Imaging Fermionic Atoms in a Quantum Gas Microscope

by

Graham Edge

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Physics
University of Toronto

Abstract

Imaging Fermionic Atoms in a Quantum Gas Microscope

Graham Edge
Doctor of Philosophy
Graduate Department of Physics
University of Toronto
2017

This thesis describes the methods used to achieve single-atom and single-site resolved fluorescence imaging of a quantum gas of fermionic atoms held in an optical lattice. This kind of quantum gas microscopy has been previously applied to bosonic atoms in two experiments, and has allowed for extensive studies of many-body quantum systems. Here we detail methods to extend the techniques of quantum gas microscopy to fermionic species of atoms, opening the door to quantum simulation of electronic models of condensed matter physics.

We prepare quantum degenerate gases of fermionic ^{40}K through sympathetic cooling with a gas of ^{87}Rb . The gas prepared is trapped less than $1\,\mathrm{mm}$ away from a thin vacuum window, an experimental geometry which permits interrogation with a large Numerical Aperture microscope objective. The cloud is loaded into a $527\,\mathrm{nm}$-period cubic optical lattice, ideal for quantum simulation of many-body models of fermionic particles such as the Hubbard model. The optical lattice is designed to allow for potential depths up to $V/k_B = 250\,\mu\mathrm{K}$, sufficient to strongly confine ^{40}K atoms into individual lattice sites.

EIT cooling — a technique previously demonstrated in ion-trapping experiments — is used to cool the fermionic atoms deeply towards the ground motional state of the lattice wells. We collect the photons scattered during continuous application of EIT cooling, while atoms are kept pinned in the lattice for many seconds. We strongly suppress all other sources of light with the use of a modulated cooling scheme and a spinning chopper wheel. Collecting the scattered light produced during $2\,\mathrm{s}$ of EIT cooling allows us to resolve single atoms in a fluorescence image. We demonstrate that the high resolution of our imaging apparatus also permits atom positions to be precisely identified on the length scale of the optical lattice. Thus we demonstrate the realisation of a quantum gas microscope for fermionic atoms, allowing for precise investigation of the ground states and dynamics of correlated quantum states of fermionic particles.

Preface

The results described in this thesis use an experimental apparatus which was already partially constructed at the time I joined the team. Dave McKay and Dylan Jervis had together already constructed the vacuum system, control electronics, MOT laser system and optics, and the magnetic transport system. Additionally Karl Pilch had designed a current stabilization system for the our Feshbach coils. I was involved in one final bake-out of the vacuum system in 2012 when it became necessary to replace the atomic sources for ^{40}K , and I constructed the ^{40}K sources which are currently in use in both experiments in the Thywissen group. Over my first years in the group I assisted Dylan and Dave in producing the first degenerate clouds of ^{87}Rb and ^{40}K in this apparatus.

The incorporation of all remaining aspects of the experimental apparatus was performed with the help of several Master's students and visitors. Dan Fine worked on many early aspects of the experiment, including the laser system operating at 405 nm that we attempted to use initially for fluorescence imaging. Carolyn Keirans and Thomas Maier performed microscope objective testing and assessed the stability of the apparatus we currently use to mount the objective. Ryan Day helped with our earliest attempts to perform laser cooling of ^{40}K using the D1 transition. Will Cairncross, Simon Heun, and Ian Kivlichan all worked on the construction of lasers for use in this experiment. Daniel Nino set up a laser system for the $4S_{1/2} \rightarrow 5P_{1/2}$ transition of ^{40}K , which has not yet been put to use in the experiment. Christian Veit built the broadband laser system currently being installed in the experiment to form a dimple trap.

Stefan Trotzky has been an unstoppable force in our lab, present for almost all of research presented here. His oversight and leadership have kept the lab moving constantly forward, and he has made large individual contributions to the development of lattice beam stabilization electronics, a stabilization system for ambient magnetic field drifts, and image analysis code used to examine our fluorescence images - all of which are crucial to the results of this thesis.

Rhys Anderson joined the lab near the end of my Ph.D, and was instrumental to our final successful efforts to achieve single-atom sensitive imaging. We worked together to implement plane selection of the ^{40}K atoms in our optical lattice, optimize EIT cooling for fluorescence imaging, and track down and destroy sources of background light in the fluorescence imaging path. He has now taken charge of the experiment working alongside new doctoral students Peihang Xu and Vijin Venu.

Above all, Joseph Thywissen has guided this project from its infancy to the current state of functional quantum simulator. His oversight has kept effort in the lab focussed, and several crucial suggestions - such as the decision to explore laser cooling on the D1 transition - have helped guide this experiment to success.

The main publication resulting from the work described in this thesis is:

- G.J.A. Edge, R. Anderson, D. Jervis, D.C. McKay, R. Day, S. Trotzky, and J.H. Thywissen (2015). Imaging and addressing of individual fermionic atoms in an optical lattice. *Physical Review A.* **92** 063406

Additionally, my early contributions to this experiment as well as a second experiment in our group resulted in the following publications which are not discussed in this thesis:

- A.B. Bardon, S. Beattie, C. Luciuk, W. Cairncross, D. Fine, N.S. Cheng, G.J.A. Edge, E. Taylor, S. Zhang, S. Trotzky, and J.H. Thywissen (2014). Transverse Demagnetization Dynamics of a Unitary Fermi Gas. *Science* **344** 722

- D.C. McKay, D. Jervis, D.J. Fine, J.W. Simpson-Porco, G.J.A. Edge, and J.H. Thywissen (2011). Low-temperature high-density magneto-optical trapping of potassium using the open 4S→5P transition at 405 nm. *Physical Review A.* **84** 063420

To my wife Caoimhe

Contents

Chapter 1

Introduction

"... nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy."

— Richard Feynman

This thesis describes the construction of an experiment which uses an ultracold quantum gas of fermionic atoms to simulate the behaviour of electrons moving throughout the crystal lattice inside a material. The same basic rules which govern the movement of electrons inside a material can lead to strikingly different behaviours, leading to different classes of common materials like metals, insulators, and semiconductors which are enormously important technologically. Particular electron behaviour can also give rise to exotic materials with very interesting properties such as such as high-temperature superconductors, graphene, and topological insulators. For any particular material structure, it can be difficult to predict whether any interesting behaviour might be supported, and it is challenging to produce mathematical methods or computational simulations which can fully capture the range of possible behaviours of the electrons in a material, so that it is complicated to produce predictions of what the behaviour of any given material might be.

By trapping an ultracold gas of ^{40}K atoms in an optical lattice trap which approximates the landscape seen by the electrons moving throughout a material, we construct a simulator for the behaviour of electrons in materials. Our simulator is highly tunable, allowing the simulated material to be investigated over a wide range of relevant parameters. The key feature of our particular experiment is the implementation of a high-resolution fluorescence imaging system and laser cooling techniques which allow us to image the result of our simulation with extreme precision — a technique currently referred to as quantum gas microscopy. In this thesis I will expound on the design constraints, experimental features, and analytical techniques which we use to construct and calibrate our quantum gas microscope.

1.1 Ultracold Atoms

The field of ultracold atomic physics studies dilute gases of atoms, cooled to extremely low temperatures near absolute zero. Although any gas should condense into a solid or liquid phase at these low temperatures, this process happens only over extremely long timescales if the gas has a low density. In the time that elapses between the cooling of the gas and its eventual collection into a solid or liquid, the

gas is considered to be 'meta-stable'. In a similar manner, experiments can be performed on radioactive atoms so long as the time scale for the experiment is short compared to the radioactive lifetime of the particle being studied.

A Dilute Metastable Gas: Ultracold atom experiments occur inside a vacuum chamber, in the presence of extremely few atoms of any kind of background gas — typical vacuum densities are around one million (10^6) particles per cubic centimeter or lower. Into this vacuum an extremely dilute gas of atoms is introduced, and is levitated and trapped inside the vacuum chamber using either magnetic or optical forces. The highest density that the trapped gas reaches in ultracold atom experiments is less than 10^{15} particles per cubic centimeter, far more dilute than the density of the air we breathe which is closer to $10^{19}\,\text{cm}^{-3}$. The low density of the gas means that pairs of atoms collide with one another relatively rarely (perhaps only hundreds of times per second), and groupings of three atoms coming together in a collision is rarer still (less than once per second in the entire gas). It is only the simultaneous collisions of three particles which lead a cold gas to pair off into molecules on its path towards forming a solid, and so it is the rarity of these collisions that prevents the gas from quickly turning to solid even when it is cooled to extremely low temperatures. Entire experiments can be performed before even a small fraction of the gas is able to take the step towards molecular formation and thermodynamic equilibrium.

An Extremely Cold, Dilute Metastable Gas: Every element of a typical ultracold atom experiment is held at room temperature, except for the atoms themselves. While it is trapped inside the vacuum system, the gas can be cooled to extremely low temperatures around $10^{-6}\,\text{K}$ through the techniques of laser cooling. After the lowest temperatures accessible through laser cooling are reached, even further cooling can be performed by 'evaporatively' cooling the cloud — allowing or forcing the hottest fraction of the gas to escape so that the atoms which remain can form a thermal distribution with an even lower temperature. Evaporative cooling can lead to the production of extremely cold gases with temperatures as low as $10^{-9}\,\text{K}$ [1], although the number of atoms in the gas must be correspondingly be reduced to smaller and smaller values. In our experiment we use evaporative cooling to produce ultracold gases of around 10^5 atoms at temperatures near $10^{-7}\,\text{K}$.

An Extremely Cold, Dilute Metastable Quantum Mechanical Gas: When a gas is cooled to ultracold temperatures, the quantum mechanical laws which determine its behaviour begin to be seen more prominently. At high temperatures, much of the wave-like behavior of the gas predicted by quantum mechanics is invisible, and the classical rules for the motion of particles or the thermodynamical properties of gases can be used to determine the expected behaviour. This is because at high temperatures the thermal de Broglie wavelength λ_{th}, which sets the length scale for the wave-like behaviour of quantum mechanics, becomes incredibly small. As the gas is cooled λ_{th} increases, and at the ultracold temperatures reached in modern experiments λ_{th} can become larger than typical distance between atoms in the gas. At this point, the way that atoms in the gas arrange themselves with respect to one another can becomes profoundly different, and the cold gas of atoms is described as being *quantum degenerate*.

An Extremely Cold, Dilute Metastable Quantum Mechanical Gas of Bosons or Fermions: Exactly how the quantum behavior of the individual atoms manifests itself in terms of the overall collective behaviour of the gas depends on the quantum statistical properties of the particles involved. In nature, the elementary particles from which comprise all matter in the universe come in two flavours, fermions and bosons. The behaviour of fermions is fundamentally predicted by Fermi-Dirac statistics, while bosons follow Bose-Einstein statistics. These two possibilities lead to strikingly different behaviour

in a quantum mechanical system, yet predict identical 'classical' behaviour as the temperature is increased. Electrons, protons, and neutrons are all fermions, while photons and the recently discovered Higgs particle are examples of bosons. The Cooper pairs which form in superconductors consist of two fermionic electrons that are bound together, which results in a single quasiparticle with the overall character of a boson. Similarly, the atoms that make up an ultracold gas are a collection of many electrons, neutrons, and protons, and a total atom can be treated as a single particle with a statistical character determined by the addition of all the fundamental particles it contains. An atom will act like a fermion or a boson, depending on whether the total number of constituent protons, neutrons, and electrons is odd or even. This means that isotopes of the same atomic species that differ only by the presence or absence of a single neutron in the nucleus will have completely different quantum statistics and behaviour. When cooled to to quantum degeneracy, bosonic atoms undergo a dramatic phase transition and become a Bose-Einstein condensate while fermionic atoms fill up all the lowest energy quantum states available to form a Fermi sea.

1.2 Quantum Simulation

Given the effort required to produce an ultracold dilute gas of bosons or fermions, it is natural to ask what purpose such a gas may serve. Since few other objects in the universe are as cold as 10^{-6} K, it is natural to suspect that ultracold gases behave like nothing else on Earth. This is only partially true. Although gases at ultracold temperatures do not occur naturally anywhere, the quantum mechanical rules that describe the bahaviour of such a gas are the same ones used to describe any quantum object. Thus studying a quantum gas allows insight into the behaviour of other quantum systems which do naturally occur in nature. The dilute gases used in ultracold atom research only become quantum degenerate at extremely low temperatures, but objects which are lighter or more dense than a gas of atoms will behave according to quantum mechanics at a higher temperature range. For instance, the electrons moving throughout a metal or semiconductor are best understood using quantum mechanics even at room temperature (300 K). An extreme case is that of neutron stars, which are so dense that they are quantum mechanical objects even at temperatures of 10^5 K. The process of using one quantum system (like an ultracold gas) to study another quantum system (like the electrons in a metal) was first proposed by Richard Feynman in 1981 [2], and is referred to as *quantum simulation*.

Studing an ultracold dilute gas to learn about the quantum mechanical underpinnings of an electron gas or a neutron star is only a productive idea if there is information to be gained that was not already accessible by studying the electron gas or the neutron star directly. In the case of a neutron star this is clearly the case, since only limited measurements can be made of these astrophysical objects. However in this thesis we are concerned with quantum simulation of electrons in the crystal lattice of a solid material, and so we should be explicit what advantage might be gained over a direct observation of the material using the full wealth of tools common in the field of condensed matter physics. In a quantum simulation of this sort, ultracold atoms are used as a substitute for electrons and are placed into an trap which has the periodic structure of a crystal lattice seen by electrons inside a real material. The lattice trap for the ultracold atoms is usually generated by laser beams, and can be free of any structural irregularities or defects that might occur in real materials. The strength of the periodic trap is fully controllable in an experiment, and might be tuned to allow the atoms to move in only one or two spatial directions to mimic — for instance — the two-dimensional confinement of electrons in cuprate superconductors.

The interaction between electrons due to their Coloumb repulsion can be mirrored in ultracold gases by their low-energy scattering properties. At the low temperatures used in experiments, atoms interact with one another only via a short-range *contact* interaction with a strength parameterized by the scattering length a. Through the use of Feshbach resonances [3] it is possible to tune the scattering length and hence the inter-particle interaction strength for ultracold atoms over a wide range, allowing a simulation to be performed with particles that are attractive, repulsive, or non-interacting.

The kinds of elementary particles which are simulated can also be tuned in an ultracold gas quantum simulation. The electrons in a material are fermionic particles with two possible states of their intrinsic angular momentum or *spin*, different species or isotopes of atoms may be used to simulate the different kinds of behaviour arising from a system of fermions, bosons, or even mixtures of the two kinds of particles. Atoms have many possible angular momentum states resulting from the combined spin of the protons and neutrons in the nucleus with the angular momentum and spin of the electronic orbitals that surround it. Two of these spin states can be isolated in order to simulate the two spin states of electrons, or one can state can be isolated in analogy with electrons which are polarized by a strong electric field. Further though, more than two spin states can be used to simulate the effect that this would have on an electronic system.

In sum, a system of ultracold atoms can be controlled and tuned in many ways, and such systems have proven to be versatile systems in which to explore quantum mechanical models. These benefits have been used to simulate the physics of graphene [4], explore lattice physics in the presence of extremely large magnetic flux [5], and to explore the topological edge states that appear in a system that realizes the quantum Hall effect [6]. In addition to these explorations of topics in condensed matter physics, the high degree of control afforded with laser-cooled atoms has enabled tests of the universality of free fall [7], produced extremely accurate atomic clocks [8], and allowed precision measurements to constrain certain models for dark energy [9].

1.3 ^{40}K and ^{87}Rb

Alkali metals are the most common atomic species used for ultracold atom experiments, as the presence of only a single valence electron leads to a rather simple electronic structure and so requires low experimental complexity. Of all the alkali metals, the bosonic atom ^{87}Rb is by far the most popular. A combination of atomic properties make it amenable to the laser cooling and evaporative cooling techniques which are essential to producing an ultracold gas. This led ^{87}Rb to be used in one of the first demonstrations of Bose-Einstein condensation in 1995 [10]. The production of the first quantum-degenerate gases of fermionic atoms came later in 1999 [11] with the cooling of ^{40}K (a rare isotope of potassium, which is radioactive with a half-life over one-billion years). Fermi gases are generally harder to bring to quantum degeneracy because Fermi-Dirac quantum statistics forbids collisions between identical fermionic atoms at ultracold temperatures. Collisions between atoms are necessary for evaporative cooling to succeed, and so suppression of these collisions can effectively prevent cold temperatures from being reached. This problem can be solved by cooling a gas of fermions divided into multiple spin states, so that not all atoms are identical and collisions may still occur.

Another solution to allow efficient evaporative cooling of fermionic atoms is to cool two different kinds of atoms together in the same experiment, with one kind of atom which is easy to cool acting as a refrigerant for another type of atom. This process of using one ultracold gas to cool another is
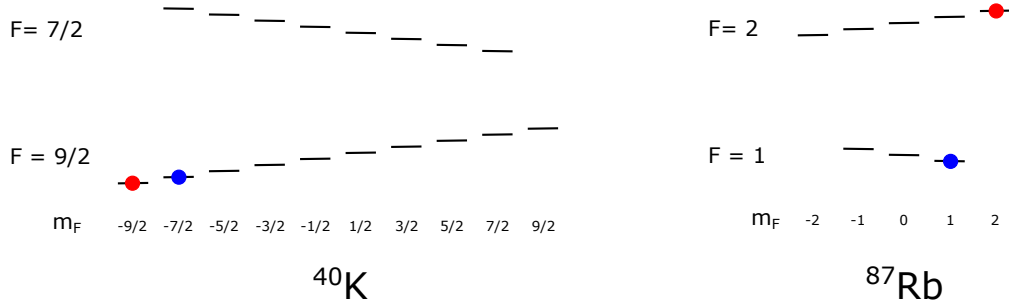
Figure 1.1: Illustration of the different internal spin states of a ^{40}K or ^{87}Rb atom in its lowest-energy electronic state. The states are denoted by quantum numbers F and m_F. F describes the total angular momentum resulting from a addition of the nuclear spin to the electron spin and the electron orbit, while m_F describes the projection of the total angular momentum along the direction of an applied magnetic field. We prepare ^{40}K in the $|F = 9/2, m_F = 9/2\rangle$ state which can be magnetically trapped, before transferring the cold gas to the $|F = 9/2, m_F = -9/2\rangle$ state denoted by the red circle. For quantum simulation of the two spin states of electrons we can transfer part of the cloud to the $|F = 9/2, m_F = -7/2\rangle$ state indicated by the blue circle. The two important states of ^{87}Rb are the $|F = 2, m_F = 2\rangle$ state (red) that we trap magnetically and the $|F = 1, m_F = 1\rangle$ state (blue) which in general provides better sympathetic cooling.

known as sympathetic cooling. In our experiment we prepare quantum degenerate gases of fermionic ^{40}K through sympathetic cooling with bosonic ^{87}Rb . These two species have been shown previously to form a good combination for sympathetic cooling, with large gases of ^{87}Rb being easy to prepare, and the inter-atomic scattering properties of the two species well suited for evaporative cooling. ^{87}Rb and ^{40}K have been used together in this way in many ultracold gas experiments [12–15], and a mixture of these two gases has even been combined to form an ultracold gas of heteronuclear diatomic molecules [16].

The different spin states of ^{40}K and ^{87}Rb atoms are shown in Figure 1.1. We begin our experiment by using a magnetic trap to confine ^{40}K atoms in the $|F = 9/2, m_F = 9/2\rangle$ state and ^{87}Rb atoms in the $|F = 2, m_F = 2\rangle$ state. In the final stages of cooling we adjust the spin state of ^{40}K in order to prepare degenerate Fermi gases in the $|F = 9/2, m_F = -9/2\rangle$ spin state which has the lowest overall energy in a magnetic field. To simulate the two spin states of an electron we can create a mixture of the $|F = 9/2, m_F = -9/2\rangle$ and $|F = 9/2, m_F = -7/2\rangle$ states, and the interaction between atoms in these two spin states may be tuned experimentally using an Feshbach resonance at 202 G. For the final evaporative cooling stages we transfer the ^{87}Rb atoms in the experiment into the $|F = 1, m_F = 1\rangle$ state to ensure collisional stability with ^{40}K atoms in $|F = 9/2, m_F = -9/2\rangle$.

1.4 Quantum Gas Microscopy

All ultracold gas experiments occur in the presence of a trap to contain the gas, most commonly a smooth harmonic potential generated by focussed laser beams. The presence of the trapping potential typically creates a fundamental difference between the ultracold gas and the system which it is meant to simulate. This difference is often taken into account through the use of a local density approximation, treating the gas as simulating an untrapped system while possessing a spatially varying chemical potential and density. Under this approximation, a trapped gas might realize several different quantum phases of a simulated system simultaneously, in different regions of the trapping potential. This poses a problem

for experimental measurements which probe the entire cloud at once, since the measurable signatures of any one phase may be buried by the signal of other coexisting phases. Imaging an ultracold gas in the trap with high resolution allows for measurements to be made locally so that different coexisting phases in the trap can be characterized in parallel.

In recent years several experiments have developed techniques to obtain high-resolution in-situ images of bosonic quantum gases held in a periodic trap [17–20]. These techniques have enabled a wealth of new information to be gained from quantum simulation of the bosonic Hubbard model, a prototypical quantum many-body system. Specifically, looking at the ultracold gas in-situ allowed for observation of the 'wedding-cake' structure of incompressible Mott-insulating plateaus that is a result of the external confinement of the trapping potential [19–21]. Of these four high-resolution experiments, two groups working with ^{87}Rb achieved fluorescence imaging of the gas with enough sensitivity to make single atoms of the gas visible, and with a high enough resolution to distinguish the atomic positions on each site of the periodic potential. This combination of technical achievements allows a complete reconstruction of the site by site density of a quantum gas in a strongly correlated many-body state, and has been termed *quantum gas microscopy*. The power of this single-atom and single-site resolving measurement technique has led to an overwhelming number of scientific results regarding the physics of bosonic quantum gases [22–30].

The work in this thesis describes efforts to extend the techniques of quantum gas microscopy to a fermionic atomic species. This would allow the control and precision of this measurement technique to be applied to quantum simulation of technologically relevant models in condensed matter physics. This research direction was pursued in parallel by several experimental groups around the world. Developing this kind of experiment for fermionic atoms was not a straightforward task even starting with the results of the two ^{87}Rb quantum gas microscopes, as the specific laser cooling techniques used to collect single-atom sensitive fluorescence images for ^{87}Rb do not directly carry over to the commonly used fermionic atoms ^{6}Li and ^{40}K . Instead along with the construction efforts for a new quantum gas experiment, a large investigative effort was required to explore nonstandard laser cooling techniques which could provide the necessary combination of cooling power and fluorescence generation. Five groups including ours published results showing single-atom sensitive and lattice-site resolved fluorescence imaging of fermionic atoms in 2015: three using ^{40}K [31–33] and two using ^{6}Li [34, 35]. The specific laser cooling techniques and imaging geometry used in each experiment differs, but all have been able to extend the techniques of quantum gas microscopy to fermionic atoms.

1.5   Outline

This thesis is structured into 5 chapters, with each containing a mixture of theoretical calculations and experimental techniques. First the basic theory of degenerate fermions in optical traps is presented. The next 4 chapters are organized roughly by the chronological order which we follow in each experimental cycle to produce site-resolved fluorescence images, first detailing how the apparatus was built to accommodate a high-resolution microscope, then explaining the methods used to produce cold two-dimensional gases in the microscope focus, subsequently detailing the laser cooling techniques used to produce a strong fluorescence signal, and finally discussing the achieved site-resolved fluorescence images.

Chapter 2 explores the mathematical foundations for the various optical traps used in this experiment. The behaviour of an ultracold gas of fermionic atoms held by such a trap is discussed, including the signatures used to determine the temperature of the degenerate gas. Various theoretical aspects of optical lattice physics are detailed, including the limits in which fermions trapped in an optical lattice can be used to perform a quantum simulation of the Hubbard model.

Chapter 3 contains specific experimental details of our quantum gas microscope apparatus, including details of the incorporation of a large-NA microscope objective into a cold atom experiment. We explore the vacuum chamber design used to allow a high degree of optical access, as well as the microscope and other imaging optics used to collect the images. Also included are details of the high-power optical lattice we use to trap atoms in the microscope focus.

Chapter 4 details the steps taken to prepare a two-dimensional ultracold gas in the focus of the microscope objective. Sympathetic cooling of ^{40}K by ^{87}Rb prepares a cold gas, which is then loaded into the optical lattice. We show how a single two-dimensional slice of the optical lattice can be isolated spectroscopically before fluorescence imaging.

Chapter 5 contains all theoretical and experimental details of the EIT cooling technique used to cool atoms in the microscope focus while simultaneously producing steady atomic fluorescence. First we compare EIT cooling to the alternative technique of Raman Sideband cooling which has also been used in quantum gas microscope experiments. Next we detail the experimental implementation of EIT cooling with particular care being given to shielding of the fluorescence imaging system from the laser cooling beams. Experimental data is shown which confirms that EIT cooling can cool ^{40}K atoms in our optical lattice close to their motional ground state.

