin polarisation of the sample with more than 95% of atoms in the $|2, 2\rangle$ state.

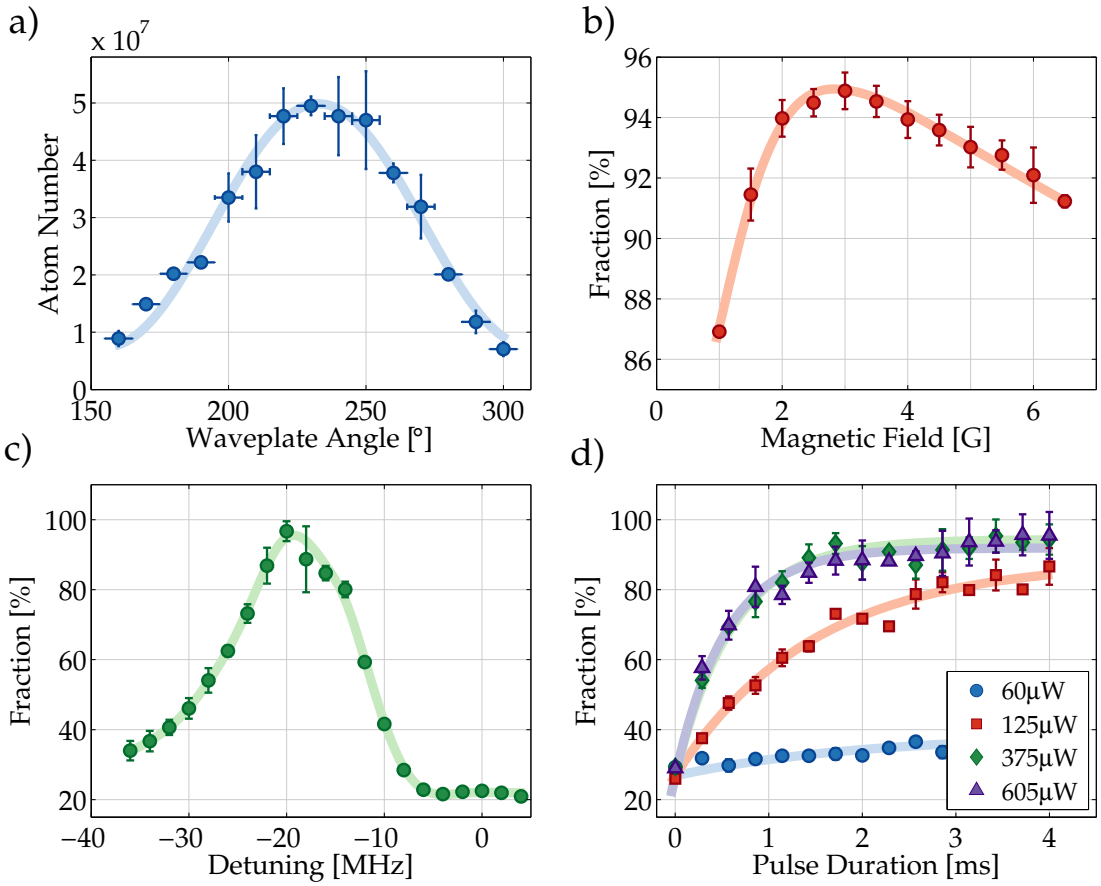

FIGURE 4.19: Optical pumping optimisation: a) The angle of the quarter waveplate is adjusted whilst observing the number of atoms recaptured in the magnetic trap. The largest number of atoms occurs with an angle of 230°. A rotation from this point by 90° flips the handedness of the light and instead produces σ^- polarisation, which optically pumps atoms into the anti-trapped state $|2,-2\rangle$. b) The fraction of atoms in the $|2,+2\rangle$ state, determined using the Stern-Gerlach type experiment, as a function of the applied external quantisation magnetic field. The optimum magnetic field magnitude is 3 G. c) The fraction of the total atom number pumped into the $|2,+2\rangle$ state as a function of the beam detuning from the transition $|F=2\rangle \rightarrow |F'=2\rangle$. The process is most effective with 20 MHz detuning from resonance, as described in the text. d) The fraction of the total atom number pumped into the $|2,+2\rangle$ state as a function of the duration of the laser pulse. The measurement is performed for four different total laser powers, 60 μW (●), 125 μW (■), 375 μW (◆), 605 μW (▲). With smaller laser powers, longer times are required to reach saturation. In the experiment, 400 μW of power is used for 3 ms.

4.5 Macroscopic Copper Trap

After being positioned $\sim 1.5\,\mathrm{mm}$ from the surface of the atom chip, the cloud is transferred into a purely magnetic trap for further cooling. An Ioffe-Pritchard trap, as described in Sec. 2.5.4, is created by a combination of a current though the Z-wire in the copper understructure along with an homogeneous external field using the X-Bias coils. Although it is important to achieve a high elastic collision rate to facilitate efficient evaporative cooling, more importantly at this stage we wish to transfer as many atoms as possible from the MOT whilst minimising any non-adiabatic heating. For this reason, atoms are first transferred into a "mode-matched" magnetic trap, and then adiabatically compressed afterwards. Perfect mode-matching means that the phase-space density is maximised - or even preserved - during the transfer, and the process must be carefully optimised to achieve this. If the trap is too tight the cloud will be heated. On the other hand, the gas will spatially expand if the trap is too weak. Both of these effects are non-adiabatic, and will lead to loss in phase-space density.

The MOT is well-described by a Gaussian density distribution, and the phase-space density will be maximised when loading into a trap with harmonic frequencies

$$\omega_i = 2\pi \times f_i = \sqrt{\frac{k_B T}{m\sigma_i^2}}, \tag{4.8}$$

where σ_i is the RMS width of the MOT in the i^{th} dimension, and T is its temperature [230]. Our MOT has dimensions $\sigma_x = 800\,\mu\mathrm{m}$ and $\sigma_y = \sigma_z = 300\,\mu\mathrm{m}$, and a temperature of $50\,\mu\mathrm{K}$, after compression and sub-Doppler cooling, which means a magnetic potential with trap frequencies of $\{\omega_\perp, \omega_{/\!/}\} \approx 2\pi \times \{37\,\mathrm{Hz}, 14\,\mathrm{Hz}\}$ should be designed. In order to suppress plain evaporation, it should also be ensured that the trap has a depth of at least $\sim 500\,\mu\mathrm{K}$ (assuming an η parameter of ~ 10, as described in Sec. 2.6).

To obtain such low frequencies, the trap is weakened by setting the value of the magnetic field at the trap bottom fairly high. However, this also reduces the trap depth for a given Z-wire current, and so we must compromise slightly in terms of frequency in order to ensure a trap which is sufficiently deep to capture all the atoms. The initial magnetic trap is produced by $80\,\mathrm{A}$ through the copper Z-wire, and $28\,\mathrm{A}(46\,\mathrm{G})$ through the external X-Bias coils. There is no anti-Ioffe field (Y-Bias) applied during loading, in order to have a high field at the trap bottom, with the value of this field being given by the component produced by the legs of the Z-wire structure. A calculation of the resulting trapping potential is shown in Fig. 4.20. This trap has harmonic frequencies of $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{76\,\mathrm{Hz}, 17.8\,\mathrm{Hz}\}$, a residual Ioffe field of $15.1\,\mathrm{G}$, a depth of $1.1\,\mathrm{mK}$, and is located at $1.1\,\mathrm{mm}$ from the surface of the atom chip. Using a Z-wire current of $80\,\mathrm{A}$, it is not possible to achieve perfect mode-matching whilst maintaining depth in the presence of gravitational sag, and we find the trap in Fig. 4.20 to be a good compromise.

With a magnetic trap having the correct properties, the transfer process itself is

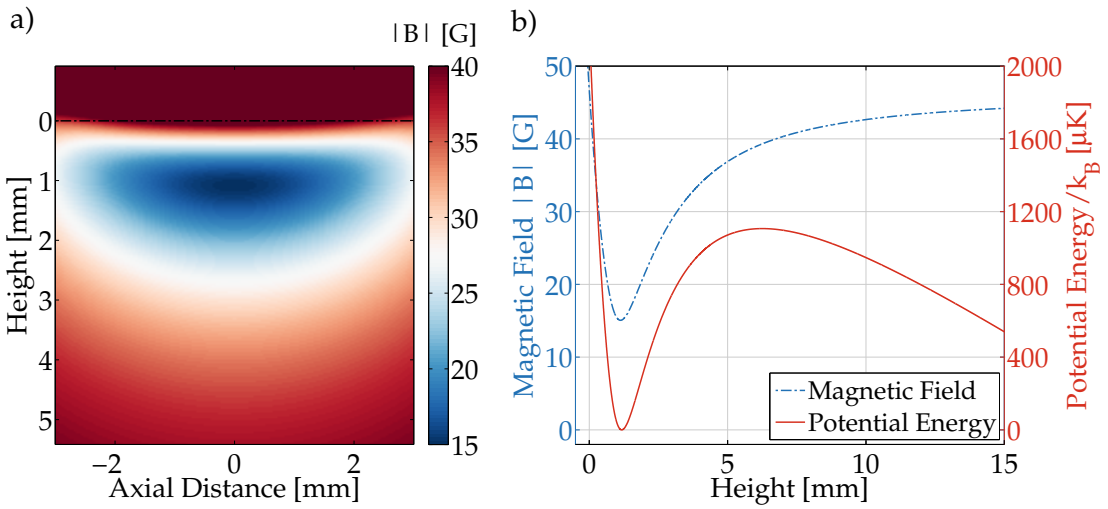

FIGURE 4.20: Calculation of the initial magnetic trap: a) The modulus of the magnetic field |B| produced by $80\,\mathrm{A}$ in the Z-wire structure and an external bias field of $B_x = 46\,\mathrm{G}$, calculated using finite element methods. The value of the magnetic field at the minimum is $|\mathrm{B}_0| = 15\,\mathrm{G}$, and the broken line indicates the surface of the atom chip. b) A line cut at the centre of the axial coordinate, as a function of height from the atom chip surface. The plot shows both the modulus of the magnetic field and the total energy including the gravitational potential, the zero of which has been set to the position of the minimum. The addition of gravity lowers the trap depth to $1.1\,\mathrm{mK}$, and induces a vertical sag of $44\,\mu\mathrm{m}$. This potential has harmonic frequencies $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{76\,\mathrm{Hz}, 17.8\,\mathrm{Hz}\}$, and is located at $1.1\,\mathrm{mm}$ from the chip surface.

optimised. After extinguishing the MOT laser light and fields, and performing optical pumping, the cloud is not confined and begins to expand. The magnetic fields must be switched on as fast as possible, but still slow compared with the Larmor frequency in order to maintain atoms in the correct spin state. The switch on is done by pre-charging the power supplies with a voltage, whilst blocking the current using the home-made current shutters described in Sec. 3.6. When the shutters are opened, the stored energy allows the magnetic field to build up faster, overcoming the inductance of the bias field coil. The current through the less inductive Z-wire is slowed down in order to maintain a constant ratio of currents throughout the transfer, which keeps the position of the magnetic minimum fixed in space. The time taken for the currents to rise is depicted in Fig. 4.21 (a), showing that the magnetic potential is built around the atoms to greater than 90% depth in $7\,\mathrm{ms}$. An optical density image of the atoms $1\,\mathrm{s}$ after switching on the magnetic fields is shown in Fig. 4.21 (b), and agrees very well with the calculated potential in Fig. 4.20.

Special care is also taken to align the spatial positions of the two traps during the transfer. This is usually done by adjusting the position of the laser cooled cloud using small external bias fields, since the position of the magnetic trap is fixed by the other criteria described above. Any mismatch results in excitation of a sloshing motion of the centre-of-mass of the cloud, leading to detrimental heating. In the axial direction, the

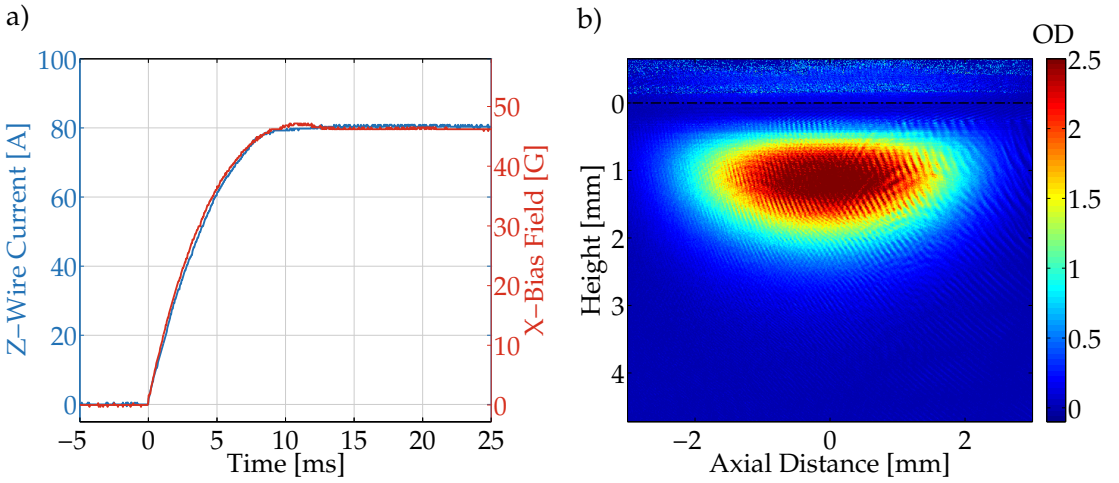

FIGURE 4.21: Switch-on of the initial magnetic trap: a) The rise time for the Z-wire current and X-Bias field is shown, reaching 90% of the maximum after $7\,\mathrm{ms}$. Pre-charging the driving power-supplies beforehand allows overcoming the inductance of the coil. b) Optical density absorption image of the cloud of $\sim 1 \times 10^8$ atoms $1\,\mathrm{s}$ after switching on the fields, measured with $2\,\mathrm{ms}$ time-of-flight. The broken line indicates the surface of the atom chip.

amplitude of the sloshing after loading the magnetic trap provides an indicator of the mismatch, as shown in Fig. 4.22 (b), with the two traps being well-aligned when this amplitude is minimised. For matching in the radial direction, the temperature after loading is minimised whilst adjusting the vertical position of the MOT.

The sloshing motion also allows measurement of the trap frequencies, and by applying sinusoidal fits to the data in Fig. 4.22 (b) the axial frequency is found to be $\omega_{/\!/} = 2\pi \times (18.1 \pm 0.2)$ Hz. To measure the radial frequency, we apply a small non-adiabatic change to the position of the trapping potential by pulsing on the current through a chip wire. This gives a small kick to the cloud, which then oscillates inside the potential at the trap frequency. Fig. 4.22 (a) shows the result of this measurement, indicating a frequency of $\omega_\perp = 2\pi \times (71 \pm 3)$ Hz. Both measured frequencies are compatible with the calculated values in Fig. 4.20.

The absolute value of the magnetic field at the trap minimum is also measured using the so-called *RF knife* technique, shown in Fig. 4.23 (a). Here, an RF field at constant frequency is applied using the copper Z-wire for $5\,\mathrm{s}$, which causes outcoupling of atoms from the trap due to spin flips depending on their energy above the trap bottom, as described in Sec. 3.8. As the RF knife frequency is scanned closer to the trap bottom, more and more atoms are removed, until the frequency creates a hole in bottom of the trap removing all atoms. An RF frequency below the trap bottom is not resonant with any atoms, and so the number remaining is not affected by the knife. The measurement shows a bottom resonance frequency of $\nu_0 = (10.3 \pm 0.5)\,\mathrm{MHz}$, which corresponds to a magnetic field of $|\mathrm{B}_0| = (14.7 \pm 0.7)\,\mathrm{G}$ by using Eq. 3.10, agreeing with the expected value from the simulation.

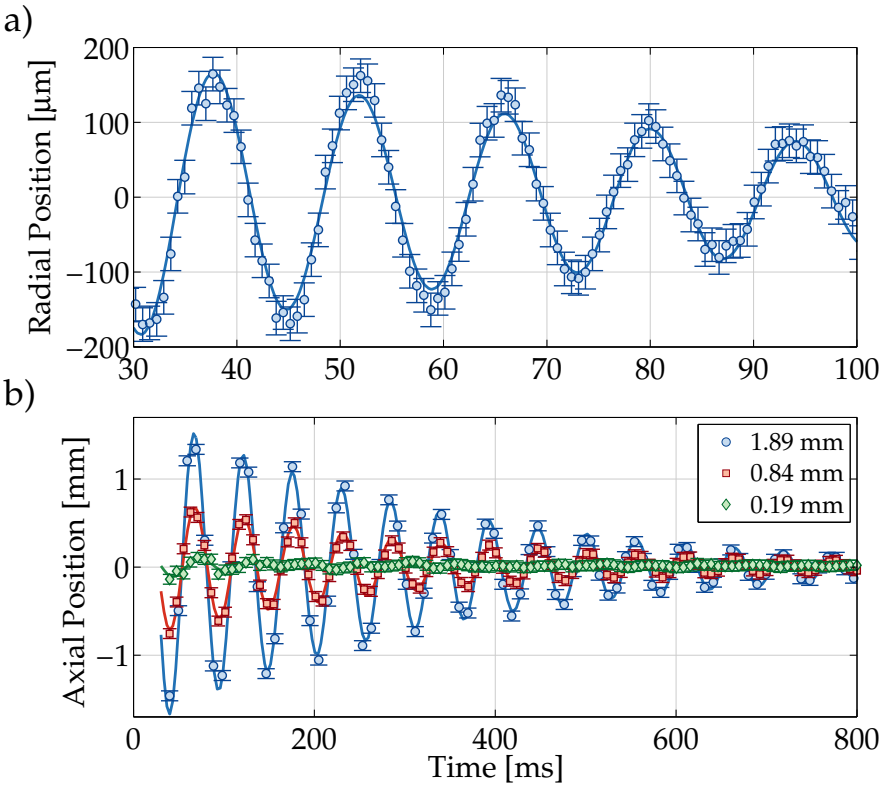

FIGURE 4.22: Initial magnetic trap frequencies: a) The radial trap frequency is measured by applying a magnetic kick to displace the cloud inside the trapping potential, which is done by pulsing a 300 mA current through an atom chip wire for 300 µs. The measurement is performed after evaporative cooling the cloud to $\sim$ 8 µK to ensure that the oscillation samples only the harmonic part of the potential. A sinusoidal fit determines the frequency to be $\omega_\perp = 2\pi \times (71 \pm 3)$ Hz. b) The axial frequency is measured by observing the centre-of-mass sloshing motion of the cloud immediately after transfer. The oscillation is measured for three different relative offsets in axial position, and the two traps can be aligned by minimising the overall amplitude. A sinusoidal fit gives an axial trap frequency of $\omega_{/\!/} = 2\pi \times (18.1 \pm 0.2)$ Hz.

Finally, it is important to know the lifetime of the magnetically trapped atoms, as this influences the efficiency of the evaporative cooling. For the relatively low densities in this initial trap, the loss is dominated by collisions with hot, untrapped atoms, which depends on the background vacuum pressure. This loss rate is independent of the cloud density, but is proportional to the total number of trapped atoms, N, which should therefore decay exponentially according to

$$N(t) = N(0) \exp\left(-\frac{t}{\tau_{\mathrm{bg}}}\right), \tag{4.9}$$

where $N(0)$ is the initial number of atoms trapped and τ_{bg} is the lifetime. The lifetime can be determined by measuring the number of atoms remaining after a variable holding time in the magnetic trap, as shown in Fig. 4.23 (b). With careful optimisation,

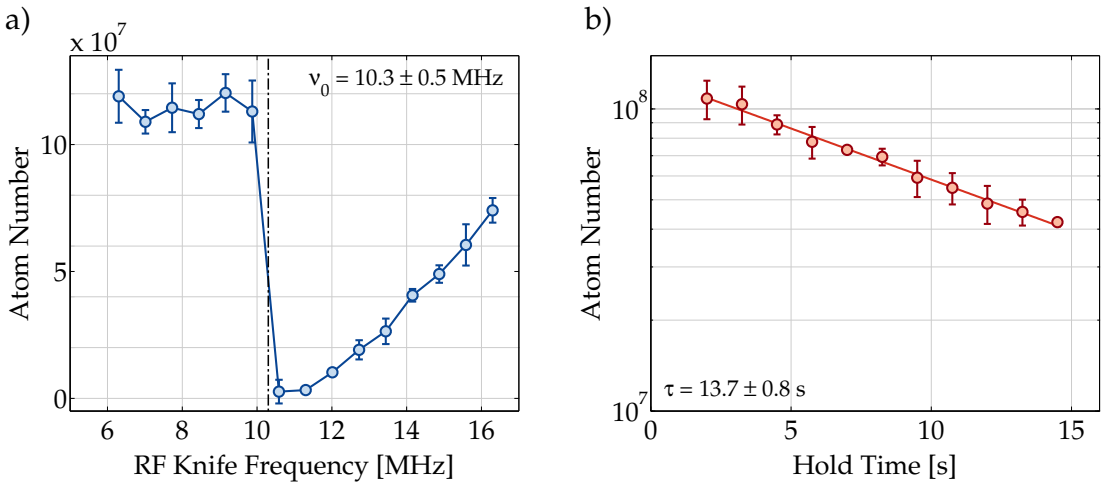

FIGURE 4.23: a) Measurement of the absolute value of the magnetic field at the trap minimum for the initial mode-matched trap, using the RF knife technique described in the text. The frequency which is resonant with the trap bottom is measured to be $\nu_0 = (10.3 \pm 0.5)$ MHz, corresponding to a magnetic field of $|B_0| = (14.7 \pm 0.7)$ G. b) Example of a measurement of the trap lifetime, by measuring the number of atoms remaining after a variable hold time, with an exponential fit giving a lifetime of (13.7 ± 0.8) s

we are routinely able to transfer 1.5×10^8 atoms with an equilibrium temperature of $120\,\mu$K and a lifetime in the trap of $\tau_{bg} \sim 15\,$s, depending on dispenser settings and present vacuum condition. These conditions are comparable to - and in many cases exceed - those achieved in similar experimental setups utilising atom chips (see, for example, [174, 231–233]).

4.5.1 Adiabatic Compression

After transferring atoms into the mode-matched magnetic trap, it is then compressed to increase the density, and therefore the elastic collision rate, to facilitate efficient evaporative cooling. The elastic collision rate determines the time required for the sample to rethermalise, and therefore the rate at which evaporative cooling can be performed, as detailed in Sec. 2.6.

The trap is compressed by increasing the strength of the confining fields, which increases the gradients. There are two practical limits on how much compression can be applied. The first is set by the amount of current that can safely be passed through the copper Z-wire, which is typically kept below $\sim 100\,$A since it can expand through resistive heating for excessive currents (we have even observed a degradation of the vacuum quality when passing high currents for prolonged periods of time). Secondly, increasing the X-Bias field moves the cloud closer to the atom chip surface, which will result in inefficient surface-induced evaporation if it is brought too close [61]. Our strategy is therefore to set the Z-wire current to its safest maximum at $80\,$A, and then

to gradually compress the cloud with the bias fields as it cools and reduces in size, keeping the atoms as close to the chip surface as possible without loss.

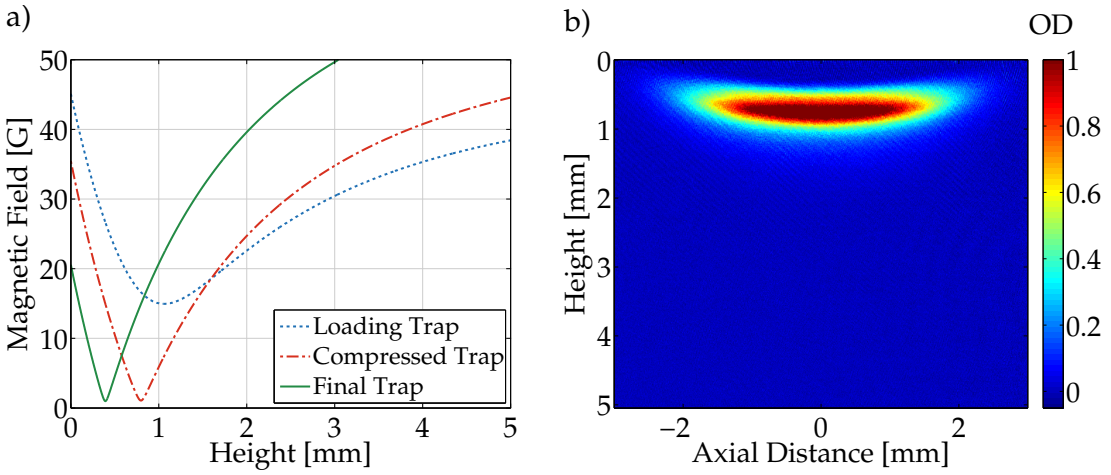

FIGURE 4.24: Magnetic trap compression: a) The calculated total magnetic field modulus as a function of height above the atom chip surface for the mode-matched (dotted), initial compressed (dash-dot line), and final (solid) traps. The first stage pushes the trap bottom down to $\sim 1\,\mathrm{G}$, changing the potential from harmonic to linear. Ultimately, the trap frequencies are increased from $2\pi \times \{76\,\mathrm{Hz}, 17.8\,\mathrm{Hz}\}$ to $2\pi \times \{460\,\mathrm{Hz}, 22\,\mathrm{Hz}\}$, and the final gradient is $420\,\mathrm{G\,cm^{-1}}$. b) An optical density absorption image of the cloud *in-situ* after the first compression stage - the "Compressed Trap" in (a). The height is $800\,\mu\mathrm{m}$, and cannot be reduced at this stage in order to prevent atoms hitting the chip surface.

After loading from the MOT, the cloud has a temperature of $120\,\mu\mathrm{K}$, with trap frequencies $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{76\,\mathrm{Hz}, 17.8\,\mathrm{Hz}\}$ in the mode-matched trap. The trap bottom is $15\,\mathrm{G}$ and the minimum is $1.1\,\mathrm{mm}$ from the chip surface, as described in Sec. 4.5 above. After allowing $100\,\mathrm{ms}$ settling time in this trap, the main X-Bias field is ramped up over $500\,\mathrm{ms}$ from $46\,\mathrm{G}$ to $55\,\mathrm{G}$, moving the trap minimum to $800\,\mu\mathrm{m}$ from the chip surface. Simultaneously, the anti-Ioffe Y-Bias field is ramped up to $13.3\,\mathrm{G}$, which cancels the field produced by the legs of the Z-wire structure and pushes the trap bottom down to $\sim 1\,\mathrm{G}$, and is enough to suppress Majorana spin-flip losses. A comparison of the potential change into this "Compressed Trap" is shown in Fig. 4.24 (a), which shows that the atoms now experience a linear potential, instead of harmonic, with a gradient of $295\,\mathrm{G\,cm^{-1}}$.

Together, these two changes to the trapping potential adiabatically increase the temperature to $300\,\mu\mathrm{K}$. The peak density of the cloud is increased to $n_0 = 1.5 \times 10^{11}\mathrm{cm^{-3}}$, which boosts the elastic collision rate to $\Gamma_{\mathrm{el}} \sim 30\,\mathrm{s^{-1}}$, according to Eq. 2.36. The dominant loss rate is one-body collisions with the background gas, given by the reciprocal of the trap lifetime, which in our case is $\Gamma_{\mathrm{bg}} = 1/\tau_{\mathrm{bg}} \sim 0.07\,\mathrm{s^{-1}}$. The so-called "ratio of good to bad collisions" after compresion is therefore $\Gamma_{\mathrm{el}}/\Gamma_{\mathrm{bg}} \sim 400$, which is within the regime in which it is possible to achieve runaway evaporation described in Sec. 2.6.2.

After the initial compression stage, the depth of the trap is given by the magnetic

potential energy barrier on the chip surface, which is calculated to be $1.3\,\mathrm{mK}$. The barrier to the gravity side is much higher at $2.6\,\mathrm{mK}$, and so atoms that are too hot will be lost to the chip surface. For this reason, the trap is not positioned any closer to the surface at this point, as shown in Fig. 4.24 (b). However, as the cloud is cooled through evaporative cooling, described later in Sec. 4.6, its spatial extension is reduced and more compression can be applied. The trap is linearly compressed throughout the ramp by increasing the X-Bias from $55\,\mathrm{G}$ to $70\,\mathrm{G}$, and reducing the anti-Ioffe Y-Bias to $12.1\,\mathrm{G}$ accordingly to maintain a $\sim 1\,\mathrm{G}$ minimum field. This second stage increases the trap frequencies to $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{460\,\mathrm{Hz}, 22\,\mathrm{Hz}\}$, the gradient to $420\,\mathrm{G\,cm^{-1}}$, and brings the centre-of-mass to a height of $360\,\mu\mathrm{m}$ above the chip surface.

A measurement of the final radial trap frequencies is given in Fig. 4.25, and was obtained by two methods. In Fig. 4.25 (a), the position of the cloud was modulated by applying an AC current through an atom chip wire, which leads to atom loss as a result of paramterically-induced heating when the driving frequency matches the harmonic trap frequencies. For comparison, Fig. 4.25 (b) shows a frequency measurement by observing a centre-of-mass oscillation, with both methods agreeing and being consistent with the calculated value.

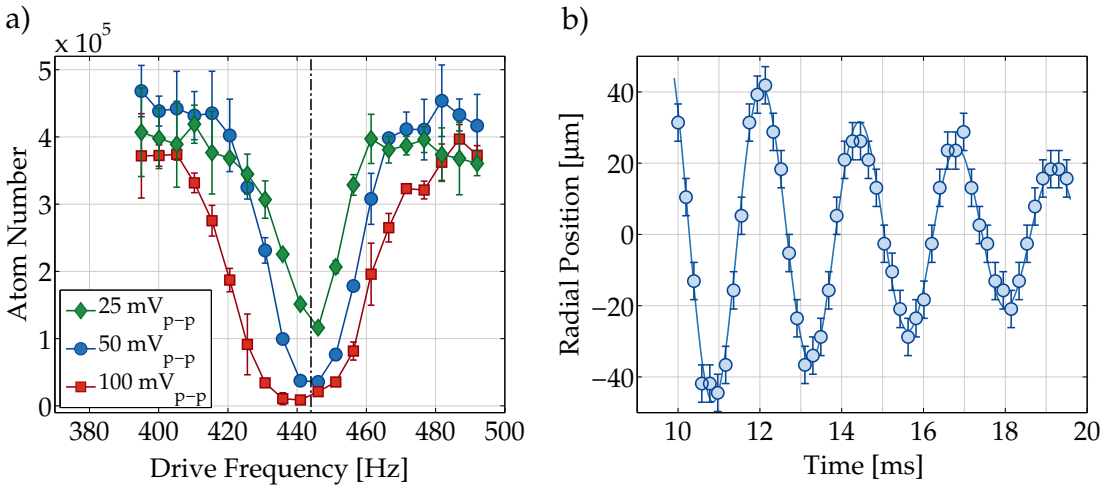

FIGURE 4.25: Radial trap frequencies of the final compressed trap: a) The number of atoms remaining in a cloud pre-cooled to $1\,\mu\mathrm{K}$ is measured as a function of the driving frequency of an AC current pulsed for $1\,\mathrm{s}$ using one of the atom chip Z-shaped wires. When the drive frequency becomes resonant with the harmonic trap frequency, atoms are parametrically excited and heated out of the trap. Three different peak-to-peak voltage amplitudes for the driving signal were tested - $25\,\mathrm{mV_{p\text{-}p}}$ (◆), $50\,\mathrm{mV_{p\text{-}p}}$ (●), $100\,\mathrm{mV_{p\text{-}p}}$ (■) - with larger values leading to a broadening of the resonance. The frequency is measured to be $\omega_\perp = 2\pi \times (443 \pm 5)$ Hz. b) The radial trap frequency as measured by pulsing a $300\,\mathrm{mA}$ current through an atom chip wire for $300\,\mu\mathrm{s}$ and observing the centre-of-mass oscillation of a $\sim 5\,\mu\mathrm{K}$ induced by the magnetic kick. A sinusoidal fit yields a frequency of $\omega_\perp = 2\pi \times (420 \pm 3)$ Hz, consistent with both the value measured in (a) and the expected result of the simulation.

4.6 Evaporative Cooling

After having compressed the magnetic trap to increase the elastic collision rate, the phase-space density of the gas is further increased using the technique of forced evaporative cooling. This is done by energy-selectively outcoupling atoms from the tail of the Maxwell-Boltzmann velocity distribution, by applying a radio-frequency magnetic field to drive transitions between trapped and anti-trapped Zeeman sublevels of the $|F = 2\rangle$ hyperfine ground state, as described in Sec. 3.8. For a constant RF knife, the evaporation rate will reduce exponentially according to Eq. 2.35 [153] until eventually becoming effectively suppressed completely when $\eta \gtrsim 10$. Therefore, the RF frequency, and consequently the trap depth, must be lowered iteratively in order to continue the evaporation. In practice, the frequency is ramped continuously from $20\,\mathrm{MHz}$ to $\sim 1\,\mathrm{MHz}$, depending on the set value of the trap bottom.