Chapter 6 shows the fluorescence images we obtain with single-atom sensitivity and single-site resolution. We experimentally extract the point spread function of the imaging system and compare it both with theoretical expectations and also other quantum gas microscope experiments. We detail our algorithm to extract the optical lattice occupancy from a single fluorescence image. We show the optimization of the EIT cooling process based on observed atom movement during during fluorescence images.

Chapter 2

Quantum Gases in Optical Traps

The ultracold gases that are produced in our experiment are passed between several different trapping geometries before they arrive finally in the optical lattice for quantum simulation. For the initial collection of cooled atoms from a room temperature vapour we use a magneto-optical trap which leverages the Doppler shift to reduce the atom velocities while confining them with a combination of magnetic fields and radiation pressure. We subsequently transport the laser-cooled atoms into a second chamber using a purely magnetic trap. In the second chamber the atoms are evaporatively cooled in a hybrid trap which combines the magnetic trap with a repulsive optical 'plug'. The penultimate trap is a crossed-beam optical trap in which the final cooling to quantum degeneracy is performed. Finally, the quantum gas is loaded into the optical lattice for quantum simulation experiments. The purely optical trapping potentials in the last two steps of this sequence are vitally important for the results of this thesis. In this chapter we explore the dipole potential used to generate these two trapping configurations, and detail the expected behaviour when they are used to confine a quantum gas of fermions.

2.1   Dipole Potential

For precision studies and manipulation of quantum objects like atoms or ions, it is preferable that they are isolated from their environment and confined in space. For an electrically charged ion, the necessary trapping can be provided with an external electric field. For a neutral alkali atom, unaffected by a static electric field, magnetic fields can instead be used to produce a confining force by interacting with the magnetic dipole associated with the valence electron's spin and angular momentum. Magnetic trapping arrangements such as the quadrupole trap or Ioffe-Pritchard trap are very common in ultracold atom experiments, and were used in the first demonstrations of Bose-Einstein condensation [10, 36].

When trapping a gas of neutral atoms with low temperature, an alternative trapping configuration is possible, created by the interaction between an oscillating electric field and the electric dipole moment that it induces in a neutral atom. Such a trap offers much more control over the potential shape, more bandwidth for trap dynamics, and more freedom over the internal atomic states which may be trapped. However within common experimental constraints the depth of a trap created using this electric dipole interaction is typically much weaker than can be achieved with magnetic trapping, and so that in many experimental both kinds of traps are used in different circumstances.

The dipole potential V_{dip} for a two-level atom (with ground and excited states $|1\rangle$ and $|2\rangle$) interacting

with an oscillating electric field with intensity $I(\vec{r})$ is [37]

$$V_{\text{dip}}(\vec{r}) = -\frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) I(\vec{r})\,, \tag{2.1}$$

where ω_0 is the resonant frequency of the two-level system so that $\hbar\omega_0 = E_2 - E_1$, and Γ is the spontaneous decay rate of the excited state. When ω is not too different from ω_0, the second term of (2.1) (commonly referred to as the counter-rotating term) contributes minimally to the potential and can be disregarded. Expressing the electric field frequency ω in terms of its detuning from resonance $\Delta = \omega - \omega_0$, the dipole potential takes on the simple form

$$V_{\text{dip}}(\vec{r}) = \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I(\vec{r})\,. \tag{2.2}$$

An important feature of the dipole potential is easily visible in (2.2): for $\Delta > 0$ (blue detuning), the potential will repel the atom from regions of high intensity, and conversely for $\Delta < 0$ (red detuning) atoms are attracted to regions of high intensity. By focussing a laser beam with $\Delta < 0$, a local maximum of intensity is formed which will act as a three-dimensional trap for a neutral atom.

Although the expression (2.2) provides a simple interpretation of the dipole potential, the approximation made by dropping the second term of (2.1) can lead to significant errors for beams with very large detuning. In our experiment, atoms with $\omega_0 \approx 390\,\text{THz}$ are trapped with a laser $\omega = 284\,\text{THz}$, corresponding to a red-detuning of $\Delta \approx -100\,\text{THz}$. In this case, the second term of (2.1) constitutes a 15% correction to the overall trapping potential and cannot be ignored. Using (2.1), we calculate that with a beam intensity of $1\,\text{kW}\,\text{cm}^{-2}$ we expect a trap depth of $V_{\text{dip,K}}/k_B = 1.37\,\mu\text{K}$ for ^{40}K and $V_{\text{dip,Rb}}/k_B = 1.53\,\mu\text{K}$ for ^{87}Rb .

In addition to experiencing the potential U_{dip} created by the applied laser beam, an atom in a high intensity oscillating field will also scatter photons at a rate Γ_{sc}. In a semi-classical picture, the dipole potential arises due to the in-phase motion of the atomic dipole with the applied field, while the photon scattering corresponds to energy dissipation due to the out-of-phase motion of the dipole and the field. As a function of the applied beam parameters, Γ_{sc} is given by [37]

$$\Gamma_{\text{sc}}(\vec{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\omega}{\omega_0} \right)^3 \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right)^2 I(\vec{r})\,. \tag{2.3}$$

Both V_{dip} and Γ_{sc} scale in proportion with the applied beam intensity $I(\vec{r})$, but with a choice of large red-detuning $\omega \ll \omega_0$, the scattering rate can become very small since it scales roughly with Δ^{-2}. In our experiment, even with the highest attainable trapping beam intensities of around $500\,\text{kW}\,\text{cm}^2$, the scattering rate is only $2\,\text{Hz}$. Thus photon scattering is rare during typical experimental timescales of 10-100 ms, and the optical dipole potentials that we create can rightfully be thought of as conservative potentials.

2.1.1   Multi-Level Atoms

In expressing the dipole potential (2.1) above, the atom was assumed to behave as a two-level quantum system with resonant frequency ω_0. In this case the ground state of the two-level system feels the potential U_{dip}, and the excited state feels the opposite potential $-U_{\text{dip}}$. Thus two-level atoms in the
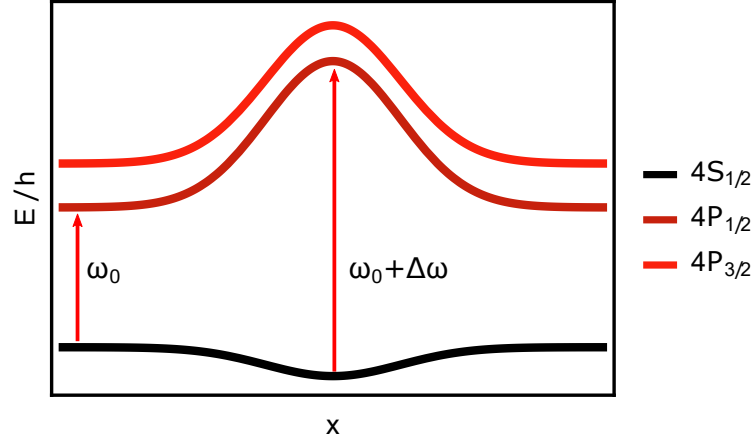
Figure 2.1: Example of the effect of the dipole potential due to a $\lambda = 1054\,\text{nm}$ beam focussed near $x = 0$. Energies for the $4S_{1/2}$, $4P_{1/2}$ and $4P_{3/2}$ states of ^{40}K are shown, and the shift in the resonance frequency $\Delta\omega$ of the $|4S_{1/2}\rangle \rightarrow |4P_{1/2}\rangle$ transition is indicated. $|4S_{1/2}\rangle \rightarrow |4P_J\rangle$ transition energies and fine structure are not to scale.

electronic ground (excited) state are attracted to (repelled from) the centre of an intense red-detuned laser beam.

For real atoms which have many excited levels, the dipole potential for a state $|i\rangle$ arising due to a beam of frequency ω is obtained by summing over the two-level shifts from each possible atomic transition $|i\rangle \rightarrow |j\rangle$ with resonant frequency $\omega_{ij} = (E_j - E_i)/\hbar$. The contributions from different transitions are weighted by the transition strength and the appropriate detuning $\Delta_j = \omega - \omega_{ij}$. For alkali atoms and red-detuned trapping beams, the shifts of the electronic ground state $|nS\rangle$ are dominated by the transition to the first electronic excited state $|nS\rangle \rightarrow |nP\rangle$. The $|nP\rangle$ state is split by the fine structure into two states, denoted $D1$ and $D2$. For calculation of the ground-state trapping potential, all higher-energy states can typically be ignored. If the detuning of the trap beam is not large compared to the atomic fine structure Δ_F or hyperfine structure Δ_{HF}, then the transition strengths will vary depending on the particular values of the angular momentum quantum numbers $|F, m_F\rangle$, creating a spin-dependent optical potential. Alternatively, if the frequency of the trapping laser lies between the $D1$ and $D2$ states, then combination of blue-detuning from one transition and red-detuning from the other can lead to an overall cancellation of the dipole potential.

For experiments with ultracold atoms trapping laser wavelengths are commonly chosen near $1064\,\text{nm}$, due to the availability of high power Nd:YAG laser sources. If the particular wavelength chosen is far from resonance with any other transitions then the simple two-level picture will hold, with an attractive ground state potential of depth V_0 and a repulsive excited state potential of the same strength. However for ^{40}K atoms a trap laser near $1064\,\text{nm}$ will have a relatively small blue-detuning from the $|4P\rangle \rightarrow |3D\rangle$ transition at $1180\,\text{nm}$. This enhances the repulsive potential of the $|4P\rangle$ excited state, creating a shift roughly 5 times greater than the trap depth for the ground state. These shifts are shown for the $4S_{1/2}$, $4P_{1/2}$ ($D1$), and $4P_{3/2}$ ($D2$) states of ^{40}K in Figure 2.1.

To image and manipulate atoms that are held in an optical trap, additional resonant beams with $\omega = \omega_0$ are commonly used. Due to the large excited-state shifts for ^{40}K , the appropriate resonant frequency is shifted by $\Delta\omega$ when the atoms are located inside the trapping laser field. In our experiment with $\lambda_{\text{trap}} = 1054\,\text{nm}$, the resonance shift $\Delta\omega/\hbar$ for transitions to the $D1$ and $D2$ states is about $170\,\text{kHz}$

when the trapping beam intensity is $1\,\mathrm{kW\,cm^{-2}}$. For low beam intensities the shift is then smaller than the excited state linewidth of $\Gamma/2\pi = 6\,\mathrm{MHz}$ and can be ignored, but at our highest trapping intensities of $500\,\mathrm{kW\,cm^{-2}}$ the shift is $> 80\,\mathrm{MHz}$. Any difference in trapping intensity seen by atoms in different trap regions then leads to inhomogeneous broadening of the optical transition, which affects our ability to image, manipulate, or cool all atoms in the trap simultaneously.

For the experiments described in this thesis, where atoms are trapped in an intense optical trap while emitting photons through repeated cycling between the $|4S_{1/2}\rangle$ and $|4P_{1/2}\rangle$ states, the large excited state shifts pose another issue. During the time the atoms spend in the exited state before spontaneously emitting a photon, they are very strongly repelled by the optical potential. This may lead atoms to move through or be ejected from the trap after several spontaneous emission cycles. For some laser cooling schemes it has been suggested that by careful choice of the detuning of the cooling laser beams, atomic dressed states can be created which admix a small fraction of the atomic excited state with a large fraction of an atomic ground state [31]. Such a state can be trapped by the optical potential, but may still produce spontaneous emission. The laser cooling technique described in this thesis makes use of a dressed state that is predominantly composed of the trapped atomic ground state, and we find that this technique is able to produce spontaneous emission using the $|4S_{1/2}\rangle \rightarrow |4P_{1/2}\rangle$ transition without large observed mobility or loss of the trapped atoms.

Although the trapping potential seen by an atom undergoing spontaneous emission can be controlled in some ways by the dressing of the atomic basis, it is not generally possible to adjust the trapping potential of the emitting state to be identical to one of the ground states. More control over the excited state trapping potential can be gained by careful choice of the wavelength used for optical trapping. A carefully chosen trap wavelength can balance the anti-trapping due to the ground-state transition with a trapping effect due to coupling to another electronic excited state [38]. In our experiment the trapping wavelength $\lambda_{\mathrm{trap}} = 1054\,\mathrm{nm}$ was specifically chosen to produce identical trapping for the $|4S_{1/2}\rangle$ and $|5P_{3/2}\rangle$ states, as described in Appendix A. For the laser cooling techniques described in this thesis which do not make use of the $5P_{3/2}$ excited state, this wavelength choice has little influence on the cooling performance.

2.1.2 Optical Dipole Traps

Optical dipole traps use the electric dipole potential (2.1) to trap atoms in space by confining them within patterns of high or low intensity laser light. With the intensity patterns that can be created by using one of the many Hermite-Gauss or Laguerre-Gauss spatial modes for a laser beam, as well as the more arbitrary intensity patterns that can be realized with a spatial light modulator [39, 40] or a time-averaged potential [41, 42], it is clear that there are many different types of trapping potentials that can be created. For instance, clouds of ultracold atoms have been squeezed into a pancake [43], confined within a hard-walled box [44], and allowed to circulate inside of a ring [45, 46]. The two trapping configurations used in this experiment are among the most simple and the most common — the approximately harmonic trap formed at the focus of a single red-detuned Gaussian laser beam, and the periodic trap formed from the interference pattern of a retro-reflected beam.

A focussed Gaussian laser beam of wavelength λ travelling along the x direction has the spatial intensity profile

$$I(x, y, z) = I_0 \left(\frac{w_{0,y} w_{0,z}}{w_y(x) w_z(x)} \right)^2 \exp\left(-\frac{2y^2}{w_y(x)^2} - \frac{2z^2}{w_z(x)^2} \right), \qquad (2.4)$$
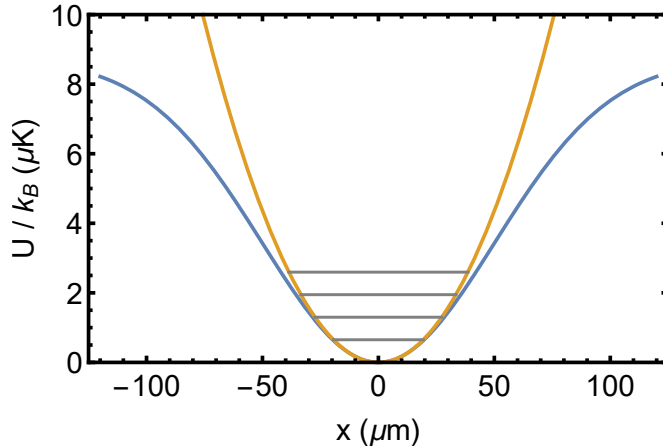
Figure 2.2: The Gaussian profile of an optical dipole trap (blue) and the corresponding harmonic approximation (orange), for ^{40}K atoms in a beam with $P = 1\,\mathrm{W}$ and $w = 100\,\mu\mathrm{m}$. Grey lines depict the quantized energies $(n + \frac{1}{2})\hbar\omega$ of the harmonic oscillator potential for $n = 100, 200, 300, 400$.

where I_0 is the total intensity, $w_{0,i}$ are the (potentially different) beam waists along each direction $i = y, z$, and as a function of the axial direction x, the beam's intensity variation is characterized by the Rayleigh lengths $x_{R,i} = \pi w_{0,i}^2/\lambda$ according to the expression

$$w_i(x) = w_{0,i}\sqrt{1 + (x/x_{R,i})^2}\,. \tag{2.5}$$

Since $x_{R,i}$ is typically much larger than $w_{0,i}$, the intensity gradients of (2.4) are much stronger along y and z than along x. Thus a red-detuned beam with the profile (2.4) typically produces an elongated, cigar-shaped trapping geometry. To generate a more symmetric trap, a second focussed Gaussian beam may be added in the perpendicular direction y to create additional confinement along x. As long as the two beams have perpendicular polarizations, or a large difference in optical frequency, then no stable interference pattern will be formed by the intersection of the two beams. The resulting intensity pattern is then simply the sum of the two intensity distributions. Denoting the parameters of the first (second) beam with the subscript A (B) we then have

$$I(x, y, z) = I_{0,A}\left(\frac{w_{0,A,y}w_{0,A,z}}{w_{A,y}(x)w_{A,z}(x)}\right)^2 e^{-\frac{2y^2}{w_{A,y}(x)^2} - \frac{2z^2}{w_{A,z}(x)^2}} + I_{0,B}\left(\frac{w_{0,B,y}w_{0,B,z}}{w_{B,y}(x)w_{B,z}(x)}\right)^2 e^{-\frac{2y^2}{w_{B,y}(x)^2} - \frac{2z^2}{w_{B,z}(x)^2}}\,. \tag{2.6}$$

In this pattern there is tight confinement along all three directions. Atoms trapped in the crossing region will not explore much of the axial direction of either beam, so that $w_{A/B,i}(x) \approx w_{0,A/B,i}$ and it is possible to ignore the dependence on the Rayleigh lengths of the beams.

If a cloud of atoms trapped by the intensity pattern (2.6) is cold enough to sit only near the bottom of the trap, then it sees an approximately harmonic potential $V(x, y, z) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)$. An example of this harmonic approximation is shown in Figure 2.2. The trap frequencies ω_i are found by using the intensity (2.6) to find the trapping potential (2.1), and then expanding the potential around
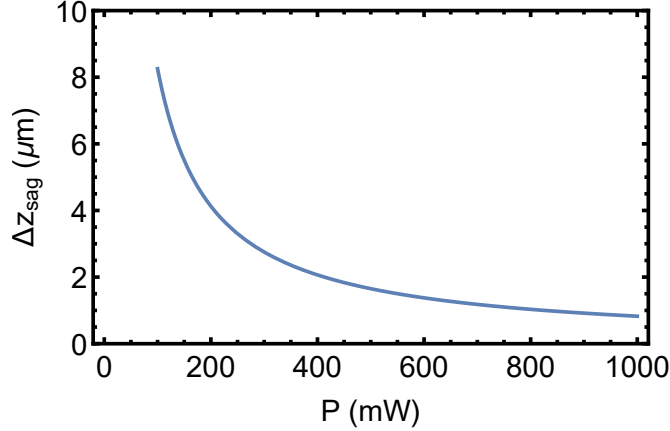
Figure 2.3: Estimated gravitational sag between ^{87}Rb and ^{40}K in our crossed optical dipole trap. The simple formula (2.10) based on the harmonic approximation has been used, which is not appropriate for very low powers.

its minimum to lowest order in x, y, and z:

$$\omega_x = \frac{4V_{0,B}}{mw_{0,B,x}^2} \tag{2.7}$$

$$\omega_y = \frac{4V_{0,A}}{mw_{0,A,y}^2} \tag{2.8}$$

$$\omega_z = \frac{4}{m}\left(\frac{V_{0,A}}{w_{0,A,z}^2} + \frac{V_{0,B}}{w_{0,B,z}^2}\right), \tag{2.9}$$

where $V_{0,A}$ denotes the maximum dipole potential depth that results from the peak beam intensity $I_{0,A}$, containing all of the atomic parameters of (2.1). For a given power P_A this peak intensity is $I_{0,A} = 2P_A/\pi w_{0,A,y}w_{0,A,z}$, and with common experimental parameters of $P = 1\,\text{W}$, $w_0 = 100\,\mu\text{m}$, and $\lambda = 1054\,\text{nm}$ the achieved trap depths are roughly $V_0/k_B \approx 8.7\,\mu\text{K}$ for ^{40}K atoms, and the mean effective harmonic trapping frequency is roughly $\bar{\omega}/2\pi = 170\,\text{Hz}$. Slight differences in the trap detuning and atomic polarizability create a larger trap depth $V_0/k_B \approx 9.7\,\mu\text{K}$ for ^{87}Rb atoms. Taking into account the differences in mass, the effective harmonic trapping frequencies for ^{40}K and ^{87}Rb differ by $\omega_K/\omega_{Rb} \approx 1.4$.

2.1.3 Gravitational Sag

In an experiment, atoms are subject to gravity in addition to the applied optical dipole potential. The additional linear potential $U_g = mgz$ tilts the harmonic trap, shifting the trap minimum position to $z_{\text{sag}} = g/\omega_z^2$. For weak traps this tilt becomes large, and the harmonic approximation around the beam intensity maximum will no longer be appropriate. For a trap formed from two crossed beams with $P = 1\,\text{W}$ and $w_z = 100\,\mu\text{m}$, this sag for ^{40}K atoms is around $7\,\mu\text{m}$. For ^{87}Rb atoms, the trap depth is similar, but the mass is much larger, leading to a sag of $13\,\mu\text{m}$. Since we use ^{87}Rb to cool ^{40}K , it is important that the two largely overlap in the trap to allow for interspecies collisions and thermalization.

An important parameter is then the differential sag

$$\Delta z_{\text{sag}} = \frac{g}{\omega_{z,K}^2 \omega_{z,Rb}^2}(\omega_{z,K}^2 - \omega_{z,Rb}^2) \approx \frac{g}{\omega_{z,K}^2} \,, \tag{2.10}$$

where the simplification is due to the approximation $\omega_K^2 = 1.95\omega_{Rb}^2 \approx 2\omega_{Rb}^2$. To minimize the effects of gravitational sag in our experiment, we form our optical dipole trap out of elliptical beams with $w_{0,A,y} \approx w_{0,B,x} \approx 200\,\mu$m and $w_{0,A,z} \approx w_{0,B,z} \approx 40\,\mu$m. This gives a similar overall trap volume and trap depth to a circularly symmetric beam with a $100\,\mu$m waist, but with a much larger vertical trapping frequencies ω_z. This reduces the gravitational sag to only a few μm for our range of beam powers. This gravitational sag will still be too large if it exceeds the size of the two ultracold clouds in the trap. We can estimate the density profile of the ^{40}K cloud using the calculations presented in Section 2.2, and we find that the vertical extent of the cloud is typically greater than $15\,\mu$m, so that a gravitational sag $< 5\,\mu$m should not significantly reduce thermalization between the two species.

2.2   Degenerate Fermi Gases

For a collection of N non-interacting quantum particles with temperature T, distributed in a quantum system with discrete energy levels E_i, the average occupation of a quantum state with energy E_i is given by the distribution

$$n(\beta, \mu, E_i) = \frac{1}{e^{\beta(E_i - \mu)} \pm 1} \,, \tag{2.11}$$

where μ is the chemical potential of the gas, and the dependence on T is included in the parameter $\beta = (k_B T)^{-1}$, with k_B the Boltzmann constant. If the particles are fermions, the positive sign in the denominator is taken, and (2.11) is the Fermi-Dirac distribution. For bosons, the negative sign is instead used, leading to the Bose-Einstein distribution. This slight difference leads to tremendously different behaviour for the two different types of particles. In particular, the Fermi-Dirac distribution predicts that any quantum state i will be occupied by no more than one identical fermion — the Pauli exclusion principle. Adding the occupation probabilities for all discrete quantum states gives the total particle number

$$N = \sum_i n(\beta, \mu, E_i) \,, \tag{2.12}$$

and this relationship provides a relationship between μ, N, and T.

If the energy levels E_i are very closely spaced, they may be treated in the semi-classical limit by writing the energy as $E(\vec{r}, \vec{p})$ — a function of the position $\vec{r}$ and momentum $\vec{p}$. In this limit of classical energies, the quantum statistics of the Fermi-Dirac distribution will still play a significant role in determining the properties of the system at low temperatures. To calculate the expected distributions of particles over $\vec{r}$ or $\vec{p}$, the classical energy $E(\vec{r}, \vec{p})$ is used in (2.11) to give the semi-classical phase space density

$$f(\beta, \mu, \vec{r}, \vec{p}) = \frac{1}{\zeta^{-1} e^{\beta E(\vec{r}, \vec{p})} + 1} \,, \tag{2.13}$$

where the dependence on μ is expressed through the fugacity $\zeta = e^{\beta \mu}$. To add the occupation probabilities for all states as in (2.12), the sum over all energy states E_i is now replaced by an integral of (2.13) over E, weighted by the appropriate density of states $g(E)$. In free space the quantum states are plane waves with momentum $\vec{p}$ and kinetic energy $E(\vec{r}, \vec{p}) = p^2/2m$, giving $g(E) = 2\pi(2m)^{3/2}\sqrt{E}$ for three

dimensions. This density of states can be used to determine the relationship between N and μ for a free Fermi gas.

The addition of a trapping potential changes $E(\vec{r}, \vec{p})$, meaning that a new density of states must be determined. In our optical dipole trap, the trapping potential is well approximated by a three-dimensional harmonic oscillator with geometric mean trapping frequency $\bar{\omega}$. The density of states for a 3D harmonic trapping potential is

$$g(E) = \frac{E^2}{2(\hbar\bar{\omega})} \,. \tag{2.14}$$

Using this $g(E)$, the number of fermions occupying states in the range of energy $(E, E + \Delta E)$ can be calculated by integrating with (2.13):

$$N(E, \Delta E) = \int_E^{E+\Delta E} g(E)f(\beta, \mu, E)\, dE \,. \tag{2.15}$$

For $T = 0$, the Fermi-Dirac distribution (2.13) behaves like a step function — predicting unity occupation of all states with $E < \mu$, and zero occupation of states with $E > \mu$. Since all states below μ are filled, a relationship between N and μ is easily obtained by integration in this case:

$$N = \int_0^\infty g(E)f(\beta, \mu, E)\, dE = \int_0^\mu g(E)\, dE = \frac{1}{6}\left(\frac{\mu}{\hbar\bar{\omega}}\right)^3 \,. \tag{2.16}$$

This expression determines the chemical potential for N fermions at $T = 0$ filling a harmonic trap characterized by the trap frequency $\bar{\omega}$. The chemical potential of a collection of fermions at zero temperature is referred to as the Fermi energy E_F, and an associated Fermi temperature T_F can be defined through $E_F = k_B T_F$. Using (2.16), we have an expression for the Fermi temperature in a harmonic trap:

$$T_F = \frac{\hbar\bar{\omega}}{k_B}\left(6N\right)^{1/3} \,. \tag{2.17}$$

It is also sometimes convenient to define a Fermi momentum p_F and a corresponding Fermi wavevector $k_F = p_F/\hbar$ so that

$$E_F = \frac{p_F^2}{2m} = \frac{\hbar^2 k_F^2}{2m} \,. \tag{2.18}$$

For non-zero temperatures, (2.16) does not hold, and for a fixed number N of fermions the chemical potential μ is defined implicitly through the integral expression

$$N = \int_0^\infty g(E)f(\beta, \mu, E)\, dE = \frac{-\mathrm{Li}_3\left(-\zeta\right)}{(\beta\hbar\bar{\omega})^3} \,, \tag{2.19}$$

where $\mathrm{Li}_n\left(x\right)$ is the polylogarithm function, defined by

$$-\mathrm{Li}_n\left(-\zeta\right) = \frac{1}{\Gamma(n)}\int_0^\infty \frac{t^{n-1}}{\zeta^{-1}e^t + 1}\, dt \,. \tag{2.20}$$

The additional function $\Gamma(n)$ which appears in (2.20) is the Gamma function.

For a collection of fermions, the quantum statistical properties will only become important when the temperature of the cloud is made comparable to or lower than the Fermi temperature T_F. Unlike a system of bosons at $T < T_c$, there is no dramatic phase transition that occurs at this temperature, but

the temperature range around T_F is where the occupation of lowest energy states becomes significant and the effects of Pauli exclusion are seen. Temperatures for which $T/T_F < 1$ are therefore the most interesting to study experimentally. By combining (2.19) with (2.17), a simple expression is found which relates the ratio T/T_F to the fugacity ζ:

$$(T/T_F)^{-3} = 6\text{Li}_3(-\zeta) . \tag{2.21}$$

For a system of electrons moving throughout a metal, T_F is very large (10^4 K), and so even at room temperature quantum statistics are important in determining the system's behaviour. On the other hand, in a system of ultracold atoms the smaller particle density and weaker confining potentials result in $T_F \approx 10^{-6}$ K. This means that extensive cooling is required to bring a trapped Fermi gas to quantum degeneracy. The first such degenerate cloud was produced using ^{40}K atoms in 1999 [11]. More recently, values of T/T_F as low as 0.04 have been reported in ultracold gases of fermions [47].

2.2.1   Density Distribution

In order to determine the value of T/T_F in an experimental setting, it would be convenient to measure the fugacity of the cloud ζ to be used in (2.21). Since ζ cannot be probed directly in our experiment, it is instead typically determined through its influence on the density profile or momentum distribution of the cloud. To find the density distribution for fermions trapped in a 3D harmonic potential $V(\vec{r}) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)$, the density distribution can be obtained by integrating (2.13) over all possible momenta, associating each quantum state with a phase space volume h^3. The resulting density distribution is succinctly expressed as

$$n(\vec{r}) = -\frac{1}{\Lambda_{\text{th}}^3}\text{Li}_{3/2}\left(-\zeta e^{-\beta V(\vec{r})}\right) , \tag{2.22}$$

where Λ_{th} is the Thermal de Broglie wavelength

$$\Lambda_{\text{th}} = \sqrt{\frac{2\pi\hbar^2\beta}{m}} . \tag{2.23}$$

In the limit $T/T_F \gg 1$, the fugacity becomes very small and the Polylogarithm can be approximated using $\text{Li}_n(\zeta) \approx \zeta$. In this case the density distribution takes the form of a Gaussian distribution. The widths of the Gaussian profile are $\sigma_i = \sqrt{k_B T/m\omega_i^2}$ along each trap axis i, just as would be expected for a classical ideal gas. At low temperatures, $T/T_F \ll 1$, the limiting behaviour is instead $\text{Li}_n(\zeta) \approx \ln^n(\zeta)/\Gamma(n+1)$, and the density profile assumes the form

$$n(\vec{r})|_{T/T_F \ll 1} = \frac{8}{\pi^2}\frac{N}{R_{F,x}R_{F,y}R_{F,z}}\left(1 - \frac{x^2}{R_{F,x}^2} - \frac{y^2}{F_{F,y}^2} - \frac{z^2}{R_{F,z}^2}\right)^{3/2} , \tag{2.24}$$

where the Fermi radii $R_{F,i}$ for each axis are $R_{F,i} = \sqrt{2E_F/m\omega_i^2}$. Rather than the smooth wings of the Gaussian distribution, the low-temperature density distribution falls abruptly to zero at $R_{F,i}$.

Density profiles for a Fermi gas in various temperature ranges are plotted in Figure 2.4 (a), showing the Gaussian shape of a cloud with $T/T_F = 1$, and the inverted parabolic shape of (2.24) for a cloud with $T/T_F = 0.01$. Temperatures in between these limits are characterized by a distribution which looks
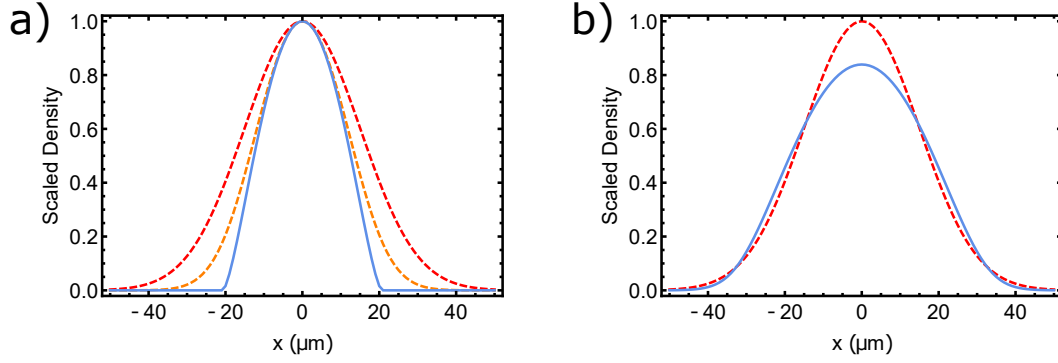
Figure 2.4: 1D density profiles of a Fermi gas in a harmonic trap. a) Profiles of $n(x)/n(0)$ rescaled by the peak density, for $T/T_F = 0.01$ (blue), $T/T_F = 0.5$ (orange), and $T/T_F = 1$ (red). A harmonic trapping frequency $\omega_x/2\pi = 100\,\mathrm{Hz}$ is assumed, and with $N = 10^5$ atoms a typical peak density is $n(0) \approx 5 \times 10^{12}\,\mathrm{cm}^{-3}$. b) Difference in density profiles for a cold cloud in a strong trapping potential (blue, $T/T_F = 0.2$, $\omega_x/2\pi = 100\,\mathrm{Hz}$), and a hot cloud in a weak trapping potential (red, $T/T_F = 1$, $\omega_x/2\pi = 35\,\mathrm{Hz}$). The same atom number is assumed in both cases, and the densities are scaled by the peak density of the hotter cloud. Only the hotter cloud can be well fit by a Gaussian distribution. For the parameters used in this plot the wings of the two density distributions are similar, but the colder cloud shows the flatter central density which is characteristic of the low temperature behaviour of (2.22).

approximately Gaussian at the edges of the cloud, but with a reduced density in the centre, as shown in Figure 2.4 (b). Because the differences between a Gaussian profile and the true form of (2.22) are so slight, it can be very difficult to distinguish between the two when examining the shape of a given cloud. In particular, if the density profile is measured with a finite spatial resolution, subtle differences in the shape of the cloud edges may be completely obscured. For this reason it is very difficult to use in-situ measurements of the cloud density in order to extract ζ.

2.2.2   Identifying a Degenerate Fermi Gas

In order to characterize the value of T/T_F in an experimental setting, the difference shapes of the density distributions at high Temperature and low Temperature can be leveraged. If one can measure the in-trap density profile of a cloud with enough precision, one can infer ζ from the shape of the distribution. An alternative method is to release the cloud from the trap and perform a fit based on the shape of the momentum distribution after a fixed time-of-flight (TOF). Measuring the momentum distribution is typically easier, since the larger expanded cloud is easier to image faithfully with a low resolution imaging system. Further, trapped degenerate gases can reach high enough density that standard absorption imaging techniques become unreliable, a problem which is mitigated by allowing the cloud to expand to a lower density. For these reasons thermometry of harmonically-trapped Fermi clouds is often performed with absorption images of the cloud after a fixed TOF. This can give accurate results for a Fermi cloud in which any other effects such as interparticle interactions do not have a large effect on the density distribution.

To determine the profile of a cloud with the initial density distribution (2.22), after release and expansion for a time t, we use the result that TOF expansion is equivalent to a rescaling of the spatial coordinates $x_i \to x_i/\sqrt{1 + \omega_i^2 t^2}$ with simultaneous rescaling of the density from $n \to n/(\prod_i \sqrt{1 + \omega_i^2 t^2})$

[48]. Thus the distribution of a Fermi gas in a TOF image will be

$$n_{\text{TOF}}(\vec{r}, t) = -\frac{1}{\Lambda_{\text{th}}^3} \frac{1}{\sqrt{1 + \omega_x^2 t^2}\sqrt{1 + \omega_y^2 t^2}\sqrt{1 + \omega_z^2 t^2}} \text{Li}_{3/2}\left(-\zeta e^{-\frac{x^2}{\sigma_x(t)^2} - \frac{y^2}{\sigma_y(t)^2} - \frac{z^2}{\sigma_z(t)^2}}\right), \qquad (2.25)$$

where the size of the cloud after time t is set by the time-dependent widths $\sigma_i(t)$, defined by

$$\sigma_i(t)^2 = \frac{2k_B T}{m\omega_i^2}\left(1 + \omega_i^2 t^2\right). \qquad (2.26)$$

To connect the TOF density distribution (2.25) with the image that will be obtained by illuminating the cloud with a resonant laser beam, we must make a connection between the cloud shape and the two-dimensional shadow that it will cast. To quantify the absorption of a resonant probe laser beam propagating along x, a comparison of the beam intensity pattern with $(I_a(y, z))$ and without $(I_0(y, z))$ atoms is used to determine the optical depth (OD) of the cloud using

$$\ln\left(\frac{I_0(y, z)}{I_a(y, z)}\right) = \text{OD}(y, z). \qquad (2.27)$$

The optical depth, the measure of the attenuation of the resonant probe beam with wavelength λ, can also be calculated from the cloud density (2.25) by integrating the atom density over the direction of beam propagation and scaling by the resonant absorption cross section $\sigma_{\text{sc}} = 3\lambda^2/2\pi$:

$$\text{OD}(y, z, t) = \sigma_{\text{sc}} \int n_{\text{TOF}}(x, y, z, t)\, dx. \qquad (2.28)$$

For the specific form of the distribution given in (2.25), this integration is easily carried out using the property of polylogarithm functions

$$\int_{-\infty}^{\infty} \text{Li}_n\left(ze^{-x^2}\right) dx = \sqrt{\pi}\, \text{Li}_{n+1/2}(z), \qquad (2.29)$$

so that the expected OD of a Fermi cloud after TOF is given by

$$\text{OD}(y, z, t) = -\frac{\sigma_{\text{sc}}(k_B T)^2 m}{2\pi\hbar^3 \omega_x} \frac{1}{\sqrt{1 + \omega_y^2 t^2}\sqrt{1 + \omega_z^2 t^2}} \text{Li}_2\left(-\zeta e^{-\frac{y^2}{\sigma_y(t)^2} - \frac{z^2}{\sigma_z(t)^2}}\right). \qquad (2.30)$$

This theoretical form can then be directly compared with the experimentally obtained OD from an absorption image.

2.2.3 Fitting Functions

In the same way that the density distribution of a Fermi cloud has a Gaussian shape at high temperature, the TOF expression for OD (2.30) simplifies when $T/T_F \gg 1$ to a Gaussian distribution. The standard deviation of the Gaussian distribution is related to the width $\sigma_i(t)$ of the density distribution by $\sigma_{\text{Gauss}} = \sigma_i(t)/\sqrt{2}$. For large expansion times (such that $\omega_i t \gg 1$), the Gaussian width becomes independent of the trap parameters and is instead only determined by the particle mass m and the temperature. Thus for hot clouds, the temperature T can be determined easily by fitting the TOF OD with a Gaussian
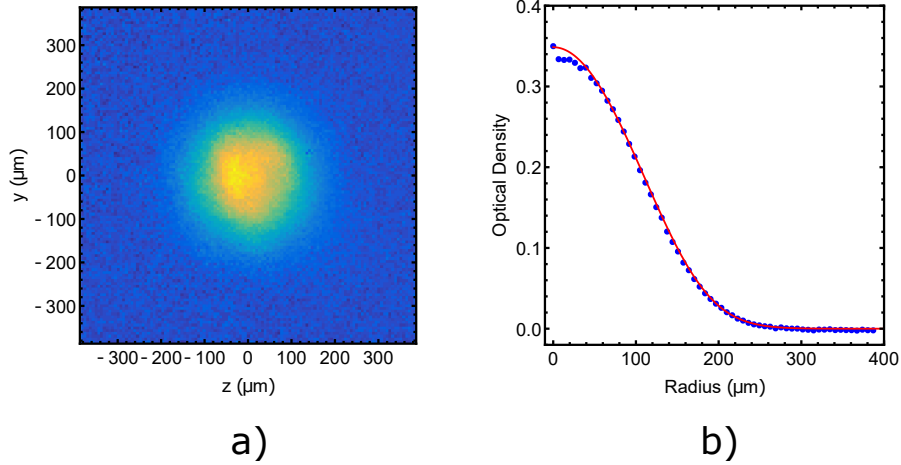
Figure 2.5: Determining T/T_F for a Fermi gas. a) A TOF OD image of a degenerate Fermi gas, released from a harmonic trap. Images of 5 separate clouds are averaged in order to improve the signal to noise. Some distortion of the radially symmetric structure is visible, and is caused by interference fringes in the absorption imaging beams. b) Radially averaged OD for the previous image (blue points) with a best-fit density distribution of the form (2.33). T/T_F for the cloud is determined to be 0.28. A reduction in scattering cross section by a factor of 3 has been assumed in the fit, as noted in the text.

function, finding the standard deviation σ_{Gauss}, and calculating

$$T_{\text{Gauss}} = \frac{m}{k_B} \left(\frac{\sigma_{\text{Gauss}}}{t} \right)^2 . \tag{2.31}$$

As the cloud temperature is reduced below $T/T_F = 1$, Gaussian fits become progressively worse, signalling the change in the TOF density distribution due to the importance of quantum statistics. In order to determine the cloud temperature, the full form of (2.30) must be used for fitting.

If the characteristics of the trap are well known, then the unknown parameters in (2.30) are T, N, and ζ. Two of these parameters will determine the fit, with the third constrained by the other two by the relationship (2.19). Since the number of atoms N can in principle be obtained trivially from an image by integration:

$$N_{\text{exp}} = \frac{1}{\sigma_{\text{sc}}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{OD}(y, z) \, dy \, dz \,, \tag{2.32}$$

it is possible to perform a one-parameter fit to determine ζ. Due to the relationship (2.21), the fugacity varies over several orders of magnitude as T/T_F is reduced towards the lowest attainable temperatures. It is better in this case to use $Q \equiv \ln(\zeta)$ as the fitting parameter to improve the fit convergence.

Experimentally, an average of several images taken for identically prepared Fermi clouds can be used to reduce the signal to noise and improve the accuracy of the fit. Further, if the trap frequencies are nearly equal in the image plane, or if the TOF is large enough that (2.26) is nearly independent of ω, then a radial average of the image may be performed in order to improve the fit accuracy. Some residual distortion is present in our experimental images due to fringes in our imaging beam, and this distortion is mitigated slightly by fitting the radial average. The fitting function required for a radially averaged

image, expressed in terms of the single fit parameter Q, is

$$n(r) = \frac{A}{Z(Q)^2} \left(-\text{Li}_2 \left(-e^{Q - \frac{r^2}{w^2} Z(Q)} \right) \right) , \tag{2.33}$$

with the function $Z(Q) = (-\text{Li}_3 \left(-e^Q \right))^{1/3}$ containing additional dependence on the fugacity, and the parameters A and w fully constrained by the values of m, $\bar{\omega}$, t, and N according to the expressions

$$A = \frac{\sigma_{\text{sc}} m \bar{\omega}^2 N^{2/3}}{2\pi \hbar \omega_x \sqrt{1 + \omega_y^2 t^2} \sqrt{1 + \omega_z^2 t^2}} , \tag{2.34}$$

and

$$w = \left(\frac{2 \hbar \bar{\omega} N^{1/3} t^2}{m} \right)^{1/2} . \tag{2.35}$$

An experimental image showing the OD of a degenerate Fermi gas is shown in Figure 2.5 (a). The cloud was released from a trap with $\omega_i = 2\pi (44, 44, 312)$ Hz, which are measured by observing the oscillation of the cloud in the trap after an initial displacement. Integrating the image according to (2.32) gives $N = 6.6 \times 10^4$, and so the Fermi temperature is $T_F = 298$ nK. Using these experimental values to fit the one parameter function (2.33) does not give good agreement. This occurs because the lone fit parameter Q simultaneously determines the width and the shape of the cloud, and cannot give good agreement with both unless all physical terms in the parameters A and w are correct. The improper fitting indicates that there may be an error in the measurement of ω_i, deviation from the calculated N, or imperfections in the imaging system. Of these three possibilities, N is most likely to be in error since saturation effects, imperfect laser linewidth, or poor polarization of the imaging beam can all lead to a smaller absorption cross section than the predicted σ_{sc}. Allowing the number of atoms in the cloud to vary in the fit, we find an increase of N by a factor of 3 leads to a good fit with the expected shape of a Fermi cloud. The resulting fit is shown in Figure 2.5 (b), with $T/T_F = 0.28$.

If we trust the quality with which we fit the Fermi momentum distribution, the fit results can be taken as a calibration of the imaging system — we can proceed by assuming all absorption images are obtained with an effective cross section that is $\sigma_{\text{eff}} = \sigma_{\text{sc}}/3$. To improve the confidence with which we determine T/T_F from absorption images, it is possible to independently calibrate the relationship between atom number N and the absorption we observe experimentally. This can be done by identifying a parameter of an ultracold cloud which is number-dependent, but which can be measured with an uncalibrated imaging system. To calibrate an imaging system for a bosonic species one can measure the critical temperature T_c for Bose-Einstein condensation, or the Thomas-Fermi radius of a trapped BEC [49]. In our experiment with an entirely separate fluorescence imaging system with single-atom sensitivity, it should be possible to calibrate our absorption imaging system by comparing images of the same cloud taken with both modalities.

2.3   Optical Lattices

2.3.1   Creating an Optical Lattice

To create a periodic potential energy landscape for a quantum gas, one can interfere multiple laser beams to create an intensity pattern which has fine features on the order of the wavelength λ. For two
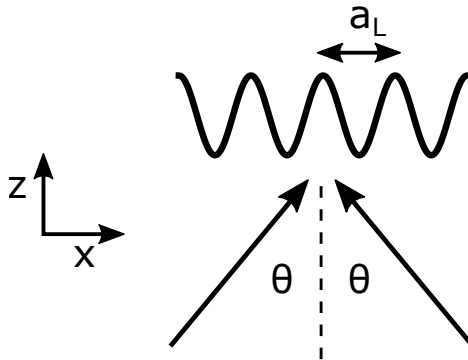
Figure 2.6: Schematic of two beam crossing to produce an interference pattern along the x direction.

laser beams with the same intensity I_0, crossing each other at an angle 2θ as depicted in Figure 2.6, an interference pattern is formed along the x direction with periodicity

$$a_L = \frac{\lambda}{2\sin(\theta)} \; . \tag{2.36}$$

If the frequencies of the two laser beams are the same, and the phase difference is constant, then this interference pattern is stationary in space. The electric field of the lasers still oscillates at the optical frequency ω, but the overall intensity varies as $I(x) = 4I_0 \cos(k_L x/2)^2$, where $k_L = 2\pi/a_L$ is the lattice wavevector. This intensity pattern together with the dipole potential (2.1) forms a periodic potential for ultracold atoms referred to as an optical lattice. For a given beam geometry, the lattice wavevector k_L is directly related to the optical wavevector $k_O = 2\pi/\lambda$ by (2.36).

To ensure that the relative phase of the two interfering beams is stable, it is common experimentally to form an optical lattice by retro-reflecting a beam and interfering it with itself. This references the phase difference between the two interfering beams to the mechanical position of the retro mirror, so that phase stability of the interference pattern can be enforced by mechanical stability of the mirror. When the beam is retro-reflected $\theta = \pi/2$, $k_L = 2k_O$, and the lattice period is simply $a_L = \lambda/2$. We form our optical lattice in this way, with $\lambda = 1054\,\text{nm}$, giving a lattice of periodicity $a_L = 527\,\text{nm}$. The optical lattice potential is $V_L(x) = V_L \cos(k_L x/2)^2$, where the lattice depth V_L is determined by the dipole potential (2.1), taking into account the factor of 4 increase in intensity over a single Gaussian beam. The behaviour of particles moving through an optical lattice is characterized by the ratio between this depth V_L and the lattice recoil energy

$$E_R = \frac{\hbar^2 k_L^2}{8m} \; . \tag{2.37}$$

For our retro-reflected lattice potential, the recoil energy of the lattice is the same as the optical recoil $\hbar^2 k_O^2/2m$ associated with absorption or emission of a photon at the lattice wavelength. For ^{40}K atoms in our lattice with $a_L = 527\,\text{nm}$, the recoil energy is $E_R = k_B(215\,\text{nK}) = h(4.49\,\text{kHz})$.

By combining multiple optical lattice potentials along different axes, almost arbitrary periodic potential landscapes can be realized. A particularly simple arrangement uses separate retro-reflected beams to create three optical lattice potentials aligned along perpendicular axes x, y, and z. If the electric fields forming the three optical lattice potentials are oriented perpendicular to one another, they will

not interfere with one another, creating a simple cubic lattice geometry with the separable potential

$$V(x, y, z) = V_{L,x} \cos(k_L x/2)^2 + V_{L,y} \cos(k_L y/2)^2 + V_{L,z} \cos(k_L z/2)^2. \qquad (2.38)$$

Even if the polarizations of the three electric fields are not perfectly perpendicular, a large difference in frequency between the lattice beams will cause any residual cross-interference to oscillate more quickly than the atoms can dynamically respond. In this case the atoms see only the time-averaged intensity of the oscillating interference, which maintains the separable form above. In our experiment we thus prevent any cross-interference between lattice beams by offsetting the optical frequency of the three different axes by $> 130\,\mathrm{MHz}$.

The periodic potential of an optical lattice can generate interesting physical behaviour similar to the delocalized movement of electrons among a lattice of ions in a typical metal or semiconductor. To emulate this kind of physical behaviour calls for lattice depths on the order of 2-20 E_R. This places only modest requirements on the required beam intensity, allowing a good deal of freedom over the focussed size of the beams used to form the optical lattice. However, for our experiment we want to obtain fluorescence images of atoms trapped in individual sites of the lattice, which requires us to reach very large lattice depths of around $1000\,E_R$. To reach the necessary intensity we focus our optical lattice beams to Gaussian waists of $w_x = w_y = 60\,\mu\mathrm{m}$ and $w_z = 85\,\mu\mathrm{m}$. With these waists, lattice beam powers $(P_x, P_y, P_z) = (1, 1, 2)\,\mathrm{W}$ give each beam a total standing wave intensity of $70\,\mathrm{kW\,cm^{-2}}$ which creates a lattice depth of $450\,E_R$. With the total available laser power we can reach powers $(P_x, P_y, P_z) = (2.5, 2.5, 5)\,\mathrm{W}$ leading to a total intensity close to $500\,\mathrm{kW\,cm^{-2}}$ and lattice depths $> 1000\,E_R$ in each axis.