In order to obtain the largest increase in phase-space density for minimal loss of atoms, there are several parameters which must optimised. The shape of the RF frequency ramp should be tailored in order to maintain the optimum value for the truncation parameter of $\eta \sim 6 - 8$. If the elastic collision rate and loss rate were constant throughout the process, then the RF frequency should be lowered exponentially in order to remove the same fraction of energy per unit time, ensuring that the energy distribution is only rescaled. In practice, both rates change as the atom number and temperature are reduced and so the ramp is adjusted experimentally.

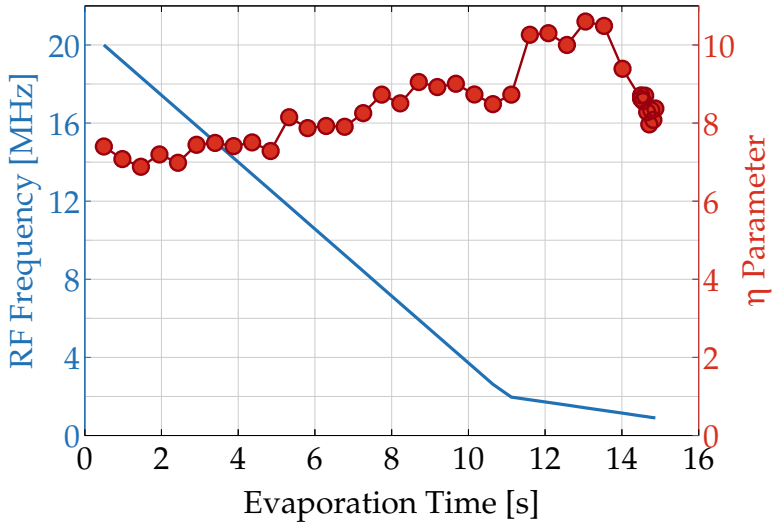

FIGURE 4.26: Truncation of the energy distribution: The trap depth is reduced by using two linear ramps of the RF frequency from $20\,\mathrm{MHz}$ to $\sim 1\,\mathrm{MHz}$ (blue solid line, left axis). By interrupting the ramp at various times and measuring the temperature, the truncation parameter η is calculated (red markers, right axis).

There have been several numerical studies into the efficiency of the evaporative cooling process, resulting in non-exponential frequency ramps in general, but which

can be well-approximated by linear functions [234, 235]. We find that the cooling efficiency is relatively insensitive to the specifics of the ramp shape, but is more influenced by the total duration of the sweep. An example of a typical frequency ramp is shown in Fig. 4.26, which uses two linear segments over $15\,\mathrm{s}$. By interrupting the cooling at various times and measuring the temperature, the truncation parameter is also calculated using Eq. 3.11 and is shown on the plot. For this particular example, η slowly grows from 7 to 10 over the first $13\,\mathrm{s}$, which is an indicator that the rethermalisation is becoming more efficient and the ramp could afford to be performed faster. In the final $2\,\mathrm{s}$ there is a sharp drop in η, which points to an increasing temperature since the truncation energy is relatively constant over this time. It has been shown using a quantum kinetic simulation that it is advantageous to *accelerate* the cooling as the gas approaches degeneracy to counteract the temperature increase arising from density-dependent three-body losses [235], but this has not as yet been tested in our setup.

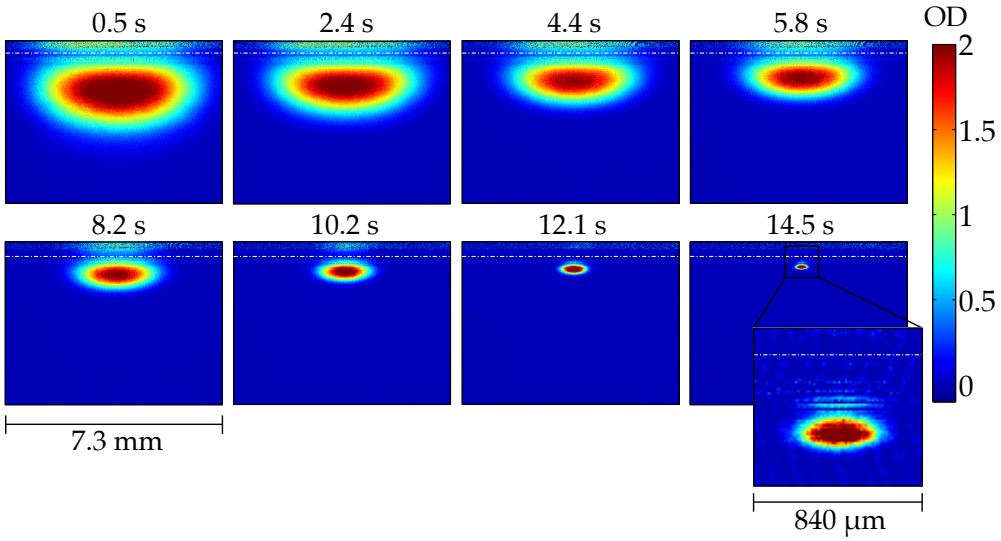

FIGURE 4.27: Images throughout the evaporation ramp: Optical density absorption images are recorded by interrupting the $15\,\mathrm{s}$ long cooling process at various times, in order to ensure that the cloud does not make contact with the chip surface (indicated by the dashed-dot line).

As described in Sec. 4.5.1, with the Z-wire current fixed at $80\,\mathrm{A}$ the X-Bias field is linearly increased over the duration of the ramp from $55\,\mathrm{G}$ to $70\,\mathrm{G}$. Figure 4.27 shows the position of the cloud minimum being moved from $800\,\mu\mathrm{m}$ to $360\,\mu\mathrm{m}$ from the atom chip surface during this process. This allows maximum compression and collision rate to be maintained throughout the sweep, whilst ensuring there is only minimal loss due to inefficient plain evaporation to the surface. These images demonstrate the different spatial scales involved throughout the process, with the cloud extension being reduced from $\sim 7\,\mathrm{mm}$ to several hundred $\mu\mathrm{m}$, and justifies switching to the second imaging "Andor System" described in Sec. 3.9.1.

The total duration of the entire evaporative cooling sequence is critical to its efficiency. The only timescales involved in the process are given by the loss rate due to

collisions with the background gas, and by the rethermalisation time required due to the elastic collision rate. If the ramp is too slow, then overwhelming losses due to background collisions will cause a reduction in phase-space density. On the other hand, if the ramp proceeds too quickly the RF knife will cut too deep into the velocity distribution (correspnding to a small η parameter), and atoms will be removed which would otherwise have rethermalised. This amounts then only to velocity selection, and again reduces the efficiency of the cooling process.

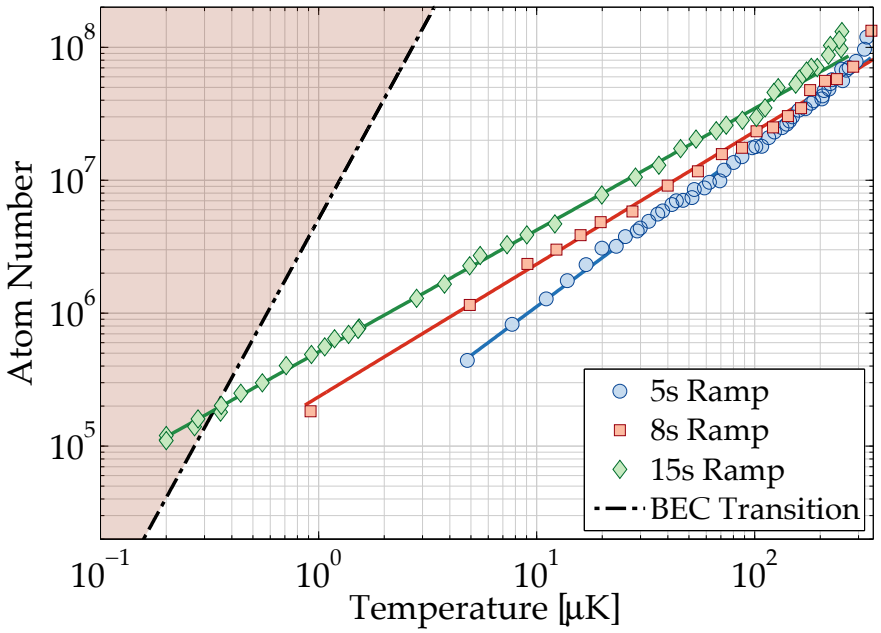

FIGURE 4.28: Efficiency of evaporation: The temperature and remaining atom number are shown for three different cooling durations - 5 s (○), 8 s (□), and 15 s (◇). To allow direct comparison between the various cooling times, the shape of the RF frequency ramp was simply chosen to be *linear* in all cases, being reduced from 20 MHz to several tens of kHz above the trap bottom. A shallower slope indicates more efficient cooling that is less costly in atom number. Only the 15 s ramp could be continued into the BEC regime (shaded region), given by the condition PSD = 2.612.

The reduction in temperature for a given loss of atoms is shown in Fig. 4.28, and serves as our main diagnostic for the efficiency of the cooling phase. In order to measure this, the sequence was interrupted at successive times during the cooling, to measure both the temperature and atom number remaining. The linear slope, s, of these data is used to assess how effective the cooling is, with a smaller slope indicating smaller atom loss for each unit reduction in temperature. Several durations of the entire ramp were tested, from 5 s to 15 s, as shown in the figure, with the 15 s ramp producing the highest number of atoms at the end of the ramp, having a slope of $s = 0.9$. The duration is not increased beyond 15 s in order to both limit the Ohmic heating of the copper Z-wire, and to maximise the repetition rate of experimental runs. Only the 15 s ramp could be extended into the quantum degeneracy regime, as for shorter times

the phase-space density began to drop upon continued cooling. This minimum time required is ultimately set by the elastic collision rate, and therefore the amount of compression applied to the trap, which in the end is given by the current that can be safely passed through the understructure wire.

The time-evolution of salient properties of the trapped gas are shown in Fig. 4.29 for this $15\,\mathrm{s}$ ramp. Again for this data series, the RF sweep is interrupted after various times, in order to measure temperature and atom number remaining. These are then used in conjunction with the trap frequencies to calculate the *in-situ* size using $\sigma_i = \sqrt{k_B T/m\omega_i^2}$, from which the phase-space density can be determined. Here, the RF frequency was ramped from $20\,\mathrm{MHz}$ to several tens of kHz above the trap bottom, reducing the temprature from $300\,\mu\mathrm{K}$ to $\sim 400\,\mathrm{nK}$, at the expense of a loss in atom number from 1×10^8 to 2×10^5. Throughout the ramp, the phase-space density steadily grows by six orders of magnitude, illustrating the power of the evaporative cooling technique.

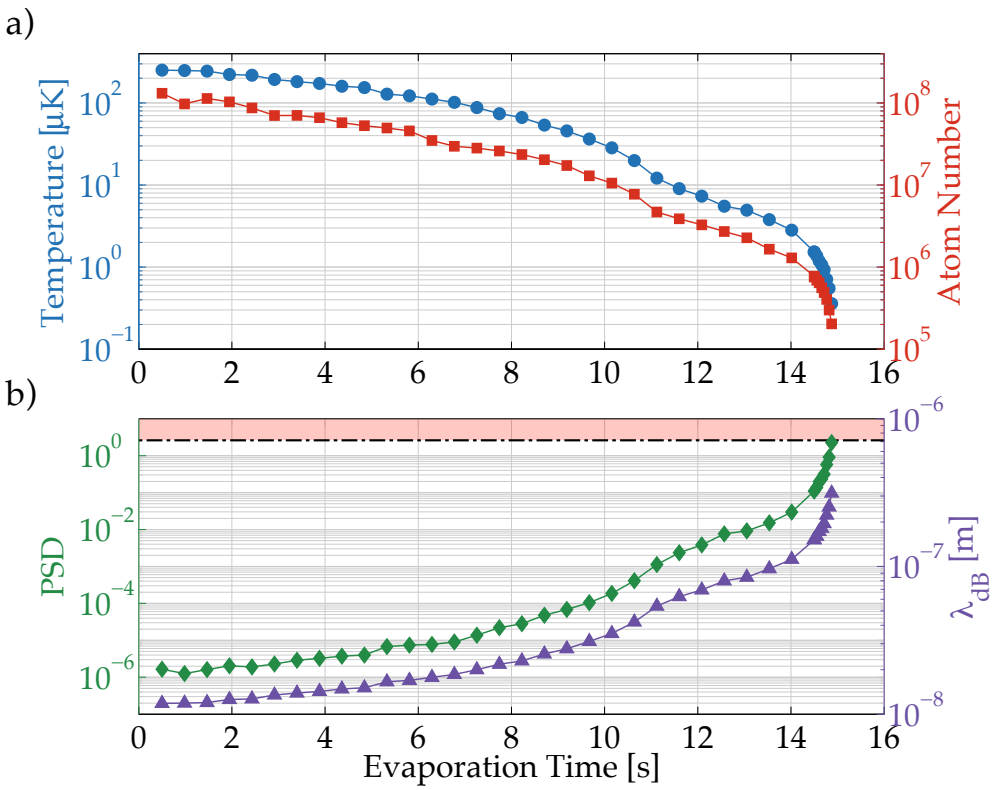

FIGURE 4.29: Evolution of the cloud properties: a) The temperature (●, left axis) is reduced from $300\,\mu\mathrm{K}$ to $400\,\mathrm{nK}$, whilst the atom number (■, right axis) is reduced from 1×10^8 to 2×10^5. b) The phase-space density (◆, left axis) increases from 10^{-6} to ~ 1, and the de Broglie wavelength (▲, right axis) is increased from $10\,\mathrm{nm}$ to $0.3\,\mu\mathrm{m}$. The shaded region denotes the BEC transition on the phase-space density axis, with PSD = 2.612.

Finally, using evaporatively cooled clouds, the trap bottom can be characterised

more thoroughly, as shown in Fig. 4.30. The magnetic field at the minimum is measured using the RF knife technique to be $(795 \pm 10\,\mathrm{kHz})$ in this particular case, corresponding to $(1.14 \pm 0.01\,\mathrm{G})$. We find that the most efficient cooling occurs when the trap bottom is placed at $800\,\mathrm{kHz} - 900\,\mathrm{kHz}$, since for lower values the RF power is greatly attenuated due to pushing the various components in the line to the limit of their bandwidth. The measurement is performed for several values of the RF field amplitude, which is controlled using the voltage-variable attenuator in the RF line, as described in Sec. 3.8. The absolute value of the magnetic field at the site of the cloud is not known, since this is given by the amount of power radiated by the copper Z-wire, which is sensitive to impedence matching at the vacuum feedthrough. However, the data allow comparison between relative amplitudes, and it can be seen that for larger values the resonance at the trap bottom is power-broadened. For this reason, the amplitude is typically ramped down in the final $1\,\mathrm{s}$ of the cooling sweep to ensure that the RF knife possesses the necessary resolution for cooling cleanly all the way down to the trap bottom.

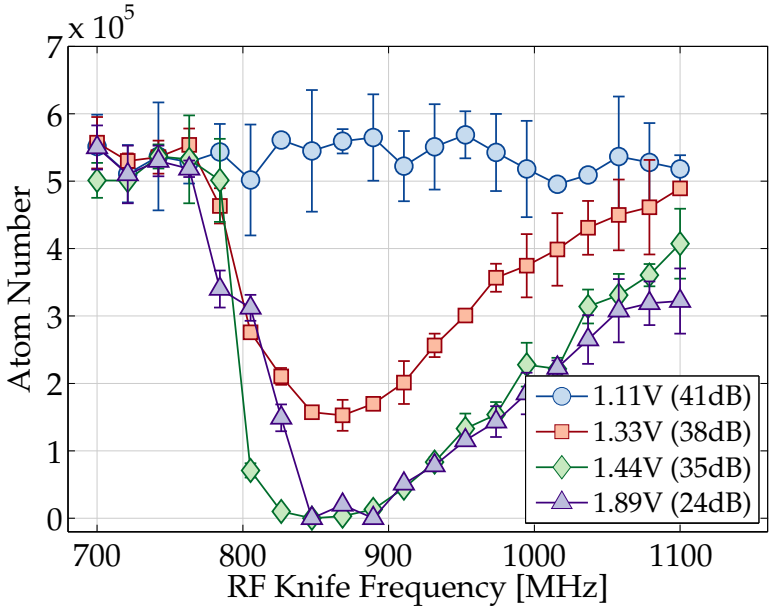

FIGURE 4.30: RF knife measurement of the trap bottom: The resonance frequency at the trap minimum is measured to be $(795 \pm 10\,\mathrm{kHz})$, corresponding to an absolute value of the magnetic field of $(1.14 \pm 0.01\,\mathrm{G})$, by performing an RF knife measurement with a $\sim 5\,\mu\mathrm{K}$ cloud. Four different amplitudes of the RF are used, corresponding to four control voltages on the voltage-controlled attenuator - $1.11\,\mathrm{V}$ (○), $1.33\,\mathrm{V}$ (□), $1.44\,\mathrm{V}$ (◇), and $1.89\,\mathrm{V}$ (△). These control voltages lead respectively to values of attenuation of $41\,\mathrm{dB}$, $38\,\mathrm{dB}$, $35\,\mathrm{dB}$ and $24\,\mathrm{dB}$. For small amplitudes, the RF has no effect on the trapped atoms, whilst for larger amplitudes the resonance is power-broadened. The control voltage required to completely just empty the trap is $\sim 1.4\,\mathrm{V}$.

With the ability to now create very cold samples using evaporative cooling, the lifetime of the atoms in the copper-Z trap can be more accurately determined. A refined

lifetime measurement is shown in Fig. 4.31, in which a cold cloud at a temperature of $1\,\mu\mathrm{K}$ was prepared by evaporative cooling in a $1.1\,\mathrm{mK}$ deep trap created using the copper-Z. The RF field was then removed and the cloud allowed to evolve in the trap for variable hold times of up to one minute. The data fit well to an exponential function, indicating density-independent one-body losses which result primarily from collisions with untrapped background atoms in the vacuum chamber. The observed time constant of $\tau = (37 \pm 1)\ \mathrm{s}$ is sufficient for our experimental timescales, and shows that the vacuum pressure of 3×10^{-11} mbar obtained in Sec. 3.1 is satisfactory. In addition, the evolution of the temperature during the hold time is shown in the inset of Fig. 4.31, from which a heating rate of $(19 \pm 1)\ \mathrm{nK\,s^{-1}}$ is deduced from the linear fit, and is sufficiently small so as not to inhibit efficient evaporation.

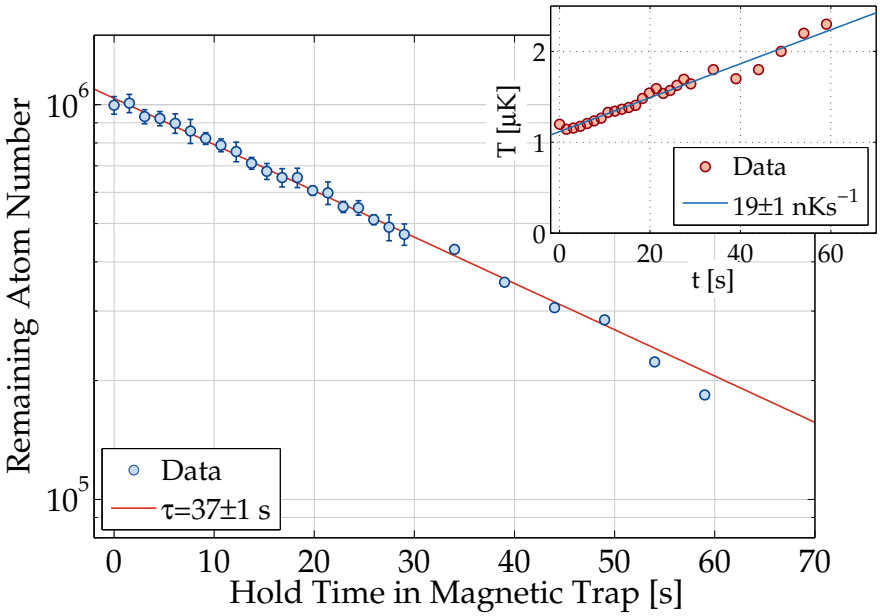

FIGURE 4.31: Lifetime and Heating Rate Measurement Using a Cold Cloud: A $1\,\mu\mathrm{K}$ cloud is prepared in an "infinitely deep" trap using the copper-Z, and the atom number remaining is plotted for variable hold times . An exponential fit indicates a vacuum-limited time constant of $\tau = (37 \pm 1)\ \mathrm{s}$. Inset: The evolution of the temperature T of the cloud during the hold time shows a heating rate of $(19 \pm 1)\ \mathrm{nK\,s^{-1}}$.

This lifetime measurement should be compared with that in Fig. 4.23 (b) in Sec. 4.5, which was also performed in a $1.1\,\mathrm{mK}$ deep trap but using a hotter $120\,\mu\mathrm{K}$ cloud, and showed a much shorter lifetime of $\tau = 14\,\mathrm{s}$. The difference between the two measurements can be attributed to the "spilling" of hot atoms out of the trap that happens when using the hotter cloud, which significantly reduces the observed lifetime. In contrast, in this refined measurement the trap is essentially infinitely deep for the $1\,\mu\mathrm{K}$ cloud. These measurements illustrate the importance of accounting for the trap depth by using sufficiently cold samples when trying to accurately determine the vacuum-limited lifetime.

4.7 Initial Observation of Bose-Einstein Condensation

An indication that the cooling process is working efficiently is a more pronounced increase in optical density after time-of-flight, since the peak optical density measured at large expansion times is directly proportional to the elastic collision rate [236], the important quantity for reaching runaway evaporation. When we detect an increasing optical density at the end of the evaporation ramp, and the system is clearly producing cold clouds with temperatures of $\sim 1\,\mu\mathrm{K}$, we begin looking for evidence of condensation taking place, whilst gradually lowering the final frequency of the RF. Figure 4.32 shows the first image of a BEC obtained in our lab, on 23rd September 2014. Each image was taken with an expansion time of $20\,\mathrm{ms}$ after fast switch-off of the trapping fields. Upon lowering the final value of the RF frequency by only $25\,\mathrm{kHz}$, a sharp peak appeaks at the centre of the density distribution, indicating a large fraction of atoms in lower momentum states, and signifying the onset of BEC. In addition, the integrated optical density profile displays a change in the distribution from Gaussian to a bimodal shape, which is characteristic of a BEC in the Thomas-Fermi regime with a residual thermal fraction.

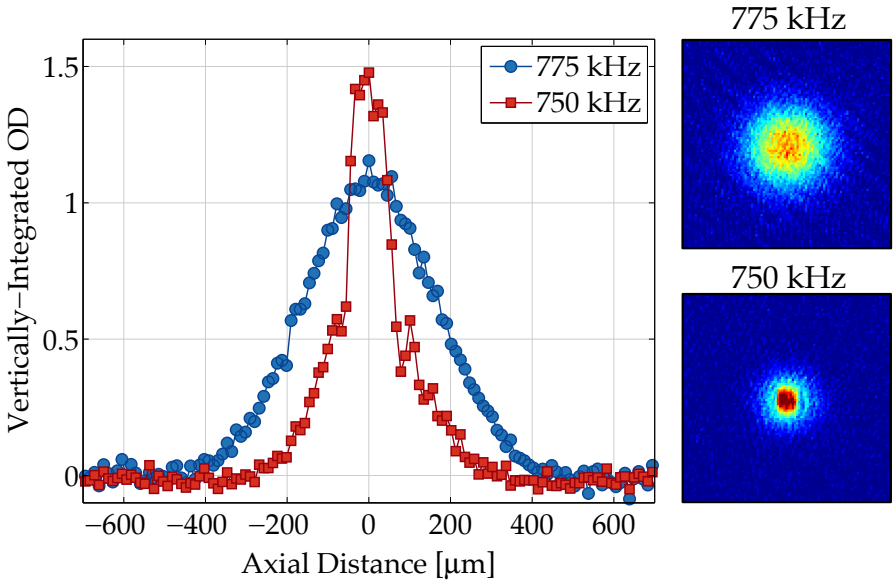

FIGURE 4.32: First observation of BEC in the experiment: A sharp peak in the optical density distribution after $20\,\mathrm{ms}$ time-of-flight indicates a macroscopic number of atoms in the lower momentum states, and is one of the defining signatures of the BEC transition. Upon lowering the RF frequency from $775\,\mathrm{kHz}$ (●) to $750\,\mathrm{kHz}$ (■) the optical density distribution changes from Gaussian to the Thomas-Fermi bimodal shape.

The first BECs produced were very unstable, and not repeatable from shot-to-shot. Aside from variations between consecutive experimental runs - most probably due to noise on the magnetic trapping fields, we observed that after approximately one hour the BEC could not be produced anymore. This was due to the fact that the background pressure in the chamber would build up over the course of the day from pulsing of the

dispensers, by which time the lifetime of the magnetic trap had degraded enough to inhibit BEC production. This problem was solved with the use of the 2D MOT, which ensures consistent background pressure throughout the day. The system stability has been improved dramatically since the first realisation - it is now possible to produce a condensate after 30 minutes from switching on the apparatus, and can be maintained throughout the working day or overnight.

4.8 Atom Chip Traps

Although the first BECs in our experiment were produced in the copper Z understructure trap, the micropotentials created by the atom chip should ultimately be used to manipulate the atomic clouds. The copper Z trap was initially used as a proof-of-principle, and did not require the atom chip wires at all - in fact, the first BEC in our lab was produced using only a high-quality gold mirror, fabricated by Dr. Jessica Maclean, necessary for the reflection MOT. However, being able to push the cooling to degeneracy in the copper trap is an indication that the process is working efficiently, and will lead to larger atom numbers on the chip eventually. It is also a reassurance that there are no other factors which may later inhibit evaporation, for example any problems with the RF line, or excessive heating rates. In the end, the copper Z is intended for use as an intermediate trap to perform initial evaporative pre-cooling, increasing the phase-space density by enough to facilitate efficient transfer from the MOT to the chip traps.

Potentials created by the atom chip are more flexible, and provide control on a finer scale than those created with the mesoscopic structures, or by external coils. The close proximity of the trapping wires to the atomic cloud allows much stronger field gradients for more modest electrical currents ($\sim 1\,\mathrm{A}$ versus $\sim 100\,\mathrm{A}$). This reduces problems that can arise from switching large currents though inductive loads - such as high voltage spikes or persistent Eddy currents - and also facilitates more rapid evaporative cooling.

4.8.1 Loading the Atom Chip

There are several considerations when designing a loading scheme from the copper Z trap into the atom chip, with the aim being to achieve a minimum loss in phase-space density throughout the transfer. The main limitation comes from the amount of current which can be safely passed through the chip trap wire. Initial capture is done using one of the largest Z-shaped wires, as shown in Fig. 4.33 (a). Based on measurements performed outside the vacuum before mounting the chip, shown in Fig. 3.19 of Sec. 3.3.3, we estimate that this $50\,\mu\mathrm{m}$ wide wire will undergo a relative resistance increase of $\sim 20\%$ for a continuous current of $1\,\mathrm{A}$. In practice however, this $1\,\mathrm{A}$ is passed in a pulsed fashion, for no more than $\sim 5\,\mathrm{s}$ out of the entire cycle of $25\,\mathrm{s}$, and so there should

be no risk of damaging the wire - even in the unlikely event of the current staying on continuously.

The required depth of the chip trap must also be large enough to prevent spilling of hot atoms, and so is directly dependent on the amount of pre-cooling applied in the copper trap. For a fixed wire current, larger depths can be achieved by increasing the bias field. However, this moves the minimum closer to the surface, and so is limited by the thermal spatial extension of the cloud in the copper trap. We find that a good compromise between all these interlinked criteria is to create the initial chip trap at a distance of $360\,\mu\mathrm{m}$ from the chip surface, which is achieved by countering the $1\,\mathrm{A}$ chip wire current with a $5.3\,\mathrm{G}$ bias field using the X-Bias coil, as shown in Fig. 4.33 (b).

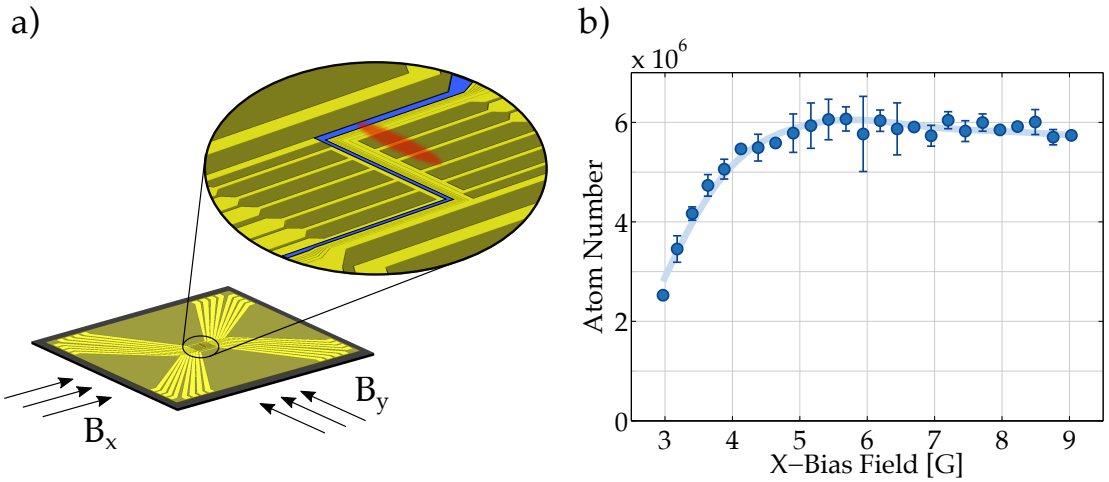

FIGURE 4.33: a) Illustration of the configuration used for initial trapping on the atom chip. The cloud is trapped by passing a current of $1\,\mathrm{A}$ through the widest $50\,\mu\mathrm{m}$ Z-shaped wire (highlighted in blue on the zoomed view). b) Atom number captured in the chip trap as a function of the applied X-Bias field, for $1\,\mathrm{A}$ wire current. The maximum atom number is achieved for $5.3\,\mathrm{G}$, which creates a trap at $360\,\mu\mathrm{m}$ from the surface.

The resulting trap has a depth of $150\,\mu\mathrm{K}$ to the gravity side (whilst the barrier to the chip surface is essentially infinite), meaning that the atoms must be pre-cooled to at least several tens of $\mu\mathrm{K}$ in order to suppress plain evaporation. The frequencies in this trap are $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{210\,\mathrm{Hz}, 22\,\mathrm{Hz}\}$, and although this radial frequency is smaller than the $450\,\mathrm{Hz}$ of the copper trap, as long as the transfer is performed slowly enough the frequency mismatch should lead only to adiabatic cooling, which does not result in a loss in phase-space density.

With the properties of the target trap selected, the transfer process itself is optimised. At the end of the cooling phase in the copper Z trap, linear ramps are applied to each of the trapping currents, shown in Fig. 4.34. The ramp process lasts for $300\,\mathrm{ms}$, after which a period of $100\,\mathrm{ms}$ settling time is allowed before any further changes are made to the chip trapping potential. The best results are obtained when ramping all currents in parallel simultaneously, except for a short delay on the copper

Z current (discussed later). This maintains the position of the trap minimum at $360\,\mu m$ from the chip surface throughout, and ensures there is always sufficient Ioffe field to suppress Majorana spin-flips. The X-Bias field is ramped to its target value of $5.3\,G$, whilst the Y-Bias is ramped down to zero. This is because the Ioffe field produced by the leads of the chip Z-wire for a current of $1\,A$ at this height was measured to be $|B_0| = (1.00 \pm 0.02)\,G$ ($\nu_0 = (700 \pm 15)\,kHz$), and so no additional external longitudinal field is needed for the initial chip trap.

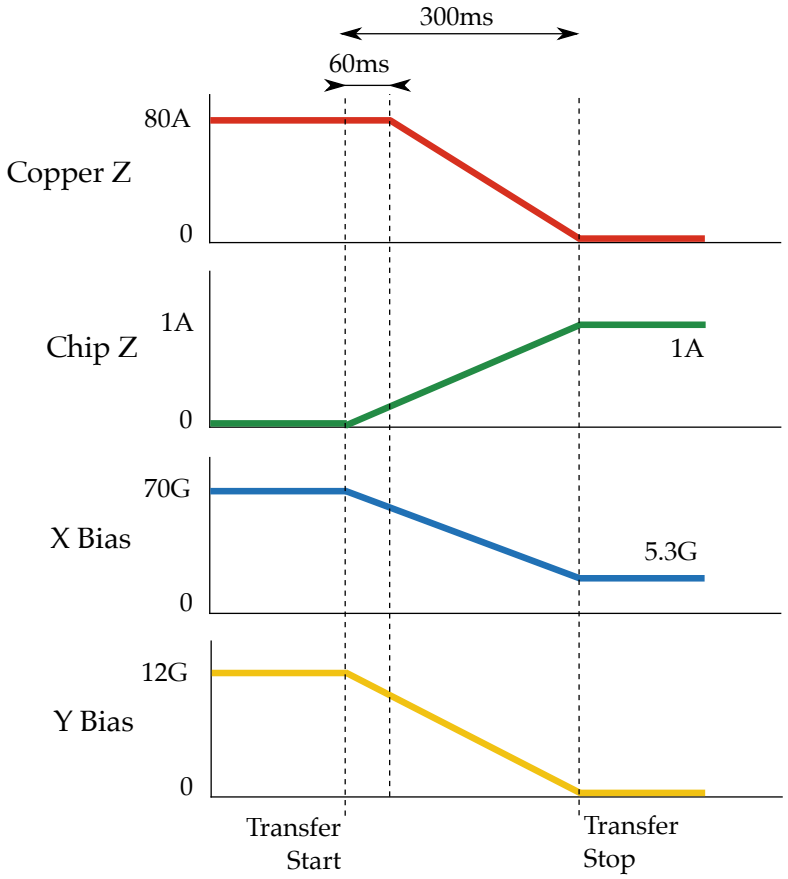

FIGURE 4.34: Ramps of the trapping fields for transfer from the copper Z trap to the atom chip trap: All fields are ramped from their respective start to final values linearly over $300\,ms$, with the exception of the copper Z-wire current. The start time for this ramp is delayed by $60\,ms$, as described in the text, to increase the transfer efficiency. With this scheme, the position of the trap minimum remains approximately fixed in space at $360\,\mu m$ throughout the transfer.