2.3.2 Bandstructure in a Periodic Potential

Replacing the harmonic confinement with the periodic potential of an optical lattice leads very different behaviour for the quantum gas. Due to the similarity between the optical lattice potential and the periodic crystal structure that electrons move through in real materials, the behaviour of the quantum gas can be easily understood in terms of textbook solid state physics. The Schrödinger equation of a particle of mass m in a 1D lattice of periodicity a_L and depth V_L is given by

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V_L \cos^2(k_L x/2)\right)\phi(x) = E\phi(x), \qquad (2.39)$$

The solutions $\phi(x)$ can be labelled by two quantum numbers q and n. The quasimomentum q is the conserved quantity arising due to the discrete translational symmetry of the lattice, in the same way that for a free particle the linear momentum p is conserved due to the continuous translational symmetry of free space. Due to the discrete symmetry, q is only unique up to multiples of the lattice recoil momentum $\hbar k_L$, and so it is customary to only consider the range $|q| < \hbar k_L/2$, referred to as the first Brillouin zone [50]. Within this range there are infinitely many solutions of (2.39) for each value of q, each identified by a particular band index n. The energy solutions in each band n are smooth functions of q, and the particular form of the energy solutions $E_n(q)$ is referred to as the bandstructure. According to Bloch's theorem, the eigenstates $\phi_q^{(n)}(x)$ for each n and q will take on the form

$$\phi_q^{(n)}(x) = e^{iqx/\hbar} u_q^{(n)}(x), \qquad (2.40)$$

where $e^{iqx/\hbar}$ is a plane wave travelling with quasimomentum q, and $u_q^{(n)}(x)$ is an arbitrary function which shares the translational periodicity of the lattice.

By solving the Schrödinger equation it is possible to determine the Bloch eigenstates (2.40) and the bandstructure $E_n(q)$. In order to understand properties of the system occurring on single lattice sites, it is helpful to construct solutions of (2.39) which are localized near a single lattice site. This is accomplished by taking the Fourier transform of Bloch wavefunctions

$$w_n(x - x_j) = \int_{k_L/2}^{k_L/2} e^{-iqx_j/\hbar} \phi_q^{(n)}(x) \, dq \,, \tag{2.41}$$

to generate the Wannier functions $w_n(x - x_j)$ each localized on a site $x_j = ja_L$. These Wannier functions are useful when considering a local interaction between two particles such as the contact interaction present in a gas of cold atoms.

2.3.3 Bandstructure Calculation

For a given lattice potential, a great deal of understanding is gained by examining the bandstructure $E_n(q)$. This is easily done numerically by setting up a differential equation for the periodic function $u_q^{(n)}(x)$ and solving it in Fourier space. By substituting (2.40) into (2.39), we obtain

$$\left(\frac{(\hat{p} + q)^2}{2m} + V(x) \right) u_q^{(n)}(x) = E_n(q) u_q^{(n)}(x) \,. \tag{2.42}$$

Since the function $u_q^{(n)}$ and the potential $V(x)$ are both periodic in a_L, they can be expanded as discrete Fourier sums

$$u_q^{(n)} = \sum_j c_j^{n,q} e^{ijk_L x} \,, \tag{2.43}$$

and

$$V(x) = \sum_{j'} V_{j'} e^{ij'k_L x} \,. \tag{2.44}$$

In a 1D optical lattice where the potential can be written $V(x) = V_L \sin^2(k_L x/2)$ [1], the Fourier components V_j' in (2.44) are trivially $V_0 = V_L/2$, $V_{\pm 1} = -V_L/4$, and $V_j' = 0$ for $|j'| > 1$. Furthermore, the kinetic term $(\hat{p} + q)^2/2m$ or (2.42) can easily be applied to each term in the expansion (2.43), yielding

$$\frac{(\hat{p} + q)^2}{2m} u_q^{(n)} = \sum_j \frac{(\hbar j k_L + q)^2}{2m} c_j^{n,q} e^{ijk_L x} \,. \tag{2.45}$$

Combining these results, we can express the differential equation (2.42) in Fourier space as

$$\sum_j \left[\left(\frac{(\hbar j k_L + q)^2}{2m} + \frac{V_L}{2} \right) c_j^{n,q} e^{ijk_L x} - \frac{V_L}{4} \left(e^{i(j+1)k_L x} + e^{i(j-1)k_L x} \right) c_j^{n,q} \right] = E(q) \sum_j c_j^{n,q} e^{ijk_L x} \,. \tag{2.46}$$

---

[1] Here $\sin^2(k_L x/2)$ is used rather than $\cos^2(k_L x/2)$ so that zero energy is conveniently referenced to the bottom of each lattice well.
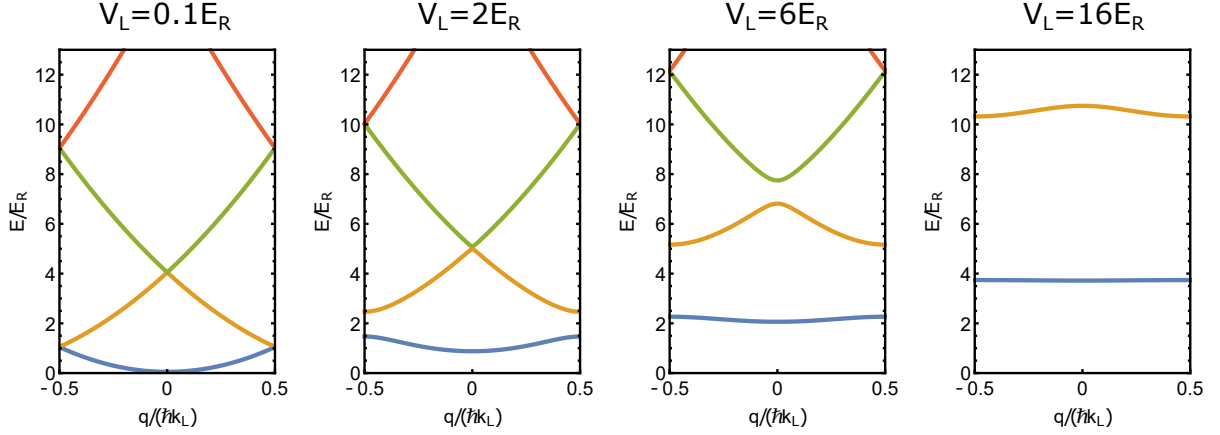
Figure 2.7: Structure of the energy bands for different lattice depths V_L. The ground band and first three excited bands are plotted in blue, orange, green, and red in increasing order. For very low depths the bandstructure is dispersion relation for a free particle, folded back into $|q| < 0.5\hbar k_L$ by the Bragg condition. In this case the energy width ΔE of the ground band is equal to E_R. As the depth increases energy gaps open between the bands. The curvature of each band decreases with increasing depth, corresponding to a decreasing ΔE and an increasing effective mass of the particles. At very high depths the bands are almost completely flat and independent of q (as shown for $V_L = 16\,E_R$) , so that $\Delta E \to 0$ and only the band index n is needed to specify the energy of the particles.

Collecting terms with the same oscillatory mode $e^{ijk_L x}$ we end up a set of coupled linear equations for the coefficients $c_j^{n,q}$:

$$\left(\frac{(\hbar j k_L + q)^2}{2m} + \frac{V_L}{2} \right) c_j^{n,q} - \frac{V_L}{4}(c_{j-1}^{n,q} + c_{j+1}^{n,q}) = E(q)c_j^{n,q} \,. \tag{2.47}$$

Using only a finite number of terms j in the expansion (2.43), the system (2.47) can be written as a matrix equation for any quasimomentum q and band index n. This matrix equation can be easily solved numerically, and for lattice depths $V_L < 40\,E_R$ good results are obtained when using only a 11×11 matrix corresponding to $|j| \leq 5$ [2]. Figure 2.7 shows the results of the above bandstructure calculation for several lattice depths. As the depth increases, the bands become flatter and the energy gaps between adjacent bands increase.

In order to solve for the Bloch functions, the coefficients $c_j^{n,q}$ need only be plugged back into the Fourier expansion (2.43) of $u_q^{(n)}$, and then multiplied with a plane wave as in (2.40). However, the full shape of the Bloch functions is not typically very illuminating. The general procedure outlined here to solve for the bandstructure of a simple optical lattice potential is easily generalized to more complicated lattice geometries, and simply adds more terms $V_{j'}$ to the Fourier transform of the potential in (2.44).

2.3.4 Mathieu Equation

For realistic materials, the potential energy landscape $V(x)$ will be periodic, but is unlikely to be as simple as the purely cosinusoidal potential generated by an optical lattice. Some understanding of a general potential $V(x)$ can be gained by analytic methods if we assume the tight-binding limit where wavefunctions are assumed to be localized on each lattice site with minimal overlap, or the nearly-free

---

[2]However, if this calculation is used to determine the energy spectrum in very deep lattices such as the $V_L \approx 1000\,E_R$ achieved experimentally, it is necessary to extend the sum up to $|j| \leq 10$
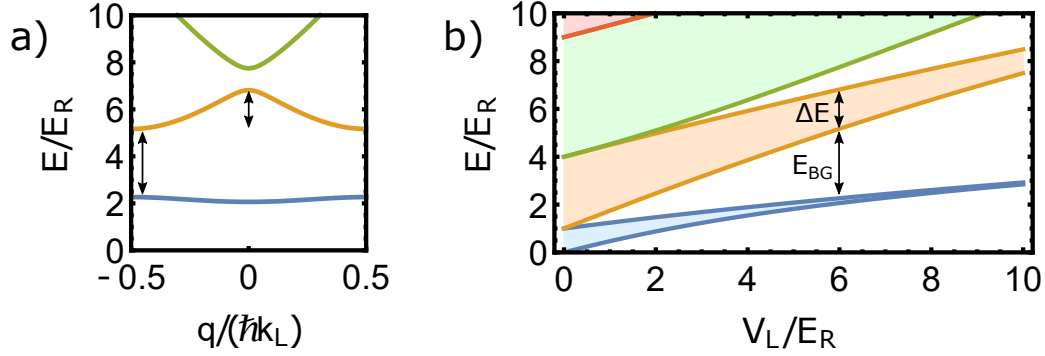
Figure 2.8: Energy bands calculated using (a) numerical bandstructure and (b) the Mathieu equation. Numerical bandstructure is evaluated for a lattice depth of 6 E_R. Bandstructure calculations provide the full structure of $E(q)$, while the Mathieu coefficients provide the range of energies for which a solution exists.

electron limit where wavefunctions are assumed to be plane-waves which are perturbed by the presence of the lattice. However, for the specific case of the cosine potential generated by an optical lattice, an alternative approach is possible, since the Schrödinger equation (2.39) can be brought into the form of the well studied Mathieu equation [51]

$$\frac{\partial^2 \phi}{\partial \tilde{x}^2} + (a - 2Q \cos(2\tilde{x})) \, \phi = 0 \,, \tag{2.48}$$

by substituting the dimensionless parameters a and Q which are related to the lattice depth by

$$a = \frac{(E - V_L/2)}{E_R} \,, \tag{2.49}$$

and

$$Q = \frac{V_L}{4E_R} \,. \tag{2.50}$$

The band energies of the cosine lattice are related to the values of a for which real solutions of the (2.48) exist. These values of a are separated into bands indexed by the integer n, just as in the bandstructure calculated previously. The lowest and highest values of a in band n are specified by the Mathieu characteristics $a_n(Q)$ and $b_{n+1}(Q)$, which can be found in tables or called by built-in functions of programs such as Wolfram Mathematica. The minimum and maximum energies in the nth band of a lattice of depth V_L can then be directly calculated based on these characteristics:

$$E_{n,\min} = E_R \left(a_n(V_L/4E_R) + \frac{V_L}{2E_R} \right) \,, \tag{2.51}$$

and

$$E_{n,\max} = E_R \left(b_{n+1}(V_L/4E_R) + \frac{V_L}{2E_R} \right) \,. \tag{2.52}$$

This also leads to a simple expression for the bandwidth ΔE_n of a given band n:

$$\Delta E_n = E_R \left(b_{n+1}(V_L/4E_R) - a_n(V_L/4E_R) \right) \,. \tag{2.53}$$
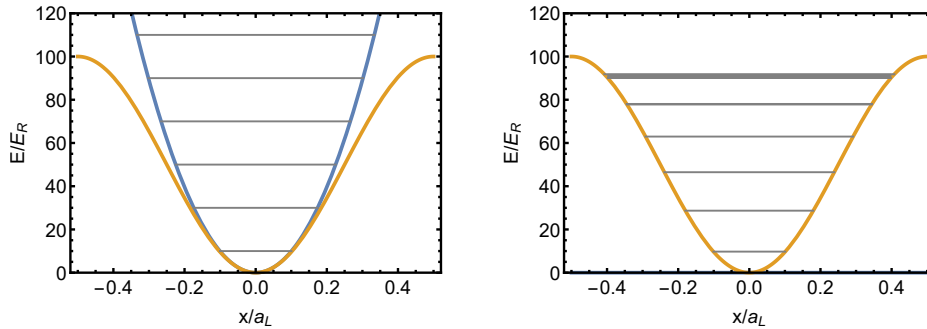
Figure 2.9: Harmonic approximation for a deep optical lattice site with $V_L = 100\,E_R$. On the left the lattice potential (orange) is overlaid with the harmonic approximation (blue), showing good agreement near the bottom of the well. The first six quantized energy levels $(n + \frac{1}{2})\hbar\omega$ of the harmonic potential are shown. On the right, the six energy bands of the lattice potential are shown, calculated using the Mathieu equation. The bandwidth increases for bands closer to the top of the lattice potential.

The energy bands calculated using the Mathieu characteristics are shown in Figure 2.8, beside a plot of the numerical bandstructure $E(q)$ generated with the methods of the previous section.

2.3.5   Harmonic Approximation for Deep Lattices

At large lattice depths, the lowest bands flatten into a single energy level which is independent of q. These flat energy bands also become evenly spaced from one another. This is a result of the atomic wavefunctions becoming more and more confined near the minimum of the sinusoidal potential, which is well approximated by a harmonic trap which features such discretely spaced energy levels. At large depths, the lattice potential can be reasonably well approximated by an array of harmonic traps, each of which may contain one or more particles, and between which tunneling is negligible. By expanding the potential $V(x) = V_L \cos^2(k_L x/2)$ around the centre of each lattice site $x = x_c$, we find $V(x) \approx \frac{1}{2}m\omega^2(x - x_c)^2$, with the harmonic trapping frequency ω is given by

$$\omega = \frac{E_R}{\hbar}\sqrt{4(V_L/E_R)}\,. \tag{2.54}$$

Due to the tight confinement in an optical lattice, only a few discrete energy levels exist within the finite trap depth. This is depicted in Figure 2.10, where there are only six bound states in the lattice with $V_L = 100\,E_R$. The higher bound states will sample enough of the lattice well that the harmonic approximation does not accurately predict their energies.

The energy gap between the ground and first excited bands is a measurable quantity which can be used to calibrate the depth of the optical lattice. Calculating this gap using (2.54) provides a good estimate, but will always overestimate the true energy gap. Using (2.54) to determine the gaps between higher bands will lead to increasingly large errors. To ensure that the correct energy gaps are chosen at all energies, the solutions to the Mathieu equation can be used rather than the simple harmonic form. Since the two methods agree closely at low lattice depths, we can also refer to the effective harmonic trapping frequency on each site in terms of the calculated band energies $\omega_{\mathrm{Mathieu}} = (E_1 - E_0)/\hbar$. At our highest accessible lattice depths of $1200\,E_R$, the trap frequency is $\omega_{\mathrm{Mathieu}}/2\pi \approx 300\,\mathrm{kHz}$.
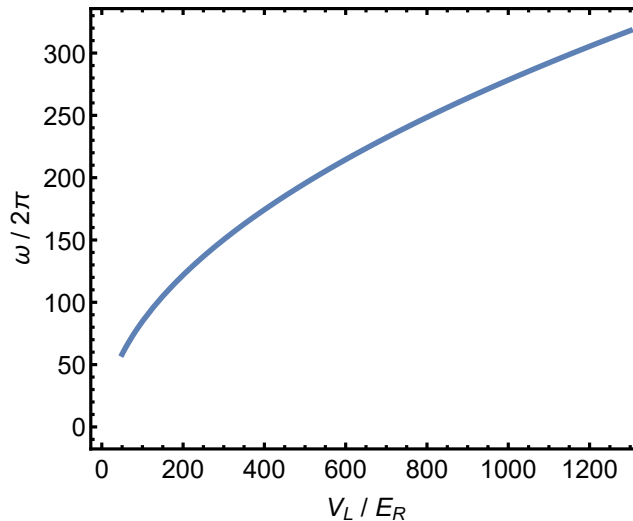
Figure 2.10: Transition frequency from the ground band to the first excited band, calculated using the Mathieu equation. The harmonic approximation (2.54) tends to overestimate this frequency by about 5 kHz.

2.3.6   Higher-Dimensional Lattices

In the above discussion we have assumed a 1D lattice potential $V(x) = V_L \cos(k_L x/2)^2$. In our experiments we generate a 3D optical lattice potential using three separate retro-reflected beams. This results in the potential $V(x, y, z) = V_{L,x} \cos(k_L x/2)^2 + V_{L,y} \cos(k_L y/2)^2 + V_{L,z} \cos(k_L z/2)^2$. The lattice depths $V_{L,i}$ for each axis need not necessarily be the same. As long as the three lattice axes are perpendicular, and the beams do not cross-interfere, this potential is fully separable in the three spatial directions and can still be understood in terms of the 1D calculations presented above. This simplification is true only because of the simple cubic lattice geometry that we implement, and will not generally be true for more complicated lattices.

Each individual lattice direction i will have a bandwidth of the ground band $W_{0,i}$ and a bandgap $E_{BG,i}$. A particle in the lowest band of each lattice is then restricted to a range of energies

$$W_{3D} = \sum_i W_{0,i} \, . \tag{2.55}$$

Two particles in the lowest band can have a total energy up to $2W_{3D}$, which may be redistributed among the particles during a collision. If $2W_{3D} > W_{0,i} + E_{BG,i}$ for some lattice axis i, then it is possible for a collision between two particles to result in excitation to the first excited band of this axis. Preventing this type of collisional population of excited bands requires that the lattice depth be increased, reducing the bandwidth and thus the available collisional energy. If all three lattice depths are equal and $V_{L,i} > 3.74 \, E_R$ scattering to exited bands is energetically forbidden, so that particles can be considered fully confined to the ground band of the system.

Our experiment is designed to obtain fluorescence images of the density distribution of atoms in a two-dimensional lattice. To realize such a two-dimensional system, we can use a very large lattice depth in the vertical direction $V_{L,z} > 50 \, E_R$, suppressing any tunneling along this axis. This leads to a vanishing bandwidth in the vertical direction $W_{0,z} \to 0$, so that the two-dimensional bandgap is
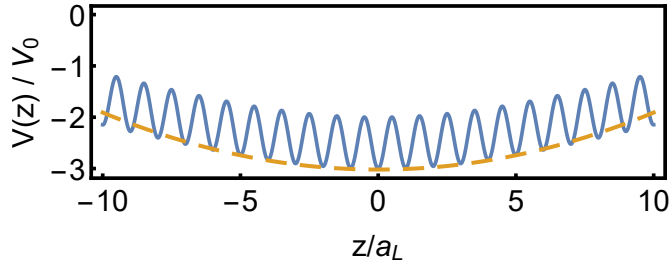
Figure 2.11: Lattice potential with additional harmonic trapping. The potential $V(x)$ (blue) is plotted along the x direction, in the presence of additional lattice beams propagating along the y and z directions. All beams are taken to have relatively small Gaussian waists of $10\,\mu$m to exaggerate the harmonic confinement. The barrier height between adjacent lattice sites is approximately independent of x because the Rayleigh length of the lattice beams is hundreds of times larger than a_L and so has little effect on this scale. The effective harmonic confinement of (2.59) is shown by the dashed orange line.

determined only by the horizontal lattice beams:

$$W_{2D} = W_{0,x} + W_{0,y}\,. \tag{2.56}$$

With such asymmetric lattice depths, the effective energy bandwidth is dominated by the remaining two horizontal lattices. The available scattering energy for two colliding particles is reduced and the horizontal lattice beams need only be stronger than $2.85\,E_R$ to forbid excitation to excited bands.

2.3.7 Harmonic Confinement for Gaussian Beams

The periodic potential landscape of an optical lattice arises due to the modulated intensity generated by interference between multiple beams. In addition to the this periodic intensity pattern which varies on the scale of λ, there will be a larger scale variation in intensity due to the profile of the beams used to produce the interference. If Gaussian beams are used to produce the optical lattice, the length scale of this larger intensity variation will be the beam waist w_0. For a one-dimensional optical lattice which is formed by interfering a Gaussian beam propagating along z with its retro-reflection, the overall potential created has the form

$$V(r,z) = -V_0 \left(\frac{w_0}{w(z)}\right)^2 e^{-2(x^2+y^2)/w(z)^2} \cos^2(k_L z/2)\,, \tag{2.57}$$

where w_0 and $w(z)$ are the usual waist parameters for the Gaussian beam. Along the z direction, this potential is the desired sinusoidal shape, with a slowly varying envelope due to the factor $(w_0/w(z))^2$ which accounts for the focussing of the beam. Since $w(z)$ varies very little on the typical experimental length scale, this factor is typically neglected and the potential can be assumed to have the form $V(z) = V_0 \cos^2(k_L z/2)$ used in the calculations of the previous sections. Along the r direction, the focussing of the beam forms a smooth Gaussian trapping potential, which can be approximated by a harmonic trap as considered previously in (2.9).

If three perpendicular Gaussian beam lattices are overlapped to produce a 3D lattice potential (with appropriate polarization and frequency to prevent cross-interference), then the overall potential is the sum of three contributions of the form (2.57). Now, when moving along the z direction, a particle will see the $\cos^2(k_L z/2)$ potential due to the beam propagating along z, as well as the Gaussian intensity

variation of the two beams propagating along x and y. Near the centre of the focussed beams the exponential factors can be expanded to first order, and the overall lattice potential for three lattice beams with individual depths V_x, V_y, and V_z takes the approximate form

$$V(x, y, z) = -V_x \left(1 - \frac{2(y^2 + z^2)}{w_{0,x}^2}\right) \cos^2(k_L x/2) - V_y \left(1 - \frac{2(x^2 + z^2)}{w_{0,y}^2}\right) \cos^2(k_L y/2)$$

$$- V_z \left(1 - \frac{2(x^2 + y^2)}{w_{0,z}^2}\right) \cos^2(k_L z/2). \quad (2.58)$$

The three terms in parentheses in (2.58) describes two important features of a 3D optical lattice potential — harmonic confinement and spatially dependent lattice depth. To visualize the harmonic confinement, consider the positions of the lattice sites $x, y, z = na_L$ so all the cosine factors are equal to one. When moving from one of these sites to the next there is still a difference in energy. This site-to-site energy difference takes the form of a harmonic trap $V(x, y, z) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)$ with trapping frequency along x given by

$$\omega_x^2 = \frac{4}{m}\left(\frac{V_y}{w_{0,y}^2} + \frac{V_z}{w_{0,z}^2}\right). \quad (2.59)$$

The other frequencies ω_y and ω_z defined in a similar way by cyclic permutation of the indices. Combining this slow energy shift between lattice sites with the lattice potential which creates barriers between lattice sites suggests a simple form for the optical lattice potential:

$$V(x, y, z) \approx \frac{1}{2}m\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right) + V_x \cos^2(k_L x/2) + V_y \cos^2(k_L y/2) + V_z \cos^2(k_L z/2), \quad (2.60)$$

which is just the sum of a 3D uniform lattice potential and 3D harmonic trap.

The potential (2.60) neglects one important feature of the real potential (2.58), which is the spatial variation of the optical lattice depth. The amplitude of the $\cos^2(k_L x/2)$ standing wave along x depends on the position $(y^2 + z^2)$, which just expresses the squared radial distance from the axis of the x lattice beam. Moving further out into the Gaussian profile of the x lattice beam leads to a lower local beam intensity and thus a smaller potential barrier for tunneling along the x direction. In a deep lattice where the individual lattice sites can be approximated by harmonic oscillator wells, we similarly expect the on-site trapping frequencies of the individual sites to be position dependent. Since the on-site harmonic trapping frequency (2.54) along the x direction is proportional to to the square root of the lattice depth V_x, it will be strongest in the centre of the x lattice beam and will be reduced away from the beam axis by a factor $(1 - 2(y^2 + z^2)/w_{0,x}^2)^{1/2}$. For our experimental lattice beam waists, sites within a radius of $10\,\mu$m from the trap centre will have all three lattice depths within 5% of the peak depth. This allows us to work with a few hundred lattice sites in a 2D plane for which the lattice depth can be assumed to be constant.