It was found that having a short delay on the start time for the ramp down of the copper Z current was crucial to achieving the best transfer efficiency. The start and end times of the copper Z ramp were adjusted empirically, and the results of the analysis given in Fig. 4.35. It can be seen that the best result is reached with a $60\,ms$ delay before starting the ramp. With this "*Start Delay*", it is then optimal to have the ramp finish simultaneously with the ramps on the other trapping fields (a "*Stop Delay*" of zero). Adding this short delay ensures sufficient potential barrier to the chip surface at

early times before the chip wire current has fully established. This leads to a four-fold increase in the number of atoms recaptured, and the process effectively reaches 100% transfer efficiency.

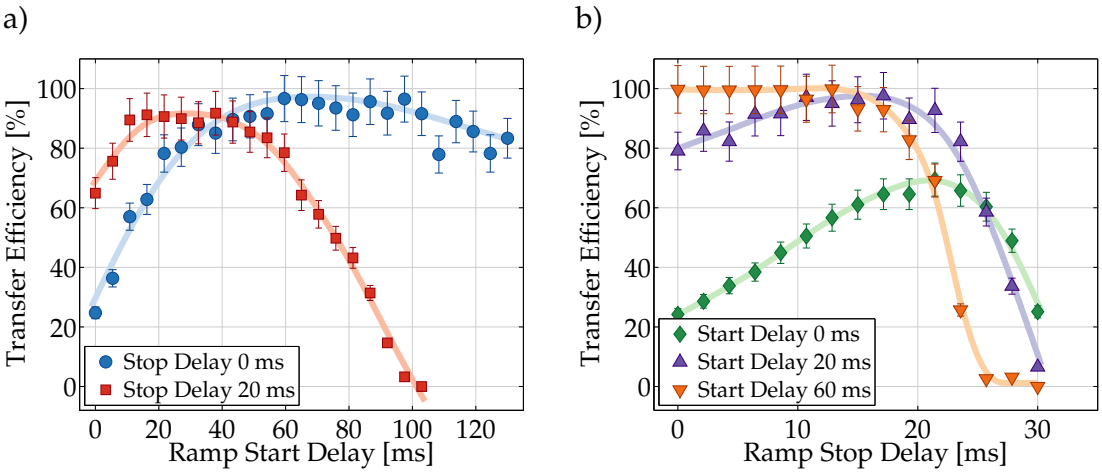

FIGURE 4.35: Delays added to ramp timings of the copper Z-wire: a) The atom number transfer efficiency between the copper and chip traps is measured as a function of the "*Start Delay*" added to the ramp start time (see Fig. 4.34 for timings). b) Transfer efficiency versus the "*Stop Delay*" added to the end of the ramp. Overall, by adding a start delay of 60 ms, and with zero stop delay, the captured number of atoms increases by a factor of four, with the transfer efficiency reaching 100%.

Determining Optimum Temperature for Loading

With the loading sequence in place, it then needs to be determined how much to cool the cloud in the copper trap before performing the transfer, because too little cooling will result in atoms spilling over the depth of the chip trap. In general, there are three different possible approaches to this:

- Cool the gas all the way to degeneracy in the copper trap, and then transfer the BEC to the atom chip. This approach is not generally used in atom chip experiments, due to the fact that the BEC is more sensitive to detrimental heating and excitation which could occur in the transfer process itself [237].

- Cool to just above the transition temperature (to $\sim 1\,\mu$K), transfer to the chip, and then apply a small RF knife to promote atoms into the BEC. This method was used to produce the first BEC on the atom chip in our lab, as it does not require any additional optimisation or compression in the atom chip.

- Reduce the temperature by just enough so that a $\sim 100\%$ transfer efficiency into the chip trap can be obtained, and then use evaporative cooling in the chip to achieve the final large increase in phase-space density.

The third approach is the most efficient to use, because it can benefit from the high trap frequencies and collision rates attainable in chip traps, and therefore reduces the cycle time. If most of the cooling is to be performed in the chip, it is advantageous to load as many atoms as possible, giving the highest initial elastic collision rate.

To determine how cold to make the cloud before transfer, the final stop RF frequency in the copper Z trap was varied, whilst looking at both the temperature and atom number before and after transferring, with the results displayed in Fig. 4.36. The final RF frequency was halted at values from $8\,\mathrm{MHz}$ to $1\,\mathrm{MHz}$, corresponding to temperatures in the copper trap from $80\,\mu\mathrm{K}$ to $8\,\mu\mathrm{K}$, as shown in Fig. 4.36 (a). The ratio between the temperature of the cloud in the copper trap to that in the chip trap is shown in the inset, and it can be seen that for higher final frequencies this ratio increases, indicating plain evaporation of the hot atoms above the chip trap depth. Below $4\,\mathrm{MHz}$, the ratio is constant at a value of ~ 1.7, consistent with the temperature decrease from adiabatic cooling due to the difference in frequencies between the two traps.

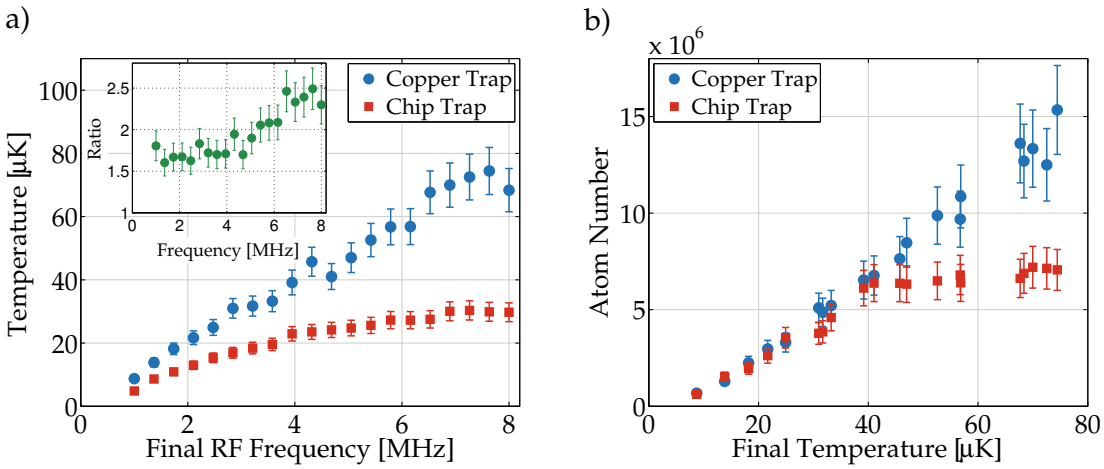

FIGURE 4.36: Properties of the transfer for various final copper trap temperatures: a) The temperature of the cloud in both the copper trap (before transfer) and the chip trap (after transfer) is plotted as a function of the final RF stop frequency of the copper evaporation ramp. The ratio of the two temperatures is shown in the inset, and is constant at ~ 1.7 up to $4\,\mathrm{MHz}$ consistent with adiabatic cooling. For higher frequencies, the ratio begins to increase, indicating plain evaporation. b) Atom number in the copper and chip traps as a function of the final temperature from pre-cooling with the copper trap evaporation ramp. If the cloud is pre-cooled below $40\,\mu\mathrm{K}$ then atoms can be transferred without loss.

The conclusion that hot atoms are being evaporated at higher frequencies is confirmed in Fig. 4.36 (b), which depicts the number of atoms remaining before and after the transfer, as a function of the final temperature in the copper trap. These results show that the cloud should be cooled to at least $40\,\mu\mathrm{K}$ (corresponding to a temperature of $20\,\mu\mathrm{K}$ in the chip after adiabatic cooling), in order to maintain a transfer efficiency of close to unity. In practice, the cloud temperature is typically reduced slightly further, to ensure that spilling from the chip trap is suppressed, increasing the robustness of

the procedure. The evaporation ramp in the copper previously described in Sec. 4.6 can now be truncated at the point where 100% of the atoms are able to be transferred, which typically means stopping the ramp at $3.5\,\mathrm{MHz}$ (after $10\,\mathrm{s}$ of pre-cooling), allowing 5×10^6 atoms to be captured in the chip trap.

Before going any further, it is important to determine the heating rate of a cloud in the chip trap, as this will act against any evaporative cooling that is applied. Possible sources of noise include mechanical vibrations (from various equipment in the lab), any near resonant light entering the chamber, or more likely current noise on the power supplies driving the chip wires. Current noise can cause both centre-of-mass vibrations and parametric excitation. This is particularly problematic if the frequency of the noise is comparable to the eigenfrequencies of the trap, or any higher harmonics. A measurement of the heating rate is shown in Fig. 4.37. This was measured by cooling the gas after loading to $\sim 550\,\mathrm{nK}$, then removing the RF knife and monitoring the evolution of the temperature. A heating rate of $(120 \pm 30)\,\mathrm{nK\,s^{-1}}$ is observed, which is comparable to other atom chip setups (see, for example, [233]), and is satisfactory for our experiment.

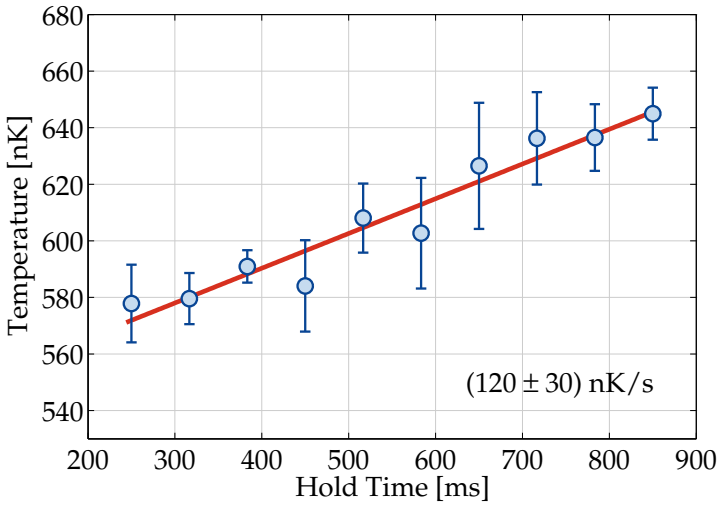

FIGURE 4.37: Heating rate in the atom chip trap: By monitoring the temperature evolution of a pre-cooled cloud in the atom chip trap, the heating rate is determined by a linear fit (solid line) to be $(120 \pm 30)\,\mathrm{nK\,s^{-1}}$.

4.8.2 Compression, Cooling and BEC on Chip

With the loading optimised, the next step was to produce an on-chip BEC. The process is qualitatively the same as that used in the copper Z trap - namely an adiabatic compression to increase the elastic collision rate, followed by forced RF evaporative cooling, as described in Sec 4.5.1 and 4.6. The trap frequencies attainable are in excess of $1\,\mathrm{kHz}$, permitting faster RF sweeps of less than several seconds. This is advantageous because it both reduces the effect of loss due to the finite trap lifetime, whilst also minimising the time needed for passing current through the chip wire.

The trap can be compressed by either increasing the external bias field, or by reducing the current through the chip wire. Both methods will increase the transverse trap frequency, but the effect of the bias field is stronger since the radial frequency scales as $\omega_\perp \propto B_{\text{bias}}^2/I$, if the bottom offset field is held constant [238]. Since the potential energy barrier to the chip surface is extremely large, the amount of compression that can be applied is ultimately limited by three-body loss processes, which involves two atoms combining to form a bound state and a third atom carrying away the excess energy. This loss rate is density dependent, and is given by Ln^2, where $L \sim 2 \times 10^{-29}\text{cm}^6\,\text{s}^{-1}$ for ^{87}Rb in the $|F = 2, m_F = +2\rangle$ ground state, making three-body losses significant for densities above $\sim 10^{14}\text{cm}^{-3}$ [239–241].

As part of an initial characterisation for the purposes of trap compression, the effect of varying the applied external bias fields on the position of the cloud was measured, with the results depicted in Fig. 4.38. After loading the atom chip as described in Sec. 4.8.1, both the perpendicular (B_x) and the vertical (B_z) bias fields were varied linearly over $500\,\text{ms}$, which forces the magnetic minimum to traverse the paths labelled ① and ② in Fig. 4.38 (a), respectively. Increasing B_x simply moves the cloud directly closer to the surface, whilst increasing B_z rotates the cloud around the central wire (around the y-axis). However, since absorption imaging is performed along the x-direction, motion along the x-axis is not detected and the effect of increasing either field is only to reduce the observed cloud height from the chip surface, as plotted in Fig. 4.38 (b) and (c). With a fixed $1\,\text{A}$ of current passing though a $50\,\mu\text{m}$ wide chip wire, either B_x is varied from $12\,\text{G}$ to $70\,\text{G}$, or B_z from zero to $30\,\text{G}$, in both cases resulting in the height, z_0, being reduced over the range from $150\,\mu\text{m}$ to $20\,\mu\text{m}$.

To compare the experimentally realised potentials with simulations, the data in Fig. 4.38 (b) and (c) are plotted along with the expectation from calculations. Firstly, the curves for a simple infinitely thin, infinitely long wire are shown (which in the case of Fig. 4.38 (b) simply follows the form $z_0 = \mu_0 I/2\pi B_x$). Also shown is the result of a full calculation for a finite dimension Z-shaped wire, assuming uniform current density in each section of the conductor, obtained using the expressions for the magnetic field components given in Eq. A.8−A.10 of Appendix A. As expected, in Fig. 4.38 (b) the infinitely thin model breaks down for heights on the order of and below the wire width (when the wire begins to no longer appear infinitely thin at the site of the atoms), whereas the full calculation shows excellent agreement over the whole data set, demonstrating its importance for close distances. In Fig. 4.38 (c), the two models give similar results, with both describing the data well, due to overall larger absolute distance from the trapping wire when the position offset along the x-direction is taken into account, making the use of the full calculation less critical in this case.

Control of the Trap Bottom

When approaching the chip surface during the trap compression, the Ioffe (y) component of the field is also reduced. This is because there is a maximum in the component

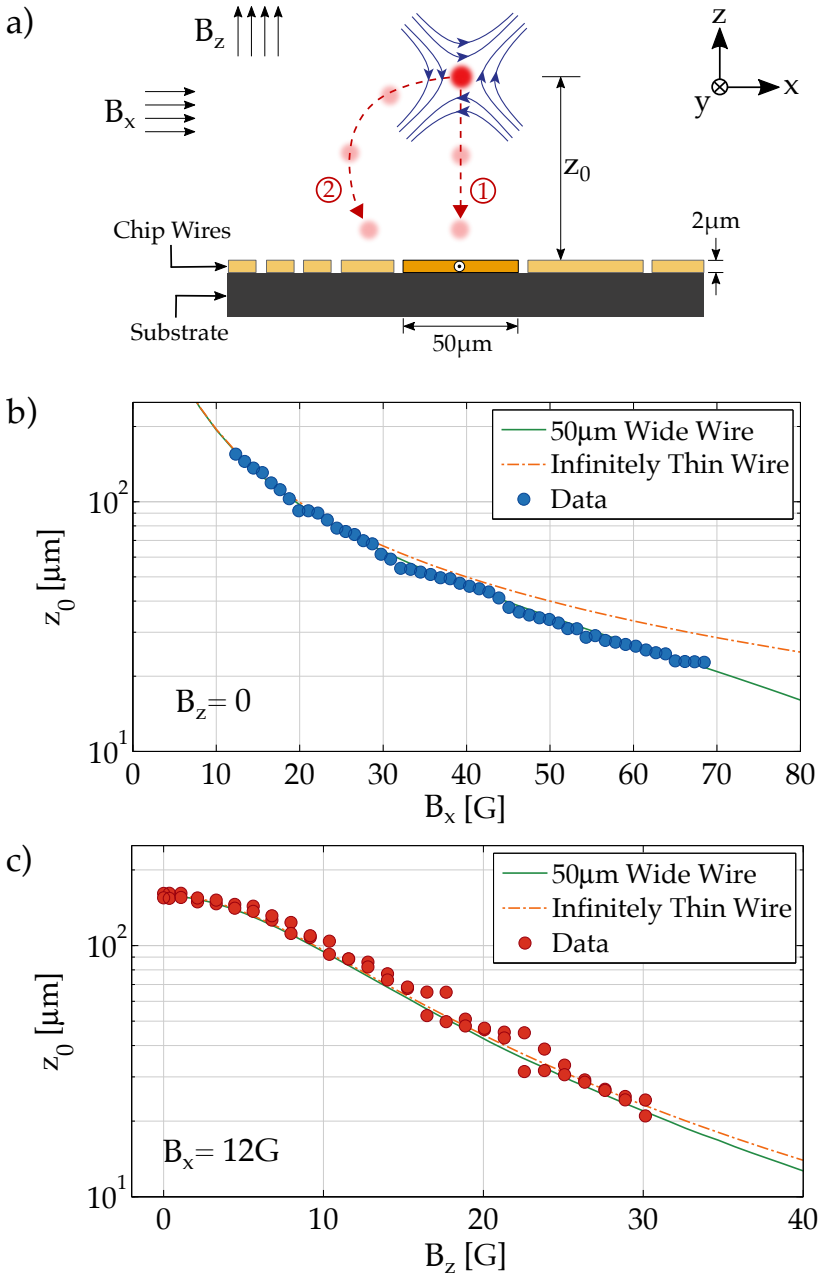

FIGURE 4.38: Characterisation of the effect on the cloud position when varying external bias fields: a) Measurement schematic: A current of 1 A through a $50\,\mu\mathrm{m}$ wide atom chip wire together with externally applied bias fields B_x and B_z creates the two-dimensional quadrupole in the x-z plane, at a height of z_0 from the chip surface. b) The trap height as a function of B_x whilst keeping $B_z = 0$, forcing the cloud to follow the path labelled ① shown in (a). c) The trap height as a function of B_z, with B_x fixed at $12\,\mathrm{G}$, resulting in the path labelled ② in (a). Shown along with the data in (b) and (c) are the results of a calculation using both an infinitely thin, infinitely long wire model (dash-dotted line), and a full Z-shaped conductor with central cross-section $50\,\mu\mathrm{m} \times 2\,\mu\mathrm{m}$ (solid line).

of the field produced by the legs of the Z wire at some distance above the chip surface. For our geometry this occurs at $880\,\mu\mathrm{m}$, and so below this distance the Ioffe field always reduces with decreasing height. This decrease in the trap bottom was measured for increasing X-bias fields (by using the RF knife-edge technique), and is shown in Fig. 4.39 (a). There is a constant discrepancy between the measured and calculated values by approximatey $0.25\,\mathrm{G}$, which can be attributed to the fact that no attempt is made in our experiment to null external fields (such as the Earth's field of $\sim 0.5\,\mathrm{G}$). It would be good to include compensation coils in the future, to account for ambient magnetic fields.

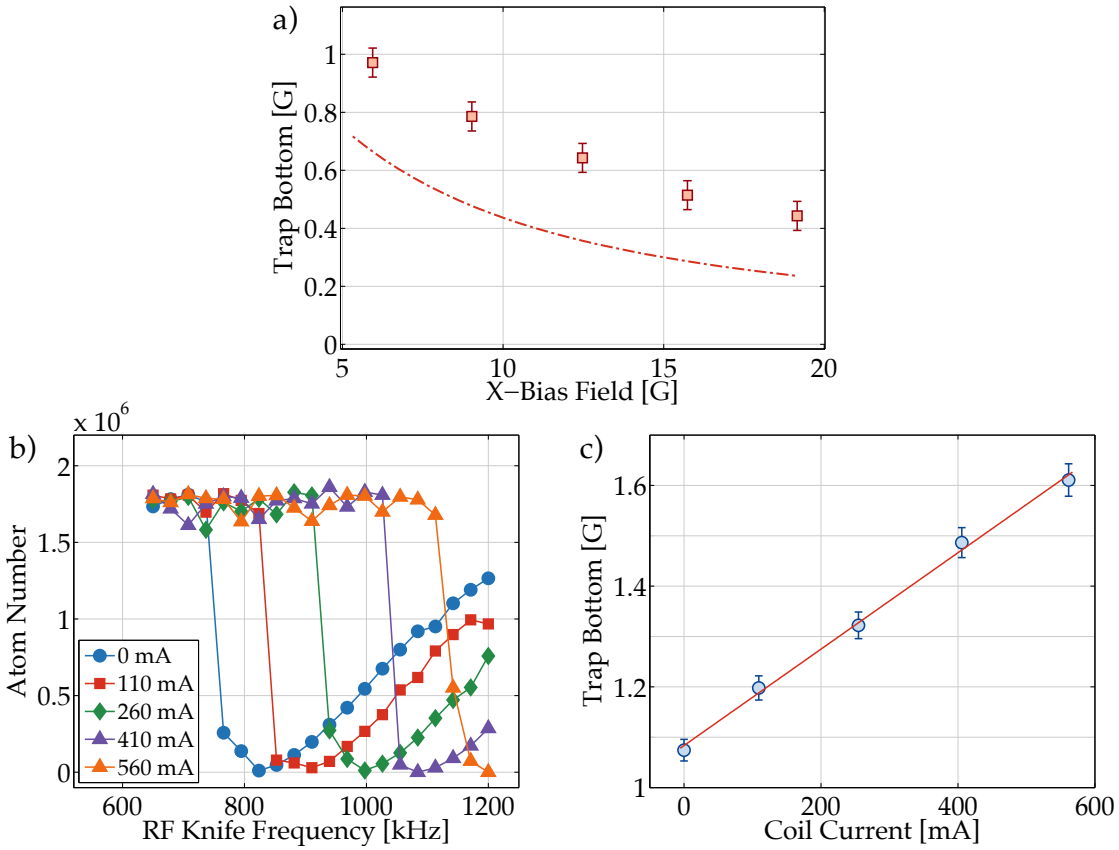

FIGURE 4.39: a) The shift in the absolute value of the magnetic field at the trap bottom when increasing the X-Bias field during trap compression. The dash-dotted line is the result of a calculation, and the discrepancy of $\sim 0.25\,\mathrm{G}$ is attributed to the influence of an ambient magnetic field. b) For calibration of an additional external field to control the trap bottom, RF knife measurements are used to measure the shift for various currents throught the external coil. c) The shift in trap bottom as a function of current through the coil calibrates the magnetic field produced by the coil per unit current to be $(0.98 \pm 0.07)\,\mathrm{G\,A^{-1}}$.

In order to maintain the trap bottom at a value of $\sim 1\,\mathrm{G}$ when applying the compression (to avoid Majorana transitions), a controllable external field was added. In contrast to the copper Z trap, where the longitudinal field was required to *cancel* the

field due to the Z legs, this field should *add* to the Ioffe field. Due to the relative orientation of the Z wires, the same pair of coils can be used in both traps provided there is some way to reverse the current flow through them in between (for example, using an H-bridge circuit). Since the field required is only on the order of $1\,\mathrm{G}$, we decided to add a second small set of coils on top of the main ones, to simply create a field in the opposite direction. The shift in the trap bottom was measured as a function of the current through these coils to calibrate the field, shown in Fig. 4.39 (b). From these data, the magnetic field produced by the small coils is determined to be $(0.98 \pm 0.07)\,\mathrm{G\,A^{-1}}$, shown in Fig. 4.39 (c), and provides a way to control the offset field at the trap minimum when applying different amounts of compression.

For BEC production on the chip, the best compression results were obtained when increasing the main X-Bias field after loading from $5.3\,\mathrm{G}$ to $12.4\,\mathrm{G}$ over $500\,\mathrm{ms}$, which positions the cloud at a distance of $160\,\mu\mathrm{m}$ from the chip surface. In addition, the Y-Bias field controlling the residual Ioffe component is increased in parallel to $640\,\mathrm{mG}$, which fixes the RF frequency at the trap bottom at $\nu_0 = (832 \pm 2)\,\mathrm{kHz}$. A measurement of the trap bottom using pulsed RF spectroscopy on a trapped BEC is shown in Fig. 4.40, which allows the frequency to be determined much more precisely than if using a thermal cloud (compare, for example, with Fig. 4.30).

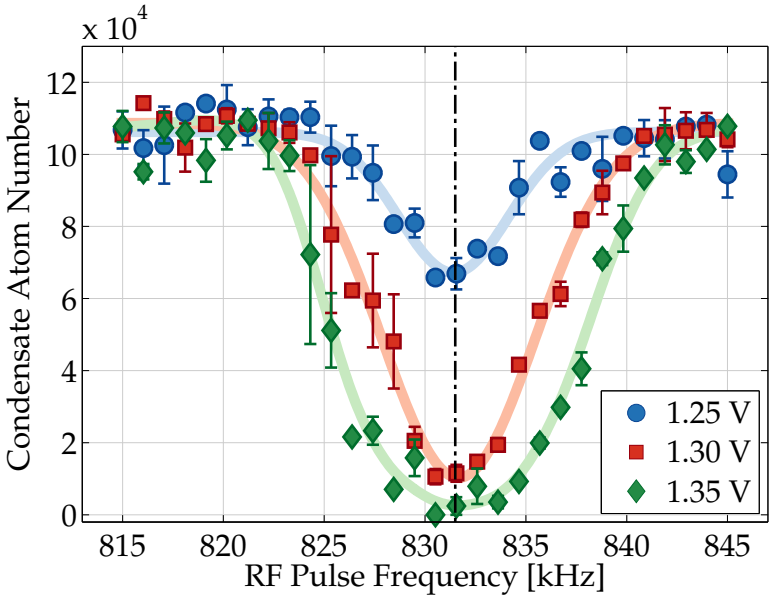

FIGURE 4.40: RF spectroscopy of the trap bottom using a BEC: A pure condensate was exposed to a $30\,\mathrm{ms}$ pulse of RF field of varying frequencies, for RF amplitudes of $1.25\,\mathrm{V}, 1.30\,\mathrm{V}$ and $1.35\,\mathrm{V}$. A trap bottom of $832\,\mathrm{kHz}$ $(1.2\,\mathrm{G})$ is observed, in addition to broadening of the resonance for increasing RF amplitudes. This method is much more precise than when using the RF knife measurement with thermal clouds.

The compression ramp modifies the harmonic trap frequencies to $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{850\,\mathrm{Hz}, 20\,\mathrm{Hz}\}$, and increases the trap depth to $500\,\mu\mathrm{K}$. During compression the cloud

is adiabatically heated to $30\,\mu\mathrm{K}$, and so the RF evaporation ramp is initiated with a frequency of $3\,\mathrm{MHz}$ corresponding to $\eta \sim 8$. Several durations for the evaporation sweep were tested, with the largest BEC atom numbers occuring with $2500\,\mathrm{ms}$ (although condensates can be produced with as little as $500\,\mathrm{ms}$ of cooling). Absorption images of the cloud are shown in Fig. 4.41, taken by interrupting at various stages throughout the RF sweep. The transition to BEC occurs at a temperature of $\sim 800\,\mathrm{nK}$, after ramping the RF frequency down to $890\,\mathrm{kHz}$, with an atom number on the chip of 8×10^5, which is in broad agreement with the prediction for an ideal gas in an harmonic trap (Eq. 2.58). When using only the copper-Z trap to produce BEC (as can be seen in Fig. 4.28 of Sec. 4.6) the number of atoms present when the phase transition occurs is a factor of four smaller than when using the atom chip (2×10^5 compared to 8×10^5). One reason for this is that the required transition temperature itself in Eq. 2.58 is simply higher when using the atom chip, because the transverse trap frequencies used are much higher ($850\,\mathrm{Hz}$ compared to $460\,\mathrm{Hz}$). In addition, even when comparing directly the number of atoms remaining at a temperature of $800\,\mathrm{nK}$, the atom chip offers approximately a factor of two improvement in atom number, due to the increased efficiency of cooling in a tighter trap.

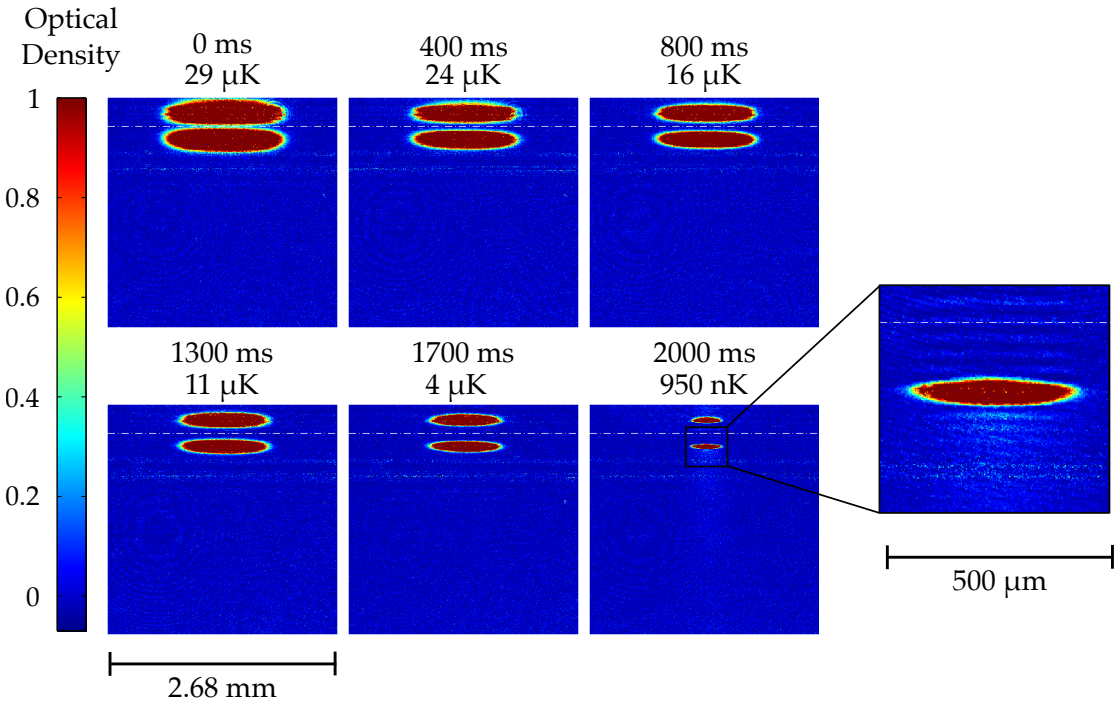

FIGURE 4.41: Atom Chip Evaporation Ramp: Optical density images of the cloud for various temperatures during cooling with $1\,\mathrm{ms}$ time-of-flight after switching off the trapping potential. The times indicated are with respect to the time at which the RF sweep is initiated. In these data, the imaging beam has been slightly inclined with respect to the plane of the atom chip surface, resulting in a second image of the cloud forming due to two beam paths traversing the cloud [198, 238]. This allows the height of the trap from the chip surface (indicated by the dotted line) to be determined as the half-separation of the two clouds.

Chapter 5

Characterisation and Experiments with BECs

This chapter presents selected data from measurements performed as part of the initial characterisation of BECs in both the copper Z-wire and atom chip traps. Section 5.1 describes the fitting routines applied in order to correctly analyse the experimentally obtained optical density distributions, and therefore retrieve accurate physical properties of the samples. Measurements of the ground state occupation as a function of temperature, the condensate expansion dynamics, and the condensate lifetime are presented in sections 5.2, 5.3, and 5.4, respectively. Section 5.5 illustrates the importance of controlled switch-off of the trapping potentials, by demonstrating the unwanted projection over magnetic states in the presence of a field zero. Finally, section 5.6 gives an example of how BECs can be produced without the aid of external coils, using only the integrated copper trapping structures which are unique to this experimental apparatus.

5.1 Analysis of Absorption Images

To accurately determine various properties of BECs, including condensate fraction, total atom number, and temperature, the obtained absorption images must be interpreted correctly. In Sec. 3.9, the analysis of the density distribution treated the gas as ideal and classical, which leads to a Gaussian distribution that works well for hotter clouds. However, near or below the transition temparature ($T \lesssim T_c$), there is a modification due to the influence of Bose-Einstein statistics, and the spatial density distribution of the thermal component, $n_{\mathrm{Th}}(x, y, z)$, is well-described by a *Bose-enhanced* Gaussian function in which the chemical potential has been set to zero [97]

$$n_{\mathrm{Th}}(x, y, z) = \frac{n_{\mathrm{Th}}(0)}{g_{3/2}(1)}\, g_{3/2}\Bigg(\exp\Bigg[-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2} - \frac{z^2}{2\sigma_z^2}\Bigg]\Bigg), \tag{5.1}$$

where $g_k(x) = \sum_{n=1}^{\infty} x^n/n^k$ is the polylogarithm function, $n_{\mathrm{Th}}(0)$ is the density at the centre of the cloud, and σ_i is the RMS width of the Gaussian in the i^{th} dimension.

Through normalising to the total number of atoms, the peak density is calculated as[1]

$$\int n_{\mathrm{Th}}(\boldsymbol{r})\, d^3\boldsymbol{r} \equiv N_{\mathrm{Th}} \implies n_{\mathrm{Th}}(0) = \frac{N_{\mathrm{Th}}}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_z}\frac{g_{3/2}(1)}{g_3(1)}. \tag{5.2}$$

For high temperatures or for distances far away from the centre, the Bose-enhanced function tends to the Gaussian of Eq. 3.24. This is because atoms in the wings have the highest velocities and will be in the lowest density regions of the cloud, and so occupation numbers of these states are much less than one, rendering the quantum statistics effects insignificant. Since the two distributions are identical away from the centre, we apply fits only to the wings of the thermal cloud. This approach allows the temperature to be determined regardless of whether the Gaussian or Bose-enhanced Gaussian function is chosen.

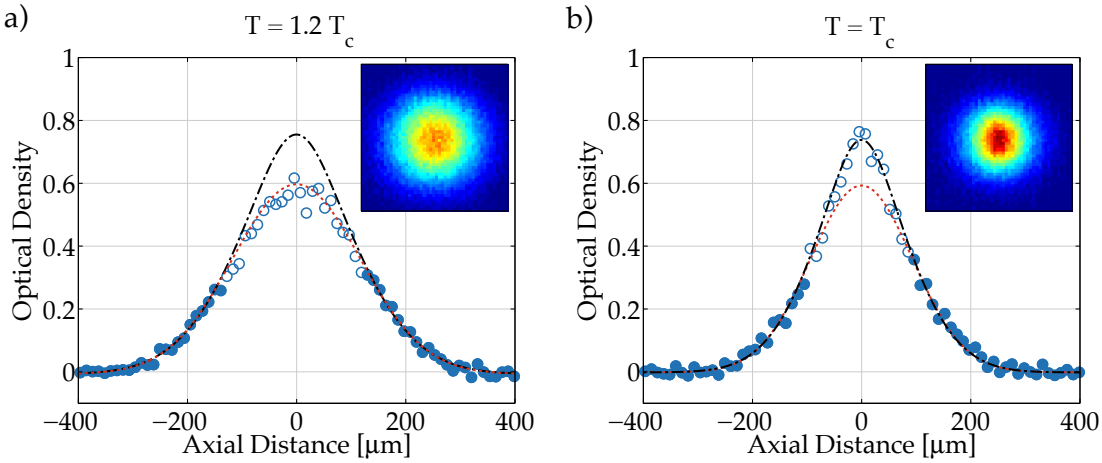

FIGURE 5.1: Effect of Bose-enhancement of thermal clouds: Images show axial cuts through the optical density distributions obtained by absorption imaging after 15 ms time-of-flight. The data in (a) show a cloud at $T = 1.2T_c$, whilst (b) is just before the onset of condensation at $T \sim T_c$. The dotted red lines are the result of purely Gaussian fits, whilst the black dash-dot lines are for the saturated Bose-enhanced function of Eq. 5.1. Filled markers denote data which have been included for the fit, whilst open circles have been masked out. The Gaussian curve describes the data well in (a), but the Bose-enhanced function is a better fit in (b), when the effect of quantum statistics becomes non-negligible.

An example of the difference between the two cases is shown in Fig. 5.1, which shows cuts through the barycentre of the fits to the distributions of two thermal clouds - one just before the onset of condensation ($\mathrm{T} \sim \mathrm{T_c}$) and a second, hotter cloud ($\mathrm{T} \sim 1.2\mathrm{T_c}$). It can be seen that for the cloud at $1.2\mathrm{T_c}$, a purely Gaussian function better fits the measured data when extrapolated over the central region of the cloud. However, for the gas just above $\mathrm{T_c}$, the Bose-enhanced function agrees well with the measured profile. There is some flexibility with the choice of which data to include as the wings in

[1]Using the result $\int_{-\infty}^{+\infty} g_k\left(\alpha \exp\left[-\frac{x^2}{\beta}\right]\right) dx = \sqrt{\pi\beta}\, g_{k+1/2}(\alpha)$ to integrate out all three spatial dimensions.

the fit. If the region is too small, the calculated temperature can be skewed depending on the choice of fitting function. If the region is too large, the temperature becomes unreliable due to a low signal-to-noise ratio. It is found that a reasonable compromise is to mask out data points that are below $\sim 1\sigma$ from the centre.

In addition to the thermal component, for a pure condensate in the Thomas-Fermi regime confined to an harmonic trap the density is given from Eq. 2.74 by the inverted parabola

$$n_{\mathrm{BEC}}(x,y,z) = n_{\mathrm{BEC}}(0)\ \max\left\{1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2}, 0\right\}, \tag{5.3}$$

where $n_{\mathrm{BEC}}(0)$ is the peak density at the centre, given from normalisation on the total number of condensate atoms by

$$n_{\mathrm{BEC}}(0) = \frac{15}{8\pi}\frac{N_{\mathrm{BEC}}}{R_x R_y R_z}, \tag{5.4}$$

and R_i is the Thomas-Fermi radius in the i^{th} dimension.

Fitting Partially Condensed Clouds

For partly condensed clouds consisting of both thermal atoms and those in the BEC state, the density distribution is assumed to be a bimodal function - specifically, a sum of the Bose-enhanced Gaussian distribution and the Thomas-Fermi inverted parabola. The fitting function applied to the measured optical density images is a sum of Eq. 5.1 and Eq. 5.3 that have been integrated over the imaging direction z. It contains nine free parameters A_i, and is given by

$$\begin{aligned}\mathrm{OD}(x,y) = \mathrm{A}_1 + \mathrm{A}_2\, g_2\left(\exp\left[-\frac{(x-\mathrm{A}_3)^2}{2{\mathrm{A}_5}^2} - \frac{(y-\mathrm{A}_4)^2}{2{\mathrm{A}_6}^2}\right]\right) \\ + \mathrm{A}_7\ \max\left\{\left(1 - \frac{(x-\mathrm{A}_3)^2}{{\mathrm{A}_8}^2} - \frac{(y-\mathrm{A}_4)^2}{{\mathrm{A}_9}^2}\right)^{3/2}, 0\right\}.\end{aligned} \tag{5.5}$$

Note that the Bose-enhanced part is modified from $g_{3/2}()$ to $g_2()$, and the exponent of the inverted parabola is modified from 1 to $3/2$, which both result from integrating out one of the spatial dimensions. The number of atoms in each of the thermal and condensed components can then be determined respectively using

$$N_{\mathrm{Th}} = \left(\frac{7}{15}\right)^{-1}\frac{M^2}{A}\frac{1}{\sigma_0}2\pi\ g_3(1)\ \mathrm{A}_2\ \mathrm{A}_5\ \mathrm{A}_6 \tag{5.6}$$

$$N_{\mathrm{BEC}} = \left(\frac{7}{15}\right)^{-1}\frac{M^2}{A}\frac{1}{\sigma_0}\frac{2\pi}{5}\ \mathrm{A}_7\ \mathrm{A}_8\ \mathrm{A}_9, \tag{5.7}$$

where the parameters $\mathrm{A}_5, \mathrm{A}_6, \mathrm{A}_8$, and A_9 are given in metres. Here, σ_0 is the on-resonance scattering cross-section for a two-level atom, the factor $7/15$ accounts for

the fact that imaging is performed in zero-field, and M^2/A is due to the finite pixel size of the CCD, all of which are described in more detail in Sec. 3.9.

Our fitting procedure for bimodal clouds follows that outlined in [242] and is essentially a three-step process, with the various stages shown in Fig. 5.2:

1. Before any fitting is applied, any artifacts or interference fringes in the measured optical density image are firstly removed to a large extent using the algorithm developed in Ref. [207], which is described in more detail in Appendix B. Then, the resultant optical density image, shown in Fig. 5.2 (a), is fit with a simple 2D bimodal function consisting of a sum of a *pure* Gaussian and an inverted parabola. A horizontal cut through the centre of the resultant fit is shown in Fig. 5.2 (b). This step neglects the effect of Bose-enhancement in the interest of minimising processing time, and is used only to retrieve initial estimates of the parameters of the cloud. Taking the results of this simple fit to calculate temperature from the Gaussian wings has been shown to underestimate the actual temperature by $\sim 13\%$ [242].

2. Next, the Thomas-Fermi radii R_i obtained in step 1 are used to mask out the central region of the image in order to exclude the contribution from the condensate, as shown in the inset of Fig. 5.2 (c). Typically, we mask out data with $x_i < 1.5R_i$ [242], ensuring that the condensate and any distortion of the thermal cloud due to the presence of the condensate is neglected. A fit using the Bose-enhanced g_2 function is then applied to the remaining image - which now contains only thermal atoms, as shown in Fig. 5.2 (c). This step fixes the parameters A_1, A_2, A_5 and A_6 in Eq. 5.5, fully defining the density distribution of the thermal atoms, from which N_{Th} and the temperature can be obtained.

3. Finally, the contribution from the thermal fraction obtained in step 2 is subtracted from the original data, leaving only atoms in the BEC state as shown in the inset of Fig. 5.2 (d). Again using the Thomas-Fermi radii estimates from step 1, a region of typically $0.8 < R_i < 1.2$ is masked out, to exclude the potentially distorted interface between the two components. The inverted parabola function is then fit to this data, and is plotted in Fig. 5.2 (d), which finally provides the remaining parameters A_3, A_4, A_7, A_8 and A_9 to fully define the BEC atom number N_{BEC}. The total atomic distribution is then given by the sum of the thermal and BEC components, and is depicted in Fig. 5.2 (e).

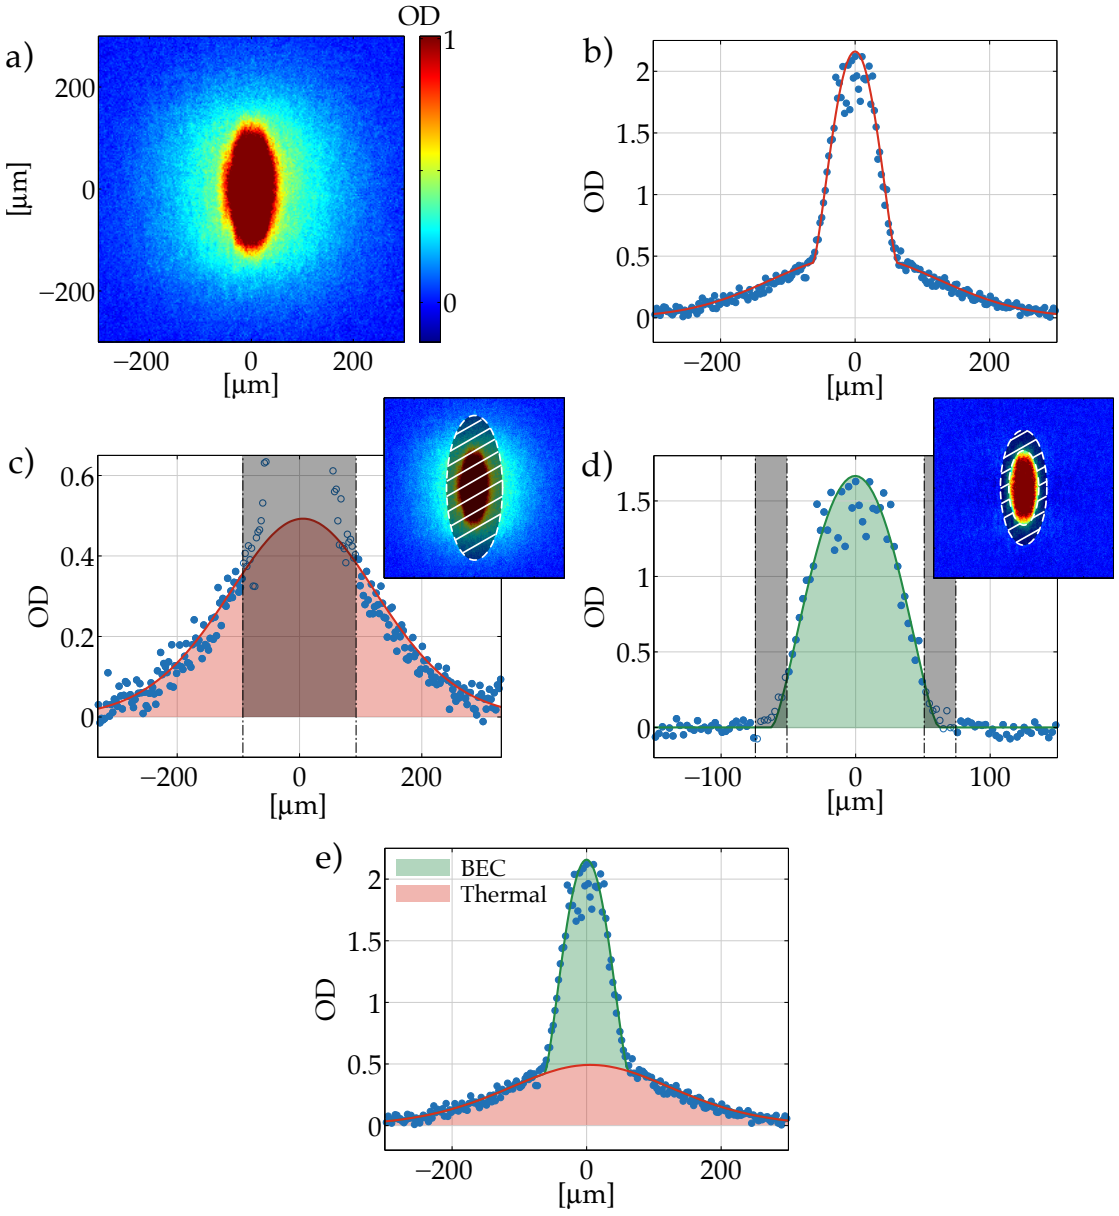

FIGURE 5.2: Procedure for Fitting Images of Partially Condensed Clouds: a) The original measured optical density image, recorded after 22 ms TOF expansion. b) The result of Step 1 described in the text, fitting a simple bimodal function (with a non-Bose-enhanced Gaussian) to retrieve initial estimates of parameters. The line plot is a horizonatal (axial) cut through the barycentre of both the original data (points) and the 2D fit (solid line). c) The result of Step 2, fitting a g_2 Bose-enhanced Gaussian to determine the thermal distribution (indicated by red shaded area). The grey shaded area (hatched on inset) depicts the BEC region which is excluded from this fitting step. d) The result of Step 3, in which the thermal component has been subtracted, leaving only BEC atoms (shown in inset). After masking out the interface region (grey shaded area), a fit using the Thomas-Fermi inverted parabola defines the distribution of BEC atoms (shown by the green shaded area). e) The final resultant fit, after applying the three step procedure. For this particular image, the extracted physical parameters are $T = 495\,\mathrm{nK}$, $N_{\mathrm{Th}} = 4.4 \times 10^5$, and $N_{\mathrm{BEC}} = 1.6 \times 10^5$ (condensate fraction = 27%).

5.2 Characterisation of the Transition

After obtaining stable BECs, one of the first basic things to measure is the critical temperature, and the subsequent dependence of the condensate fraction on the temperature of the gas. Data taken as the threshold for condensation is crossed are given in Fig. 5.3. Each image is obtained by evaporatively cooling the gas in the copper Z-trap to various lower temperatures, by further reducing the RF frequency, and then taking an absorption image after $15\,\mathrm{ms}$ of ballistic expansion. The harmonic frequencies of the trap were measured to be $\{\omega_\perp, \omega_{/\!/}\} = 2\pi \times \{450\,\mathrm{Hz}, 23\,\mathrm{Hz}\}$, and the RF frequency at the trap bottom to be $\nu_0 = 651\,\mathrm{kHz}$. The data are fit using the procedure described in Sec. 5.1, to a 2D function consisting of a Bose-enhanced Gaussian for the thermal component and an inverted parabola for the Thomas-Fermi condensate, integrated along the imaging axis to obtain column densities. The 1D profiles are then obtained by taking an axial cut through the centre of both the measured data and the fitted distributions. Temperatures are determined from the wings of the thermal component, whilst atom numbers are calculated by integrating over the fitted profiles, and the condensate fraction is defined by the ratio of the number in the condensate to the total number of atoms.

Over the data series, the atom number reduces by an order of magnitude from 3.6×10^5 to 3.9×10^4, whilst the temperature decreases accordingly from $600\,\mathrm{nK}$ to $160\,\mathrm{nK}$, corresponding to a truncation parameter of $\eta \sim 6$. For lower atom numbers in the thermal cloud, it becomes increasingly difficult to obtain reliable temperatures from the thermal wings, as the signal approaches the noise floor. The image in Fig. 5.3 (a) has been fit to a pure Gaussian, whilst in (b) it was necessary to apply the Bose-enhanced function from Eq. 5.1 as quantum statistics become more relevant. A small condensate peak appears in (c), and upon further reduction of the RF frequency, the condensate fraction increases continuously up to 62% (the largest measured fractions that can be reliably fit to a bimodal distribution are usually $\sim 70\%$). Finally, an almost pure condensate of 40,000 atoms with no discernable thermal fraction is shown in (j), to which a pure Thomas-Fermi profile has been fit.

The phase transition to BEC is taken to be between images (b) and (c), giving a critical temperature of $T_\mathrm{c} = 405 \pm 15\,\mathrm{nK}$, reached with 2.4×10^5 atoms. The theoretical critical temperature of an ideal gas confined to an harmonic trap in the thermodynamic limit is given by Eq. 2.58, reproduced here for convenience

$$T_c^0 = 0.94 \frac{\hbar}{k_B} \omega_\mathrm{ho} N^{1/3}, \tag{5.8}$$

which for these trap parameters is calculated to be $470\,\mathrm{nK}$. However, there are several effects which cause a shift in the critical temperature of a realistic system. Firstly, there is the effect of finite particle number, arising because the gas is not in the thermodynamic limit ($N \to \infty$). The shift in critical temperature due to finite size effects can be obtained by accounting for the non-zero energy of the lowest single-particle state in the

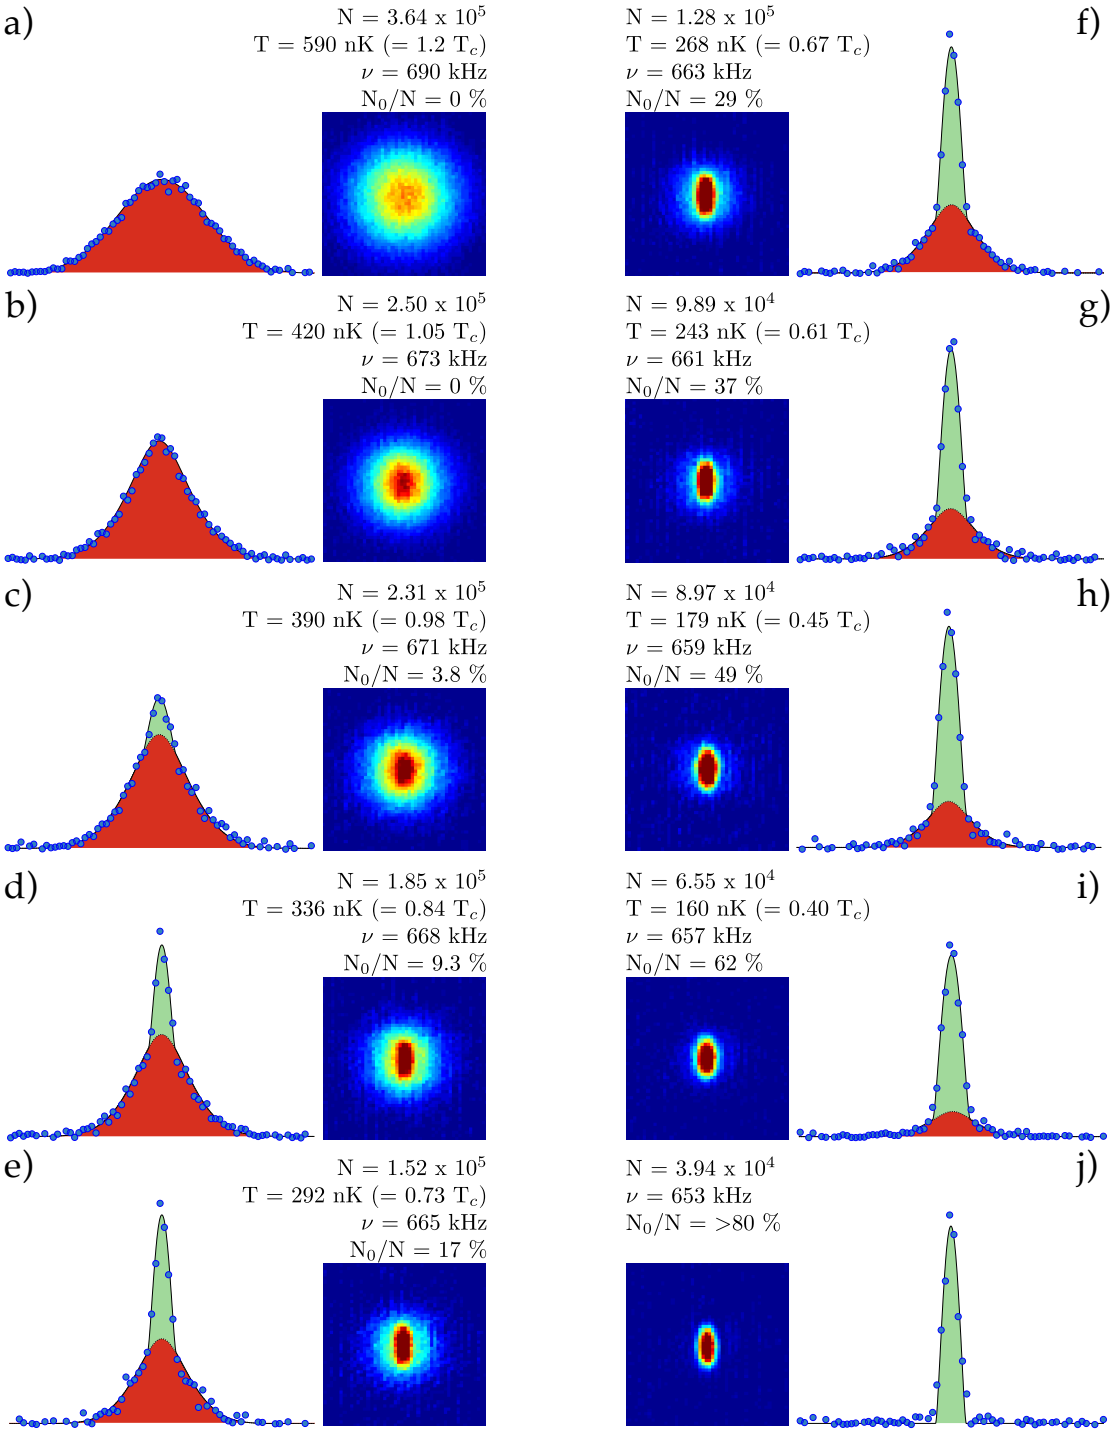

FIGURE 5.3: Images of the formation of a BEC in the copper Z-trap: The images (a) - (j) are absorption images after 15 ms time-of-flight, for varying lower final RF frequencies of evaporation. With reducing temperature, a condensate peak appears between (b) and (c), and so the critical temperature is deduced to be 405 ± 15 nK with 2.4×10^5 atoms. Upon further cooling, the condensate fraction grows until no discernable thermal fraction remains in (j), with a total of 40,000 atoms in the BEC.

calculation of Eq. 5.8 [243], instead of approximating it to zero as was done in Sec. 2.7.1. The second effect is caused by the repulsive interactions between particles [244], which tend to lower the density at the centre of the trap and hence require further cooling in order to achieve the criterion for BEC of PSD ~ 2.6. Including both of these effects, the corrected critical temperature T_c is given by [243, 244]

$$T_c = T_c^0 \left[1 - \underbrace{0.73 \frac{\bar{\omega}}{\omega_{\text{ho}}} N^{-1/3}}_{\text{finite size}} - \underbrace{1.33 \frac{a}{a_{\text{ho}}} N^{1/6}}_{\text{interactions}} \right], \tag{5.9}$$

where $a_{\text{ho}} = \sqrt{\hbar/m\omega_{\text{ho}}}$, $a = 5.5\,\text{nm}$ is the s-wave scattering length [245], and $\bar{\omega} = (2\omega_\perp + \omega_{/\!/})/3$ and $\omega_{\text{ho}} = (\omega_\perp^2 \omega_{/\!/})^{1/3}$ are the arithmetic and geometric means of the trapping frequencies, respectively. With the parameters in our trap for the data in Fig. 5.3, finite size and interaction effects reduce the expected critical temperature by 2.2% and 6.9%, respectively, leading to $T_c = 426\,\text{nK}$, which is in reasonable agreement with the measured value of $405 \pm 15\,\text{nK}$.

Using the same data set, the fraction of atoms in the BEC state as the temperature is gradually reduced is depicted in Fig. 5.4. The measured condensate fraction is plotted versus the temperature scaled by T_c^0 for an ideal gas, the value of which is recalculated for each point based on the total number of atoms remaining, accounting for the overall loss of atoms when evaporatively cooling over the course of the data set.

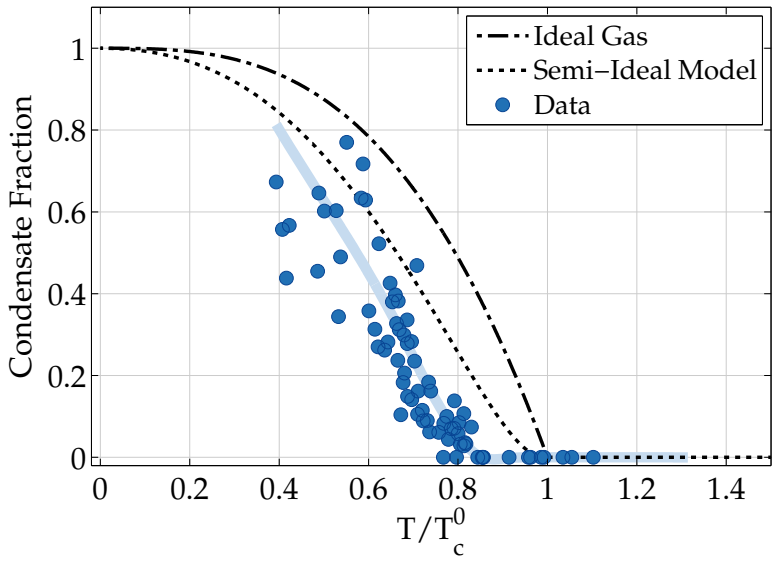

FIGURE 5.4: Growth of condensate fraction: The condensate fraction N_0/N is measured as a function of scaled temperature T/T_c^0, where T_c^0 is the theoretical critical temperature for a non-interacting ideal bose gas in the thermodynamic limit. The data are compared to both the ideal gas model (dashed-dot line) and the analytical semi-ideal model (dotted line), as described in the text. The smaller measured critical temperature of $T_c \sim 0.8 T_c^0$ may be attributed to an overestimation of $\sim 20\%$ in the detected atom number during absorption imaging, for example due to imperfect polarisation of the probe beam.

For the ideal gas in an harmonic potential, the condensate fraction should grow as $N_0/N = 1 - \left(T/T_c^0\right)^3$ [97], calculated also in Eq. 2.59, which is depicted by the dash-dotted line in Fig. 5.4. Also shown is the so-called *semi-ideal* model (dotted line), which includes interaction effects from condensed atoms but neglects those in the thermal cloud [246]. This model then considers the thermal component to be an ideal gas that experiences a combination of both the external trapping potential and the mean-field repulsion by the condensate, and is a good approximation because the density of the thermal part is so dilute that its interactions are almost negligible. The curve takes the implicit form [246]

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c^0}\right)^3 - \xi\frac{\zeta(2)}{\zeta(3)}\left(\frac{T}{T_c^0}\right)^2\left(\frac{N_0}{N}\right)^{2/5}, \tag{5.10}$$

where the parameter ξ accounts for the strength of interactions of atoms in the condensate, and is given by

$$\xi = \frac{1}{2}\zeta(3)^{1/3}\left(15N^{1/6}\frac{a}{a_{\text{ho}}}\right)^{2/5}. \tag{5.11}$$

For our trap parameters, the value of ξ is ~ 0.4, and does not depend strongly on N due to the power $1/15$ scaling. Several of the first experiments measuring the condensate fraction dependence on temperature, although clearing observing a sudden transition at $T \sim T_c^0$, already showed signs of deviating from the non-interacting model [158, 247]. Later, a more thorough study indicated clear deviations from the ideal gas model, and showed that the semi-ideal curve of Eq. 5.10 can describe the experimental data very well [248]. Our measured data in Fig. 5.4 shows an offset from the semi-ideal curve towards a reduced condensate fraction, along with a critical temperature of $T_c \sim 0.8T_c^0$. This discrepancy could occur if any of the physical parameters are incorrectly calibrated. The required correction to the harmonic trap frequencies is too large to explain it, however an overestimation of the inferred atom number of $\sim 20\%$ can account for the difference. This is quite possible, since any weak stray magnetic fields present during imaging can affect the polarisation of the probe beam, and therefore the photon scattering rate with atoms via the Clebsch-Gordan coefficients for various possible transitions. The type of data can in fact be used in the future as a means of calibrating the absolute atom number when the scattering cross-section is not known precisely.

5.3 Condensate Expansion

Another key signature of the presence of BEC comes from observing the expansion after switching off the trapping potential. An example is shown in Fig. 5.5, which depicts a cloud of 6.5×10^4 BEC atoms with 40% condensate fraction released from the copper Z-trap with frequencies $\left\{\omega_\perp, \omega_{/\!/}\right\} = 2\pi \times \{500\,\text{Hz}, 23\,\text{Hz}\}$, for times-of-flight of $3\,\text{ms} - 25\,\text{ms}$. The expansion shows the characteristic anisotropy, resulting in a reversal

of the aspect ratio at large times when compared with in-trap values. This is in contrast to a purely thermal cloud, which undergoes ballistic expansion and tends to an aspect ratio of 1 at large times due to the isotropic momentum distribution - an example of which is given in Fig. 3.41 of Sec. 3.9.

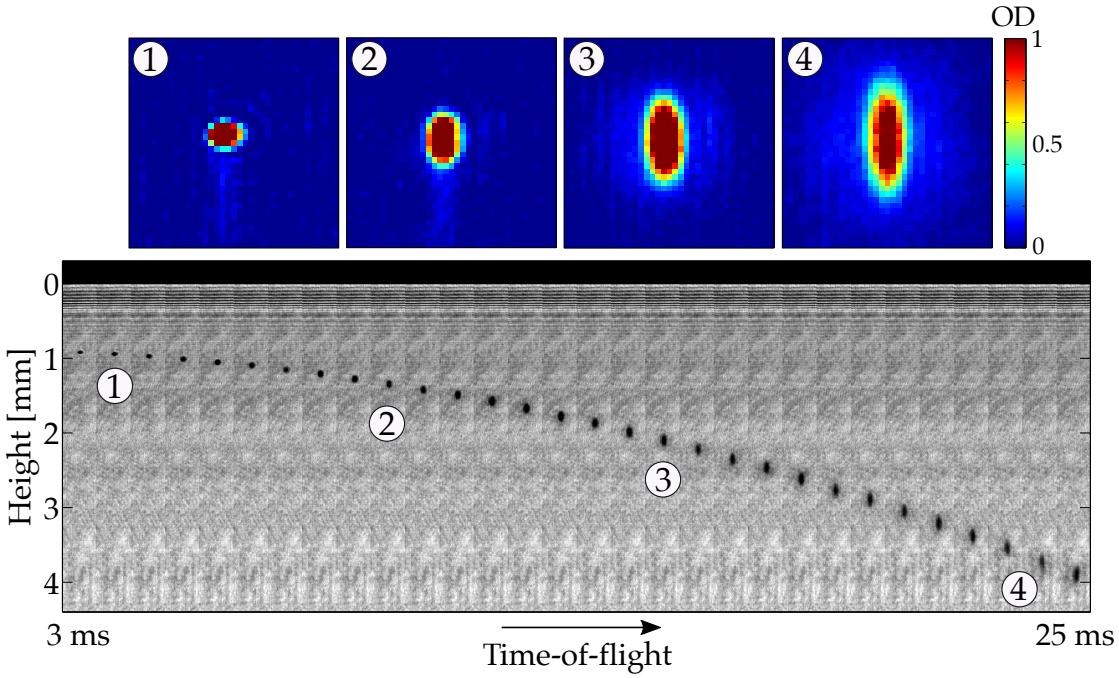

FIGURE 5.5: Expansion of a BEC during time-of-flight: A BEC of 6.5×10^4 atoms and 40% condensate fraction is prepared by directly evaporatively cooling in the copper Z-trap, initially at a distance of $800\,\mu\mathrm{m}$ from the chip surface. After fast removal of the trapping fields, the cloud expansion is recorded for times-of-flight from $3\,\mathrm{ms}$ to $25\,\mathrm{ms}$. The raw camera CCD pictures (lower panel) show the nature of the expansion during free-fall, whilst the optical density absorption images at selected times (upper panel) depict the reversal of the aspect ratio due to a release of interaction energy in the gas. Note that the centre-of-mass motion of this data was used to calibrate the imaging system magnification in Fig. 3.36 of Sec. 3.9.1

The BEC expansion is due to the anisotropy in the initial trapping potential, which has a stronger confinement in the radial direction by a factor of ~ 20 compared to the axial direction. The stronger confinement forces are balanced by repulsive interactions in the condensate, which leads to a larger gradient of the mean-field energy in the radial direction. After switching off the trapping potential, this interaction energy is rapidly converted to kinetic energy, leading to an accelerated radial expansion, whilst the axial direction barely changes. The interaction energy soon becomes negligible as the cloud becomes more dilute, at which point the expansion proceeds with constant velocity in each direction.