In a deep lattice, the spatial variation in on-site harmonic trapping leads to a variation in the ground state energy in each lattice well. Sites farther from the centre of the lattice have weaker confinement, and so a smaller zero-point energy offset $\hbar\omega/2$. This spatially varying energy offset effectively decreases the harmonic confinement of the lattice beams expressed in Equation (2.59). By using a harmonic approximation for $\cos^2(k_L x/2)$ to estimate the confinement frequency of a lattice well with barrier

height V_L, a corrected harmonic confinement frequency ω_x along the x direction can be obtained:

$$\omega_x^2 = \frac{4}{m} \left(\frac{V_y - \sqrt{V_y E_R}}{w_{0,y}^2} + \frac{V_z - \sqrt{V_z E_R}}{w_{0,z}^2} \right) , \tag{2.61}$$

which takes into account both the spatially varying AC Stark shift and the spatially varying lattice confinement. Equation (2.61) will not appropriately express the confinement of weak lattices outside of the tight-binding limit where lattice sites cannot be approximated by a simple harmonic trap.

The harmonic confinement effects and spatially dependent depths of an optical lattice can be minimized by choosing lattice beams with large Gaussian waists. Since the finite-sized cloud of ultracold atoms must still be confined in some way, such a large-beam optical lattice can be combined with an additional harmonically confinement from an optical dipole trap or magnetic trap. This arrangement provides a useful separation between control of the lattice depths and the confinement strength, and also minimizes any spatially dependent lattice depth effects. Some groups even implement blue-detuned beams to gain further control over the external lattice confinement [47, 52]. In our experiment choosing large lattice beams to minimize the spatial variation is not possible due to the large lattice depths we require for in-situ imaging. As a result, the effective harmonic trapping present in our experiment depends strongly on the lattice depths chosen.

When we increase the vertical optical lattice intensity to $V_z > 50\,E_R$ to realize a 2D lattice system with matched horizontal depths V_H, only the external confinement in the horizontal plane is relevant. This harmonic trapping is the combined trapping effect of the lattice beams given by the frequency ω_x in (2.59) and any additional horizontal trapping due to the optical dipole trap as given by the frequency ω_x in (2.9). For matched beam waists for the horizontal lattice beams ($w_{0,x} = w_{0,y} \equiv w_{Latt,H}$) and matched horizontal profiles for the dipole trap beam ($w_{0,A,y} = w_{0,B,x} \equiv w_{ODT,H}$), the combined frequency for external confinement in the lattice is given by

$$\omega_{\text{ext}}^2 = \frac{4}{m} \left(\frac{V_H - \sqrt{V_H E_R}}{w_{Latt,H}^2} + \frac{V_z - \sqrt{V_Z E_R}}{w_{Latt,z}^2} + \frac{V_{\text{ODT}}}{w_{\text{ODT},H}^2} \right) . \tag{2.62}$$

Our large horizontal dipole trap waist of $w_{\text{ODT},H} = 200\,\mu\text{m}$ means that ω_{ext} is dominated by the effect of the lattice beams. For typical parameters of $V_H = 5\,E_R$, $V_z = 50\,E_R$, and a $300\,\text{mW}$ optical trap providing $V_{\text{ODT}} \approx 16\,E_R$, the external confinement is $\omega_{\text{ext}}/2\pi = 180\,\text{Hz}$.

At first glance, the presence of the additional harmonic trapping suggested by the potential (2.60) seems to require re-evaluation of all discussion of the previous sections regarding periodic potentials. The eigenstates of the harmonically trapped optical lattice may not necessarily have the same bandstructure and Bloch states of the uniform lattice potential. Exploration of this problem has shown [53] that the eigenfunctions still have closed-form solutions, depending once more on Mathieu functions. In the limits typically explored with optical lattice systems, the solutions take on simple forms which are easy to understand. In a pure lattice potential with low depth, the dynamics are essentially those of a free particle with its mass m replaced by an effective mass m^*. Similarly, for a weak lattice potential superimposed with a harmonic trap, the lowest energy eigenstates look like harmonic oscillator wavefunctions with m replaced by m^*. These harmonic oscillator wavefunctions have a size given by the harmonic oscillator length $l_o = (\hbar/m\omega^*)^{1/2}$ with $\omega^*/\omega_{\text{ext}} = (m/m^*)^{1/2}$, and can still be delocalized over many lattice sites for small values of ω^*.
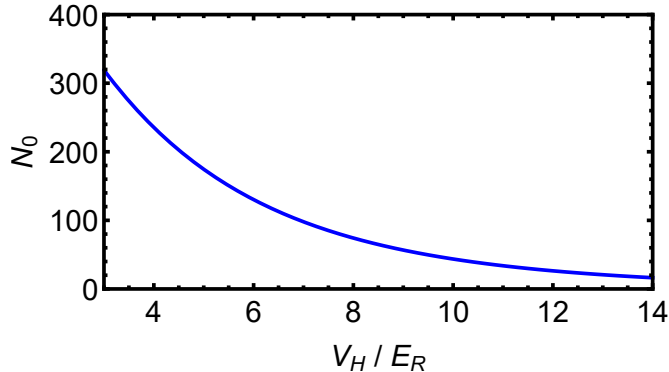
Figure 2.12: Characteristic atom number N_0 for an optical lattice realized with our experimental parameters. The vertical lattice depth is held at a depth of $50\,E_R$ to restrict hopping to two dimensions. With increasing lattice depth, the decreasing bandwidth and increasing confinement cause N_0 to drop rapidly, and for depths exceeding roughy $16\,E_R$ very few atoms can be considered delocalized.

Although the lowest energy eigenstates of the potential (2.60) are delocalized as expected for Bloch waves, this is not the case for all eigenstates. Increasing the number of fermions in the system leads to occupation of more and more harmonic-oscillator-like states with increasing energy, until the particles occupy the full range of energies one would expect for the ground band of the lattice potential alone. In a pure lattice potential, the presence of a bandgap would ensure that further increasing the number of fermions in the system would come at a large energy cost. However, in the harmonically trapped lattice there are still accessible eigenstates with energy only slightly higher than the top of the ground band. These states exist far away from the centre of the harmonic trap, and are tightly localized near individual lattice sites rather than spreading across the lattice [54]. The localization can be understood as result of the harmonic confinement creating an energy offset between different lattice sites — once the energy shift between the centre of the lattice and a site at distance x becomes larger than the bandwidth ΔE_0, tunneling between between the two regions will be cut off.

To quantify the effect of the harmonic potential in a given optical lattice, it is common to consider the length scale ξ for which

$$\tfrac{1}{2}m\omega_{\mathrm{ext}}^2(\xi a_L)^2 = \frac{W}{2}\,, \tag{2.63}$$

so that ξ is the distance from the origin (in lattice sites) for which the harmonic confinement raises the energy by half of the bandwidth W. For all lattice sites within the radius ξ the lowest energy states are not localized. In a 2D lattice the number N_0 of such sites is

$$N_0 = \pi\left(\frac{W}{m\omega_{\mathrm{ext}}^2 a_L^2}\right)\,. \tag{2.64}$$

When the number of atoms in the trapped lattice system is increased to $2N_0$, they must necessarily occupy energies greater than the bandgap in the centre of the lattice. Thus at zero temperature, $2N_0$ atoms may be placed into the lattice before localized states are filled. In our experiment $N_0 \approx 200$ for low 2D lattice depths, but drops rapidly with increasing lattice depth (due primarily to the shrinking bandwidth).

For large particle numbers N and deep lattice depths, the ratio $N_0/N \to 0$. In such a case, very few of the atoms in a gas exist in delocalized states, and the gas can be considered to be completely localized

to individual lattices sites. In this $zero$-tunneling limit it is easy to estimate the density of states and Fermi temperature of the 2D gas by ignoring the small fraction of particles in delocalized states [55]. With only one particle per lattice site, and the site j away from the origin associated with a harmonic energy $\frac{1}{2}m\omega_{\text{ext}}^2 a_L^2 j^2$, the density of states is simply

$$g(E) = \frac{2\pi E}{m\omega_{\text{ext}}^2 a_L^2} \ . \tag{2.65}$$

Integrating over energy (as in (2.16) for a harmonic trap and no lattice) we find the Fermi energy of the trapped lattice gas

$$E_F = \frac{m\omega_{\text{ext}}^2 a_L^2 N}{2\pi} \ . \tag{2.66}$$

This Fermi energy indicates the maximum energy of N fermions loaded at $T = 0$ into a harmonically trapped lattice with spacing a_L and confinement ω_{ext}.

As long as the harmonic confinement present due to the lattice (or an additional dipole trap) is slowly varying compared to the optical lattice potential, the essential lattice physics described previously is not disrupted. The effect of the additional harmonic trap can be taken into account with the use of the local density approximation, substituting the effect of the external potential by a spatially varying chemical potential

$$\mu(\vec{r}) = \mu - V(\vec{r}) \,, \tag{2.67}$$

where μ is the overall chemical potential of the gas, and $V(\vec{r})$ is the trapping potential. In our case the trapping potential is harmonic, but the local density approximation can be used for more general trap configurations. In the context of the local density approximation, preparing a gas in a trap is equivalent to preparing many different samples of the quantum system, each located at the point $\vec{r}$ with different chemical potentials $\mu(\vec{r})$. A great advantage of using an in-situ probe such as single-site fluorescence imaging is the ability to study the gas as a function of $\vec{r}$, extracting information about how system behaviour changes in response to the effective local μ.

2.4 Hubbard Model

In the previous sections we have outlined some of the properties of non-interacting fermions in an optical lattice. The Schrödinger equation considered only the kinetic energy of the particles and the confining potential, with no possibility for inter-particle interactions. This was appropriate for the experimental case of ultracold fermionic atoms which are all in the same internal state, because at low energies collisions between identical fermionic atoms become forbidden by quantum statistics. In real materials the Coulomb interactions between electrons in a crystal structure can have profound influences on the net behaviour of a material, and it is possible to incorporate inter-particle interactions with ultracold atoms by using two different internal states of an atom labelled as spin-up $| \uparrow \rangle$ and spin-down $| \downarrow \rangle$. In this case an interaction emerges between the two different spin states due to the contact interaction experienced by neutral atoms. To explore the parameter space of attractive and repulsive interactions, Feshbach resonances are commonly used to experimentally tune the strength and sign of the contact interaction.

When the inter-particle interactions are strong enough to dominate system behaviour, new models beyond the simple bandstructure of the previous sections are necessary. Since the contact interactions

between ultracold atoms occur locally, it is natural to consider the lattice system in the limit where the lattice depth is large enough to localize the wavefunctions near a single lattice site. In this tight-binding limit, the wavefunctions are best expressed in the basis of Wannier states and the interaction energy can be added to the Hamiltonian for two particles localized on the same lattice site.

Without yet incorporating interactions in the system, one can express the lattice Hamiltonian of the system in the basis of Wannier functions, resulting in the tight-binding Hamiltonian. Particles are assumed to occupy only the lowest energy band of the lattice potential, and so only the $n = 0$ Wannier functions are necessary to describe the quantum state. The Hamiltonian is written in second-quantized form as a sum over all sites of the lattice

$$H_{TB} = -t \sum_{\sigma} \sum_{\langle i,j \rangle} \left(c_{i,\sigma}^{\dagger} c_{j,\sigma} + h.c. \right) , \tag{2.68}$$

with the kinetic energy written in terms of a tunneling of particles between neighboring sites i and j. Here $c_{i,\sigma}^{\dagger}$ and $c_{i,\sigma}$ are creation and annihilation operators for a fermion of spin σ on site i, and only one particle of each spin state can exist on a given site due to Fermi statistics. The parameter t is the tunneling energy scale, and is given by

$$t = - \int w_0^*(x) \left(-\frac{\hbar^2}{2m} \nabla^2 + V(x) \right) w_1(x) \, dx , \tag{2.69}$$

where $w_0(x)$ and $w_1(x)$ are Wannier functions on adjacent sites. This expression can be simplified by expressing the Wannier functions as an inverse Fourier transform of the Bloch waves as in (2.41). This transforms (2.69) into an integral over quasimomentum q:

$$t = -\frac{1}{k} \int_{-k_L/2}^{k_L/2} E(q) e^{iqa} \, dq . \tag{2.70}$$

From this equation we see that the tunneling energy t is given by the Fourier transform of the energy band $E(q)$. In writing (2.68) we assumed only nearest-neighbor tunneling, which is appropriate in the case of deep lattices with cosine-shaped bands. In this case there is only one Fourier component for $E(q)$, so that (2.70) can be evaluated to simply equate t with the width of the energy band:

$$t = \frac{W}{4} . \tag{2.71}$$

The expression (2.71) for t reflects the behaviour we have already seen in the previous sections: in deeper lattices with small bandwidth W, tunneling is less likely and has a smaller effect on the Hamiltonian. For a given optical lattice depth, we can calculate W in each direction from band theory, immediately giving the tunneling energy. Using bandwidths calculated using the Mathieu equation (2.71), W and t are plotted versus the one-dimensional lattice depth in Figure 2.13, showing the exponential decay of the tunneling rate with increasing lattice depth.

In a symmetric two-dimensional or three-dimensional lattice potential, the tunnel coupling between sites of each lattice axis is still given by the matrix element t. However, the total range of energies available to a particle is increased by the lattice dimensionality d to $W_{\text{tot}} = dW$, as in (2.55) and (2.56). By including the dimensionality factor, the tunnel coupling can be expressed in terms of the total bandwidth of single particle states by the expression $t = W_{\text{tot}}/4d$.
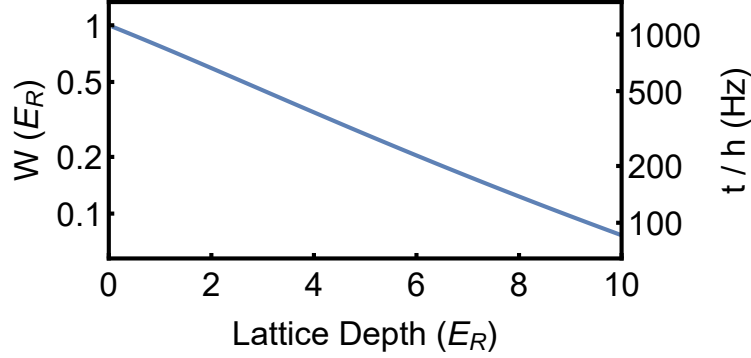
Figure 2.13: Log plot of both the bandwidth W and the corresponding tunneling rate t/h predicted from (2.71) for a one-dimensional lattice, as a function of the lattice depth. The left axis, expressed in units of E_R is general for any optical lattice and atomic species. The right axis uses the recoil energy specific to ^{40}K in a lattice of period $a = 527\,$nm to express the tunneling rate in Hz.

The full Hubbard model is written by adding the effects of inter-particle interactions to the tight-binding model above. Two atoms of opposite spin can occupy the same lattice site, and will be subject to an additional interaction energy U, which will be proportional to their density overlap. The two energy scales t and U parameterize the Hubbard Hamiltonian

$$H = -t \sum_\sigma \sum_{\langle i,j \rangle} \left(c_{i,\sigma}^\dagger c_{j,\sigma} + h.c. \right) + U \sum_i n_{i,\uparrow} n_{i,\downarrow}. \tag{2.72}$$

Two particles occupying the same lattice site must have an identical Wannier wavefunction due to the restriction to only the lowest energy band. Since the probability density of each particle is proportional to $|w_0(x)|^2$, the density overlap is proportional to the 3D spatial integral of $|w_0(x)|^4$. For cold atoms of mass m, interacting via a contact interaction parameterized by the scattering length a_s, the interaction energy U is specifically given by [56]

$$U = \frac{g}{2} \int |w_0(\mathbf{x})|^4 \, d\mathbf{x}, \tag{2.73}$$

where $g = 4\pi a_s \hbar^2/m$.

In a deep lattice for which each site is well approximated at low energies by a harmonic trapping potential, the Wannier function for the lowest band can be replaced by the ground-state wavefunction of a 3D harmonic oscillator. In this case the integral in (2.73) is easily evaluated to give

$$U = \frac{g}{2} \left(\frac{1}{2\pi \bar{l}_0^2} \right)^{3/2} = \frac{\hbar \bar{\omega}(a_s/\bar{l}_0)}{\sqrt{2\pi}}, \tag{2.74}$$

where $\bar{\omega}$ and $\bar{l}_0$ are the geometric mean values of the trap frequency and oscillator length for the 3D trap. If asymmetric lattice depths $V_z \gg V_y = V_y$ are used to realize a two-dimensional Hubbard model, the trap frequency of the z lattice axis still affects the calculation of U through $\bar{\omega}$. If the vertical lattice is kept at a large depth $V_z = 50\,E_R$, varying the horizontal lattice depths will result in a small variation of U in the range of a few kHz, as shown in Figure 2.15. Adjusting the lattice depth in this way also necessarily varies t, so that the system can be tuned from tunneling-dominated to interaction-dominated as the lattice depth varies. Further tuning of the interactions is possible by using a Feshbach resonance
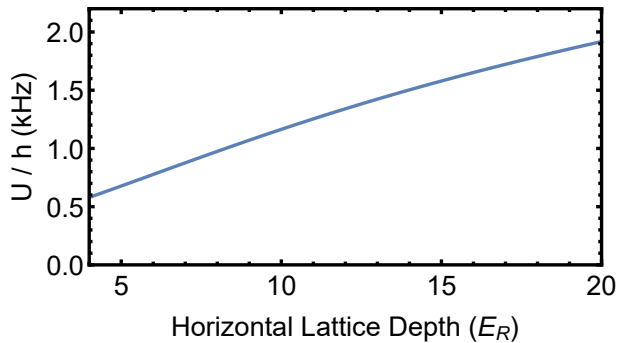
Figure 2.14: Hubbard interaction parameter U for ^{40}K atoms interacting with the background scattering length $a_s = 174\,a_0$ in an asymmetric 3D lattice with period $a = 527\,$nm. The vertical lattice depth is taken to be $V_z = 50\,E_R$, so that tunneling will occur only in the remaining horizontal lattice direction. The harmonic approximation of (2.74) is used, replacing the Wannier states with ground state harmonic oscillator wavefunctions.

to independently vary the scattering length a, which makes it possible to adjust U while keeping t fixed.

2.4.1 Hubbard Model with Harmonic Trapping

Taking into account the additional trap confinement present in an ultracold atom experiment, the Hubbard model is expressed as

$$H = -t \sum_\sigma \sum_{\langle i,j \rangle} \left(c_{i,\sigma}^\dagger c_{j,\sigma} + h.c. \right) + U \sum_i n_{i,\uparrow} n_{i,\downarrow} + \sum_i \epsilon_i n_i \,, \tag{2.75}$$

where the ϵ_i terms take into account the presence of the trapping potential shifting the energies of different lattice sites. For a 1D harmonic trap we would have $\epsilon_i \propto i^2$. As discussed previously harmonic confinement is typically taken into account using the local density approximation, unless the energy shifts between sites are so large that particles become localized. The phase diagram of the Hubbard model (2.75) is determined entirely through the competition of the three energy scales t, U, ϵ, along with the temperature T and the chemical potential μ.

When realizing the Hubbard model experimentally for a quantum simulation, three dimensionless ratios of the Hubbard parameters and T can be used to characterize the state prepared. Typically the temperature is stated as $k_B T / U$ in units of the interaction strength, or as $k_B T / t$ in units of the tunneling energy. The interaction energy is commonly stated as $U/8t$ in terms of the bandwidth for two dimensions. The total number N of particles in both spin states is best compared to the parameter N_0 from (2.64). For two spin states it takes $2N_0$ particles to fill the harmonic trap up to half of the bandwidth, suggesting a characteristic density $\rho = N/2N_0$. With this definition $\rho = 1$ corresponds to a filling of roughly one atom per site in the centre of the lattice, which is the important case of half-filling for the spin-$\frac{1}{2}$ system. These three ratios can be used to compare the results of different experiments and theoretical models.

With low density $\rho \ll 1$ the many-body state is metallic and compressible. As ρ is increased to 1 to produce half-filling, the centre of the trap will still metallic as long as the the interaction is small ($U/8t < 1$). Increasing the density further will entirely fill the ground band in the centre of the trap, producing a band-insulating incompressible state with a fixed density of two particles per site (one in each
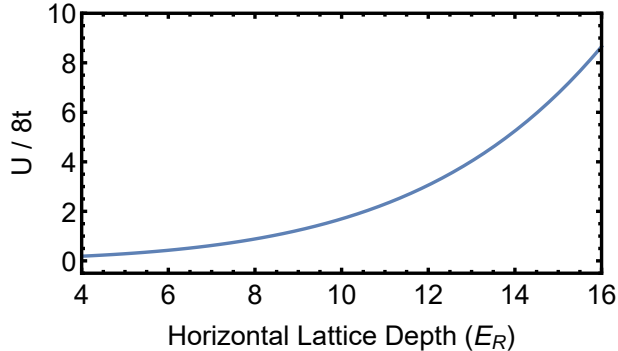
Figure 2.15: The Hubbard parameter $U/8t$ for ^{40}K atoms in a two-dimensional lattice with $a_L = 527$ nm. The lattice depth in the vertical direction is assumed to be $V_z = 50\,E_R$. The background scattering length of $a_{BG} = 174a_0$ is assumed.

spin state). Alternatively, starting with half-filling ($\rho = 1$) and increasing the interaction energy until it dominates tunneling ($U/8t \gg 1$) produces a Mott insulating state in the centre of the trap. These band-insulating and Mott-insulating states will only form at sufficiently low temperature. Temperatures much lower than the bandgap are required to realize a band insulator, and $k_B T \ll U$ is required to realize the interaction-driven Mott insulator. For the range of temperatures accessible to cold atom experiments, the interplay of these three important phases has been explored for a 3D Hubbard model [57], showing clearly that all of the phases can coexist in a trapped gas. This is easy to understand in the context of the local density approximation, where different regions of the cloud locally realize different instances of the standard Hubbard model (2.72).

Simply by tuning the lattice depth, large variation over $U/8t$ is achievable. Assuming the background scattering length for ^{40}K atoms, a range of typical lattice depths gives values of $U/8t$ ranging from much less than 1 to around 10 as plotted in Figure 2.15. This variation should be enough to tune between band-insulating and Mott-insulating phases as described above. To vary the temperature $k_B T/t$ we can adjust the amount of cooling performed on the ultracold gas prior to loading the optical lattice. Reaching extremely low temperatures requires sophisticated cooling techniques, and is a topic of current interest in the cold-atom community [58]. Tuning the density ρ is straightforward by varying the number of atoms trapped by the experiment. Recent quantum gas microscope experiments have shown that by making use of small volume dimple traps, small shot-to-shot variation in N can be achieved [59].

Previous ultracold atom experiments have examined 3D trapped Hubbard models, where large numbers of particles can coexist in the same state. This allowed for the characterization of Mott-insulating states based on global signatures such as double-occupancy and compressibility [14, 52]. Recently, quantum-gas microscope experiments have been used to explore the coexisting band-insulating, Mott-insulating, and metallic states of a Fermi gas in a 2D optical lattice [31,35,59]. Here the high-resolution imaging afforded by a quantum-gas microscope is crucial to read out the state of a relatively small number of atoms. There is a great theoretical interest in probing beyond the well understood insulating states of the Hubbard model, to see if it contains any signs of d-wave superfluidity or other exotic behaviour. Very recently, experiments with quantum-gas microscopes have been able to observe the appearance of anti-ferromagnetic correlations in a Mott insulator at low temperature [60–62], one of the last well characterized phases of the Hubbard model. It is very exciting to see what light these kinds of experiments may be able to shed on other regions of the Hubbard model phase diagram for which good

theoretical results are lacking.

Chapter 3

Microscope Apparatus

Achieving site-resolving fluorescence imaging requires considerable forethought in the design of the experiment. Typical ultracold atom experiments operate with a small cloud trapped under vacuum in a glass or metal chamber, completely surrounded with a precisely tuned jungle of optics and electronics. A significant challenge for quantum gas microscopy is the incorporation of all necessary trapping, cooling, and probing beams in the experiment while still leaving plenty of room to place a microscope objective very close to the ultracold cloud. To achieve high resolution imaging of an ultracold gas, the microscope objective should view the cloud with tremendous optical access. This can be achieved by using a glass vacuum cell, with a long-focal-length microscope outside the cell imaging the cloud within [20, 63]. Our experiment follows this basic approach, but uses an especially thin vacuum-air interface to allow the use of microscopes with a shorter focal length near 3 mm. Alternative approaches to high resolution imaging make use of a refracting element either mounted inside the vacuum system [64] or bonded to one of the windows [18, 31, 34]. The key to any high-resolution imaging arrangement is to collect the largest range of emission angles for the light interacting with the cloud. The emission directions collected are quantified by the numerical aperture (NA) of the imaging system. Achieving a large NA allows small features to be resolved in an image, and in the context of fluorecence detection also ensures efficient collection of the emitted light.