For a condensate in the Thomas-Fermi regime, it has been shown [249] that for a trap in which the frequencies depend on time, $\omega_i \rightarrow \omega_i(t)$, the density takes the form of the Thomas-Fermi profile from Eq. 2.71 in Sec. 2.7.3, but with radii which now also

depend on time

$$n(\mathbf{r},t) = n(0,t)\,\max\left\{1 - \frac{x^2}{R_x^2(t)} - \frac{y^2}{R_y^2(t)} - \frac{z^2}{R_z^2(t)}, 0\right\}. \tag{5.12}$$

The parabolic shape is therefore preserved in time, but with radii that are scaled according to $R_i(t) = R_i(0)b_i(t)$, with the time-evolution parameters governed by the three coupled differential equations

$$\ddot{b}_i + \omega_i^2 b_i - \frac{\omega_i^2}{b_i b_x b_y b_z} = 0. \tag{5.13}$$

Here, the second term is due to the trap potential and the last term to interactions, so the evolution of the gas during time-of-flight can be modelled by setting the second term equal to zero at $t = 0$. For an axially-symmetric trap such as ours, with frequencies $\omega_x \equiv \omega_z \equiv \omega_\perp$ and $\omega_y \equiv \omega_{/\!/}$, and defining a scaled time parameter $\tau = \omega_\perp t$, the equations reduce to

$$\frac{d^2 b_\perp}{d\tau^2} = \frac{1}{b_\perp^3 b_{/\!/}} \qquad\qquad \frac{d^2 b_{/\!/}}{d\tau^2} = \frac{\lambda^2}{b_\perp^2 b_{/\!/}^2}, \tag{5.14}$$

where $\lambda = \omega_{/\!/}/\omega_\perp$. These two equations have been solved analytically [249] for $\omega_\perp \gg \omega_{/\!/}$ (which is fulfilled for cigar-shaped traps), giving the equations

$$b_\perp(\tau) = \sqrt{1+\tau^2} \tag{5.15}$$

$$b_{/\!/}(\tau) = 1 + \lambda^2\left(\tau \arctan\tau - \ln\sqrt{1+\tau^2}\right), \tag{5.16}$$

highlighting the different nature of the expansion in each direction for condensates in the Thomas-Fermi regime. The radial direction undergoes accelerated expansion for $\tau < 1$, during which time the mean-field energy is released into kinetic energy. For our traps, this typically takes place in the first millisecond of expansion. The axial expansion is only significant for $\tau > \lambda^{-2}$, corresponding to $t \gtrsim 150\,\mathrm{ms}$ and so is not seen in our typical time-of-flight measurements (the falling cloud would hit the bottom viewport of the vacuum chamber after $\sim 80\,\mathrm{ms}$).

Figure 5.6 shows the measured Thomas-Fermi radii for the same data set described in Fig. 5.5. For $\tau \gg 1$, the radial expansion of Eq. 5.15 is linear, and the radial size evolves according to $R_\perp(t) = R_\perp(0)\; b_\perp(t) \simeq R_\perp(0)\;\omega_\perp t$. The chemical potential μ of the BEC in the Thomas-Fermi limit, given by Eq. 2.72, can then be written as

$$\mu = \frac{1}{2} m\,\omega_\perp^2 R_\perp^2(0) \approx \frac{1}{2} m \left(\frac{dR_\perp(t)}{dt}\right)^2 = \frac{1}{2} m v_\perp^2, \tag{5.17}$$

for $\tau \gg 1$ and $R_\perp(t) \gg R_\perp(0)$, where $v_\perp$ is the expansion velocity in the radial direction, and provides a way to experimentally measure the chemical potential. The radial expansion velocity is measured in Fig. 5.6 to be $5.38\,\mu\mathrm{m\,ms^{-1}}$, corresponding to

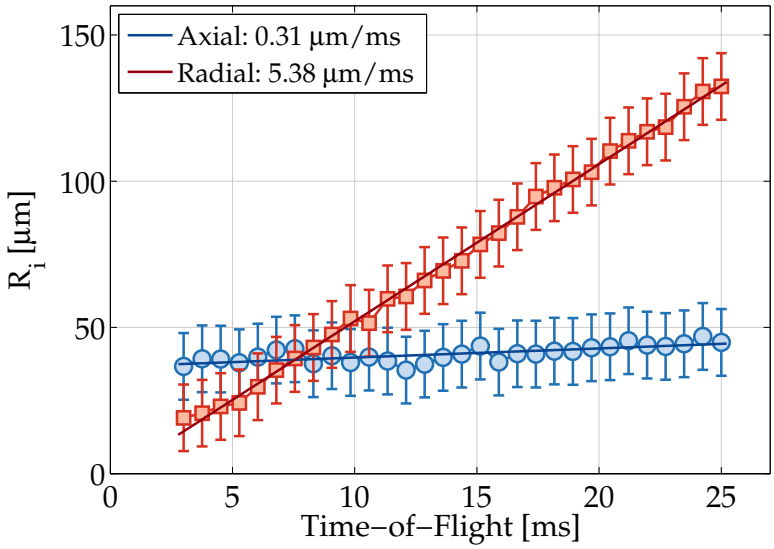

FIGURE 5.6: Thomas-Fermi radii during BEC expansion: The Thomas-Fermi radii are measured for various times-of-flight for the axial (○) and radial (□) directions. Linear fits result in expansion velocities of $0.31\,\mu\mathrm{m\,ms^{-1}}$ and $5.38\,\mu\mathrm{m\,ms^{-1}}$ for the two directions, respectively.

$\mu = k_\mathrm{B} \times 151\,\mathrm{nK} = h \times 3.14\,\mathrm{kHz}$ - consistent with the value of $145\,\mathrm{nK}$ calculated using Eq. 2.73.

5.4 Condensate Lifetime

An important quantity to know is the lifetime of the condensate, which determines the time available for performing experiments, and there are several processes which can cause decay of the BEC atom number. Firstly, there are always one-body losses due to collisions with the background gas, although this loss rate is typically orders of magnitude smaller than the BEC decay rate (the vacuum limited lifetime in the experiment is on the order of $20\,\mathrm{s}$). There are also both two and three body inelastic collisions between trapped atoms in the gas. Three-body losses occur when three atoms come together and two of them form a molecule, with the energy released going into the third atom [214]. All three atoms are removed from the trap, and the rate of the process is proportional to $n^2(\boldsymbol{r})$. Two-body spin relaxation can also occur, whereby an atom can spin flip into an anti-trapped state during a collision, releasing the Zeeman energy, and this process is proportional to $n(\boldsymbol{r})$. The dominant decay process for ^{87}Rb condensates in the $|F = 2, m_F = +2\rangle$ state are three-body losses [239], and due to the strong scaling these typically become relevant for peak densities higher than $10^{14}\ \mathrm{cm^{-3}}$ to $10^{15}\ \mathrm{cm^{-3}}$. In addition to collisional losses there is also technically-induced heating. Noise on the magnetic trapping fields below the kilohertz range leads to parametric heating of atoms if the frequency coincides with harmonics of the eigenfrequencies of the trap [250].

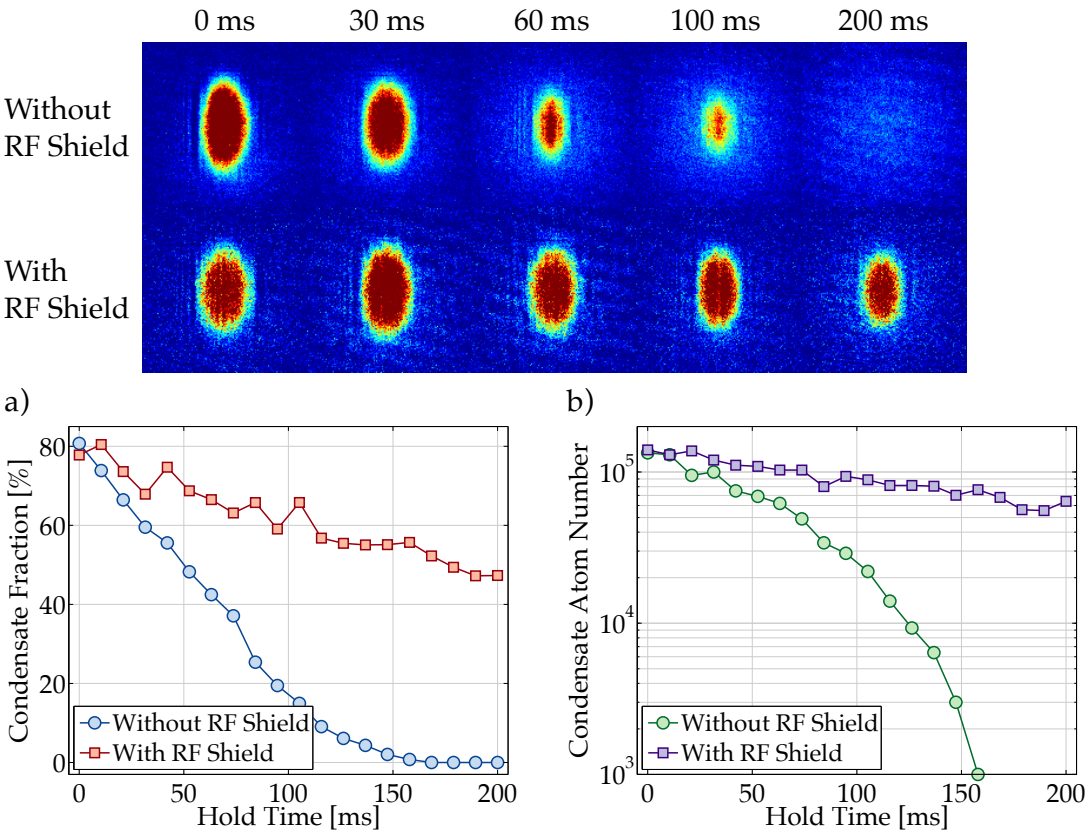

FIGURE 5.7: Lifetime of the BEC with and without an RF shield: An almost pure BEC, with 1.4×10^5 atoms and 80% condensate fraction is prepared on the atom chip by evaporatively cooling down to a RF frequency of $845\,\mathrm{kHz}$ ($15\,\mathrm{kHz}$ above the trap bottom), and imaged after $15\,\mathrm{ms}$ of ballistic expansion. The evolution is then monitored over the following $200\,\mathrm{ms}$ hold time both with and without a $900\,\mathrm{kHz}$ RF shield (lower and upper rows of optical density absorption images, respectively), as described in the text. Without the shield, the condensate fraction reduces to zero within $150\,\mathrm{ms}$, shown in (a). The application of the shield enables removal of the heated atoms and prolongs the condensate lifetime to several hundred milliseconds. This in addition eliminates the accelerated non-linear loss that is present without the shield, caused by heated atoms remaining in the trap, and is shown in (b).

In order to determine the condensate lifetime, an almost pure BEC with 1.4×10^5 atoms and 80% condensate fraction (with no discernable thermal component) was prepared on the atom chip, and its time-evolution in the trap was monitored, as shown in Fig. 5.7. The top row of optical density images show what happens to the BEC if it is simply left to evolve in the trap over time after removing the evaporative cooling RF field. It can be seen that over the data set atoms are transferred into the thermal component after being gradually excited out of the BEC, resulting overall in a decrease in the condensate fraction from 80% to 0% within $150\,\mathrm{ms}$ as shown in Fig. 5.7 (a). The atom number over the evolution time is constant, indicating that heated atoms remain confined in the trap. These heated atoms can induce further losses if they remain trapped by undergoing multiple collisions with the still condensed particles, which accelerates

the loss process and leads to the non-linear reduction in BEC atom number seen in Fig. 5.7 (b).

The lifetime of the condensate can be prolonged (lower row of optical density images) to several hundreds of milliseconds by applying a so-called "*RF shield*" [251, 252], whereby the RF field is left on constantly throughout the hold time. The frequency of the RF shield is chosen in this particular example to be $\nu_{\mathrm{sh}} = 900\,\mathrm{kHz}$, which with the trap bottom of $\nu_0 = 830\,\mathrm{kHz}$ corresponds to a truncation energy of $6.7\,\mu\mathrm{K}$. Using Eq. 2.73, the initial chemical potential of the BEC is calculated to be $267\,\mathrm{nK}$, and so this RF knife is 25 times higher than the energy of the cloud. The application of the RF knife then converts the detrimental heating rate into a loss rate by removing only the very high energy atoms, whilst not causing evaporation of the condensed component. The total number of atoms remaining at the end of the data set when using the RF shield is 90% when compared to the case without, indicating that only a small fraction of particles are being evaporated. The RF shield removes the non-linearity in the condensate number, shown in Fig. 5.7 (b), which then becomes exponential and indicates that the loss is density-independent. Finally, it is noted that although the initial in-situ peak density, given by $n(0) = \mu m/4\pi\hbar^2 a$ from Eq. 2.71, is calculated to be $7.3 \times 10^{14}\ \mathrm{cm}^{-3}$, there is no evidence of three-body losses in these data, which would typically have the signature of an initial over-exponential loss at times when the densities are highest [240].

5.5 Trap Switch-Off

When extinguishing all magnetic trapping fields to perform absorption imaging of the cloud, it is important that sufficient Ioffe field is maintained throughout the switching process. Although all the trapping currents are reduced to zero in less than $100\,\mu\mathrm{s}$, any oscillations during this time can lead to the emergence of a zero in the total field during the switching. If this happens, there will then not be a well-defined quantisation axis, and atoms can be projected from the $|F = 2, m_F = +2\rangle$ state across all five magnetic substates. This is especially important for very cold clouds and BECs, which sample the region of minimum field for a significant amount of time.

An experiment was performed to examine the effects of timings of the trapping currents, whereby a small delay was added to the switch-off time of the chip Z-wire with respect to that of the main X-Bias field. Since the Z-wire produces a strong field gradient, any projection into other spin states will translate into spatial separation due to the Stern-Gerlach effect, with the separation depending on the duration of time spent in the field gradient. The data series in Fig. 5.8 shows the results for various delays of the Z-wire switch-off time ranging from $-100\,\mu\mathrm{s}$ to $300\,\mu\mathrm{s}$, with negative delays corresponding to the Z-wire being switched off *before* the bias field. Each run was performed with a BEC, and all for the same time-of-flight of $20\,\mathrm{ms}$. If no delay is added,

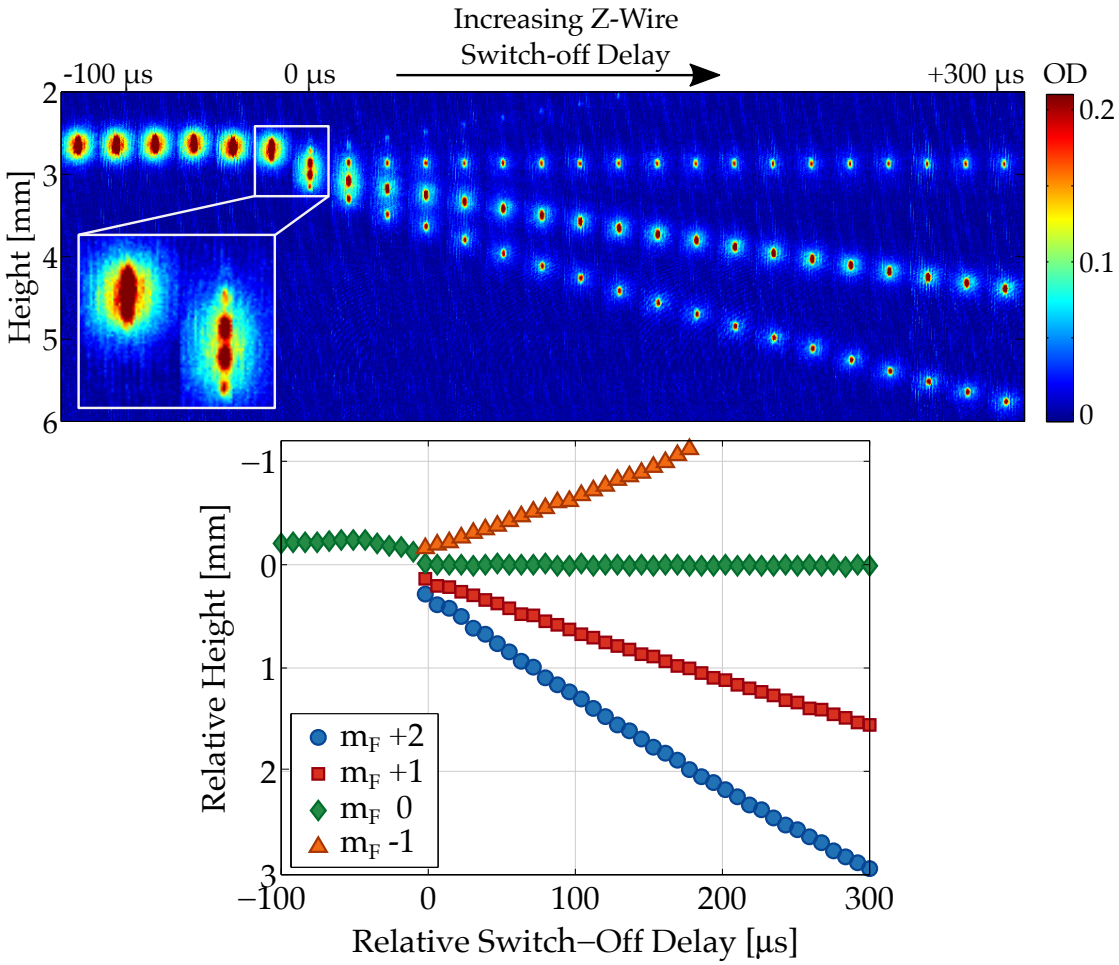

FIGURE 5.8: Spin flips during the magnetic trap switch-off process: The upper panel shows a series of optical density absorption images, taken with 20 ms ballistic expansion, whilst varying the relative delay between switch-off of the Z-wire and bias field currents. Negative delays indicate the Z-wire being switched *before* the bias field. If the two are switched simultaneously (zero delay) there is projection onto other spin states which translates into a spatial separation by the Stern-Gerlach force. The momentum imparted by the force is proportional to both the time spent in the field gradient and the m_F quantum number, as shown in the plot (lower). A delay of $-20\,\mu s$ is sufficient to suppress spin flips during the switch-off process.

i.e. the two are switched off simultaneously, it can be seen that there is indeed a projection into all spin states, but the spatial separation is small because the cloud does not spend a significant amount of time in the gradient. For larger delays, the separation increases directly proportional to the m_F quantum number through the Stern-Gerlach force, given by Eq. 2.11 in Sec. 2.2, as can be seen in Fig. 5.8 (lower). It is found that a small delay of $-20\,\mu s$ (i.e. switching the Z-wire off $20\,\mu s$ earlier than the bias field) is sufficient to prevent spin-flips in the switching process.

5.6 A BEC With Integrated Trapping Structures

Miniaturisation of cold atom setups is currently an active area of research [253–258], with the aim being to eventually create systems which can be taken out of the laboratory and be used as portable quantum sensors. The first BECs produced in the 1990s used large external coils, which consume several kW of power. The limitation came from the fact that external coils are situated far from the atoms, and so large currents are required to produce the necessary field gradients for evaporative cooling. The use of the atom chip goes some way to reducing power consumption, due to the close proximity of the cloud to the microfabricated trapping structure, and therefore allowing lower currents to be used. However, as implemented so far in this thesis, external coils are usually still needed in order to obtain homogeneous fields as required for operation of the Z-trap.

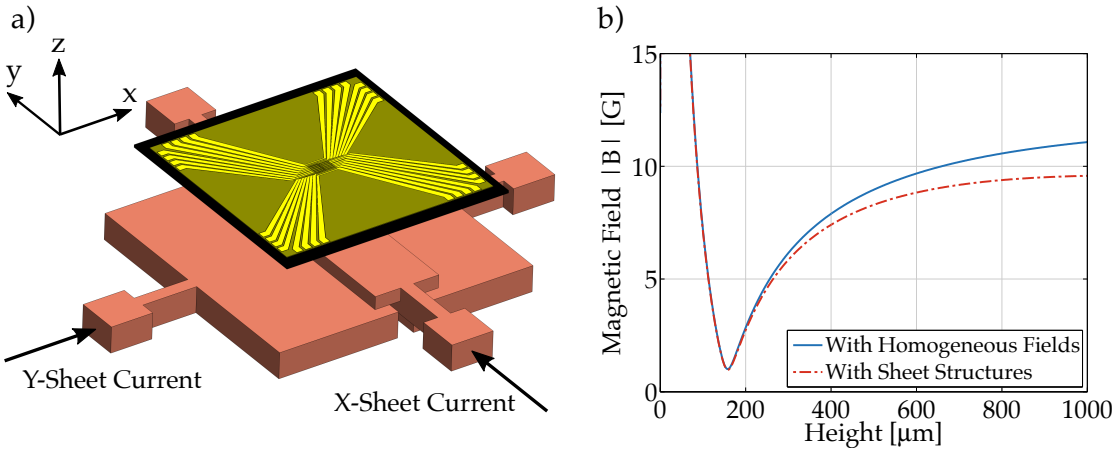

FIGURE 5.9: Arrangement for trapping without external coils: a) Two orthogonal copper sheet structures located beneath the surface of the atom chip produce nearly-homogeneous fields, and are used in conjunction with the Z-wire to trap atoms on the chip without external coils. The X-Sheet provides the main bias field for the Z-wire, whilst the Y-Sheet allows adjustment of the Ioffe field at the trap bottom. b) Calculation of the trapping field versus height from the chip surface, shown as a line cut through the trap minimum. When creating the trap using the sheet structures (dash-dot line) there is a slight additional tilt of the field as compared to the case when purely homogeneous fields are used (solid line), due to the sheet's gradient, but around the position of the minimum the potentials are almost identical. The currents used to create this trap are 1 A in the chip Z-wire, 50 A in the X-Sheet, and 7 A in the Y-Sheet.

In this experiment, there are a variety of copper structures located beneath the surface of the atom chip, as described in Sec. 3.3.1, which can be used to produce magnetic fields with low gradients. Specifically, there are two orthogonal sheet-like structures shown in Fig. 5.9 (a), labelled X-Sheet and Y-Sheet, producing magnetic fields in the x and y directions, respectively. Owing to their relatively wide widths, we have used these sheets to create the nearly homogeneous fields required for the Z-trap, and they

have been used in this setup to create a BEC after loading the atom chip *without* the use of any external coils.

As described in Sec. 4.8.2, BECs are usually produced on the atom chip with $1\,\mathrm{A}$ through the chip wire and a bias field of $12.4\,\mathrm{G}$, creating the trap minimum at a distance of $160\,\mu\mathrm{m}$ from the chip surface. From the calculations and measurements in Sec. 3.3.1, the X-Sheet structure produces a magnetic field of $0.25\,\mathrm{G\,A^{-1}}$ at this position, and so to replace the external coil we require a current in the X-Sheet of $50\,\mathrm{A}$. The Y-Sheet can then be used to adjust the residual Ioffe field at the trap bottom, and its field is measured in Fig. 5.10, giving a value of $(86 \pm 5)\,\mathrm{mG\,A^{-1}}$ and agreeing exactly with the finite element calculation in Sec. 3.3.1. Therefore, a current of $7\,\mathrm{A}$ through the Y-Sheet sets the trap bottom to be $1.2\,\mathrm{G}$. With these values the trapping potential is almost identical to one created using homogeneous external fields, as shown in Fig. 5.9 (b). Due to the narrower $10\,\mathrm{mm}$ width of the X-Sheet, the field it produces has a residual gradient, and the only difference is a slight tilt of the potential for large distances from the minimum.

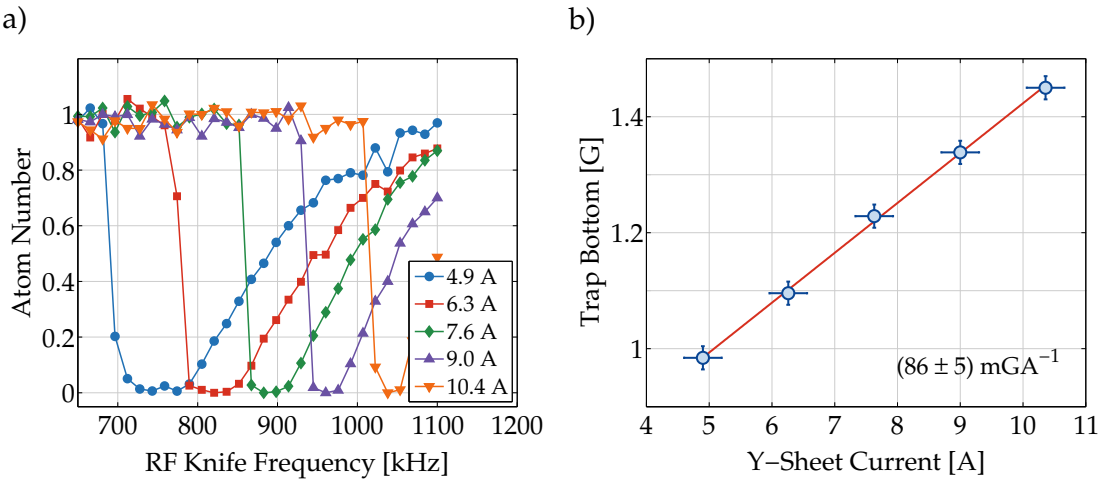

FIGURE 5.10: Calibration of the field produced by the Y-Sheet structure: a) The trap bottom shift is measured for various currents through the Y-Sheet using the RF knife technique. b) The shift in trap bottom as a function of applied current calibrates the field to be $(86 \pm 5)\,\mathrm{mG\,A^{-1}}$, agreeing exactly with a finite element calculation.

Atoms are captured in the chip trap at the end of the copper evaporative cooling phase by ramping down the copper Z-wire current to zero whilst simultaneously increasing a $50\,\mu\mathrm{m}$ atom chip wire current to $1\,\mathrm{A}$, in the same manner as that described in Sec. 4.8.1. In parallel with this, the currents through the external coils are reduced to zero, whilst those in the X and Y sheet structures are increased to $50\,\mathrm{A}$ and $7\,\mathrm{A}$, respectively. At this point, the cloud is now confined using only the integrated trapping structures, and all external coil currents have ceased. Application of a linear RF sweep from $3.5\,\mathrm{MHz}$ to $880\,\mathrm{kHz}$ over $2\,\mathrm{s}$ then produces a BEC with 10^5 atoms - the same number as in the case when purely homogeneous external fields are used.

The use of the sheet structures has several advantages. The more direct, immediate benefit for this experimental setup is that the structures are essentially non-inductive, allowing their currents to be controlled much faster. Since the trap frequencies of these typically elongated atom chip geometries are on the order of $\{\omega_\perp, \omega_{/\!/}\} \sim 2\pi \times \{1\,\mathrm{kHz}, 20\,\mathrm{Hz}\}$, the ability to apply current ramps on a sub-millisecond timescale means that non-adiabatic changes with respect to all directions of the external potential can be made. Such rapid modifications if chosen correctly constitute a *quench* of various parameters of the system (such as density, interaction strength, temperature, or dimensionality), inducing excitations [259], and offers the possibility of studying the interconnected roles of dimensionality and thermalisation in out-of-equilibrium many-body quantum gases [114], which are not yet fully theoretically understood [260, 261].

A second significance of the sheet structures described above relates to the vision of eventually bringing cold atom technology beyond the laboratory setting, with a view to fulfilling a variety of practical applications, such as magnetic field sensing [262, 263], measurements of gravity [264, 265], inertial navigation [266, 267], and atomic timekeeping devices [268], to name a few. For these systems to become portable and be taken "into the field", they will need to eventually become more rugged and robust. A large amount of progress has already been made, but most work has focussed so far mainly on the magneto-optical trap used to initially collect and cool clouds of atoms. Specifically, pyramidal reflector MOTs have been implemented, in which a single input beam can be used to provide all trapping and cooling light [269–271]. These systems then soon became miniaturised and integrated onto chips [272, 273], and eventually were fully planarised using diffraction gratings to capture a large number of atoms [274, 275] - a technique which seems to be the most promising for the future applications. Indeed, several companies are already developing portable cold atom sensor systems, which are now on their way to becoming a commercial reality.

With regards to magnetic trapping, to a large extent miniaturisation and integration has already somewhat taken place with the use of both the atom chip [54, 59], and with the now standard "U" and "Z" shaped structures located beneath chip surface [61, 212] to create the inhomogeneous fields. However, all experiments producing BECs with atom chips (for example, see Refs. [276–278] for comparable experimental setups) require at least one external homogeneous field produced by macroscopic coils. The *iSense* project [254] has managed to demonstrate magnetic trapping using innovative integrated coil-like structures beneath the chip surface [279], but that system has not pushed cooling all the way to BEC. The unique structures demonstrated here are capable of bridging the gap between length scales - on the one hand retaining something relatively small and portable, whilst also being able to provide the large magnetic fields necessary to produce quantum degenerate atomic samples. The large cross-sectional areas of the conductors lead to a more modest total power consumption for all the bias fields of $750\,\mathrm{mW}$ including all electrical contacts (compared with $5\,\mathrm{W}$ when using external coils), making the chip Z-wire itself the dominant factor with a consumption of

$3.6\,\mathrm{W}$. In addition, the trapping structures are geometrically aligned by construction, keeping the tuning of experimental parameters to a minimum, which is crucial when aiming for robust, rigid devices within a single assembly.

In summary, after describing in detail the experimental sequence in Chapter 4, this chapter has then focussed on measurements and characterisation of BECs. Section 5.1 details how the measured optical density images are analysed in order to extract the pertinent physical quantities, such as temperature, atom number, and condensate fraction. The fitting functions and algorithms applied were outlined in detail, along with the importance of accounting for the density enhancement due to Bose statistics. In Sec. 5.2 and 5.3, two of the most basic BEC measurements were performed, respectively - the growth of the condensate with reducing temperature, and the signature anisotropic expansion of the condensate after being released from the trap, with reasonable agreement being found when compared with the theoretical descriptions in each case. The lifetime of the condensate in the trap was measured in Sec. 5.4 - an important quantity which determines the amount of time available for performing experiments with BECs. It was demonstrated how the lifetime could be significantly extended from $150\,\mathrm{ms}$ to $\sim 500\,\mathrm{ms}$ with the use of a so-called "RF shield". Section 5.5 illustrated the importance of the relative timings of trapping currents during switch-off, and depicted the projection of atoms across all magnetic Zeeman levels that can happen unless care is taken to ensure a sufficient quantisation field at all times. Finally, in Sec. 5.6 it was shown how BECs can be produced using the novel sheet-like trapping structures unique to this experimental system, which are able to act as a "drop-in" replacement for external macroscopic coils when producing the homogeneous bias fields.

Chapter 6

Summary and Outlook

This thesis has described the implementation and development of an experimental system for routinely producing quantum degenerate Bose gases of rubidium-87, trapped using the micropotentials generated by an atom chip for the first time in Nottingham. After giving an overview of the current state of the relatively young field of ultracold atoms in Chapter 1, key theoretical formulations were presented in Chapter 2, with particular attention paid to the relevant underlying physics directly related to measurements in later chapters.

Chapter 3 presented the most important aspects of the experimental apparatus, including the atom chip, and systems for vacuum, lasers, precise timing control, and optical imaging. Described in addition is the unique array of millimeter-sized copper structures located beneath the atom chip surface, for producing flexible trapping potentials, the magnetic fields produced by which were characterised and compared with full finite element calculations.

A detailed account of the procedure developed for reliably producing BECs on the atom chip was given in Chapter 4. Specifically in summary, the process begins with a loading rate of $5 \times 10^7\ \mathrm{s}^{-1}$ from the 2D MOT into the 3D mirror-MOT, which captures typically 1.5×10^8 atoms within $10\,\mathrm{s}$. This cloud is then sub-Doppler cooled, spin-polarised with a purity of 96%, and transferred with $\sim 70\%$ efficiency to a purely magnetic trap generated by the copper Z-wire. After adiabatic compression followed by a $10\,\mathrm{s}$ forced evaporative cooling phase, a cloud with typically 6×10^6 atoms and a temperature of $20\,\mu\mathrm{K}$ can be loaded onto the atom chip with 100% efficiency. A final $2\,\mathrm{s}$ cooling sweep results in the transition to BEC at $T_c \sim 800\,\mathrm{nK}$, with almost-pure on-chip condensates of 10^5 atoms.

Measurements and characterisation of a selection of the salient properties of BECs were presented in Chapter 5. This included analysis of the ground state occupation as a function of temperature, the expansion dynamics of the condensate after release from the trap, and the BEC lifetime - which could be prolonged to several hundreds of milliseconds with the aid of an RF shield. Finally, it was demonstrated how BECs can be produced on the atom chip using only the integrated trapping structures, without the use of external coils. This significantly reduces power consumption required for the bias fields from $5\,\mathrm{W}$ to $750\,\mathrm{mW}$, and importantly increases the controllability of the trapping potentials.

The achievement of BEC is a significant milestone in an ultracold atom experiment, the pursuit of which was the focus of the majority of this PhD project. After producing the first condensates, approximately a further year was devoted to greatly improving the apparatus stability. This was a lengthy process, since considerable effort is required to overcome the many obstacles and technical challenges faced. Much of the equipment utilised cannot be bought as standard, and as such often needs to be custom designed and then bulit in-house. BEC production requires the orchestration of many operations both in series and in parallel to a high degree of precision, as has been presented throughout this thesis, and now after a ~ 20 minute warm-up period from switching on the system, condensates can be generated and maintained throughout the working day and overnight. This level of stability and repeatability, together with the thorough characterisation of the system carried out during this work, is important for being able to perform well-controlled, reliable studies in the future, and in particular is crucial at times when the data collection process requires repetition of the experiment several hundreds of times to build statistics.

Experimental progress in the research field of ultracold atomic gases has developed with rapid pace over the last few decades, and the number of experimental systems operating in laboratories around the world is continually growing. These experiments have evolved to the point where they can be used to realise in the lab a wide variety of simple model systems, thereby allowing direct comparison with theoretical calculations, with a view to casting new insights and obtaining deeper understanding into the field of condensed matter physics. Although the equilibrium behaviour is reasonably well-understood [280], currently a particularly active research area is the study of the dynamics of closed quantum many-body systems that are *out of equilibrium* [260, 261, 281, 282], for which there remains no general, systematic theoretical framework. Whilst much progress has been made, there are still many difficult and significant open questions. For instance, the exact mechanism and timescale on which an isolated interacting quantum many-body system can relax and evolve to thermal equilibrium - a process which is described with statistical physics in the classical case, but it is not yet clear as to how far the ideas of statistical physics can extend in the description of these problems. In addition, the question of how the finite temperature present in experiments, and the corresponding thermal fluctuations to which it leads, will affect the results. It is also difficult to say whether the generalisation of simple models will be able to accurately describe the results of more complex experimental systems, or if perhaps there exist any universal properties or scaling laws which can be applied to predict the behaviour of other systems [259]. Experiments are now possible which are capable of shedding light on some of these issues, and it will therefore be crucial for there to be a close collaboration with theory in the future.

Cold atomic gases provide an excellent testing ground for simple models due to

their high degree of controllability, and the ability to modify various properties dynamically. Owing to them being exceptionally well-isolated from the surrounding environment, trapped gases can be treated as closed systems to a good approximation on the timescales of experiments. In addition to this, their extremely low temperatures and diluteness mean that dynamics occurs on the relatively long timescales of milliseconds, which is ideal for the study of non-equilibrium phenomena. State-of-the-art technological innovations have made it possible to now probe clouds at the sensitivity level of a single atom - both *in-situ*, using high-resolution optical systems [283–285], and in time-of-flight expansion, using the so-called "light sheet" fluorescence imaging technique [286]. This technique was used, for example, to follow the complex dynamical non-equilibrium decay of atoms in a one-dimensional Bose gas back to the ground state of the trap after inducing an excitation to the first radially excited state [287]. In addition, the phase information of a gas can be directly accessed via matter-wave interference [71], and together these aforementioned methods allow the atomic density, momentum and fluctuation distributions to measured with unprecedented precision.

One of the most conceptually simple operations that can be performed to bring a system out of equilibrium is a so-called *quench*, in which an abrupt (on the relevant timescales) change is made to a parameter, such as the confining potential, interaction strength, temperature, or density, and the subsequent time-evolution is monitored. Several landmark experiments along this vein have already been conducted using ultracold gases. For example, the periodic collapse and revival of a matter-wave interference pattern was observed by Greiner et al. [288], by performing a quench of their optical lattice potential depth and following the subsequent dynamics. A particularly ideal system to study is the one-dimensional Bose gas, which is a complex system but one that is reasonably well-understood theoretically [99, 100, 289]. The apparent absence of any thermalisation of strongly-correlated one-dimensional Bose gases was observed by Kinoshita et al. [114], by trapping the clouds in an optical lattice and then quenching from equilibrium, highlighting the interplay between thermalisation and integrability of the system. The large range of extreme aspect ratios afforded by the micropotentials created using atom chips makes them an ideal tool for studying one-dimensional clouds, whilst at the same time offering the complementary advantage of being able to create a single realisation of the system of interest, and therefore to measure local phase fluctuations. This is in contrast to optical lattices, in which certain information inherently becomes averaged out in time-of-flight measurements as a result of there being many instances per shot [101]. Atom chips have been used, for example, to investigate the non-equilibrium coherence dynamics of coupled one-dimensional Bose gases [111].

The atom chip is a perfect tool for these types of studies, and the experimental system outlined in this work stands out due to the flexibility provided by the combination of the microfabricated trapping structures and the integrated sheet structures detailed in Chapter 5, which are unique to this setup. The trapping potentials can be tuned independently along the spatial dimensions over a large parameter range and

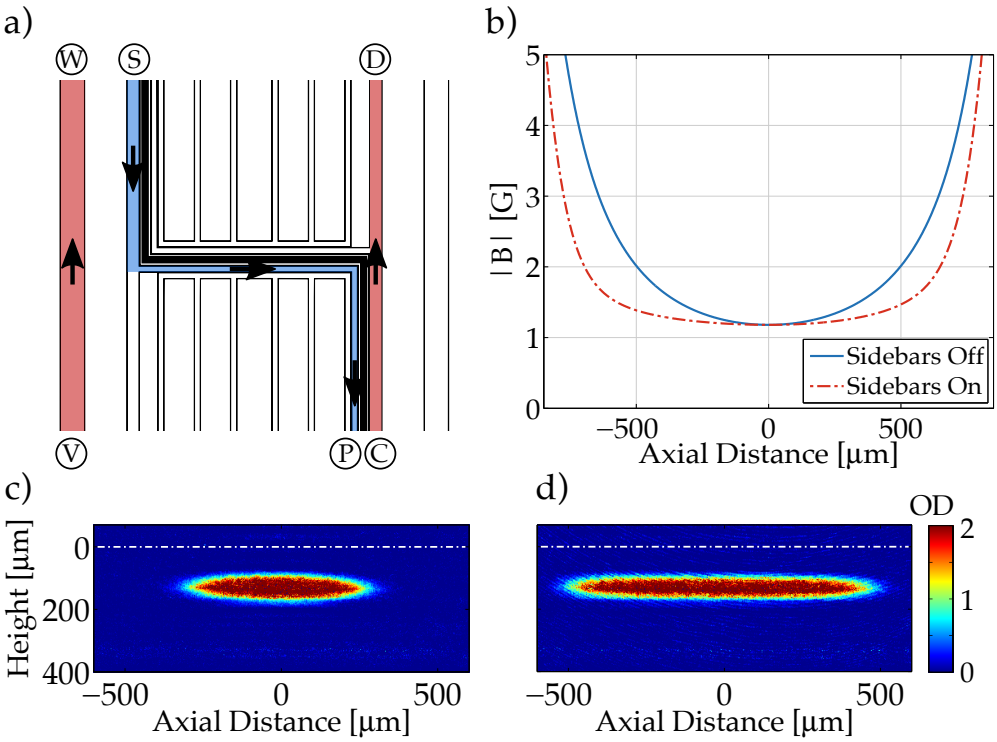

FIGURE 6.1: Tuning the axial confinement in an atom chip trap: a) Trapping on the chip is achieved by running a current of $800\,\mathrm{mA}$ through the $50\,\mu\mathrm{m}$ wide Z-shaped conductor S→P, along with a bias field of $-12.3\,\mathrm{G}$ to create a minimum at $130\,\mu\mathrm{m}$ from the chip surface. By running currents through the sidebars structures V→W ($200\,\mu\mathrm{m}$ wide) and C→D ($100\,\mu\mathrm{m}$ wide), the field produced by the legs of the Z-wire can be partially cancelled, thereby reducing the axial confinement. b) The effect of the sidebars is shown by comparing the magnetic field along the axial eigen-direction of the trap, showing a cancelling of the harmonic part of the trap and leading to a more box-like potential. In this example, $I_{\mathrm{VW}} = 2.05\,\mathrm{A}$ and $I_{\mathrm{CD}} = 0.94\,\mathrm{A}$, which modifies the axial trap frequency from $\omega_{/\!/} = 2\pi \times (17\,\mathrm{Hz} \rightarrow 3.5\,\mathrm{Hz})$, whilst leaving the radial frequency unchanged at $\omega_{\perp} = 2\pi \times 1.2\,\mathrm{kHz}$ in each case. The lower absorption images (c) and (d) show the experimental realisation of these potentials by the measured optical density of $2.4\,\mu\mathrm{K}$ clouds after $1\,\mathrm{ms}$ time-of-flight. The image in c) shows the configuration without the sidebars, and b) is the case in which they have been added.

with sub-millisecond timescales, giving purer access to quench protocols. We have recently begun to learn how to precisely manipulate clouds with the atom chip, exploring a variety of different trapping geometries and comparing them to calculations. As mentioned, due to the strong transverse confinement offered by atom chips, they are particularly suited to preparing clouds with effectively reduced dimensionality and extreme aspect ratios, and a wide range of structures are available with the chip design currently in use in the experiment, which can be used to modify the trapping potentials accordingly. An example of tuning the geometry is shown in Fig. 6.1. Here, two additional sidebar structures have been used to modify the axial confinement of the cloud. The sidebars carry currents anti-aligned with the legs of the Z-wire, which partially compensates the harmonic term in the axial trapping field, and leads to a more box-like potential in which the quartic term becomes more significant [290, 291]. In this

particular case, the axial trap frequency was relaxed from $\omega_{/\!/} = 2\pi \times (17\,\mathrm{Hz} \rightarrow 3.5\,\mathrm{Hz})$, with an unchanged radial frequency of $\omega_\perp = 2\pi \times 1.2\,\mathrm{kHz}$, corresponding to a modification of aspect ratio ($\omega_\perp/\omega_{/\!/}$) from 70 to 340. The ability of these patterned wire structures to create in a simple way such tunable potentials allows controlled experimental access to elongated Bose gases over a large parameter space.

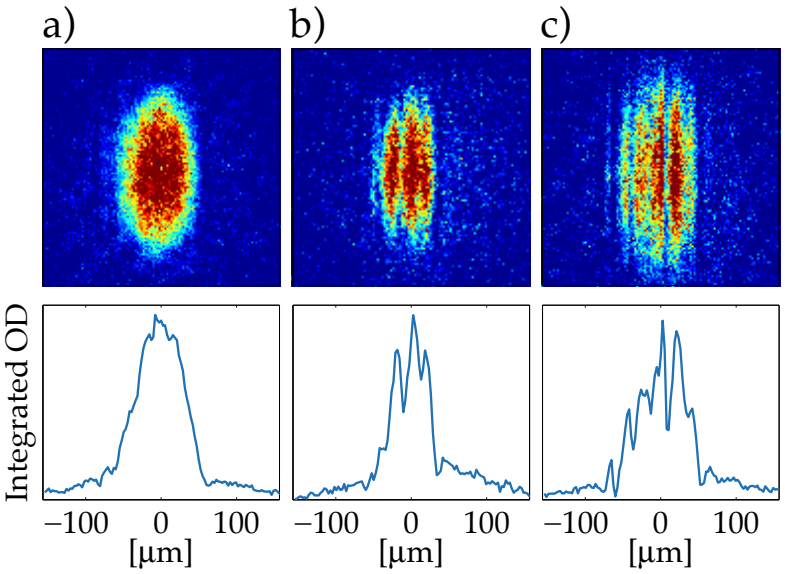

FIGURE 6.2: Examples of the emergence of phase fluctuations in elongated geometries: The optical density absorption images - along with their corresponding line profiles integrated along the vertical direction - depict condensates formed in a variety of trapping geometries, with an expansion time of $20\,\mathrm{ms}$ after switching off the trapping potential. The images correspond to trapping frequencies $1/2\pi\,\{\omega_{/\!/}, \omega_\perp\}$ of: a) $\{20\,\mathrm{Hz}, 140\,\mathrm{Hz}\}$; b) $\{16\,\mathrm{Hz}, 600\,\mathrm{Hz}\}$; c) $\{16\,\mathrm{Hz}, 1.3\,\mathrm{kHz}\}$. For more isotropic geometries such as (a), there is phase coherence over the whole gas, whilst for elongated clouds such as (b) and (c) there are axial phase fluctuations which translate to density fringes with sufficient time-of-flight.

As described in Chapter 1, when working with these elongated geometries it is possible for the gas to enter the quasi-condensate regime, whereby the equilibrium state is described by a single macroscopic wavefunction but with a fluctuating phase profile [99]. Since the fluctuations are driven by thermal excitations, only for temperatures lower than the characteristic temperature T_ϕ does the phase coherence length become longer than the length of the gas, effectively establishing a common phase [105]. Whilst recently exploring various different trapping geometries as described above, a strong indication that the gas enters the quasi-condensate regime in our experiment is the appearance of interference fringes in the BEC density distribution after time-of-flight, and selected examples of typical data obtained are given in Fig. 6.2. Since a gradient of phase in a BEC gives rise to a velocity potential [155], the fluctuating phase profile along the length of the gas effectively gives rise to an initial longitudinal velocity distribution. When the trapping potential is removed, these various parts of the gas with different initial positions can self-interfere when overlapping after a long enough expansion time, giving rise to the vertical fringes evident in Fig. 6.2 (b) and (c). It

can be seen that the significance of these fringes depends directly on the geometry in which the condensate is formed, as was observed by Dettmer, et al. [105], and for more isotropic traps such as Fig. 6.2 (a) there is essentially full phase coherence over the whole gas.

Whilst a rigorous characterisation of the interference fringes seen in Fig. 6.2 has not yet been fully carried out, preliminary analysis indicates that the aspect ratio of the trap influences their spatial frequency, as is possibly seen by the examples in (b) and (c). This is entirely plausible since the regions of constructive interference are dependent on the interplay between the phase coherence length and the overall length of the gas. If this indeed turns out to be the case, we plan to perform experiments in which a non-adiabatic quench between geometries is performed, and the dynamics by which the system reaches the new state is observed. It is an interesting question to determine the timescale on which the gas will react when quenching the trapping potential from 3D to 1D. If, for example, the interference fringes will melt away, and whether this process is symmetric when reversing the quench direction. The flexible potentials generated by atom chips are perfect for exactly these types of quench experiments, due to their high spatial and temporal controllability, and the system developed over the course of this thesis will serve well as a tool for studying out-of-equilibrium quantum gases in the near future.

Appendix A

Magnetic Field of a Finite Rectangular Conductor

The magnetic field $\boldsymbol{B}(\boldsymbol{r})$ generated by a steady current flowing through a conductor is given by the Biot-Savart law,

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_V \boldsymbol{j}(\boldsymbol{r}') \times \frac{\boldsymbol{r} - \boldsymbol{r}'}{|\boldsymbol{r} - \boldsymbol{r}'|^3} \, \mathrm{d}V', \tag{A.1}$$

where $\boldsymbol{r}'$ is the vector to an element of volume $\mathrm{d}V'$ which carries a current density vector $\boldsymbol{j}(\boldsymbol{r}')$, and the total field is given by the integral of these elements over the total conductor volume V. This equation can be applied to calculate the magnetic field at position $\boldsymbol{r}$ resulting from a constant current flowing through a finite rectangular conductor, as shown in Fig. A.1.

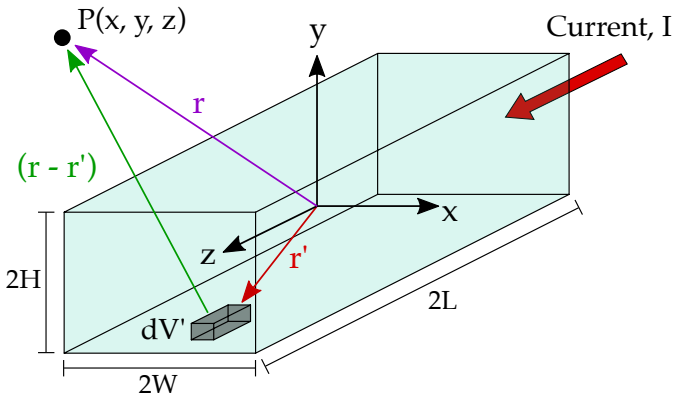

FIGURE A.1: Coordinate system for calculating the magnetic field from a rectangular conductor: An element of volume dV' located at position $\boldsymbol{r}'$ carrying an elemental current density $\boldsymbol{j}(\boldsymbol{r}')$ produces a field at position $\boldsymbol{r}$ which depends on the relative distance vector $(\boldsymbol{r} - \boldsymbol{r}')$, given by the Biot-Savart law. The total field produced by a rectangular conductor of dimensions $2W \times 2H \times 2L$ carrying a constant current I can be calculated by integrating over the total conductor volume.

Taking the coordinates of $\boldsymbol{r}$ to be (x, y, z) and those of $\boldsymbol{r}'$ to be (x', y', z'), and with the centre of the coordinate system at the centre of the conductor, we can write

$$\frac{\boldsymbol{r}-\boldsymbol{r}'}{|\boldsymbol{r}-\boldsymbol{r}'|^3} = \frac{(x-x')\hat{\mathbf{i}} + (y-y')\hat{\mathbf{j}} + (z-z')\hat{\mathbf{k}}}{\left(\sqrt{(x-x')^2+(y-y')^2+(z-z')^2}\right)^3}, \tag{A.2}$$

where $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ are unit vectors in the x, y and z directions, respectively. Assuming that the current density is *constant* throughout the conductor, and flowing in the positive z-direction, we can write

$$\boldsymbol{j}(\boldsymbol{r}') = \boldsymbol{j} = \frac{I}{4WH}\hat{\mathbf{k}}, \tag{A.3}$$

where I is the current flowing through the cross-sectional area $4WH$. Inserting Eq. A.2 and A.3 into Eq. A.1, and evaluating the cross-product, we obtain in Cartesian coordinates an expression for the total magnetic field at position $\boldsymbol{r}$

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0 I}{16\pi WH}\int_{-L}^{L}\int_{-H}^{H}\int_{-W}^{W} \frac{-(y-y')\hat{\mathbf{i}} + (x-x')\hat{\mathbf{j}}}{\left(\sqrt{(x-x')^2+(y-y')^2+(z-z')^2}\right)^3}\,\mathrm{d}x'\,\mathrm{d}y'\,\mathrm{d}z', \tag{A.4}$$

with field components present in only the x and y directions, and no component parallel to the current flow.

Evaluating the Integral

Concentrating for now on only the x component of the magnetic field, integrating out the y' dimension first using substitution results in

$$\begin{aligned} B_x(x,y,z) = \frac{\mu_0 I}{16\pi WH}\int_{-L}^{L}\int_{-W}^{W} &-\frac{1}{\sqrt{(x-x')^2+(y-H)^2+(z-z')^2}} \\ &+\frac{1}{\sqrt{(x-x')^2+(y+H)^2+(z-z')^2}}\,\mathrm{d}x'\,\mathrm{d}z'. \end{aligned} \tag{A.5}$$

Secondly, integrating out the x' direction gives

$$\begin{aligned} B_x(x,y,z) = \frac{\mu_0 I}{16\pi WH}\int_{-L}^{L} &-\ln\left[\frac{x+W+\sqrt{(x+W)^2+(y-H)^2+(z-z')^2}}{x-W+\sqrt{(x-W)^2+(y-H)^2+(z-z')^2}}\right] \\ &+\ln\left[\frac{x+W+\sqrt{(x+W)^2+(y+H)^2+(z-z')^2}}{x-W+\sqrt{(x-W)^2+(y+H)^2+(z-z')^2}}\right]\,\mathrm{d}z'. \end{aligned} \tag{A.6}$$

To perform the final integral, *Mathematica* was used to evaluate the general form

$$\begin{aligned}\int \ln\left[a+\sqrt{a^2+b^2+(x_0-x)^2}\right]\,\mathrm{d}x = x &+ b\arctan\left[\frac{x_0-x}{b}\right] \\ &- b\arctan\left[\frac{a(x_0-x)}{b\sqrt{a^2+b^2+(x_0-x)^2}}\right] \\ &+ (x_0-x)\ln\left[a+\sqrt{a^2+b^2+(x_0-x)^2}\right] \\ &+ a\ln\left[x_0-x+\sqrt{a^2+b^2+(x_0-x)^2}\right].\end{aligned} \tag{A.7}$$

When using this equation to evaluate the integral in Eq. A.6, the first two terms cancel throughout, and so can be omitted. The final field components can be collected and written in the compact form

$$\begin{cases} B_x(x,y,z) = \dfrac{\mu_0 I}{16\pi WH}\displaystyle\sum_{k,l,m=0}^{1}(-1)^{k+l+m}\,\mathcal{F}(\tilde{x}_k,\tilde{y}_l,\tilde{z}_m) & \text{(A.8)} \\ B_y(x,y,z) = -\dfrac{\mu_0 I}{16\pi WH}\displaystyle\sum_{k,l,m=0}^{1}(-1)^{k+l+m}\,\mathcal{F}(\tilde{y}_l,\tilde{x}_k,\tilde{z}_m) & \text{(A.9)} \\ B_z(x,y,z) = 0, & \text{(A.10)} \end{cases}$$

where the definitions have been used that $\tilde{x}_k = x - W(2k-1)$, $\tilde{y}_l = y - H(2l-1)$, $\tilde{z}_m = z - L(2m-1)$, and

$$\begin{aligned}\mathcal{F}(x,y,z) = y\arctan\left[\frac{xz}{y\sqrt{x^2+y^2+z^2}}\right] &- x\ln\left[z+\sqrt{x^2+y^2+z^2}\right] \\ &- z\ln\left[x+\sqrt{x^2+y^2+z^2}\right].\end{aligned} \tag{A.11}$$

These equations have been used to give an example of the characteristics of the field produced by a finite dimensioned wire, shown in Fig. A.2. For distances from the wire approximately greater than its width, the dimensions of the conductor become less significant, and the calculation for an infinitely thin wire is a good estimation. For distances close to the wire, the field components are reduced and there is a finite value of magnetic field at the conductor's surface.

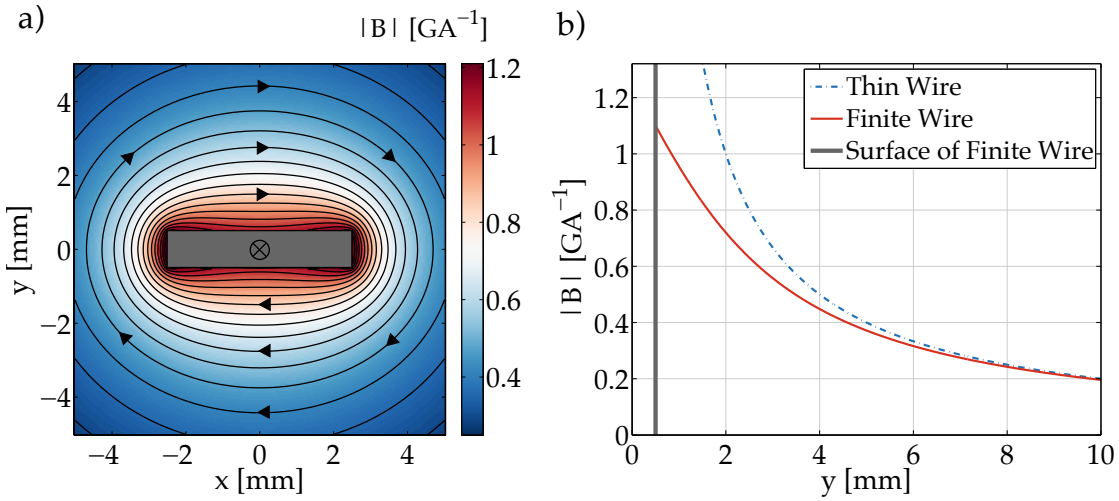

FIGURE A.2: Magnetic field $|\boldsymbol{B}|(\boldsymbol{r})$ produced in the exterior by a rectangular conductor of finite width $5\,\mathrm{mm}$ and height $1\,\mathrm{mm}$, and infinite length: a) The field lines in the xy plane perpendicular to the current flow become circular at large distances when the wire appears infinitely thin, but flatten close to the surface of the conductor. The direction of current flow is into the page. b) The magnetic field as a function of height above the wire surface. For comparision, the field of an infinitely thin wire is plotted, and the two are in agreement for distances approximately greater than the finite wire width. Importantly, whilst the infinitely thin wire field diverges for small distances, the rectangular wire has a finite value at its surface.

Appendix B

Removing Interference Fringes in Absorption Images

In cold atom experiments, quantitative information on trapped clouds is commonly gained using absorption imaging, whereby the attenuation of a probe beam due to photon scattering is measured when passing through the sample, as described in Sec. 3.9 of Chapter 3. Assuming a weak probe beam (much less than the saturation intensity for the transition), which is tuned to the atomic resonance, the intensity attenuation caused by traversing the cloud along the z direction is given from the Beer-Lambert Law [230]

$$I_{\text{out}}(x,y) = I_{\text{in}} e^{-\text{OD}(x,y)}, \tag{B.1}$$

where the optical density $\text{OD}(x,y)$ is related to the two-dimensional column density $n_{\text{2D}}(x,y)$, which is an integration along the z direction, via the scattering cross section σ_0 by the relation

$$\text{OD}(x,y) = \sigma_0 \, n_{\text{2D}}(x,y) = \ln\left(\frac{I_{\text{in}}(x,y)}{I_{\text{out}}(x,y)}\right). \tag{B.2}$$

In practice, the two images are captured using a CCD camera with a finite pixel size, and the optical density distribution is calculated on a pixel-wise basis

$$\text{OD}_{ij} = -\ln\left(\frac{A_{ij}}{R_{ij}}\right), \tag{B.3}$$

where A_{ij} and R_{ij} are the absorption and reference images received from the CCD camera, respectively. The two images will inevitably contain distinct features such as the spatial profile of the probe beam, and reflections and diffraction patterns due to elements in the optical system (for example, mirrors, waveplates, glass viewports, particles of dust, etc). Particularly obvious is often the appearance of interference fringes on the images, due to the high degree of coherence of the probe laser light. Ultimately, if all of these artifacts remain the same on both the absorption and reference images, their effect will be cancelled by the normalisation ratio in Eq. B.3. However, the position of the interference fringes are sensitive to temporal changes in path length on the order of the optical laser wavelength ($\sim 1\,\mu\text{m}$), and can easily move in the time between

taking the absorption and reference images due to vibrations of optical elements or the surrounding air (from cooling fans or switching of coils/shutters, for example). A shift of only a few pixels can be enough to cause the fringes to remain prominent on the final optical density image, an example of which is shown in Fig. B.1. This can therefore affect the measured density distribution of the cloud, thereby altering the inferred physical properties of the gas, such as atom number and temperature.

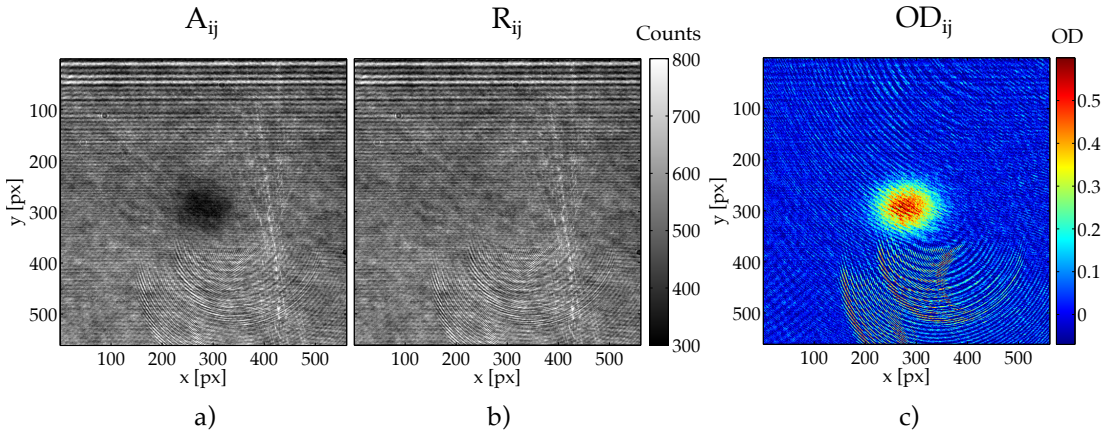

FIGURE B.1: Interference fringes in absorption imaging: The raw CCD images from both the absorption (a) and reference (b) images contain fringes and various artifacts from the beam passage through a series of optical elements. Due to the movement of these fringes between taking the two pictures, they do not cancel out fully, and remain present on the optical density distribution (c), contaminating the image.

One possible strategy for minimising the fringes is to try to reduce vibrations, such as by ensuring cooling fans are switched off, but this is difficult to achieve due to the required stability level in practice. It is also helpful to keep the time between taking the two sucessive pictures as short as possible, but for this to solve the problem it is usually necessary to reduce the time to below the frequency of the vibrations ($< 1\,\mathrm{ms}$), which requires more complex camera triggering and timing schemes (see, for example, [72]). It can be possible to remove fringes that have well-defined shapes by applying a notch filter in the Fourier transformed space of the image, and then applying the inverse Fourier transform afterwards [242]. However, when the size of the cloud becomes on the order of the spatial frequency of the fringes, this method will begin to affect the atomic density distribution, and result in a loss of information.

For the atom chip lab at Nottingham, I have implemented directly the fringe removal algorithm described by Ockeloen et al. from the University of Amsterdam [207, 292]. The method is acknowledged to have its origins at the University of Hamburg [293, 294] (which is in turn said to have been inspired by work in the group of W. Ertmer at the University of Hannover). A similar fringe-reduction algorithm was also published by the Chinese Academy of Sciences, Shanghai [295], which is based on the related so-called "*eigenface*" method of facial recognition [296, 297].

The concept of the Amsterdam method is as follows. For each experimental cycle, a single absorption image A_{ij} and a single reference image R_{ij} are acquired. Typically, the experiment is repeated many times, for example when scanning a parameter, resulting in many pairs of absorption and reference images, and interference fringes will appear on the resultant optical density distributions whenever the two images in a particular pair do not exactly match up. The idea is that for each absorption image, it could be that a different reference image in the set matches more closely - or even better in general, a weighted superposition of *all* images in the set. The algorithm therefore generates for each absorption image a brand new corresponding *optimal reference image*, Q_{ij}, which is best able to normalise out any interference fringes in an atom-free region, and can then be used instead of the original reference image to calculate the optical density. The nice thing about this method, as opposed to a Fourier filtering technique, is that the absorption images containing the atomic shadows are left completely untouched, and so the measured physical properties of the cloud cannot be affected, meaning that any features in the density distribution cannot be lost.

Algorithm

In order to work with the images, they are first unwrapped from $m \times n$ dimensional matrices into $p \times 1$ column vectors, which allows each image to be denoted by only one index x, where $x = \{1, 2, 3, \ ... \ p\}$ and $p = mn$ is the total number of pixels

$$A_{ij} = \begin{pmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{21} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{m1} & A_{m2} & \dots & A_{mn} \end{pmatrix} \Longrightarrow A_x = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{pmatrix}$$

$$R_{ij} = \begin{pmatrix} R_{11} & R_{12} & \dots & R_{1n} \\ R_{21} & R_{22} & \dots & R_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ R_{m1} & R_{m2} & \dots & R_{mn} \end{pmatrix} \Longrightarrow R_x = \begin{pmatrix} R_1 \\ R_2 \\ \vdots \\ R_p \end{pmatrix}$$

For example, a 480×640 pixels CCD picture becomes a $307,200 \times 1$ vector. Over a set of experimental cycles, there are many reference image vectors obtained, which can then be concatenated to form the matrix

$$R_{xk} = \begin{pmatrix} R_{11} & R_{12} & \dots & R_{1N} \\ R_{21} & R_{22} & \dots & R_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ R_{p1} & R_{p2} & \dots & R_{pN} \end{pmatrix}$$

where the index $k = \{1, 2, 3, \ ... \ N\}$ denotes the cycle number and N is the total number of reference images, e.g. 50. Each column in the matrix R_{xk} represents a different image

in the set.