3.1   High NA Optical Access

To prepare an ultracold gas, we laser cool both ^{87}Rb and ^{40}K in a dual-species magneto-optical trap. The two species are loaded into a magnetic trap and transported horizontally and vertically into a stainless steel and titanium *s*cience chamber, pictured in Figure 3.1. At the end of the vertical magnetic transport, atoms are held in a magnetic quadrupole trap in the centre of the science chamber, surrounded by 12 windows in the horizontal plane and 2 vertical viewports. Using the 12 horizontal windows we implement several optical beams to trap and probe the atoms: a blue-detuned beam at 760 nm optically plugs the quadrupole trap, two red-detuned beams at 1054 nm form a crossed optical dipole trap, two retro-reflected 1054 nm beams form the horizontal arms of the optical lattice, two near-resonant 770 nm beams are used to perform EIT cooling of atoms trapped in the optical lattice, and finally three resonant 767 nm beams are used to optically pump the atoms and to record absorption images of the cloud's column density along either of the horizontal lattice axes. Each of the 12 horizontal viewports is used
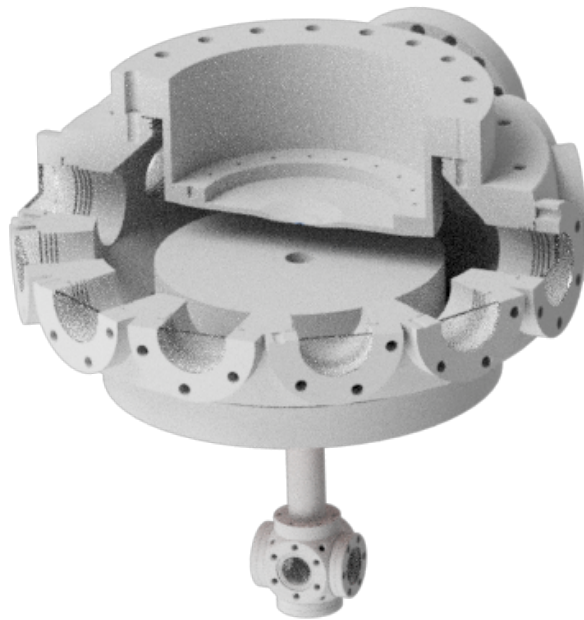
Figure 3.1: Three-quarter cut out view of the stainless steel and titanium vacuum chamber. The main chamber has a radius of 143 mm and features twelve 38 mm diameter windows equally spaced around the horizontal plane. Two windows at the top and bottom of the provide optical access in the vertical axis. The lower window is attached to a 1 inch cube (connected to the main chamber by a 17.5 mm ID tube with a length of 149 mm) and is used to couple in a vertical lattice beam. The upper window is very close to the position of the ultracold cloud, and is used to capture fluorescence.
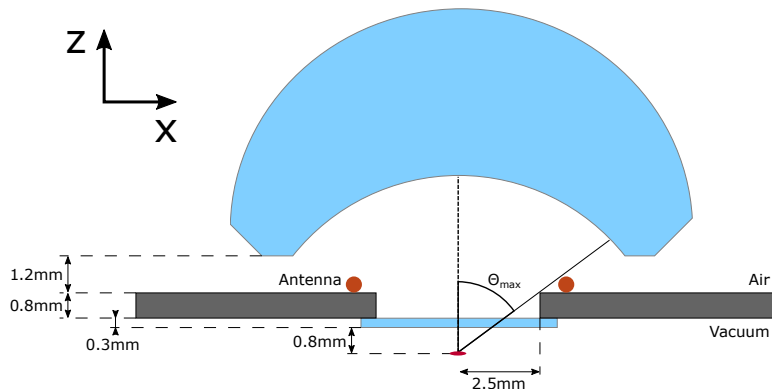
Figure 3.2: The imaging window is $0.2\,\text{mm}$ thick, with an additional $0.1\,\text{mm}$ thickness due to bonding, leading to an overall extent of $0.3\,\text{mm}$ inside the vacuum system. With the ultracold cloud located $0.8\,\text{mm}$ from the imaging window, the microscope acceptance angle θ_{max} is limited to $0.92\,\text{rad}$, giving a maximum NA of 0.8. The microscope objective is supported $1.2\,\text{mm}$ above the titanium flange, leaving room for a thin RF antenna (shown in brown) to encircle the window.

to accommodate at least one ingoing beam.

In the vertical axis, a third optical lattice beam enters through the bottom of the vacuum system and propagates upward, retro-reflecting inside the vacuum system to form the vertical standing wave. The top flange of the science chamber is recessed, placing the top window of the vertical axis very close to the horizontal plane defined by the centre of the 12 horizontal windows. This top window is used to provide high-NA optical access for fluorescence imaging, as described in this chapter.

The experimental apparatus used for the results of this thesis has been assembled over several years through the contributions of several graduate students and visitors. Extensive details on the vacuum system layout can be found in the thesis of David McKay [65]. Details of the laser system, magneto-optical trap, magnetic transport system, and plugged quadrupole trap are presented in the thesis of Dylan Jervis [66]. Below I elaborate on those aspects of the apparatus which are of central interest for the high-resolution imaging results presented in this thesis.

3.1.1   Imaging Window

To achieve a high imaging resolution, our experiment is designed with great attention given to the main imaging window in order to maximize the achievable NA. The window itself is made of sapphire and is $200\,\mu\text{m}$ thick, and was affixed to the titanium vacuum flange by UKAEA using electron diffusion bonding. The sapphire window is thick enough to maintain vacuum pressure, and but is still thin enough to allow the use of a $3\,\text{mm}$ working distance microscope objective outside the vacuum system to collect light from the atoms. An important consequence of the window's thin profile is a reduced image sensitivity to imperfections in the surface quality such as curvature or stress. The titanium flange which holds the sapphire imaging window is tapered to $0.8\,\text{mm}$ thickness near the window, and the bonding with the sapphire window adds $100\,\mu\text{m}$ thickness to the interface, leading to an overall vacuum interface thickness of $1.1\,\text{mm}$ as shown in Figure 3.2.

With the atomic cloud trapped as close as possible to the sapphire window, we collect the maximum possible range of fluorescence emission angles in the microscope. With the cloud positioned within

$10\,\mu$m of the imaging window, imaging with NA = 0.9 would be possible. This cloud position could be accomplished using evanescent traps or optical accordion lattices which are reflected off of the imaging window [31, 34, 67]. Instead, we trap the cloud $800\,\mu$m away from the window surface, simplifying the trapping configuration while still allowing for acceptance angles up to 0.92 rad and NA = 0.8.

Ultracold clouds of ^{87}Rb and ^{40}K are initially prepared in a separate chamber of the vacuum system in a magneto-optical trap (MOT). They are transferred to the main chamber by a magnetic transport system made of several sets of overlapping anti-Helmholtz coils. The last set of coils forms a quadrupole magnetic trap in the imaging chamber, positioned to trap the cloud in a final position $800\,\mu$m from the inside of the imaging window. The exact distance was adjusted by using thin machined Garolite spacers to adjust the height when mounting the top coil of the pair. Even with careful adjustment of the top coil position, in-situ images of a trapped cloud indicate that the quadrupole trap is not exactly centred in the imaging window. In order to precisely center the atoms under the window in the x-y plane, we use additional shim coils to displace the centre of the quadrupole trap using bias magnetic fields along x and y.

The titanium vacuum flange extends away from the imaging window to a radius of $74\,$mm, and at a radius of $10\,$mm a sloped region increases the vertical distance between atoms and flange from $800\,\mu$m to $1500\,\mu$m. The overall cross section of the flange and window are shown in Figure 3.3. The close proximity of the atoms and the imaging window constrains all beam paths for trapping and manipulation of the atoms to the x-y plane perpendicular to the optic axis of the imaging system. If a beam propagates too close to the flange clipping of the intensity profile will cause distortions in the desired trap. To reduce any such distortion, we constrain all trapping and cooling beam waists $w(z)$ so that the beams are at least $2w(z)$ away from the flange at all points. This ensures that clipping only at most on the $1/e^4$ level of the intensity profile, producing minimal distortion. This constraint means that beams can only tilt out of the $\hat{x}$-$\hat{y}$ plane into the $\hat{z}$ direction by roughly $25\,$mrad before unacceptable clipping would be expected. Similarly, a Gaussian beam which is focused to a small waist w_0 will have a large divergence angle $\theta_D = \lambda/\pi w_0$ which can also lead to clipping. For a $1064\,$nm beam, beam divergence constrains us to a minimum focal size of $15\,\mu$m. The most tightly focused horizontal beam used in the experiment is the elliptical optical dipole trap, for which each of the two beams has a $40\,\mu$m vertical profile.

3.1.2   RF Antenna

Near the imaging window, we also place a small loop of wire to act as an antenna to couple radiofrequency (RF) and microwave frequency radiation to the cloud. Placing the antenna near the imaging window permits a small distance between the antenna and the cloud, which helps to achieve a large coupling Rabi frequency Ω_R. However, the atoms are strongly shielded from the antenna by the titanium flange. Large Ω_R can still be achieved if the antenna partially overlaps the imaging window, providing direct 'line-of-sight' to the atoms, but this would interfere with light collection by the microscope. Instead, we form the antenna into a loop that encircles the imaging window without obstructing the clear aperture, as pictured in Figure 3.2. The antenna design is shown in Figure 3.4, along with a previous iteration which restricted the imaging window to roughly NA = 0.6 but delivered an order of magnitude more RF power to the atoms.

The RF antenna is formed from a single loop of Kapton coated AWG30 copper wire (OD = $0.54\,$mm). The loop is secured around the window by taping it to the vacuum flange, and the coaxial cable which is connected to the loop is supported by taping it to a rigid arm made of Garolite. At the other end of
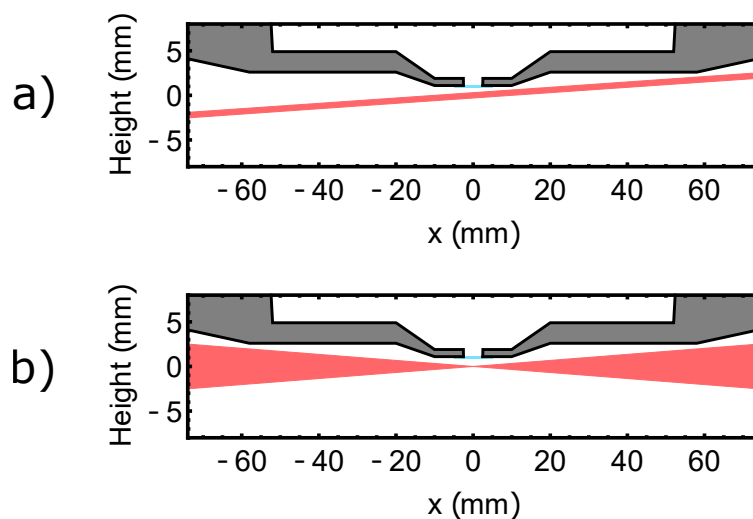
Figure 3.3: Cross section of the vacuum flange, showing the allowable beam angles and focussing. a) A Gaussian beam with $w_0 = 300\,\mu$m entering the chamber $25\,$mrad out of the plane perpendicular to the imaging axis. b) A Gaussian beam entering the chamber in the plane perpendicular to the imaging axis, focusing to a $20\,\mu$m waist. All beam profiles are calculated for $\lambda = 1064\,$nm and show the width at which the intensity drops to $1/e^4$ of its peak value.
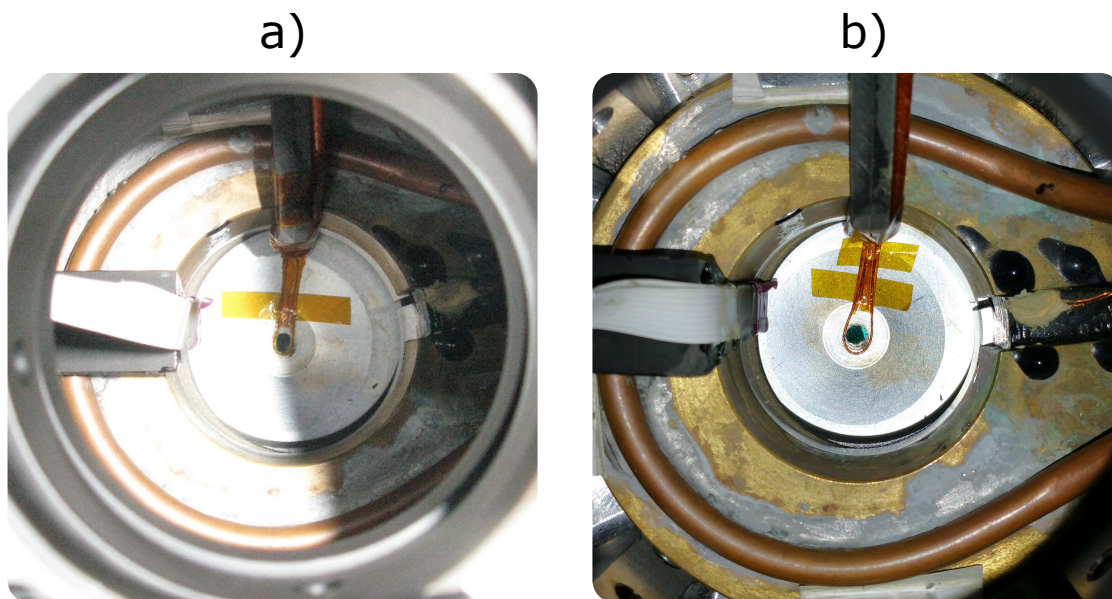
Figure 3.4: RF/Microwave antenna position. a) An old antenna design slightly overlapped the imaging window, giving good coupling to the atoms but reducing the achievable NA. b) The new antenna formed from a thinner wire, and shaped around the window to allow optical access with NA = 0.8. The atoms are effectively shielded from the antenna by the titanium flange, resulting in low RF/microwave Rabi frequencies. Also visible are the black Garolite supports for the antenna cable (top) and a 3-axis magnetic field sensor (left). The antenna, field sensor, and Garolite supports all sit within the inner radius of the top coil of the magnetic transport system.
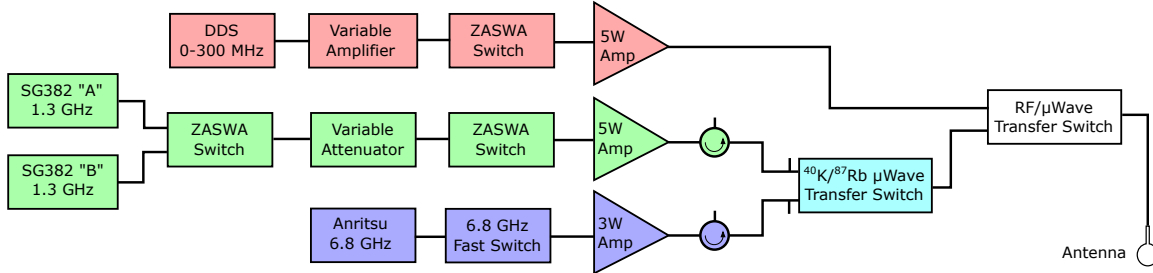
Figure 3.5: Layout of the RF and microwave electronics which are used with the single loop antenna. Microwave signals are stub-tuned before combining with a transfer switch. RF source is a DDS based on Analog Devices AD9910, built by Alan Stummer. Microwaves sources are two SRS SG382 generators for 1.3 GHz and Anritsu MG37022A for 6.8 GHz. Amplitude control is provided by a Texas Instruments VCA822 variable gain amplifier for RF, and by a Mincircuits ZX73-2500 voltage variable attenuator for 1.3 GHz. Fast switches are Minicircuits ZASWA-2-50DR for RF and 1.3 GHz, and NARDA 50222 for 6.8 GHz. Amplifiers are Minicircuits TIA-1000-1R8 for RF, Minicircuits ZHL-5W-2G+ for 1.3 GHz, and Minicircuits ZVE-3W-83+ for 6.8 GHz. Circulators are Wenteq F2568 for 1.3 GHz and Alcatel 40B86 for 6.8 GHz. Transfer switches are both Minicircuits MTS-18-12B+.

| Transition	Frequency	Amplifier Power	$\Omega_R/2\pi$
$^{87}\text{Rb} : F = 1 \rightarrow F' = 2$	6.8 GHz	3W	600 Hz
$^{40}\text{K} : F = 9/2 \rightarrow F' = 7/2$	1.3 GHz	5W	5 kHz
$^{40}\text{K} , \,^{87}\text{Rb} : |F, m_F\rangle \rightarrow |F, m'_F\rangle$	$10 - 200\,\text{MHz}$	4W	12 kHz

Table 3.1: RF and microwave frequencies and associated Rabi frequencies

the coaxial cable we use three transfer switches to connect the antenna to different frequency sources, as shown in Figure 3.5. The two microwave frequencies are separately stub tuned to the antenna using short lengths of coaxial cable. Despite similar amplifier output power at each of the three frequency ranges, we find that the achieved Ω_R is lower at higher frequencies, as shown in Table 3.1.

3.2 Microscope Objective

Our thin imaging window, tapered vacuum flange, and low-profile RF antenna combine for an overall thickness of 1.7 mm. To keep the atom cloud 800 μm away from the window it then suffices to use a microscope objective with a working distance of at least 2.5 mm. This presents several off-the-shelf options, and we initially used a Zeiss LD-Plan-Neofluar 63x/0.75. Unfortunately, the housing of this objective was slightly magnetic, so that adjusting and focusing it in such close proximity to the atom cloud caused significant changes to the magnetic field. The field shifts were large enough to interfere with the alignment of our plugged quadrupole trap, hindering the production of ultracold clouds. This alignment could be compensated with additional magnetic fields, but this compensation would need to be performed each time the microscope was adjusted, a time-consuming process. We instead decided to use a custom objective made from nonmagnetic aluminum.

With an initial goal of imaging ^{40}K on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition at 405 nm, we purchased a custom objective from Special Optics which attained NA = 0.6. When we were unable to achieve single-atom sensitive fluorescence imaging using this transition (see Appendix A), we purchased a second custom objective for the 770 nm $4S_{1/2} \rightarrow 4P_{1/2}$ transition, designed to reach NA = 0.8. Both the 405 nm and

770 nm objectives were infinity-corrected, and used with a tube lens of 200 mm focal length they had magnifications of 60x and 40x, respectively. The 770 nm objective was used for all results in this thesis apart from the results described in Appendix A. Both objectives are made with an aluminum body, and no effects on the magnetic field inside the chamber were detected when positioning and focusing either objective.

The 770 nm objective has a working distance of 3.11 mm. When focused on the atoms inside the chamber, the distance between the final lens of the objective and the outer wall of the titanium flange is 1.2 mm, as shown in Figure 3.2. To precisely align the objective over the cloud, we use three positioning stages, as shown in Figure 3.6: a Newport TTN80 goniometer for tilt in the $\hat{x}$ and $\hat{y}$ directions, a PI P-733.ZCD piezo translation stage providing fine adjustments along the $\hat{z}$ direction for focusing, and a Newport M406 two-axis linear stage with micrometer screws for centering of the microscope in the $\hat{x}$-$\hat{y}$ plane. The goniometer sits on the piezo stage, which in turn is supported by the micrometer stage. The micrometer stage is supported 12 cm above the imaging window with a custom aluminum frame. The objective is connected to the goniometer stage by a 175 mm long threaded tube, which hangs down through the centre of the other stages. The threads of the long tube provide a means of coarsely adjusting the microscope in the $\hat{z}$ direction.

This method of mounting the objective at the end of a long tube provides excellent access to all positioning stages while leaving room for the RF antenna, a three-axis field sensor, and the top coil of the magnetic transport system to be placed inside the recessed vacuum flange. However, any play in the threads which attach the tube to the goniometer stage will result in large displacements of the objective over the atoms. To reduce this freedom the tube is threaded into a 20 mm high brass cylinder to provide a large contact area. An additional lock ring was machined to lock the rotation of the tube using a set screw, but tightening this introduces large motion of the objective and so it is typically not used.

To test the long-term stability of the microscope positioning stack before integrating it in the experiment, a 25 μm pinhole was imaged using the fully mounted objective, and the position of the image was tracked over time. Over a 12 h observation period, drifts were less than 3 μm. When the aluminum support legs were heated to increase their temperature by 4° C, the image drifted by 6 μm but returned to its initial position when the heat was removed. A 0.5 μm pinhole was also imaged to gain information about the short term stability of the microscope mount, and observed fluctuations are on the order of ±50 nm on a time scale of 300 ms. Since the measurement did not confirm whether this was a fluctuation due to movement of the microscope or the pinhole, we take this as an upper bound on the possible vibrational movement of the microscope. Even if this movement was entirely due to the microscope objective, a jitter of 50 nm will only add a small correction to the overall imaging resolution, which is dominated by the resolution limit of the microscope objective. This suggests that the mounting of the objective was sufficiently stable for our purposes of fluorescence imaging in a lattice with a spacing near 500 nm.

When mounting all of the positioning stages to the vacuum flange, the magnetic field in the position of the atoms changed by roughly 1 G. This is mainly due to the micrometer stage with some contribution from the goniometer stage. The generated field does not change greatly when using the stages to adjust the microscope position, and so this does not pose a large inconvenience in the experiment. The field is characterized and removed in the experiment by applying compensating fields with extra shim coils.
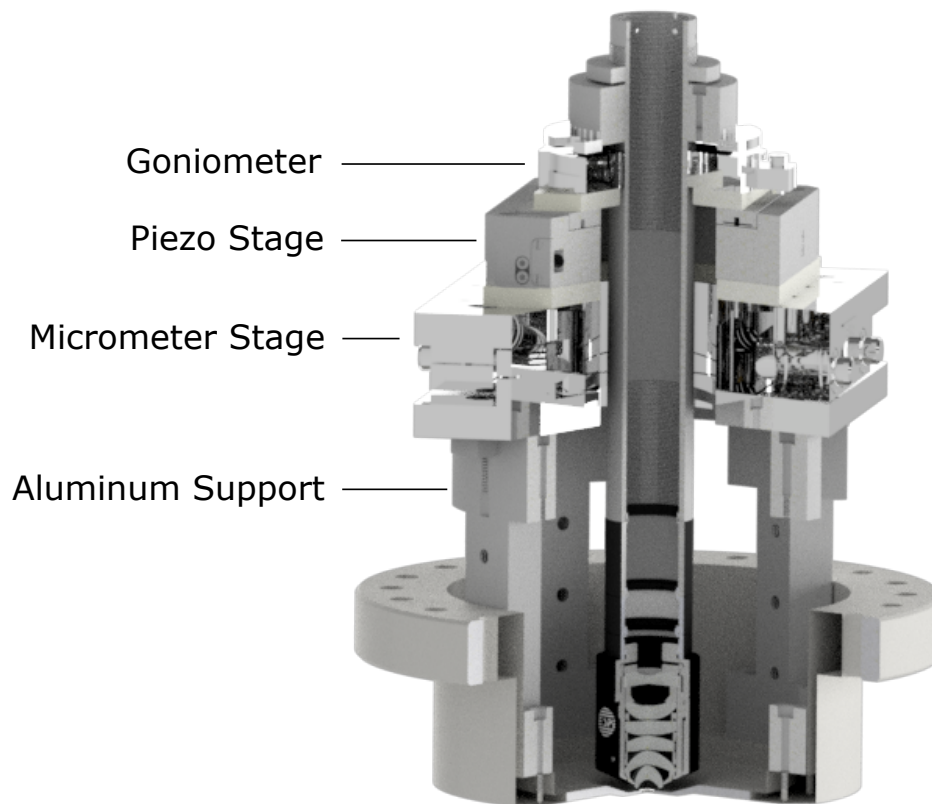
Figure 3.6: Mounting scheme for the microscope objective. For high maneuverability, the microscope objective is mounted to a goniometer, which is mounted on a piezo stage for $\hat{z}$ translation, which is mounted on a micrometer stage for $\hat{x}$ and $\hat{y}$ translation. The micrometer stage is attached to the vacuum system via aluminum legs, which raise all the mechanical stages far enough to mount the top coil of the Helmholtz pair used to access Feshbach resonances.
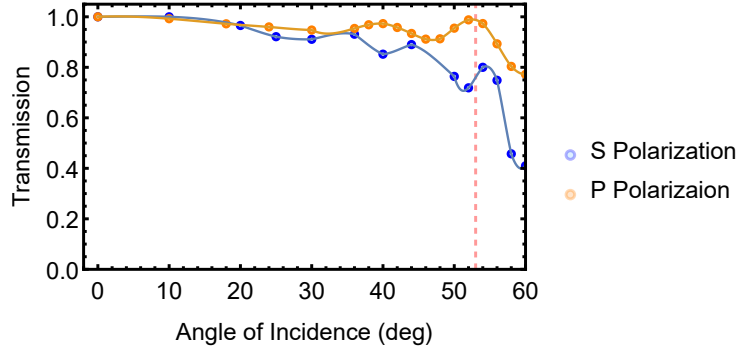
Figure 3.7: Measured transmission of the imaging window at $\lambda = 770\,\text{nm}$ as a function of the incident angle. The dashed red line indicates the maximum angle of incidence which is collected by the microscope objective. Solid orange and blue lines are a guide to the eye.