The optimal reference image Q_x for a given absorption image A_x can be expressed as a weighted sum of all the individual reference images

$$Q_x = \sum_{k=1}^{N} c_k R_{xk}, \tag{B.4}$$

where c_k are the weighting coefficients. For example, Q_x could be composed of 10% of image $\#1$, 30% of image $\#2$, -5% of image $\#3$, etc. In order to determine these coefficients, an area of the image in which there are no atoms present should be matched as well as possible between the absorption and optimal reference images. Therefore the quantity δ can be used, which looks at the total difference over all pixels between the absorption image and the new reference image for a given set of weighting coefficients

$$\begin{aligned}\delta &= \sum_{x=1}^{p}\left(A_x - Q_x\right)^2 = \sum_{x=1}^{p}\left(A_x - \sum_{k=1}^{N} c_k R_{xk}\right)^2 \\ &= \left(A_1 - \left[c_1 R_{11} + c_2 R_{12} + ... \, c_N R_{N1}\right]\right)^2 + \left(A_2 - \left[c_1 R_{21} + c_2 R_{22} + ... \, c_N R_{2N}\right]\right)^2 \\ &\quad + ... + \left(A_p - \left[c_1 R_{p1} + c_2 R_{p2} + ... \, c_N R_{pN}\right]\right)^2 .\end{aligned} \tag{B.5}$$

The difference in the squares δ is minimised when the partial derivatives with respect to all of the coefficients are simultaneously equal to zero

$$\frac{\partial \delta}{\partial c_1} = \frac{\partial \delta}{\partial c_2} = \; ... \; = \frac{\partial \delta}{\partial c_N} = 0, \tag{B.6}$$

which leads to a system of N equations that must be solved. Calculating explicitly the partial derivatives and writing the equations out in full leads to

$$
\begin{aligned}
R_{11}A_1 + R_{21}A_2 + ... + R_{p1}A_p = \; & c_1\Big(R_{11}R_{11} + R_{21}R_{21} + ... + R_{p1}R_{p1}\Big) \\
& + c_2\Big(R_{11}R_{12} + R_{21}R_{22} + ... + R_{p1}R_{p2}\Big) \\
& + ... \\
& + c_N\Big(R_{11}R_{1N} + R_{21}R_{2N} + ... + R_{p1}R_{pN}\Big) \\
R_{12}A_1 + R_{22}A_2 + ... + R_{p2}A_p = \; & c_1\Big(R_{12}R_{11} + R_{22}R_{21} + ... + R_{p2}R_{p1}\Big) \\
& + c_2\Big(R_{12}R_{12} + R_{22}R_{22} + ... + R_{p2}R_{p2}\Big) \\
& + ... \\
& + c_N\Big(R_{12}R_{1N} + R_{22}R_{2N} + ... + R_{p2}R_{pN}\Big) \\
& \vdots \\
& \vdots \\
R_{1N}A_1 + R_{2N}A_2 + ... + R_{pN}A_p = \; & c_1\Big(R_{1N}R_{11} + R_{2N}R_{21} + ... + R_{pN}R_{p1}\Big) \\
& + c_2\Big(R_{1N}R_{12} + R_{2N}R_{22} + ... + R_{pN}R_{p2}\Big) \\
& + ... \\
& + c_N\Big(R_{1N}R_{1N} + R_{2N}R_{2N} + ... + R_{pN}R_{pN}\Big).
\end{aligned}
$$

These equations can be rewritten in terms of the sums over x and k, and expressed in the vector form

$$
\begin{pmatrix}
\sum_{x=1}^{p} R_{x1}A_x &=& \sum_{k=1}^{N} c_k\left(\sum_{x=1}^{p} R_{x1}R_{xk}\right) \\
\sum_{x=1}^{p} R_{x2}A_x &=& \sum_{k=1}^{N} c_k\left(\sum_{x=1}^{p} R_{x2}R_{xk}\right) \\
& \vdots & \\
\sum_{x=1}^{p} R_{xN}A_x &=& \sum_{k=1}^{N} c_k\left(\sum_{x=1}^{p} R_{xN}R_{xk}\right)
\end{pmatrix}.
$$

Even more compactly, the entire set of N equations can be written in matrix form

$$
\begin{pmatrix}
R_{11} & R_{21} & \dots & R_{p1} \\
R_{12} & R_{22} & \dots & R_{p2} \\
\vdots & \vdots & \ddots & \vdots \\
R_{1N} & R_{2N} & \dots & R_{pN}
\end{pmatrix}
\begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{pmatrix}
=
\begin{pmatrix}
R_{11} & R_{21} & \dots & R_{p1} \\
R_{12} & R_{22} & \dots & R_{p2} \\
\vdots & \vdots & \ddots & \vdots \\
R_{1N} & R_{2N} & \dots & R_{pN}
\end{pmatrix}
\begin{pmatrix}
R_{11} & R_{12} & \dots & R_{1N} \\
R_{21} & R_{22} & \dots & R_{2N} \\
\vdots & \vdots & \ddots & \vdots \\
R_{p1} & R_{p2} & \dots & R_{pN}
\end{pmatrix}
\begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_p \end{pmatrix},
$$

or

$$\mathbf{R}^{\mathrm{T}}\mathbf{A} = \mathbf{R}^{\mathrm{T}}\mathbf{R}\,\mathbf{C}. \tag{B.7}$$

This is a linear matrix equation of the form

$$\mathbf{Y} = \mathbf{B}\mathbf{X}, \tag{B.8}$$

and can be solved by lower-upper (LU) decomposition, which is useful when the same matrix $\mathbf{B}$ must be used many times for several vectors $\mathbf{Y}$. The method consists of first writing the matrix $\mathbf{B}$ as a product of two lower and upper *triangular* matrices, such that $\mathbf{B} = \mathbf{L}\mathbf{U}$, and therefore

$$\mathbf{Y} = \mathbf{L}\mathbf{U}\mathbf{X}. \tag{B.9}$$

Then let $\mathbf{Z} = \mathbf{U}\mathbf{X}$, and solve $\mathbf{Y} = \mathbf{L}\mathbf{Z}$ to find $\mathbf{Z}$. Once $\mathbf{Z}$ is found, then the equation $\mathbf{Z} = \mathbf{U}\mathbf{X}$ can be solved to finally find $\mathbf{X}$. The advantage of this method is that triangular matrices are used in all of the equations, which are then easier to solve (the cost however is that it is required to solve two of them instead of only one). The decomposition of $\mathbf{B}$ needs only to be performed *once*, and then the triangular equations must be solved for every absorption image $\mathbf{A}$ in the set to find the corresponding coefficients $\mathbf{C}$. Once the vector $\mathbf{C}$ has been found, the optimal light vector for the particular absorption picture is easily calculated through

$$\mathbf{Q} = \sum_{k=1}^{N} c_k R_{xk} = \mathbf{R}\mathbf{C}, \tag{B.10}$$

which can then simply be reshaped back into the dimensions of the original light image.

Matlab Function

A Matlab script for solving Eq. B.7 is given the PhD thesis of Caspar Ockeloen [292], written by S. Whitlock, and has been reproduced here for our reference. After reshaping all of the absorption and reference images as decribed above, the inbuilt Matlab function *[L,U] = lu(A)* is used to first decompose $\mathbf{R}^{\mathrm{T}}\mathbf{R}$ into lower and upper triangular matrices. Finally, the corresponding coefficients vectors $\mathbf{C}$ for each absorption image is found using the Matlab function *X = linsolve(A,B)*, and the optimal light images are returned by using Eq. B.10.

```
function optrefimages = fringeRemoval(atomsimages,lightimages,bgmask)

nimgsA  = length(atomsimages); % Number of "atoms" images
nimgsR = length(lightimages);  % Number of "light" images
xdim = size(atomsimages{1},2); % Number of pixels in x dir
ydim = size(atomsimages{1},1); % Number of pixels in y dir

%%% Reshape images from matrices to vectors %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
R = double(reshape(cat(3,lightimages{:}),xdim*ydim,nimgsR));
A = double(reshape(cat(3,atomsimages{:}),xdim*ydim,nimgsA));

%%% Masking %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
k = find(bgmask(:)==1);

%%% Lower-Upper Decomposition %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
[L,U,p] = lu(R(k,:)'*R(k,:),'vector');

%%% Solve For Coefficients and Create Optimal Images %%%%%%%%%%%%%%%%%%%%%%%%
lower.LT = true; upper.UT = true;
optrefimages=cell(nimgsA,1);
for j=1:nimgsA
    b=R(k,:)'*A(k,j);
    c = linsolve(U,linsolve(L,b(p,:),lower),upper);
    optrefimages{j}=reshape(R*c,[ydim xdim]);
end

end
```

Performance

For a typical set of 50 images recorded with a 1024×1024 pixel CCD camera, on a standard desktop PC the algorithm takes approximately 3 s to decompose the matrix $\mathbf{R}^{\mathrm{T}}\mathbf{R}$ into lower and upper triangular matrices. Once the decomposition has been performed, it takes approximately a further 20 s to calculate and save the 50 new optimal light images. These times are reduced to 0.2 s and 3.3 s, respectively, for a smaller 640×480 image size.

An important quantity to know is the number of reference images that are required in a given set in order to effectively remove any imaging artifacts. Figure B.2 (c) shows the variance of pixel values, Var(A_x/Q_x), for an example image in several regions of the corrected optical density distribution, indicated in Fig. B.2 (b). The larger the variance, the more interference fringes are present in the region. It can be seen that the largest gain in image quality is obtained when using at least ∼50 reference images from the set.

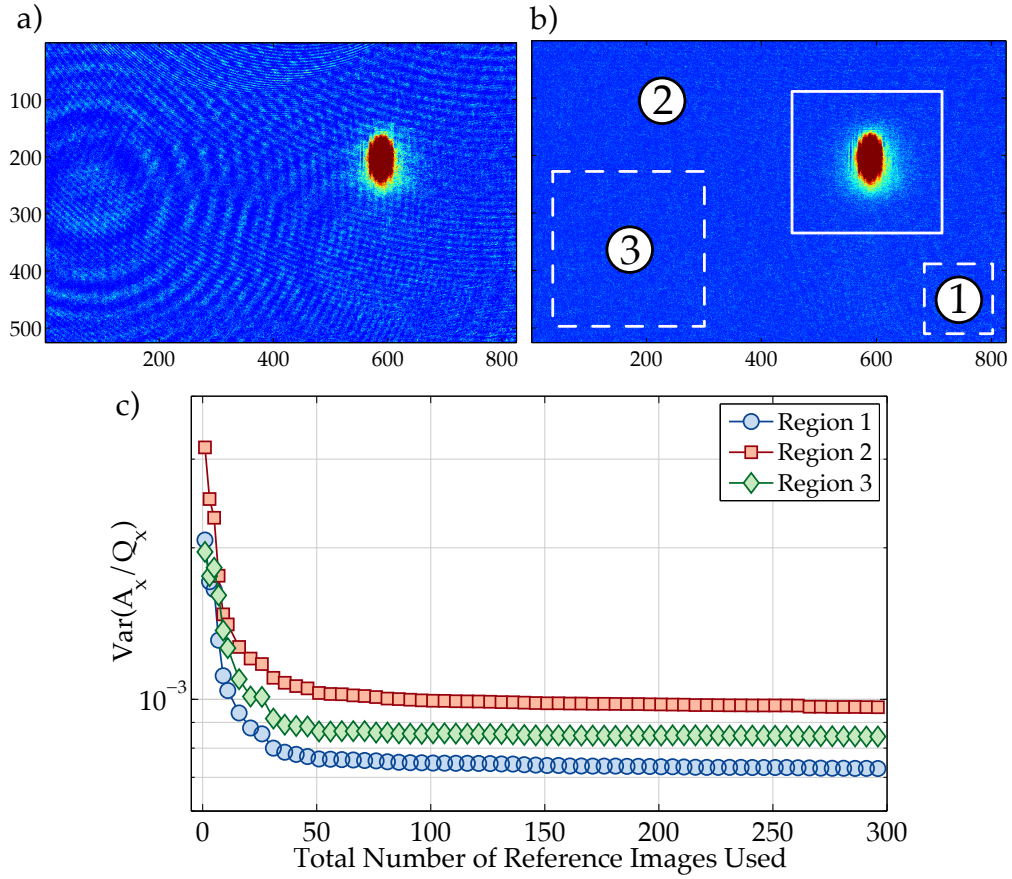

FIGURE B.2: Number of reference images requried to effectively remove fringes: The optical density of an amost pure BEC is shown both without (a) and with fringe correction (b), by using all 300 images in the set. The solid line region in (b) containing the atomic cloud was excluded from the matching algorithm. c) The variance of the regions indicated by the dashed regions in (b), plotted as a function of the number of images used in the algorithm. It can be seen that using ∼50 images already significantly improves the image quality, and that this number is not significantly dependent on the size of the region of interest.

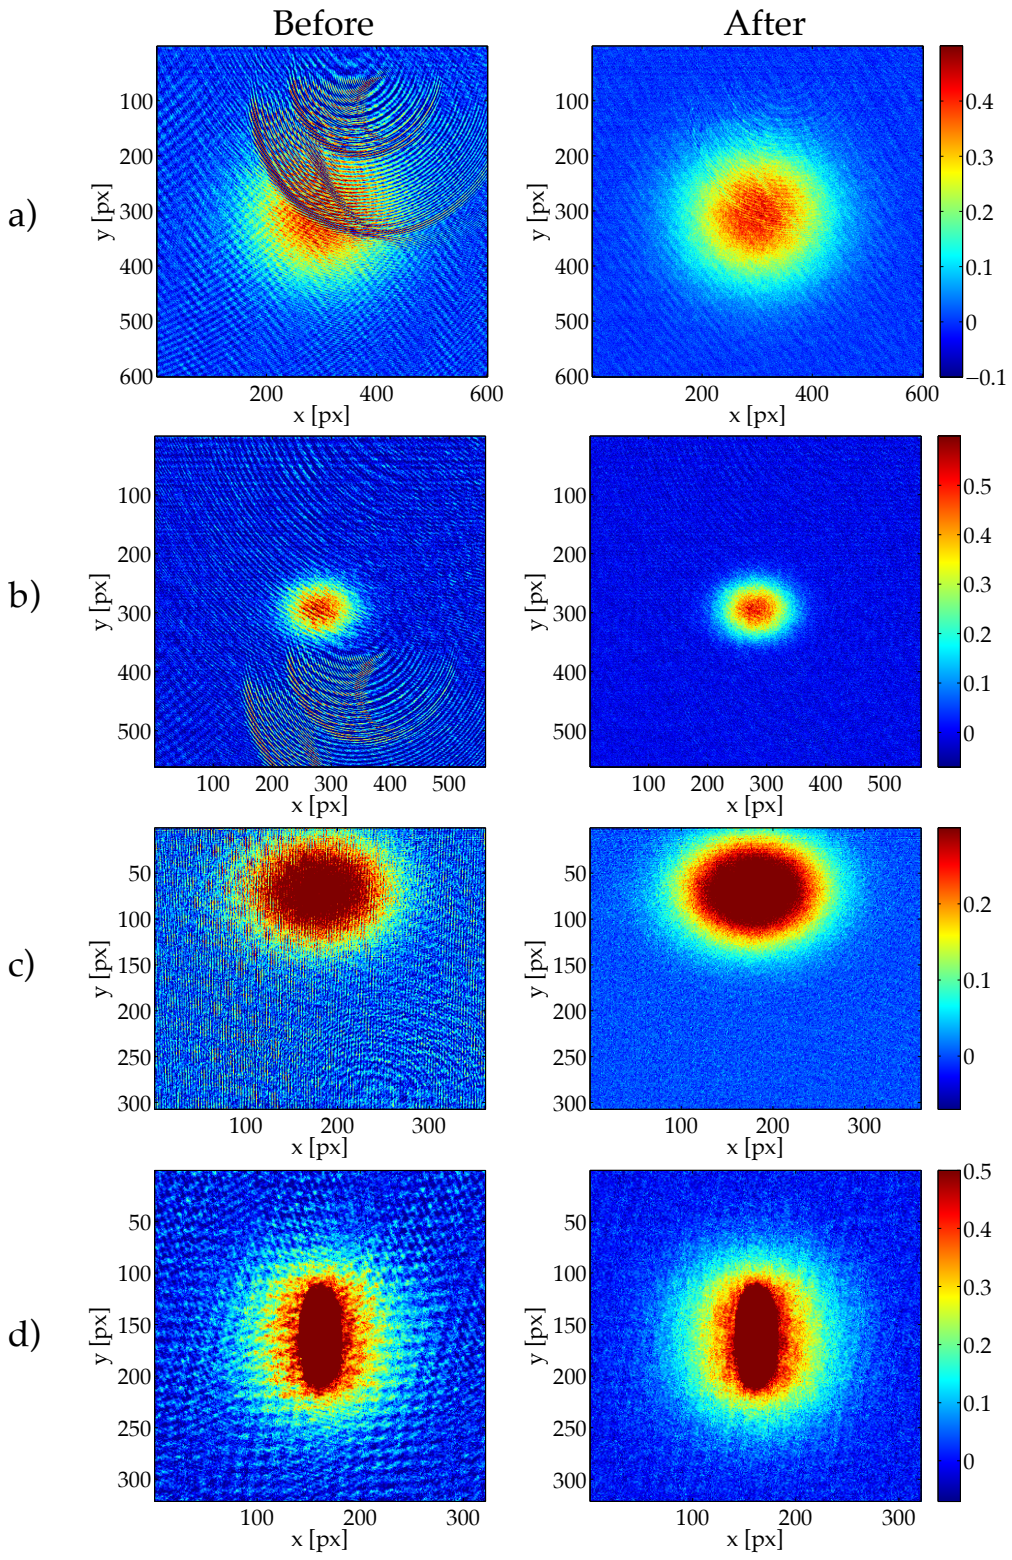

FIGURE B.3: Some examples of optical density images before (left column) and after (right column) applying the fringe removal algorithm. Images in (a), (b) and (c) depict thermal clouds at temperatures $10\,\mu K$, $1\,\mu K$ and $100\,\mu K$, respectively. The image in (d) is a BEC with $\sim 30\%$ condensate fraction.

Bibliography

[1] G. Cook and R. H. Dickerson. "Understanding the chemical potential". *Am. J. Phys.* **63** (1995), pp. 737–742.

[2] W. D. Phillips and H. Metcalf. "Laser Deceleration of an Atomic Beam". *Phys. Rev. Lett.* **48** (1982), pp. 596–599.

[3] S. Chu et al. "Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure". *Phys. Rev. Lett.* **55** (1985), pp. 48–51.

[4] E. L. Raab et al. "Trapping of Neutral Sodium Atoms with Radiation Pressure". *Phys. Rev. Lett.* **59** (1987), pp. 2631–2634.

[5] S. Chu. "Nobel Lecture: The manipulation of neutral particles". *Rev. Mod. Phys.* **70** (1998), pp. 685–706.

[6] C. N. Cohen-Tannoudji. "Nobel Lecture: Manipulating atoms with photons". *Rev. Mod. Phys.* **70** (1998), pp. 707–719.

[7] W. D. Phillips. "Nobel Lecture: Laser cooling and trapping of neutral atoms". *Rev. Mod. Phys.* **70** (1998), pp. 721–741.

[8] S. N. Bose. "Plancks Gesetz und Lichtquantenhypothese". *Z. phys* **26** (1924), p. 178.

[9] A. Einstein. "Quantentheorie des einatomigen idealen Gases". *Sitzungsberichte der Preussischen Akademie der Wissenschaften* **3** (1924), p. 18.

[10] K. B. Davis et al. "Evaporative Cooling of Sodium Atoms". *Phys. Rev. Lett.* **74** (1995), pp. 5202–5205.

[11] M. H. Anderson et al. "Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor". *Science* **269** (1995), pp. 198–201.

[12] K. B. Davis et al. "Bose-Einstein Condensation in a Gas of Sodium Atoms". *Phys. Rev. Lett.* **75** (1995), pp. 3969–3973.

[13] C. C. Bradley et al. "Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions". *Phys. Rev. Lett.* **75** (1995), pp. 1687–1690.

[14] W. Ketterle. "Nobel lecture: When atoms behave as waves: Bose-Einstein condensation and the atom laser". *Rev. Mod. Phys.* **74** (2002), pp. 1131–1151.

[15] E. A. Cornell and C. E. Wieman. "Nobel Lecture: Bose-Einstein condensation in a dilute gas, the first 70 years and some recent experiments". *Rev. Mod. Phys.* **74** (2002), pp. 875–893.

[16] M. Andrews et al. "Observation of interference between two Bose condensates". *Science* **275** (1997), pp. 637–641.

[17] M.-O. Mewes et al. "Output Coupler for Bose-Einstein Condensed Atoms". *Phys. Rev. Lett.* **78** (1997), pp. 582–585.

[18] I Bloch, T. W. Hansch, and T Esslinger. "Measurement of the spatial coherence of a trapped Bose gas at the phase transition". *Nature* **403** (2000), pp. 166–170.

[19] M. R. Matthews et al. "Vortices in a Bose-Einstein Condensate". *Phys. Rev. Lett.* **83** (1999), pp. 2498–2501.

[20] K. W. Madison et al. "Vortex Formation in a Stirred Bose-Einstein Condensate". *Phys. Rev. Lett.* **84** (2000), pp. 806–809.

[21] J. Abo-Shaeer et al. "Observation of vortex lattices in Bose-Einstein condensates". *Science* **292** (2001), pp. 476–479.

[22] S. Burger et al. "Dark Solitons in Bose-Einstein Condensates". *Phys. Rev. Lett.* **83** (1999), pp. 5198–5201.

[23] J Denschlag et al. "Generating solitons by phase engineering of a Bose-Einstein condensate". *Science* **287** (2000), pp. 97–101.

[24] D. S. Jin et al. "Collective Excitations of a Bose-Einstein Condensate in a Dilute Gas". *Phys. Rev. Lett.* **77** (1996), pp. 420–423.

[25] M.-O. Mewes et al. "Collective Excitations of a Bose-Einstein Condensate in a Magnetic Trap". *Phys. Rev. Lett.* **77** (1996), pp. 988–991.

[26] O. M. Maragò et al. "Observation of the Scissors Mode and Evidence for Superfluidity of a Trapped Bose-Einstein Condensed Gas". *Phys. Rev. Lett.* **84** (2000), pp. 2056–2059.

[27] J. E. Simsarian et al. "Magneto-Optic Trapping of ^{210}Fr". *Phys. Rev. Lett.* **76** (1996), pp. 3522–3525.

[28] M. W. Zwierlein et al. "Observation of Bose-Einstein Condensation of Molecules". *Phys. Rev. Lett.* **91** (2003), p. 250401.

[29] M. Greiner, C. A. Regal, and D. S. Jin. "Emergence of a molecular Bose–Einstein condensate from a Fermi gas". *Nature* **426** (2003), pp. 537–540.

[30] S. Jochim et al. "Bose-Einstein condensation of molecules". *Science* **302** (2003), pp. 2101–2103.

[31] J. Klaers et al. "Bose-Einstein condensation of photons in an optical microcavity". *Nature* **468** (2010), pp. 545–548.

[32] D. M. Stamper-Kurn et al. "Optical Confinement of a Bose-Einstein Condensate". *Phys. Rev. Lett.* **80** (1998), pp. 2027–2030.

[33] M. D. Barrett, J. A. Sauer, and M. S. Chapman. "All-Optical Formation of an Atomic Bose-Einstein Condensate". *Phys. Rev. Lett.* **87** (2001), p. 010404.

[34] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. "Optical Dipole Traps for Neutral Atoms". *Adv. At. Mol. Opt. Phy.* **42** (2000), pp. 95–170.

[35] S Inouye et al. "Observation of Feshbach resonances in a Bose–Einstein condensate". *Nature* **392** (1998), pp. 151–154.

[36] C. Chin et al. "Feshbach resonances in ultracold gases". *Rev. Mod. Phys.* **82** (2010), pp. 1225–1286.

[37] M. Bartenstein et al. "Collective Excitations of a Degenerate Gas at the BEC-BCS Crossover". *Phys. Rev. Lett.* **92** (2004), p. 203201.

[38] T. Bourdel et al. "Experimental Study of the BEC-BCS Crossover Region in Lithium 6". *Phys. Rev. Lett.* **93** (2004), p. 050401.

[39] M. W. Zwierlein et al. "Vortices and superfluidity in a strongly interacting Fermi gas". *Nature* **435** (2005), pp. 1047–1051.

[40] B. P. Anderson and M. A. Kasevich. "Macroscopic quantum interference from atomic tunnel arrays". *Science* **282** (1998), pp. 1686–1689.

[41] I. Bloch. "Ultracold quantum gases in optical lattices". *Nature Phys.* **1** (2005), pp. 23–30.

[42] D. Jaksch et al. "Cold Bosonic Atoms in Optical Lattices". *Phys. Rev. Lett.* **81** (1998), pp. 3108–3111.

[43] M. Greiner et al. "Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms". *Nature* **415** (2002), pp. 39–44.

[44] G. Roati et al. "Anderson localization of a non-interacting Bose–Einstein condensate". *Nature* **453** (2008), pp. 895–898.

[45] J. Billy et al. "Direct observation of Anderson localization of matter waves in a controlled disorder". *Nature* **453** (2008), pp. 891–894.

[46] J. Chabé et al. "Experimental Observation of the Anderson Metal-Insulator Transition with Atomic Matter Waves". *Phys. Rev. Lett.* **101** (2008), p. 255702.

[47] M. Albiez et al. "Direct Observation of Tunneling and Nonlinear Self-Trapping in a Single Bosonic Josephson Junction". *Phys. Rev. Lett.* **95** (2005), p. 010402.

[48] S Levy et al. "The a.c. and d.c. Josephson effects in a Bose-Einstein condensate". *Nature* **449** (2007), pp. 579–583.

[49] C.-L. Hung, V. Gurarie, and C. Chin. "From cosmology to cold atoms: observation of Sakharov oscillations in a quenched atomic superfluid". *Science* **341** (2013), pp. 1213–1215.

[50] L. J. Garay et al. "Sonic Analog of Gravitational Black Holes in Bose-Einstein Condensates". *Phys. Rev. Lett.* **85** (2000), pp. 4643–4647.

[51] R. Balbinot et al. "Nonlocal density correlations as a signature of Hawking radiation from acoustic black holes". *Phys. Rev. A* **78** (2008), p. 021603.

[52] O. Lahav et al. "Realization of a Sonic Black Hole Analog in a Bose-Einstein Condensate". *Phys. Rev. Lett.* **105** (2010), p. 240401.

[53] J.-C. Jaskula et al. "Acoustic Analog to the Dynamical Casimir Effect in a Bose-Einstein Condensate". *Phys. Rev. Lett.* **109** (2012), p. 220401.

[54] R. Folman et al. "Microscopic atom optics: From wires to an atom chip". *Adv. At. Mol. Opt. Phy.* **48** (2002), pp. 263–356.

[55] J. Reichel. "Microchip traps and Bose–Einstein condensation". *Appl. Phys. B* **74** (2002), pp. 469–487.

[56] J. Fortágh and C. Zimmermann. "Magnetic microtraps for ultracold atoms". *Rev. Mod. Phys.* **79** (2007), pp. 235–289.

[57] E. Hinds and I. Hughes. "Magnetic atom optics: mirrors, guides, traps, and chips for atoms". *J. Phys. D: Appl. Phys.* **32** (1999), R119.

[58] R. Folman. "Material science for quantum computing with atom chips". *Quantum Inf. Process.* **10** (2011), pp. 995–1036.

[59] J. Reichel and V. Vuletic. *Atom Chips*. John Wiley & Sons, 2011.

[60] M. Keil et al. "Fifteen years of cold matter on the atom chip: promise, realizations, and prospects". *J. Mod. Opt.* (2016), pp. 1–46.

[61] J. Reichel, W. Hänsel, and T. W. Hänsch. "Atomic Micromanipulation with Magnetic Surface Traps". *Phys. Rev. Lett.* **83** (1999), pp. 3398–3401.

[62] R. Folman et al. "Controlling Cold Atoms using Nanofabricated Surfaces: Atom Chips". *Phys. Rev. Lett.* **84** (2000), pp. 4749–4752.

[63] J. D. Weinstein and K. G. Libbrecht. "Microscopic magnetic traps for neutral atoms". *Phys. Rev. A* **52** (1995), pp. 4004–4009.

[64] J. Fortagh et al. "Miniaturized Wire Trap for Neutral Atoms". *Phys. Rev. Lett.* **81** (1998), pp. 5310–5313.

[65] J. Denschlag, D. Cassettari, and J. Schmiedmayer. "Guiding Neutral Atoms with a Wire". *Phys. Rev. Lett.* **82** (1999), pp. 2014–2017.

[66] A. Haase et al. "Trapping neutral atoms with a wire". *Phys. Rev. A* **64** (2001), p. 043405.

[67] W. Hänsel et al. "Bose–Einstein condensation on a microelectronic chip". *Nature* **413** (2001), pp. 498–501.

[68] H. Ott et al. "Bose-Einstein Condensation in a Surface Microtrap". *Phys. Rev. Lett.* **87** (2001), p. 230401.

[69] P. Hommelhoff et al. "Transporting, splitting and merging of atomic ensembles in a chip trap". *New J. Phys.* **7** (2005), p. 3.

[70] A. Günther et al. "Combined chips for atom optics". *Phys. Rev. A* **71** (2005), p. 063619.

[71] T. Schumm et al. "Matter-wave interferometry in a double well on an atom chip". *Nature Phys.* **1** (2005), pp. 57–62.

[72] D. M. Farkas, E. A. Salim, and J. Ramirez-Serrano. "Production of rubidium Bose-Einstein condensates at a 1 Hz rate". *arXiv preprint arXiv:1403.4641* (2014).

[73] J. Rudolph et al. "A high-flux BEC source for mobile atom interferometers". *New J. Phys.* **17** (2015), p. 065001.

[74] Y.-j. Lin et al. "Impact of the Casimir-Polder Potential and Johnson Noise on Bose-Einstein Condensate Stability Near Surfaces". *Phys. Rev. Lett.* **92** (2004), p. 050404.

[75] D. M. Harber et al. "Measurement of the Casimir-Polder force through center-of-mass oscillations of a Bose-Einstein condensate". *Phys. Rev. A* **72** (2005), p. 033610.

[76] J. M. Obrecht et al. "Measurement of the Temperature Dependence of the Casimir-Polder Force". *Phys. Rev. Lett.* **98** (2007), p. 063201.

[77] J. Fortágh et al. "Surface effects in magnetic microtraps". *Phys. Rev. A* **66** (2002), p. 041604.

[78] A. E. Leanhardt et al. "Propagation of Bose-Einstein Condensates in a Magnetic Waveguide". *Phys. Rev. Lett.* **89** (2002), p. 040401.

[79] A. E. Leanhardt et al. "Bose-Einstein Condensates near a Microfabricated Surface". *Phys. Rev. Lett.* **90** (2003), p. 100404.

[80] S. Kraft et al. "Anomalous longitudinal magnetic field near the surface of copper conductors". *J. Phys. B: At. Mol. Opt. Phys.* **35** (2002), p. L469.

[81] J. Estève et al. "Role of wire imperfections in micromagnetic traps for atoms". *Phys. Rev. A* **70** (2004), p. 043629.

[82] D.-W. Wang, M. D. Lukin, and E. Demler. "Disordered Bose-Einstein Condensates in Quasi-One-Dimensional Magnetic Microtraps". *Phys. Rev. Lett.* **92** (2004), p. 076802.

[83] P. Krüger et al. "Potential roughness near lithographically fabricated atom chips". *Phys. Rev. A* **76** (2007), p. 063621.

[84] S. Wildermuth et al. "Bose–Einstein condensates: microscopic magnetic-field imaging". *Nature* **435** (2005), pp. 440–440.

[85] S Wildermuth et al. "Sensing electric and magnetic fields with Bose-Einstein condensates". *Appl. Phys. Lett.* **88** (2006), p. 264103.

[86] S Aigner et al. "Long-range order in electronic transport through disordered metal films". *Science* **319** (2008), pp. 1226–1229.

[87] C. Henkel, S. Pötting, and M. Wilkens. "Loss and heating of particles in small and noisy traps". *Appl. Phys. B* **69** (1999), pp. 379–387.

[88] C. Henkel et al. "Fundamental limits for coherent manipulation on atom chips". *Appl. Phys. B* **76** (2003), pp. 173–182.

[89] M. P. A. Jones et al. "Spin Coupling between Cold Atoms and the Thermal Fluctuations of a Metal Surface". *Phys. Rev. Lett.* **91** (2003), p. 080401.

[90] T David et al. "Magnetic interactions of cold atoms with anisotropic conductors". *EPJ D* **48** (2008), pp. 321–332.

[91] A. Emmert et al. "Measurement of the trapping lifetime close to a cold metallic surface on a cryogenic atom-chip". *EPJ D* **51** (2009), pp. 173–177.

[92] J.-B. Trebbia et al. "Experimental Evidence for the Breakdown of a Hartree-Fock Approach in a Weakly Interacting Bose Gas". *Phys. Rev. Lett.* **97** (2006), p. 250403.

[93] J. Esteve et al. "Observations of Density Fluctuations in an Elongated Bose Gas: Ideal Gas and Quasicondensate Regimes". *Phys. Rev. Lett.* **96** (2006), p. 130403.

[94] A. H. van Amerongen et al. "Yang-Yang Thermodynamics on an Atom Chip". *Phys. Rev. Lett.* **100** (2008), p. 090402.

[95] Z. Hadzibabic et al. "Berezinskii–Kosterlitz–Thouless crossover in a trapped atomic gas". *Nature* **441** (2006), pp. 1118–1121.