3.3   Image Capture

3.3.1   Photon Collection

When collecting a fluorescence image of a cloud with single-atom sensitivity, the number of photons incident on any individual camera pixel is very low. The scattering rate during imaging is limited by several factors, and even with long exposure times of around 1 s, less than 10^5 photons are typically scattered by each atom. This makes it crucial to collect, transmit, and record the scattered light with high efficiency. The microscope will collect only a fraction of light determined by the acceptance angle, and if the atoms are assumed to fluoresce equally in all directions this collection efficiency (CE) is given by

$$\text{CE} = \int_0^\alpha \sin\theta\, d\theta = \sin^2(\alpha/2)\,, \tag{3.1}$$

where α is the admission angle of the optical system, and is related to the NA by $\text{NA} = \sin(\alpha)$ when the microscope operates in air or vacuum (with unity index of refraction). The NA enforced by our vacuum geometry and our microscope corresponds to $\alpha = 0.92$ radians, allowing us to collect up to 20% of the light scattered by the cloud.

In addition to providing a view port for fluorescence imaging, we also use the imaging window to retro-reflect one of our lattice beams, forming a vertical standing wave which is perpendicular to the window surface. For this dual purpose, the imaging window is coated to have high transmission at the fluorescence wavelength of 770 nm, and a high reflectivity at the optical lattice wavelength of 1054 nm. We measure the window transmission at normal incidence to be 8×10^{-5} for 1054 nm and 0.99 for 770 nm. For high-NA fluorescence imaging, the 770 nm light will have incidence angles up to $\alpha = 0.92$ radians, and so we additionally determine the window transmission as a function of angle. The results are shown in Figure 3.7. The transmission is seen to decrease at large angles of incidence, which is not unexpected from the coating specifications. As one would expect for a dielectric interface, the two possible electric field polarizations exhibit very different behaviour at large angles of incidence. The reduced transmission of the window at large incident angles will lower the collection efficiency of the imaging system, and can be taken into account by introducing an angle-dependent transmission $T(\theta)$ to (3.1), giving the formula

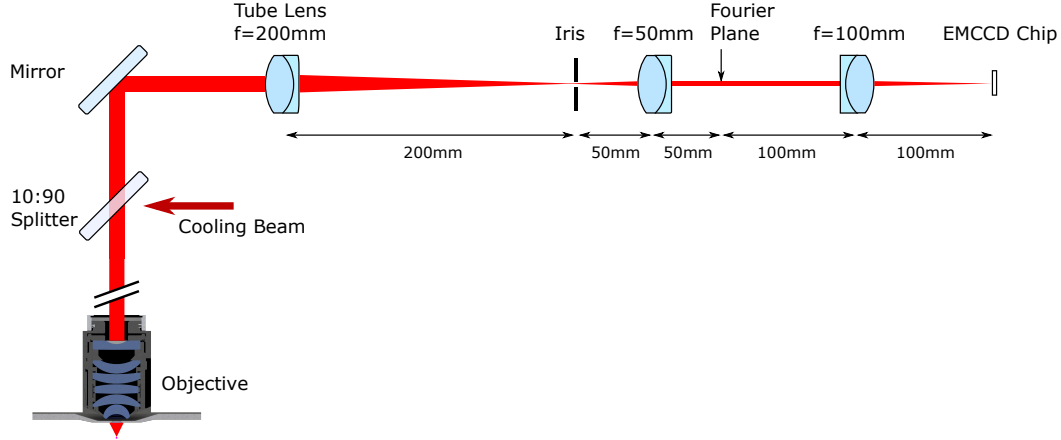$$\text{CE} = \int_0^\alpha T(\theta) \sin\theta\, d\theta\,. \tag{3.2}$$

Figure 3.8: Optical layout of the microscope imaging system. A 90% transmissive beamsplitter placed near the microscope allows cooling beams to be reflected into the imaging path and sent through the microscope towards the atoms. The objective and tube lens create an image inside the iris with a magnification of $40\times$. Two achromatic doublets of $50\,\mathrm{mm}$ and $100\,\mathrm{mm}$ focal length increase to total magnification to $80\times$ for the image on the EMCCD.

We average over both polarization states, and then use an interpolation to the measured transmission data to give an approximate functional form for $T(\theta)$, allowing us to evaluate (3.2). We calculate that the finite transmission of the imaging window reduces the collection efficiency of the optical system to 18.5%.

3.3.2   Imaging Path

The collected light is transmitted to our camera using the optical system depicted in Figure 3.8. Ensuring that all optics are properly coated ensures good transmission through the system, and so we need only ensure that the camera efficiently detects the weak light signal. A camera with high quantum efficiency ensures that the majority of photons will be converted into photoelectrons, but it is still important to detect the small number of electrons with low noise. For this purpose an electron multiplying CCD camera is ideal, where solid state amplification of the liberated photoelectrons amplifies the signal before it is read out from the chip, reducing the effect of electronic noise on the readout process. We use an Andor iXon Ultra 897, which has a quantum efficiency $> 80\,\%$ at a wavelength of $770\,\mathrm{nm}$.

The pixel size of the iXon camera is $16 \times 16\,\mu\mathrm{m}$. We use an additional 2x magnification stage after the microscope objective and tube lens to increase the overall magnification of the imaging system to 80x. This gives us an effective pixel size of roughly $200\,\mathrm{nm}$, spreading the light from each site in a $527\,\mathrm{nm}$ period lattice over an area of $2.7 \times 2.7\,\mathrm{pixels}$. This magnification is intended to be large enough to properly sample the periodicity of the lattice, without spreading the dim fluorescence signal from a single atom over too many pixels.

The use of a secondary magnification stage in the imaging system means that we have access to the first image plane between the tube lens and the magnification stage, in addition to the second image plane located at the camera sensor. In the first image plane, we place an iris in order to shield the camera from stray light not originating from the atom cloud. Additionally, inside the magnification stage we have access to a Fourier plane of the optical system. When trying to filter out unwanted light

from the imaging path, we performed spatial filtering of different light modes by placing a dark spot inside the Fourier plane. This setup for blocking unwanted light in an image is similar to the technique of dark-ground imaging [68]. This spatial filtering was unable to remove enough of the unwanted light from the imaging path, and later a spinning chopper wheel was used in the Fourier plane to periodically block all light from reaching the camera.

To shield the iXon camera from ambient light, we enclose the optical assembly in Figure 3.8 in a lens tube. A small gap in the tube is created in the position of the Fourier plane, to allow for the placement of elements such as a dark spot or a chopper wheel. To prevent ambient light from entering the lens tube through this gap, we covered the entire assembly from the mirror to the camera in a large box. Later, a narrow bandpass filter (Alluxa 770-1 OD4) placed between the last $100\,\mathrm{mm}$ lens and the EMCCD chip was used to effectively block all ambient sources of light, obviating the need for the additional box surrounding the lens tube.

In addition to using the microscope to collect light *from* the atoms, we also use the imaging path to couple light *towards* the atoms. Although at the moment this is only used to illuminate the atoms with large near-resonant beams for laser cooling and removal of atoms, the microscope optical path will be used in the future to create a tightly focussed 'dimple' trap or to project arbitrary optical potentials using a spatial light modulator. Beams to be projected down through the microscope are overlapped with the imaging path using a 10% reflective beamsplitter. Although any low-reflectance optical surface like a glass plate would appear to work for this purpose, it is important to minimize the appearance of multiple reflections in beams coupled down towards the microscope. Roughly equal strength reflections from the two parallel faces of a glass plate will form a stable interference pattern at the position of the atoms when projected through the microscope. We use a beamsplitter with a $30\,\mathrm{arcmin}$ wedge between the two glass surfaces, ensuring that only one reflection reaches the atoms and removing any issues of interference. From the perspective of high-resolution imaging, it is not ideal for a wedged optic to be placed inside the imaging path, and so after obtaining the results of this thesis the 10% reflective beamsplitter was removed and the mirror in the optical setup was replaced with a 90% reflective beamsplitter. This allows the same access to send beams through the objective, while removing unwanted transmissive optical elements from the imaging path.

3.4   Point Spread Function

The high NA of our imaging system will allow us to collect fluorescence images with high spatial resolution. The spatial resolution of an imaging system is best understood by considering the point spread function (PSF), which is the image generated for a light source which is an infinitesimally small point. Since there is no spatial structure to the point source of light, the shape of the resulting image is determined entirely by diffraction effects in the imaging system. Larger NA imaging systems collect light from a point source with a greater range of emission angles, which leads to narrower spatial distribution of intensity when an image of the source is formed. It is the effective aperture of the imaging system, transmitting light within some range of emission angles and rejecting light at larger angles, which ultimately determines the minimum possible imaging resolution. This minimum resolution may not be achieved in practice due to the presence of optical aberrations.

To calculate the size of the PSF for an imaging system with a given aperture, scalar diffraction theory is commonly used. In this method we disregard the vector nature of the electric field by assuming that

the vector direction of the electric field is the same at all points of the imaging system, so that we can consider only the field magnitude to determine the interference effects. Further we assume that the imaging system collects the spherical light waves emitted by the point object at only a small range of emission angles. In this approximate case, the light arriving at the entrance of the imaging system will look approximately like a plane wave. These approximations are obviously not generally applicable to imaging systems with large NA, but they simplify the steps of the calculation dramatically and allow for general insights to be obtained.

For a plane wave passing through a uniform aperture of radius a and focused by a lens of focal length f, the electric field distribution in the focal plane is found by describing the lens surface in polar coordinates (r, θ) and integrating the electric field arising from each point:

$$E(\rho, \phi) = \int_0^a \int_0^{2\pi} r \exp^{ikr\rho \cos \theta / f} d\theta dr\,, \tag{3.3}$$

where (ρ, ϕ) is the location of a given point in the image plane represented in polar coordinates, k is the wavevector of the light, and the assumption of small angles θ has been used to write the electric field contribution from each point as a plane wave rather than a spherical wave. Integrating over the angle θ results in a Bessel function of order zero, and further integration over r yields a Bessel function of order one. Finally squaring the electric field to find the intensity distribution, we have

$$I(\rho) = I(0) \left(\frac{2J_1(ka\rho/f)}{ka\rho/f} \right)^2 \,, \tag{3.4}$$

where $I(0)$ is the peak intensity of the distribution, and the approximation $\rho/\sqrt{\rho^2 + f^2} \approx \rho/f$ has been used (which is typically valid since ρ is bounded by the microscope field of view, which is much smaller than f). For imaging through a medium with index of refraction $n \approx 1$ the lens' Numerical Aperture (NA) is given by $\mathrm{NA} = a/f$, and so (3.4) can be re-expressed as:

$$I(\rho) = I(0) \left(\frac{2J_1(k\rho\mathrm{NA})}{k\rho\mathrm{NA}} \right)^2 \,. \tag{3.5}$$

The intensity pattern (3.5) is known as an Airy pattern, and is plotted in Figure 3.9 for our experimental values of k and NA. The image of the zero-size point source has a circular region of high intensity, which smoothly decays to a ring of zero intensity at a radius of about 500 nm. Outside of this dark band there is a series of bright rings with decaying brightness, associated with the hard edge of the imaging aperture.

Several definitions of imaging resolution can be defined based on the calculated shape of the Airy pattern. For instance the Rayleigh criterion states that two point objects are resolved if their separation is larger than the distance from the centre of the Airy pattern to the minimum of the first dark band. This gives a minimum resolvable distance of

$$d_{\mathrm{Rayleigh}} = \frac{0.61\lambda}{\mathrm{NA}}\,. \tag{3.6}$$

For our experimental NA of 0.8 and an imaging wavelength of 770 nm, this gives a resolution limit of 587 nm, which is comparable to the typical spacing of atoms in an optical lattice. Alternative definitions of what is required for two point sources to be resolved produce different definitions of the minimum
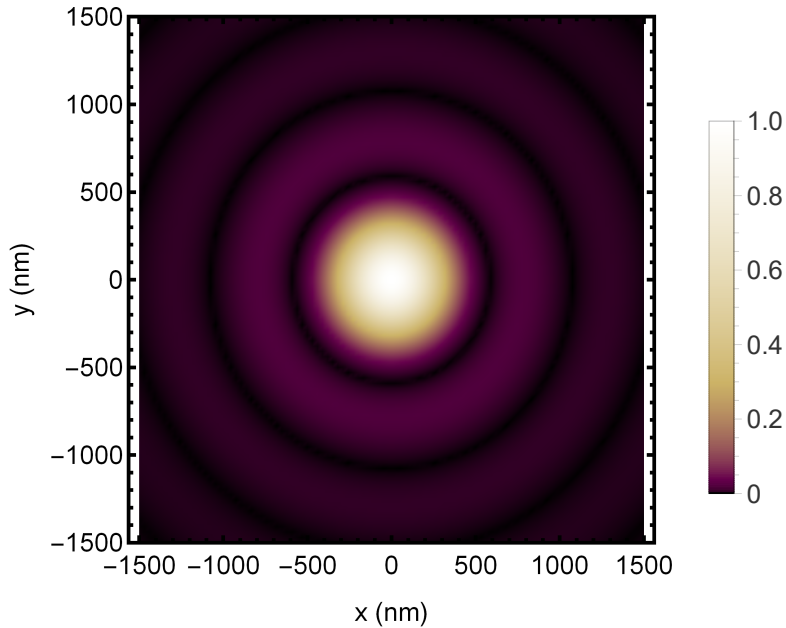
Figure 3.9: An Airy pattern plotted using (3.5) and our experimental values of k and NA. The bright central region is the Airy disc, and is surrounded by concentric bright rings. The colourmap has been chosen to accentuate the ring structure.

resolvable distance, such as the Abbe limit ($d_{\text{Abbe}} = 0.5\lambda/\text{NA}$) and the Sparrow limit ($d_{\text{Sparrow}} = 0.47\lambda/\text{NA}$). In a quantum gas microscope, our task is not to resolve atoms with arbitrary spacings, but rather to determine a distribution of atoms fixed to a periodic array of sites. The additional constraint on the interparticle spacings changes the nature of the imaging problem, and so these common definitions of imaging resolution are not exactly applicable. For instance, a quantum gas microscope working with ^{87}Rb was able to successfully reconstruct the distribution of atoms in a lattice with spacing $d = 532\,\text{nm}$ by imaging $780\,\text{nm}$ light with NA $= 0.68$ [20], a configuration which satisfies none of the above criteria.

In the presence of imaging aberrations and a small signal-to-noise ratio, it may not be possible to resolve all of the ringlike features of the Airy pattern shown in Figure 3.9. For this reason it can be instructive to compare the full-width at half-maximum (FWHM) of the experimental PSF to the theoretical value for an Airy pattern, to quantify how closely we achieve diffraction limited performance. The intensity of the Airy pattern (3.5) falls to half of its maximum value when $k\rho\text{NA} \approx 1.62$, so that the FWHM of the intensity distribution for a given λ and NA is

$$\text{FWHM}(k, \text{NA}) \approx \frac{1.62\lambda}{\pi\text{NA}} \approx \frac{\lambda}{2\text{NA}}. \tag{3.7}$$

For our experimental parameters, we would thus expect the FWHM of the diffraction-limited PSF to be $496\,\text{nm}$ based on scalar diffraction theory. As we will see later in Section 6.2, we experimentally measure the PSF of our imaging system to have a FWHM close to $600\,\text{nm}$, not far above this theoretical estimate.

3.5   Optical Lattice

The fluorescence imaging system described above allows for high-resolution detection of atoms located at the microscope focus, $800\,\mu$m inside the vacuum system. Using three retro-reflected beams, we form a cubic optical lattice in the focal position. This optical lattice is used to apply a periodic potential to our ultracold cloud for quantum simulation of the Hubbard model. After each instance of the simulation, the lattice potential is increased dramatically to localize atoms at individual sites of the lattice. High resolution fluorescence imaging can then be performed to determine the occupancy of all lattice sites, allowing precise measurement of the atomic density and various correlation functions.

The optical lattice is formed by three retro-reflected beams at $\lambda_L = 1053.6$ nm, crossing each other at roughly $90°$ and shifted in frequency with respect to one another by > 100 MHz to prevent stable cross-interference. This configuration results in a simple three-dimensional cubic lattice potential, with a lattice constant $a = 526.8$ nm. All beams are delivered to the experiment through separate polarization maintaining fibers (Oz Optics PMJ-3AHPC,3AHPC-1064-6/125) to ensure good pointing stability, and an optical isolator is placed at the output of each fiber to prevent the retro-reflected beams from being coupled back towards the laser source. The x-lattice beam is linearly polarized along $\hat{y}$, the y-lattice beam is linearly polarized along $\hat{z}$, and the z-lattice beam is linearly polarized along $(\sqrt{3}\hat{x}+\hat{y})/2$. Ideally the z beam would be polarized along $\hat{x}$, but this is difficult with the geometry of our experiment. Cross-interference between the z and x beams thus occurs in the experiment, but is time averaged to zero due to the large frequency difference between the beams. Additionally, if the two beams of the optical dipole trap are kept on during lattice loading further cross-interference is possible. This problem is similarly solved by using a large frequency difference.

The beams for the x-lattice and y-lattice are derived from an NP Photonics RFLM master laser, which seeds a fiber amplifier (Nufern SFA) to increase the power to a maximum of 40 W. The light from the amplifier is also used to form the optical dipole trap in the experiment, and a motorized waveplate (Pacific Laser Equipment) and polarizing beam splitter are used to divert power either to the dipole trapping beams or optical lattice beams as needed during the cycle. To generate the z-lattice beam, we use a second master laser (NP Photonics RFLSA) and Nufern amplifier, operating at nearly the same wavelength of 1053.6 nm but differing by several GHz in frequency. Using two separate fiber amplifiers allows us to operate each amplifier at 50% of its maximum output power while still obtaining lattice beam powers $(P_x, P_y, P_z) = (2.5, 2.5, 5)$ W. Running two amplifiers at 50% output is preferable to running one amplifier at 100% output, since this appears to improve the overall lifetime of the expensive equipment. Two optical isolators are used at the output of each fiber amplifier to protect backscattered beams from damaging the amplifier.

The lattice beams are focussed to waists $(w_x, w_y, w_z) = (60, 60, 85)\,\mu$m, providing maximum lattice depth $V_0 = 1200\,E_R$ for each axis, where $E_R = k_B \times 215$ nK is the recoil energy of ^{40}K for a photon with $\lambda_L = 1053.6$ nm. At this depth atoms in an ultracold gas are strongly confined to each site of the lattice, with negligible wavefunction overlap and essentially zero tunneling. Each lattice site is well approximated by a three-dimensional harmonic trap with oscillation frequency $\omega_0/2\pi = 300$ kHz for all three directions. For an atom harmonically trapped in such a lattice site and absorbing or emitting a photon at $\lambda_{D1} = 770$ nm, the trap energy level spacing is 33 times larger than the photon recoil energy. This large ratio between the trap energy and recoil energy is important for our efforts to perform laser cooling of the strongly trapped atoms, and is characterized by the Lamb-Dicke parameter $\eta = (E_R/\hbar\omega_0)^{1/2}$. For our experiment the Lamb-Dicke parameter for emission of a 770 nm photon is
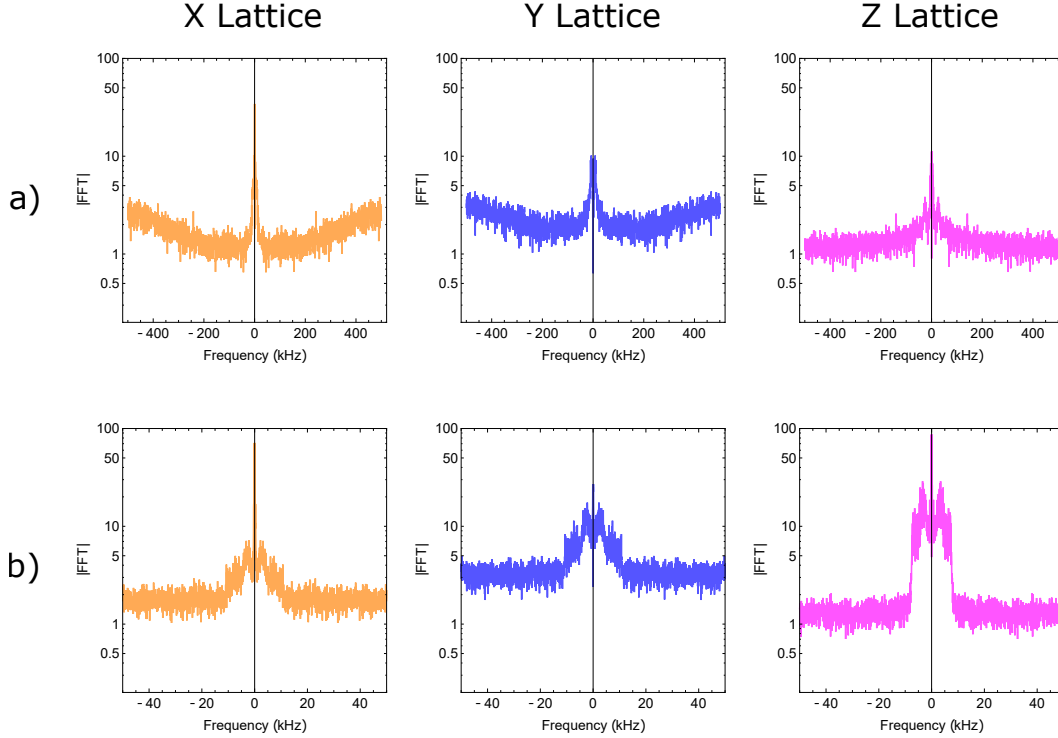
Figure 3.10: Lattice intensity noise measured with a photodiode. a) Frequency noise in the range of 100-500 kHz, which would resonantly excite harmonically trapped atoms in our deep optical lattice. b) Frequency noise in the range 10-50 kHz, comparable to the bandgap of of the lattice depths used for quantum simulation.

$\eta = 0.17$.

The power in each optical lattice beam is stabilized by reflecting a small fraction of the beam power onto a logarithmic photodiode. The logarithmic response of the photodiode is compared to the desired setpoint to generate an error signal, and PI feedback is applied to stabilize the power at the set point. Although the photodiode signal is logarithmic, the feedback signal is exponentially amplified before acting on the AOM drive amplitude, creating a feedback loop that is linear overall. Taking the logarithm before applying proportional-integral (PI) feedback allows for a large dynamic range, with similar fractional power stability at the powers relevant for Hubbard model simulation (a few tens of mW) to the powers relevant for fluorescence imaging (a few W).

The analog PI feedback has a bandwidth of roughly 1 kHz, which is smaller than the highest attainable on-site trapping frequencies of $\nu_{max} \approx 300$ kHz. It is therefore possible for relevant frequencies of amplitude noise present in the laser source to pass unfiltered into the optical lattice, leading to amplitude modulation of the potential and heating of the trapped atoms. We search for such amplitude noise by measuring the stabilized lattice beams with an independent linear photodiode. The power spectra of the optical lattice amplitude noise are shown in Figure 3.10. The spectra of the two horizontal lattice beams are indicative of the laser source they share, with amplitude noise growing steadily with increasing frequency due to a large noise peak at 1.5 MHz present in the master laser due to carrier relaxation. This high frequency behaviour is not seen in the vertical lattice beam, but increased noise is seen at frequencies < 15 kHz. The similar noise structure seen in all beams at frequencies < 15 kHz may be due to known electrical and vibrational noise sources in the fiber amplifiers [69].
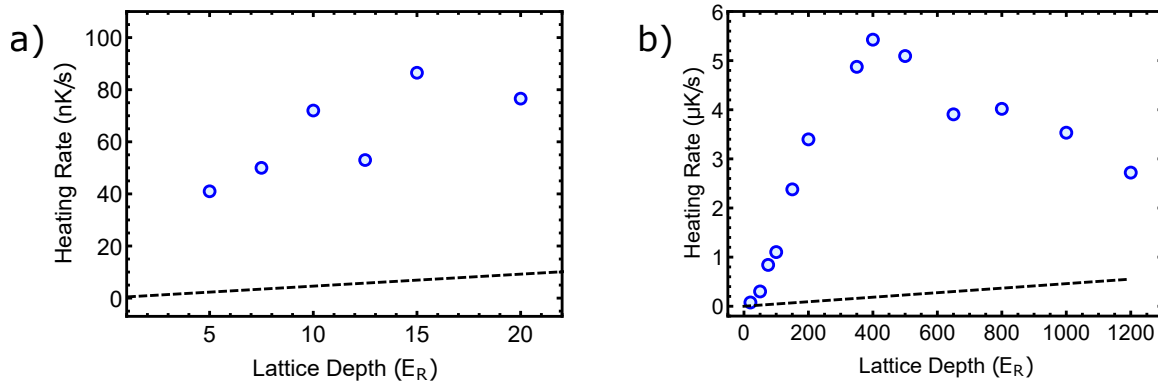
Figure 3.11: Measured heating rates in the 3D optical lattice. a) Heating rates in nK/s for low lattice depths appropriate for Hubbard model simulation. b) Heating rates in μK/s for the large lattice depths necessary for site-resolved fluorescence imaging. In both plots the black dashed line indicates the amount of photon recoil heating that would be expected due to the lattice scattering rate predicted by (2.3).