[96] P. Krüger, Z. Hadzibabic, and J. Dalibard. "Critical Point of an Interacting Two-Dimensional Atomic Bose Gas". *Phys. Rev. Lett.* **99** (2007), p. 040402.

[97] V. Bagnato, D. E. Pritchard, and D. Kleppner. "Bose-Einstein condensation in an external potential". *Phys. Rev. A* **35** (1987), pp. 4354–4358.

[98] A. Görlitz et al. "Realization of Bose-Einstein Condensates in Lower Dimensions". *Phys. Rev. Lett.* **87** (2001), p. 130402.

[99] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. "Regimes of Quantum Degeneracy in Trapped 1D Gases". *Phys. Rev. Lett.* **85** (2000), pp. 3745–3749.

[100] M. A. Cazalilla et al. "One dimensional bosons: From condensed matter systems to ultracold gases". *Rev. Mod. Phys.* **83** (2011), pp. 1405–1466.

[101] B. Paredes et al. "Tonks–Girardeau gas of ultracold atoms in an optical lattice". *Nature* **429** (2004), pp. 277–281.

[102] T. Kinoshita, T. Wenger, and D. S. Weiss. "Observation of a one-dimensional Tonks-Girardeau gas". *Science* **305** (2004), pp. 1125–1128.

[103] B. L. Tolra et al. "Observation of Reduced Three-Body Recombination in a Correlated 1D Degenerate Bose Gas". *Phys. Rev. Lett.* **92** (2004), p. 190401.

[104] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. "Phase-Fluctuating 3D Bose-Einstein Condensates in Elongated Traps". *Phys. Rev. Lett.* **87** (2001), p. 050404.

[105] S. Dettmer et al. "Observation of Phase Fluctuations in Elongated Bose-Einstein Condensates". *Phys. Rev. Lett.* **87** (2001), p. 160406.

[106] S. Richard et al. "Momentum Spectroscopy of 1D Phase Fluctuations in Bose-Einstein Condensates". *Phys. Rev. Lett.* **91** (2003), p. 010405.

[107] D. Hellweg et al. "Phase fluctuations in Bose–Einstein condensates". *Appl. Phys. B* **73** (2001), pp. 781–789.

[108] H. Kreutzmann et al. "Characterization and control of phase fluctuations in elongated Bose–Einstein condensates". *Appl. Phys. B* **76** (2003), pp. 165–172.

[109] L. Sadler et al. "Spontaneous symmetry breaking in a quenched ferromagnetic spinor Bose–Einstein condensate". *Nature* **443** (2006), pp. 312–315.

[110] S. Ritter et al. "Observing the Formation of Long-Range Order during Bose-Einstein Condensation". *Phys. Rev. Lett.* **98** (2007), p. 090402.

[111] S Hofferberth et al. "Non-equilibrium coherence dynamics in one-dimensional Bose gases". *Nature* **449** (2007), pp. 324–327.

[112] A. Widera et al. "Quantum Spin Dynamics of Mode-Squeezed Luttinger Liquids in Two-Component Atomic Gases". *Phys. Rev. Lett.* **100** (2008), p. 140401.

[113] S. Trotzky et al. "Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional Bose gas". *Nature Phys.* **8** (2012), pp. 325–330.

[114] T. Kinoshita, T. Wenger, and D. S. Weiss. "A quantum Newton's cradle". *Nature* **440** (2006), pp. 900–903.

[115] T. Kitagawa et al. "The dynamics and prethermalization of one-dimensional quantum systems probed through the full distributions of quantum noise". *New J. Phys.* **13** (2011), p. 073018.

[116] M. Gring et al. "Relaxation and prethermalization in an isolated quantum system". *Science* **337** (2012), pp. 1318–1322.

[117] D. A. Smith et al. "Prethermalization revealed by the relaxation dynamics of full distribution functions". *New J. Phys.* **15** (2013), p. 075011.

[118] A. Piccardo-Selg. "Cold Atom Sources for Portable Quantum Sensors". PhD thesis. University of Nottingham, 2013.

[119] G. Aviv. "Dynamic Manipulations of Interacting 1d Bose Gases". PhD thesis. University of Nottingham, 2014.

[120] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics, Volume 1*. John Wiley & Sons, 1977.

[121] C. J. Foot. *Atomic Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2005.

[122] D. A. Steck. *Rubidium 87 D Line Data*. (2001). URL: http://steck.us/alkalidata (visited on 09/06/2016).

[123] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics, Volume 2*. John Wiley & Sons, 1978.

[124] G. Breit and I. I. Rabi. "Measurement of Nuclear Spin". *Phys. Rev.* **38** (1931), pp. 2082–2083.

[125] J. C. Maxwell. *A Treatise on Electricity and Magnetism*. 1st. Vol. II. Clarendon Press Series. Oxford, 1873.

[126] P. N. Lebedev. "Experimental Examination of Light Pressure". *Ann. der Physik* **6** (1901), p. 433.

[127] E. F. Nichols and G. F. Hull. "The Pressure Due to Radiation. (Second Paper.)" *Phys. Rev. (Series I)* **17** (1903), pp. 26–50.

[128] R. Frisch. "Molecule Recoil Associated with Optical Radiation". *Z. Physik* **86** (1933), p. 42.

[129] A. Ashkin. "Acceleration and Trapping of Particles by Radiation Pressure". *Phys. Rev. Lett.* **24** (1970), pp. 156–159.

[130] A. Ashkin. "Atomic-Beam Deflection by Resonance-Radiation Pressure". *Phys. Rev. Lett.* **25** (1970), pp. 1321–1324.

[131] H. J. Metcalf and P. Straten. *Laser Cooling and Trapping of Neutral Atoms*. Graduate Texts in Contemporary Physics. Springer, 1999.

[132] D. A. Steck. *Quantum and Atom Optics*. (2007). URL: http://steck.us/teaching (visited on 09/06/2016).

[133] T. Hänsch and A. Schawlow. "Cooling of gases by laser radiation". *Opt. Commun.* **13** (1975), pp. 68 –69.

[134] D. J. Wineland, R. E. Drullinger, and F. L. Walls. "Radiation-Pressure Cooling of Bound Resonant Absorbers". *Phys. Rev. Lett.* **40** (1978), pp. 1639–1642.

[135] W. Neuhauser et al. "Optical-Sideband Cooling of Visible Atom Cloud Confined in Parabolic Well". *Phys. Rev. Lett.* **41** (1978), pp. 233–236.

[136] C. Monroe et al. "Very cold trapped atoms in a vapor cell". *Phys. Rev. Lett.* **65** (1990), pp. 1571–1574.

[137] W. H. Wing. "On neutral particle trapping in quasistatic electromagnetic fields". *Prog. Quant. Electron.* **8** (1984), pp. 181 –199.

[138] S. Earnshaw. "On the nature of the molecular forces which regulate the constitution of the luminiferous ether". *Trans. Camb. Phil. Soc* **7** (1842), pp. 97–112.

[139] E. Majorana. "Atomi orientati in campo magnetico variabile". *Il Nuovo Cimento* **9** (1932), pp. 43–50.

[140] C. V. Sukumar and D. M. Brink. "Spin-flip transitions in a magnetic trap". *Phys. Rev. A* **56** (1997), pp. 2451–2454.

[141] T. Bergeman, G. Erez, and H. J. Metcalf. "Magnetostatic trapping fields for neutral atoms". *Phys. Rev. A* **35** (1987), pp. 1535–1546.

[142] A. L. Migdall et al. "First Observation of Magnetically Trapped Neutral Atoms". *Phys. Rev. Lett.* **54** (1985), pp. 2596–2599.

[143] A. Piccardo-Selg. "Degenerate Quantum Gases: Towards Bose-Einstein Condensation of Sodium". Diploma Thesis. University of Heidelberg, 2008.

[144] W. Petrich et al. "Stable, Tightly Confining Magnetic Trap for Evaporative Cooling of Neutral Atoms". *Phys. Rev. Lett.* **74** (1995), pp. 3352–3355.

[145] D. E. Pritchard. "Cooling Neutral Atoms in a Magnetic Trap for Precision Spectroscopy". *Phys. Rev. Lett.* **51** (1983), pp. 1336–1339.

[146] C. F. Matta, R. J. Boyd, and A. Becke. *The Quantum Theory of Atoms in Molecules: From Solid State to DNA and Drug Design*. John Wiley & Sons, 2007, p. 3.

[147] H. F. Hess. "Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen". *Phys. Rev. B* **34** (1986), pp. 3476–3479.

[148] T. Tommila. "Cooling of Spin-Polarized Hydrogen Atoms Trapped in Magnetic-Field Minima". *EPL (Europhysics Letters)* **2** (1986), p. 789.

[149] H. F. Hess et al. "Magnetic trapping of spin-polarized atomic hydrogen". *Phys. Rev. Lett.* **59** (1987), pp. 672–675.

[150] W. Ketterle and N. V. Druten. "Evaporative Cooling of Trapped Atoms". *Adv. At. Mol. Opt. Phy.* **37** (1996), pp. 181 –236.

[151] J. Walraven. "Atomic hydrogen in magnetostatic traps". *Quantum Dynamics of Simple Systems*. Ed. by G. L. Oppo et al. Vol. 44. SUSSP Proceedings. 1996.

[152] K. B. Davis, M. O. Mewes, and W. Ketterle. "An analytical model for evaporative cooling of atoms". *Appl. Phys. B* **60** (1995), pp. 155–159.

[153] 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** (1996), pp. 381–389.

[154] F. Dalfovo et al. "Theory of Bose-Einstein condensation in trapped gases". *Rev. Mod. Phys.* **71** (1999), pp. 463–512.

[155] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. International Series of Monographs on Physics. Clarendon Press, 2003.

[156] C. J. Pethick and H. Smith. *Bose–Einstein Condensation in Dilute Gases*. Cambridge University Press, 2008.

[157] R. Bowley and M. Sanchez. *Introductory Statistical Mechanics*. Clarendon Press Oxford, 1999.

[158] J. R. Ensher et al. "Bose-Einstein Condensation in a Dilute Gas: Measurement of Energy and Ground-State Occupation". *Phys. Rev. Lett.* **77** (1996), pp. 4984–4987.

[159] E. G. M. van Kempen et al. "Interisotope Determination of Ultracold Rubidium Interactions from Three High-Precision Experiments". *Phys. Rev. Lett.* **88** (2002), p. 093201.

[160] N. N. Bogolyubov. "On the theory of superfluidity". *J. Phys. (USSR)* **11** (1947), pp. 23–32.

[161] L. Vestergaard Hau et al. "Near-resonant spatial images of confined Bose-Einstein condensates in a 4-Dee magnetic bottle". *Phys. Rev. A* **58** (1998), R54–R57.

[162] K. Dieckmann et al. "Two-dimensional magneto-optical trap as a source of slow atoms". *Phys. Rev. A* **58** (1998), pp. 3891–3895.

[163] J. Schoser et al. "Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap". *Phys. Rev. A* **66** (2002), p. 023410.

[164] S. Chaudhuri, S. Roy, and C. S. Unnikrishnan. "Realization of an intense cold Rb atomic beam based on a two-dimensional magneto-optical trap: Experiments and comparison with simulations". *Phys. Rev. A* **74** (2006), p. 023406.

[165] A. Bonucci et al. "A Novel Approach in UHV Pumping of Accelerators: the NEXTorr Pump". *Proceedings of the 2nd International Particle Accelerator Conference*. (San Sebastian, Spain). 2011, pp. 1536–1538.

[166] W. Umrath. *Fundamentals of Vacuum Technology*. Vacuum Handbook. Oerlikon Leybold Vacuum, 2007.

[167] V. Dugrain, P. Rosenbusch, and J. Reichel. "Alkali vapor pressure modulation on the 100 ms scale in a single-cell vacuum system for cold atom experiments". *Rev. Sci. Instrum.* **85**, 083112 (2014).

[168] H. Busche. "Efficient Loading of a Magneto-Optical Trap for Experiments with Dense Ultracold Rydberg Gases". Diploma Thesis. University of Heidelberg, 2011.

[169] B. Höltkemeier. "2D MOT as a Source of a Cold Atom Target". Diploma Thesis. University of Heidelberg, 2011.

[170] J. Ferreras and M. Corrigan. "Design and Assembly of a Source of Atoms at a Few Hundred Microkelvin". Diploma Thesis. University of Nottingham, 2014.

[171] T. B. Swanson et al. "Efficient transfer in a double magneto-optical trap system". *J. Opt. Soc. Am. B* **15** (1998), pp. 2641–2645.

[172] R. Engel-Herbert and T. Hesjedal. "Calculation of the magnetic stray field of a uniaxial magnetic domain". *J. Appl. Phys.* **97**, 074504 (2005).

[173] E. Fleminger. "Magnetic Simulations for Trapping and Cooling of Neutral Atoms". Masters Thesis. Ben Gurion University of the Negev, 2006.

[174] M. Extavour. "Fermions and Bosons on an Atom Chip". PhD thesis. University of Toronto, 2009.

[175] P. S. Chauhan et al. "Copper wire bonding". *Copper Wire Bonding*. Springer, 2014, pp. 1–9.

[176] R. Zaouk, B. Y. Park, and M. J. Madou. "Introduction to microfabrication techniques". *Microfluidic Techniques: Reviews and Protocols* (2006), pp. 5–15.

[177] S Groth et al. "Atom chips: Fabrication and thermal properties". *Appl. Phys. Lett.* **85** (2004), pp. 2980–2982.

[178] S. Groth. "Development, Fabrication and Characterisation of Atom Chips". PhD thesis. University of Heidelberg, 2006.

[179] C. E. Wieman and L. Hollberg. "Using diode lasers for atomic physics". *Rev. Sci. Instrum.* **62** (1991), pp. 1–20.

[180] A. Truscott, N. Heckenberg, and H Rubinsztein-Dunlop. "Frequency stabilised grating feedback laser diode for atom cooling applications". *Opt. Quant. Electron.* **31** (1999), pp. 417–430.

[181] S Saito and Y Yamamoto. "Direct observation of Lorentzian lineshape of semiconductor laser and linewidth reduction with external grating feedback". *Electron. Lett.* **9** (1981), pp. 325–327.

[182] M Fleming and A. Mooradian. "Spectral Characteristics of External-Cavity Controlled Semiconductor Lasers". *IEEE J. Quant. Electron.* **17** (1981), pp. 44–59.

[183] G. Bjorklund et al. "Frequency modulation (FM) spectroscopy". *Appl. Phys. B* **32** (1983), pp. 145–152.

[184] D. W. Preston, C. E. Wieman, and K. M. Siegbahn. "Doppler-free saturated absorption spectroscopy: Laser spectroscopy". *Am. J. Phys* **64** (1996), pp. 1432–1436.

[185] J. H. Scofield. "Frequency-domain description of a lock-in amplifier". *Am. J. Phys.* **62** (1994), pp. 129–132.

[186] J. Appel, A. MacRae, and A. I. Lvovsky. "A versatile digital GHz phase lock for external cavity diode lasers". *Meas. Sci. Technol.* **20** (2009), p. 055302.

[187] U Schünemann et al. "Simple scheme for tunable frequency offset locking of two lasers". *Rev. Sci. Instrum.* **70** (1999), pp. 242–243.

[188] E. A. Donley et al. "Double-pass acousto-optic modulator system". *Rev. Sci. Instrum.* **76**, 063112 (2005).

[189] P. Horowitz and W. Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015.

[190] M. Jones. "Preparation of Cold Samples of Caesium-133 and Lithium-6 in a Single Apparatus for Bose-Fermi Mixture Experiments". PhD thesis. University of Nottingham, 2014.

[191] S. Stellmer et al. "Laser Cooling to Quantum Degeneracy". *Phys. Rev. Lett.* **110** (2013), p. 263003.

[192] C. S. Adams et al. "Evaporative Cooling in a Crossed Dipole Trap". *Phys. Rev. Lett.* **74** (1995), pp. 3577–3580.

[193] T. Weber et al. "Bose-Einstein condensation of cesium". *Science* **299** (2003), pp. 232–235.

[194] N. Masuhara et al. "Evaporative Cooling of Spin-Polarized Atomic Hydrogen". *Phys. Rev. Lett.* **61** (1988), pp. 935–938.

[195] C.-L. Hung et al. "Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps". *Phys. Rev. A* **78** (2008), p. 011604.

[196] J.-F. Clément et al. "All-optical runaway evaporation to Bose-Einstein condensation". *Phys. Rev. A* **79** (2009), p. 061406.

[197] D. M. Farkas et al. "Efficient direct evaporative cooling in an atom-chip magnetic trap". *Phys. Rev. A* **87** (2013), p. 053417.

[198] D. A. Smith et al. "Absorption imaging of ultracold atoms on atom chips". *Opt. Express* **19** (2011), pp. 8471–8485.

[199] A. Putra et al. "Optimally focused cold atom systems obtained using density-density correlations". *Rev. Sci. Instrum.* **85** (2014), p. 013110.

[200] H. K. Anderson. "Bose-Einstein condensates in optical lattices". PhD thesis. University of Aarhus, Denmark, 2008.

[201] W. Smith. *Modern Optical Engineering*. McGraw-Hill Education, 2007.

[202] M. Andrews et al. "Direct, nondestructive observation of a Bose condensate". *Science* **273** (1996), p. 84.

[203] A. Paris-Mandoki. "A Single Apparatus for the Production of Ultracold Fermionic Lithium and Cold Bosonic Caesium Gases". PhD thesis. University of Nottingham, 2015.

[204] J. Chabé et al. "Coherent and incoherent multiple scattering". *Phys. Rev. A* **89** (2014), p. 043833.

[205] G Reinaudi et al. "Strong saturation absorption imaging of dense clouds of ultracold atoms". *Opt. Lett.* **32** (2007), pp. 3143–3145.

[206] W. Muessel et al. "Optimized absorption imaging of mesoscopic atomic clouds". *Appl. Phys. B* **113** (2013), pp. 69–73.

[207] C. F. Ockeloen et al. "Detection of small atom numbers through image processing". *Phys. Rev. A* **82** (2010), p. 061606.

[208] J. Walraven. "Thermodynamic and Collisional Properties of Trapped Atomic Gases". *Lecture Notes of the Ecole de Physique des Houches* (2008).

[209] W. Petrich et al. "Behavior of atoms in a compressed magneto-optical trap". *J. Opt. Soc. Am. B* **11** (1994), pp. 1332–1335.

[210] L. You and M. Holland. "Ballistic expansion of trapped thermal atoms". *Phys. Rev. A* **53** (1996), R1–R4.

[211] D. Stamper-Kurn. "Peeking and Poking at a New Quantum Fluid: Studies of gaseous Bose-Einstein condensates in magnetic and optical traps". PhD thesis. Massachusetts Institute of Technology, 2000.

[212] S. Wildermuth et al. "Optimized magneto-optical trap for experiments with ultracold atoms near surfaces". *Phys. Rev. A* **69** (2004), p. 030901.

[213] A. Steane, M Chowdhury, and C. Foot. "Radiation force in the magneto-optical trap". *J. Opt. Soc. Am. B* **9** (1992), pp. 2142–2158.

[214] J. Weiner et al. "Experiments and theory in cold and ultracold collisions". *Rev. Mod. Phys.* **71** (1999), pp. 1–85.

[215] M. Haw et al. "Magneto-optical trap loading rate dependence on trap depth and vapor density". *J. Opt. Soc. Am. B* **29** (2012), pp. 475–483.

[216] T. Arpornthip, C. A. Sackett, and K. J. Hughes. "Vacuum-pressure measurement using a magneto-optical trap". *Phys. Rev. A* **85** (2012), p. 033420.

[217] D. W. Sesko, T. G. Walker, and C. E. Wieman. "Behavior of neutral atoms in a spontaneous force trap". *J. Opt. Soc. Am. B* **8** (1991), pp. 946–958.

[218] K. R. Overstreet et al. "Multiple scattering and the density distribution of a Cs MOT". *Opt. Express* **13** (2005), pp. 9672–9682.

[219] J. Ramirez-Serrano et al. "Multistage two-dimensional magneto-optical trap as a compact cold atom beam source". *Opt. Lett.* **31** (2006), pp. 682–684.

[220] T. Müller et al. "Versatile compact atomic source for high-resolution dual atom interferometry". *Phys. Rev. A* **76** (2007), p. 063611.

[221] P. D. Lett et al. "Observation of Atoms Laser Cooled below the Doppler Limit". *Phys. Rev. Lett.* **61** (1988), pp. 169–172.

[222] J. Dalibard and C. Cohen-Tannoudji. "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models". *J. Opt. Soc. Am. B* **6** (1989), pp. 2023–2045.

[223] P. J. Ungar et al. "Optical molasses and multilevel atoms: theory". *J. Opt. Soc. Am. B* **6** (1989), pp. 2058–2071.

[224] C. Adams and E Riis. "Laser cooling and trapping of neutral atoms". *Prog. Quant. Electron.* **21** (1997), pp. 1–79.

[225] C. Salomon et al. "Laser cooling of cesium atoms below 3 μK". *EPL (Europhysics Letters)* **12** (1990), p. 683.

[226] A. Steane and C. Foot. "Laser cooling below the Doppler limit in a magneto-optical trap". *EPL (Europhysics Letters)* **14** (1991), p. 231.

[227] M Drewsen et al. "Investigation of sub-Doppler cooling effects in a cesium magneto-optical trap". *Appl. Phys. B* **59** (1994), pp. 283–298.

[228] W. Franzen and A. G. Emslie. "Atomic Orientation by Optical Pumping". *Phys. Rev.* **108** (1957), pp. 1453–1458.

[229] R. Benumof. "Optical pumping theory and experiments". *Am. J. Phys.* **33** (1965), pp. 151–160.

[230] W. Ketterle, D. S. Durfee, and D. M. Stamper-kurn. "Making, probing and understanding Bose-Einstein condensates". *In Proceedings of the 1998 Enrico Fermi school on Bose-Einstein condensation in Varenna, Italy.* (1999), p. 67.

[231] S. Machluf. "Coherent Splitting of Matter-Waves on an Atom Chip Using a State-Dependent Magnetic Potential". PhD thesis. Ben Gurion University of the Negev, 2013.

[232] M. Gring. "Prethermalization in an Isolated Many Body System". PhD thesis. Technische Universität Wien, 2012.

[233] B. Yuen. "Production and Oscillations of a Bose-Einstein Condensate on an Atom Chip". PhD thesis. Imperial College London, 2013.

[234] C. A. Sackett, C. C. Bradley, and R. G. Hulet. "Optimization of evaporative cooling". *Phys. Rev. A* **55** (1997), pp. 3797–3801.

[235] M. Yamashita et al. "Optimization of evaporative cooling towards a large number of Bose-Einstein-condensed atoms". *Phys. Rev. A* **67** (2003), p. 023601.

[236] J. E. Lye et al. "Images of evaporative cooling to Bose-Einstein condensation". *J. Opt. B: Quantum Semiclassical Opt.* **4** (2002), p. 57.

[237] S. Wildermuth. "One-dimensional Bose-Einstein condensates in micro-traps". PhD thesis. University of Heidelberg, 2005.

[238] S. Schneider et al. "Bose-Einstein condensation in a simple microtrap". *Phys. Rev. A* **67** (2003), p. 023612.

[239] J Söding et al. "Three-body decay of a rubidium Bose–Einstein condensate". *Appl. Phys. B* **69** (1999), pp. 257–261.

[240] E. A. Burt et al. "Coherence, Correlations, and Collisions: What One Learns about Bose-Einstein Condensates from Their Decay". *Phys. Rev. Lett.* **79** (1997), pp. 337–340.

[241] D. Jervis and J. H. Thywissen. "Making an Ultracold Gas". *Quantum Gas Experiments: Exploring Many-Body States*. Ed. by P. S. Törmä Klaus. World Scientific Publishing Co., 2015. Chap. 2, pp. 5–32.

[242] J. Szczepkowski et al. "Analysis and calibration of absorptive images of Bose–Einstein condensate at nonzero temperatures". *Rev. Sci. Instrum.* **80**, 053103 (2009).

[243] S. Grossmann and M. Holthaus. "On Bose-Einstein condensation in harmonic traps". *Phys. Lett. A* **208** (1995), pp. 188–192.

[244] S. Giorgini, L. P. Pitaevskii, and S. Stringari. "Condensate fraction and critical temperature of a trapped interacting Bose gas". *Phys. Rev. A* **54** (1996), R4633–R4636.

[245] J. P. Burke et al. "Prospects for Mixed-Isotope Bose-Einstein Condensates in Rubidium". *Phys. Rev. Lett.* **80** (1998), pp. 2097–2100.

[246] M. Naraschewski and D. M. Stamper-Kurn. "Analytical description of a trapped semi-ideal Bose gas at finite temperature". *Phys. Rev. A* **58** (1998), pp. 2423–2426.

[247] B. P. Anderson and M. A. Kasevich. "Spatial observation of Bose-Einstein condensation of $^{87}\mathrm{Rb}$ in a confining potential". *Phys. Rev. A* **59** (1999), R938–R941.

[248] F. Gerbier et al. "Experimental study of the thermodynamics of an interacting trapped Bose-Einstein condensed gas". *Phys. Rev. A* **70** (2004), p. 013607.

[249] Y. Castin and R. Dum. "Bose-Einstein Condensates in Time Dependent Traps". *Phys. Rev. Lett.* **77** (1996), pp. 5315–5319.

[250] M. E. Gehm et al. "Dynamics of noise-induced heating in atom traps". *Phys. Rev. A* **58** (1998), pp. 3914–3921.

[251] C. Myatt. "Bose-Einstein Condensation Experiments in a Dilute Vapor of Rubidium". PhD thesis. University of Colorado, 1997.

[252] H. C. W. Beijerinck. "Heating rates in collisionally opaque alkali-metal atom traps: Role of secondary collisions". *Phys. Rev. A* **62** (2000), p. 063614.

[253] J. Rushton, M. Aldous, and M. Himsworth. "Contributed review: the feasibility of a fully miniaturized magneto-optical trap for portable ultracold quantum technology". *Rev. Sci. Instrum.* **85** (2014), p. 121501.

[254] M. de Angelis et al. "iSense: a portable ultracold-atom-based gravimeter". *Procedia Comput. Sci.* **7** (2011), pp. 334–336.

[255] E. A. Salim et al. "Compact, microchip-based systems for practical applications of ultracold atoms". *Quantum Inf. Process.* **10** (2011), pp. 975–994.

[256] T Könemann et al. "A freely falling magneto-optical trap drop tower experiment". *Appl. Phys. B* **89** (2007), pp. 431–438.

[257] G. Stern et al. "Light-pulse atom interferometry in microgravity". *EPJ D* **53** (2009), pp. 353–357.

[258] R. Geiger et al. "Detecting inertial effects with airborne matter-wave interferometry". *Nat. Commun.* **2** (2011), p. 474.

[259] W. Rohringer et al. "Non-equilibrium scale invariance and shortcuts to adiabaticity in a one-dimensional Bose gas". *Sci. Rep.* **5** (2015), p. 9820.

[260] A. Polkovnikov et al. "Colloquium: Nonequilibrium dynamics of closed interacting quantum systems". *Rev. Mod. Phys.* **83** (2011), pp. 863–883.

[261] T. Langen, R. Geiger, and J. Schmiedmayer. "Ultracold Atoms Out of Equilibrium". *Annu. Rev. Condens. Matter Phys.* **6** (2015), pp. 201–217.

[262] M. Vengalattore et al. "High-Resolution Magnetometry with a Spinor Bose-Einstein Condensate". *Phys. Rev. Lett.* **98** (2007), p. 200801.

[263] M. L. Terraciano, M. Bashkansky, and F. K. Fatemi. "A single-shot imaging magnetometer using cold atoms". *Opt. Express* **16** (2008), pp. 13062–13069.

[264] A. Peters, K. Y. Chung, and S. Chu. "Measurement of gravitational acceleration by dropping atoms". *Nature* **400** (1999), pp. 849–852.

[265] J. B. Fixler et al. "Atom Interferometer Measurement of the Newtonian Constant of Gravity". *Science* **315** (2007), pp. 74–77.

[266] B. Canuel et al. "Six-Axis Inertial Sensor Using Cold-Atom Interferometry". *Phys. Rev. Lett.* **97** (2006), p. 010402.

[267] T. Müller et al. "A compact dual atom interferometer gyroscope based on laser-cooled rubidium". *EPJ D* **53** (2009), pp. 273–281.

[268] D. M. Farkas, A. Zozulya, and D. Z. Anderson. "A compact microchip atomic clock based on all-optical interrogation of ultra-cold trapped Rb atoms". *Appl. Phys. B* **101** (2010), pp. 705–721.

[269] K. I. Lee et al. "Single-beam atom trap in a pyramidal and conical hollow mirror". *Opt. Lett.* **21** (1996), pp. 1177–1179.

[270] J. Arlt et al. "A pyramidal magneto-optical trap as a source of slow atoms". *Opt. Commun.* **157** (1998), pp. 303 –309.

[271] M. Vangeleyn et al. "Single-laser, one beam, tetrahedral magneto-optical trap". *Opt. Express* **17** (2009), pp. 13601–13608.

[272] M. Trupke et al. "Pyramidal micromirrors for microsystems and atom chips". *Appl. Phys. Lett.* **88** (2006), p. 071116.

[273] S. Pollock et al. "Integrated magneto-optical traps on a chip using silicon pyramid structures". *Opt. Express* **17** (2009), pp. 14109–14114.

[274] M. Vangeleyn et al. "Laser cooling with a single laser beam and a planar diffractor". *Opt. Lett.* **35** (2010), pp. 3453–3455.

[275] C. Nshii et al. "A surface-patterned chip as a strong source of ultracold atoms for quantum technologies". *Nat. Nanotechnol.* **8** (2013), pp. 321–324.

[276] J. J. P. van Es et al. "Longitudinal character of atom-chip-based rf-dressed potentials". *Phys. Rev. A* **77** (2008), p. 063623.

[277] S. Machluf et al. "Coupling between Internal Spin Dynamics and External Degrees of Freedom in the Presence of Colored Noise". *Phys. Rev. Lett.* **105** (2010), p. 203002.

[278] S van Frank et al. "Optimal control of complex atomic quantum systems". *Sci. Rep.* **6** (2016).

[279] J. Malcolm. "Construction of a Portable Platform for Cold Atom Interferometry". PhD thesis. University of Birmingham, 2015.

[280] I. Bloch, J. Dalibard, and W. Zwerger. "Many-body physics with ultracold gases". *Rev. Mod. Phys.* **80** (2008), pp. 885–964.

[281] J. Eisert, M. Friesdorf, and C. Gogolin. "Quantum many-body systems out of equilibrium". *Nature Phys.* **11** (2015), pp. 124–130.

[282] K. L. Lee and N. P. Proukakis. "Non-equilibrium atomic condensates and mixtures: collective modes, condensate growth and thermalisation". *J. Phys. B: At. Mol. Opt. Phys.* **49** (2016), p. 214003.

[283] J. F. Sherson et al. "Single-atom-resolved fluorescence imaging of an atomic Mott insulator". *Nature* **467** (2010), pp. 68–72.

[284] W. S. Bakr et al. "Probing the superfluid–to–mott insulator transition at the single-atom level". *Science* **329** (2010), pp. 547–550.

[285] M. Wilzbach et al. "Simple integrated single-atom detector". *Opt. Lett.* **34** (2009), pp. 259–261.

[286] R Bücker et al. "Single-particle-sensitive imaging of freely propagating ultracold atoms". *New J. Phys.* **11** (2009), p. 103039.

[287] R. Bücker et al. "Twin-atom beams". *Nature Phys.* **7** (2011), pp. 608–611.

[288] M. Greiner et al. "Collapse and revival of the matter wave field of a Bose–Einstein condensate". *Nature* **419** (2002), pp. 51–54.

[289] V. Dunjko, V. Lorent, and M. Olshanii. "Bosons in Cigar-Shaped Traps: Thomas-Fermi Regime, Tonks-Girardeau Regime, and In Between". *Phys. Rev. Lett.* **86** (2001), pp. 5413–5416.

[290] J. Reichel and J.H. Thywissen. "Using magnetic chip traps to study Tonks-Girardeau quantum gases". *J. Phys. IV France* **116** (2004), pp. 265–274.

[291] J. Van Es et al. "Box traps on an atom chip for one-dimensional quantum gases". *J. Phys. B: At. Mol. Opt. Phys.* **43** (2010), p. 155002.

[292] C. Ockeloen. "Probing fluctuations in a lattice of mesoscopic atomic ensembles". Master's Thesis. University of Amsterdam, 2010.

[293] M. Erhard. "Experimente mit mehrkomponentigen Bose-Einstein-Kondensaten". PhD thesis. University of Hamburg, 2004.

[294] J. Kronjäger. "Coherent Dynamics of Spinor Bose-Einstein Condensates". PhD thesis. University of Hamburg, 2007.

[295] X. Li et al. "Reduction of interference fringes in absorption imaging of cold atom cloud using eigenface method". *Chin. Opt. Lett.* **5** (2007), pp. 128–130.

[296] L. Sirovich and M. Kirby. "Low-dimensional procedure for the characterization of human faces". *J. Opt. Soc. Am. A* **4** (1987), pp. 519–524.

[297] M. A. Turk and A. P. Pentland. "Face recognition using eigenfaces". *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR '91., IEEE Computer Society Conference on*. 1991, pp. 586–591.