The amplitude noise present in the laser system can cause heating of the cloud by driving transitions between different motional states of the lattice potential. In a deep lattice characterized by harmonic trapping with frequency ω_0 in each lattice well, the symmetric perturbation introduced by amplitude modulation at a frequency ω_{mod} will give rise to heating primarily when $\omega_{\mathrm{mod}} = 2\omega_0$. Alternatively, translational shifts of the optical lattice caused by frequency noise or pointing instabilities of the beams result in a spatially antisymmetric perturbation of the harmonic trap, and cause heating when $\omega_{\mathrm{mod}} = \omega_0$. We can probe for the presence of all types of lattice heating by holding a cloud in a deep optical lattice for various hold times, and probing the resulting temperature with TOF absorption images. Heating rates for atoms held in the optical lattice are shown in Figure 3.11.

For quantum simulation we will use a shallow optical lattice of around $10\,E_R$ in the horizontal plane, with a slightly deeper lattice around $50\,E_R$ in the vertical axis to confine the system to 2D. For these lower lattice depths, the heating rates are around $60\,\mathrm{nK/s}$ as shown in Figure 3.11(a). These heating rates are not large enough to cause heating out of the ground band of the lattice during realistic experimental timescales of $< 500\,\mathrm{ms}$, since the bandgap in the optical lattice is typically on the order of $500\,\mathrm{nK}$. However, the many-body quantum states (such as a Mott insulator) which we would like to simulate in the experiment will have much smaller excitation energy scales. For instance, a recent experiment has shown the melting of a Mott insulating state of ^{40}K when the temperature increased from $k_BT/U = 0.18$ to $k_BT/U = 0.55$, roughly equal to a temperature increase of $15\,\mathrm{nK}$. If we then needed to keep the overall lattice heating lower than $5\,\mathrm{nK}$ for a quantum simulation experiment, our heating rates at $10\,E_R$ would limit us to experimental timescales $< 100\,\mathrm{ms}$.

After realising the Hubbard model in a low-depth lattice, we ramp up the lattice to a large depth $> 1000\,E_R$. We use this deep lattice to pin the atoms while we obtain a fluorescence image, and so heating rates must be small enough to keep atoms from heating out of their initial sites. At large lattice depths the measured heating rates are a few $\mu\mathrm{K/s}$, as shown in Figure 3.11(b), which could cause non-negligible heating out of the ground state during the $2\,\mathrm{s}$ exposure time needed for single-atom fluorescence imaging. However, during fluorescence imaging it is already necessary to apply laser cooling to the cloud to suppress heating due to the scattering of fluorescence photons. Thus the additional heating of the optical lattice does not concern us unless the combined heating effects from all sources

are greater than the cooling power of our laser cooling mechanism.

3.5.1   Alignment

To achieve adiabatic transfer from the optical dipole trap into the lattice, it is important that the 6 lattice beams are precisely aligned to the position of the dipole-trapped cloud. The x and y-lattice are aligned by first aligning the forward-propagating beam to the position of the atom cloud, then tuning the retro-reflection of the beam in order to maximize the contrast of the interference pattern. Since the z-lattice is retro-reflected inside the vacuum system, it is not possible to independently align the retro-reflected beam and the focussing and alignment procedure is more complicated.

Horizontal Lattice Alignment

We use dichroic mirrors in the optical path of the two horizontal lattice beams, allowing a resonant probe beam to be sent along the x or y axis for absorption imaging. This arrangement, shown in Figure 3.12, allows for a very simple coarse alignment of the forward travelling beams of the two horizontal lattice arms: the cloud position is determined by taking an absorption image, then the lattice beam is adjusted to the same location in the image plane by observing on the CCD camera the small fraction of power transmitted through the dichroic mirror of the retro path. The image of the lattice beam on the CCD camera can also be used to focus the forward beam. This alignment technique can still lead to errors due to small chromatic shifts in the imaging system, but reliably gets the lattice beam close enough to see some trapping effect on the cloud. The retro beam can also be aligned coarsely by adjusting its positioning by eye to pass back through all optical elements returning to the fiber, before the retro beam is rejected by the optical isolator.

To more finely align the forward beam we measure the force imparted to the cloud when quickly pulsing the beam while the cloud is in free-fall. Blocking the retro-reflection arm ensures that the cloud reacts only to the forward beam, which then has the effect of an optical dipole trap. To probe the alignment of the forward lattice beam we releasing a cold cloud from our crossed-beam dipole trap, and after $200\,\mu$s of free-fall we pulse the forward lattice beam for $0.5\,$ms with a power of $100\,$mW. The force F imparted on the centre of mass of the cloud at position x by the beam centered on position x_c is easily calculated from (2.1) and (2.4):

$$F(x) = -\frac{4V_0}{w_0^2}(x - x_c)e^{-2(x-x_c)^2/w_0^2}\,, \tag{3.8}$$

where V_0 is the peak potential depth. Letting the cloud fall for some time-of-flight after the lattice pulse maps this force onto the final cloud position. When the beam is aligned to the cloud with $x_c = x$ this force vanishes, and the cloud centre of mass falls to the same position with or without the lattice beam pulse — although the overall shape of the cloud after expansion may be perturbed by the pulse. We scan the beam alignment and monitor the cloud position after $15\,$ms time-of-flight, with typical results as shown in Figure 3.13. The large slope of the linear region near x_c allows us to align the beam more precisely than the length scale set by the beam waist. Although small changes in cloud position are easier to discern when probing a ^{87}Rb BEC, we always align using the ^{40}K cloud position which is slightly different due to gravitational sag (predicted by equation (2.10)).

If we hold the forward lattice beam at high intensity and completely ramp off the dipole trap, the
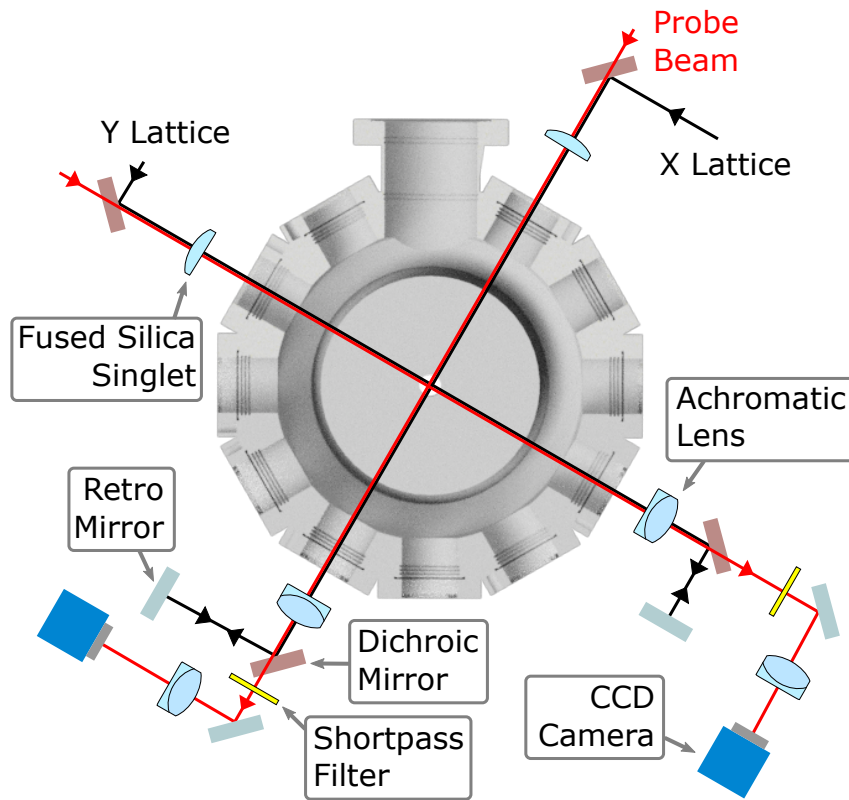
Figure 3.12: Layout of the horizontal lattice beams. Absorption imaging can be used to image the cloud along either of the lattice axes, providing a convenient alignment tool. Probe and lattice light are combined and separated using dichroic mirrors (Thorlabs DMSP805). Lattice beams are focussed on the atoms by a 200 mm fused silica singlet lens and collimated after exiting the chamber by an achromatic lens with 200 mm focal length (Thorlabs ACA254-200-B). The achromat has a $< 50\,\mu$m focal shift between the probe and lattice wavelengths, so that aligning the lens for lattice collimation simultaneously focuses the absorption imaging system. Using achromatic lenses in the entire imaging system ensures that the small fraction of lattice light transmitted by the dichroic mirror can be used to faithfully overlap the lattice beam with the cloud position. After the lattice has been aligned, the residual 1054 nm light transmitted by the dichroic mirror is blocked with a shortpass dielectric filter (Semrock FF01-945/SP). The CCD cameras used for absorption imaging are both PCO pixelfly qe with a pixel size of $6.45\,\mu$m.

Figure 3.13: Dipole Force Curve used to align the optical lattice beams. The solid line is a fit of Eq. (3.8).

cloud will be loaded into the elongated trapping potential formed by the single forward lattice beam. An in-situ absorption image taken along the axis of the lattice beam can then determine whether the cloud is pulled horizontally or vertically as it moves to the new trap centre, giving information about errors in beam alignment. Such an image of the trapped cloud position is a good indication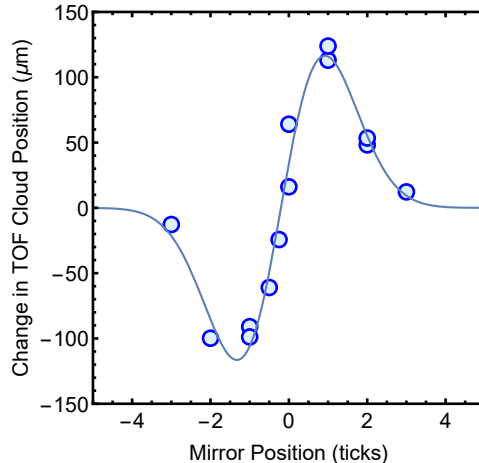 of the trapping beam centre only if the gravitational sag of the cloud is small. To minimize the effect of gravity, we use a power of at least $200\,\mathrm{mW}$, ensuring that the sag is $< 10\,\mu\mathrm{m}$. Further errors can occur if the cloud moves along the axis of the trapping beam and out of the imaging plane. We confirm with absorption images taken along a perpendicular direction that such cloud movement is small and these effects are negligible. Looking at in-situ images of this sort allows for a quick diagnosis over any lattice misalignment, but the full force curve is always used to perform alignment of the beams.

In order to align the retro-reflected beam of the horizontal lattice, we wish to maximize the depth of the interference pattern seen by the atoms. This can be done by performing Raman-Nath diffraction of a cloud to measure the energy gap between ground and second-excited bands [70], and adjusting the alignment in order to maximize the gap. In our system this is easily done with a ^{87}Rb BEC, but if the initial trap is weak this may align the lattice to the wrong vertical position due to the differential gravitational sag of ^{87}Rb and ^{40}K . Alternatively, diffraction measurements can be performed using the ^{40}K cloud so that gravitational sag presents no obstacle. However, unless a small cloud is used these measurements suffer from inhomogeneous broadening as the fermions spread out over a large range of lattice sites and quasimomenta, and therefore experience different lattice depths and energy gaps.

As an alternative method to finely align the lattice retro beams, we load the cloud into the combined potential of the well-aligned forward beam and the misaligned retro beam. As the retro beam position is adjusted, three cases are possible: (a) the retro beam misses the atoms, and the cloud is trapped only in the forward beam, (b) the retro beam is nearly aligned, and the cloud is pulled to a new trap minimum between the two beam centres, (c) the retro is completely collinear with the forward beam and the cloud position is unchanged from the position of the original dipole trap. An example of the measured cloud position versus retro beam position is shown in Fig. 3.14. Such a measurement is again subject to errors due to the effect of gravity, but these are again small when $> 200\,\mathrm{mW}$ is used in the lattice beam.

Since we align using a retro mirror at a distance of roughly f from the focusing lens, we primarily
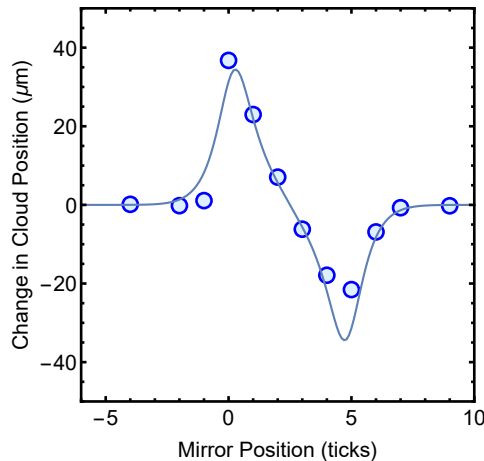
Figure 3.14: In-situ trap position measurement used to align the retro beam of the optical lattice.

align the position and not the angle of the retro beam at the location of the atoms. To consider the effect of aligning the retro beam while the atoms constantly feel the presence of the well-aligned forward beam, we can calculate the intensity distribution due to two Gaussian beams offset in one dimension by a distance a, at an axial coordinate for which the electric fields add constructively:

$$
\begin{aligned}
I(x) &\propto \left(E_0 e^{-x^2/w_0^2} + E_0 e^{-(x-a)^2/w_0^2}\right)^2 \\
&= E_0^2 e^{-2x^2/w_0^2} + E_0^2 e^{-2(x-a)^2/w_0^2} + 2E_0^2 e^{-x^2/w_0^2} e^{-(x-a)^2/w_0^2} .
\end{aligned}
\tag{3.9}
$$

For $a \gg w_0$, the constructive interference represented by the third term vanishes, and (3.9) is just the sum of Gaussian intensities for two beams. This results in two potential minima, and the atoms will be loaded into the minimum corresponding to the well-aligned forward beam. When when the distance a is comparable to w_0, the cloud will be pulled to a new point of maximum intensity. Finding the location of the trap minimum as a function of the retro alignment results in a curve quite similar to the force curve of (3.8), but with a reduced slope in the linear region. This curve is overlaid with the data in Fig. 3.14, showing good agreement.

The retro alignment based on in-situ cloud positions is convenient, but may result in errors due to gravitational sag or imperfections in the imaging system. To confirm that we have properly aligned the beams we measure the on-site trapping frequencies for ^{40}K atoms pinned in the lattice and compare to the theoretical expectation from equation (2.1) based on the known power, waist, and atomic polarizability. We can measure the on-site trapping frequencies by modulating the amplitude of the lattice beams, exciting atoms from the ground band of the lattice to excited bands. Population in the excited bands is easily counted in a band mapped image, so that the fraction of atoms excited can be reliably determined as a function of modulation frequency. In the limit of large lattice depths where atoms are harmonically trapped at each lattice site, the perturbation due to amplitude modulation is symmetric around the trap minimum, and so cannot excite atoms from the ground band to the first excited band. Instead we look for excitation from the ground band to the second excited band, at a modulation frequency very close to twice the trapping frequency of the lattice. An example of the excited fraction detected versus modulation frequency is shown in Fig. 3.15. The modulation frequencies used are larger than the feedback bandwidth for our lattice beams, and so the appropriate modulation is instead added in a
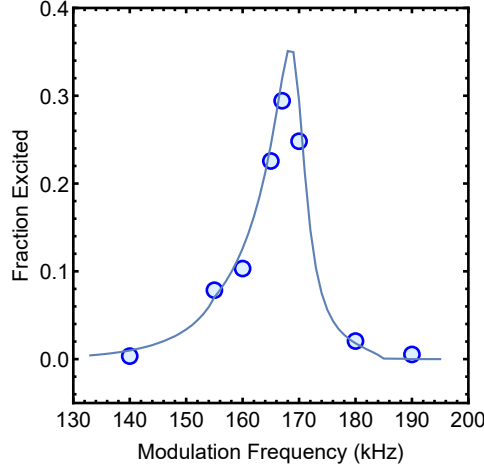
Figure 3.15: Excitation spectrum from amplitude modulation of the y-lattice. Expected lattice depth is $105\,E_R$, corresponding to a harmonic trapping frequency of $\nu_0 = 87\,\mathrm{kHz}$, and an excitation peak at $\nu = 170\,\mathrm{kHz}$ for amplitude modulation. The solid line shows the expected range of excitation frequencies for a cloud with $T/T_F = 0.6$ prior to lattice loading.

feed-forward configuration using a transformer to couple the modulation into the control circuit.

At large lattice depths, the energy bands of the lattice are quite flat and narrow, so that only a narrow range of frequencies would be expected to yield excitation to the second excited band. The excitation spectrum of Fig 3.15 is however homogeneously broadened by the Rabi frequency of the excitation and inhomogeneously broadened by the range of trap frequencies sampled by the ^{40}K cloud. By observing Rabi flopping of atoms between the ground and second excited bands, we estimate the Rabi frequency for the excitation to be $5\,\mathrm{kHz}$ for a modulation depth of 5% of the total beam power. To estimate the amount of inhomogeneous broadening, we can use the lattice density of states (the three-dimensional analogue to (2.65)) to calculate the density distribution we expect for a gas with given T/T_F. Then for a given modulation frequency we can calculate the number of sites we expect to be resonant, assuming the probability of excitation is given by a Lorentzian distribution with FWHM set by the Rabi frequency. The resulting curve is shown in Fig. 3.15 for $T/T_F = 0.6$ and a peak lattice depth of $100\,E_R$, and nicely reproduces the shape of the observed spectrum. This suggests that we have properly aligned the lattice to achieve the maximum depth expected for the given beam power. To overcome the limitations of inhomogeneous broadening, colder clouds are a benefit due to their increased density. Alternatively, with high-resolution in-situ imaging it is possible to measure the trapping frequencies for different regions of the lattice separately, as explained in the next section and shown in Figure 3.16.

Vertical Lattice Alignment

Alignment of the vertical lattice beam is more complicated, because the beam reflects off of the immovable sapphire window inside the vacuum system. The forward beam and its retro-reflection cannot be aligned separately and instead two mirrors must be walked to align simultaneously the beam position and angle to overlap with the cloud and create a strong interference pattern. This alignment is made easier because the distance from the atoms to the reflecting surface of the window is only around ten times larger than the beam waist of $85\,\mu$m. This allows a large range of incident angles to support at least partial overlap of the forward and retro beams. To coarsely align the beam to the position of the cloud we can remove the
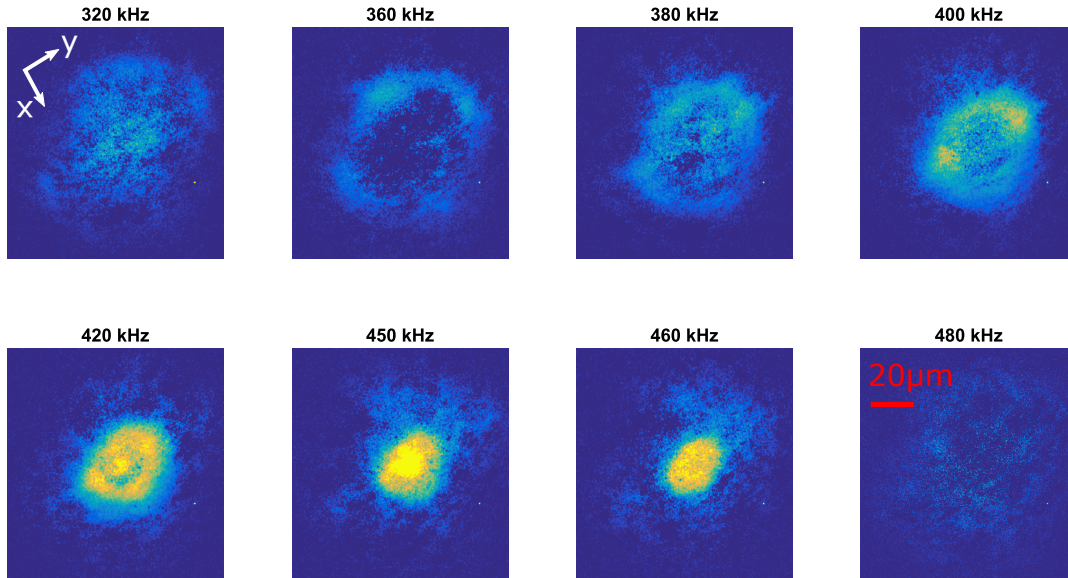
Figure 3.16: Examining the variation in on-site oscillation frequencies by taking in-situ images after amplitude modulation of the z-lattice. These images are produced by subtracting an exposure with modulation from an exposure of a similar cloud without modulation, producing an image of the cloud density that is excited by the modulation. With increasing frequency a circularly symmetric modulation pattern emerges, exciting atoms closer and closer to the centre of the trap. At 450 kHz modulation, the centre of the lattice is strongly excited, indicating that $\omega_z/2\pi = 225$ kHz in the centre of the lattice. Thus we can calibrate that the lattice depth in the centre of the trap is $680\,E_R$. The large range of possible excitation frequencies in the lattice poses a problem if we wish to site-resolve fluorescence from every lattice site visible in this field-of-view. However, the central region up to a radius of $15\,\mu$m has a very uniform on-site trapping frequency and contains more than enough lattice sites for quantum simulation of the Hubbard model.

microscope objective and instead use a low-magnification imaging system to acquire absorption images which show the cloud position in the x-y plane. A CCD camera can easily see the $< 0.01\%$ of vertical beam power which is transmitted through the sapphire window, and this signal can be used to overlap the vertical beam to the atom's position. Further, the transmitted light may be measured with the CCD at several distances from the window, giving a measurement of waist versus distance which can be used to ensure proper focussing of the lattice light.

To align finely the vertical lattice, we first use pulses with power > 100 mW to look for diffraction of a ^{87}Rb BEC during time-of-flight. Once some diffraction is observed, the power is reduced and the beam is carefully walked while using the diffraction signal to probe the lattice depth created. Unlike the horizontal lattice beams gravitational sag plays little role, so that alignment to the ^{87}Rb BEC gives a good alignment to the ^{40}K cloud as well. Any final alignment of the lattice beam to ensure spatial overlap with the ^{40}K cloud is done by adjusting a mirror located f away from the focusing lens, thereby changing the position of the beam without a large effect on the retro-reflection angle.

Calibration of the vertical lattice depth can be performed using amplitude modulation, as with the horizontal beams. These measurements still suffer from inhomogeneous broadening, but we can determine that the alignment procedure outlined above routinely produces an optical lattice with the

expected depth. Since we do not independently align the forward and retro-reflected beams, we must be careful to ensure that the interference pattern formed retains the circular symmetry of the individual beams. To probe this alignment of the vertical lattice, we use in-situ fluorescence images[1] to examine the cloud after amplitude modulation has strongly excited atoms at a given on-site trap frequency. Strong modulation causes atoms to disappear from the image — either because they were removed from the trap or because they cannot be efficiently imaged after excitation. By subtracting the resulting image from a reference image without any modulation, the spatial distribution of atoms resonant with a given modulation frequency is obtained. An example of this mapping is shown in Fig. 3.16 for a lattice depth of $700\, E_R$, and shows that the regions of constant trap frequency do possess the desired circular symmetry.

The in-situ measurement of the on-site trapping frequency of the vertical lattice gives us confidence in the lattice alignment, and also informs us of the effects of trap inhomogeneity in our system. The same technique can be used to probe the on-site trapping frequencies of the two horizontal lattice beams, to create a full picture of the trap confinement felt by atoms in a particular region of the lattice. An accurate measurement of the trapping frequency for the confinement at each site is crucial for our attempts to obtain site-resolved fluorescence images, since we prevent atom loss by cooling the bound atoms towards the ground state of the confining potential. The different trap frequencies experienced at different lattice sites mean that not all atoms in a cloud can be cooled effectively, leading to loss of atoms during fluorescence image acquisition. This is equivalent to a spatially dependent scattering rate for the atoms during the fluorescence image, so that the fluorescence image that we acquire has a different structure from the density distribution of the atoms we wish to measure. As seen in Figure 3.16 the on-site trapping frequency variation between the centre and the visible edges of the lattice can be as large as $100\,\text{kHz}$ when we use the large lattice depths appropriate for fluorescence imaging. However, this effect does not prevent us from faithfully imaging a cloud prepared in a simulation of the Hubbard model. As discussed in Section 2.3.7, only the central region of the optical lattice for Hubbard model simulations. In the centre of the lattice with a diameter of $15\,\mu$m, the variation in on-site trapping frequencies is $< 20\,\text{kHz}$ for our highest lattice depths – small enough for all atoms in this region to be cooled and imaged efficiently.

With a well aligned optical lattice to tightly confine atoms in the focus of a high-NA microscope objective, our experimental apparatus is capable of collecting any emitted fluorescence photons efficiently and with high spatial resolution. The remaining challenges are to load the appropriate ultracold gas of atoms into the lattice, and to implement a laser cooling scheme which can generate a steady stream of fluorescence photons while keeping the atoms confined to the individual lattice sites. Our solutions to these two challenges will be outlined separately in the following chapters.

---

[1]The details of our in-situ fluorescence images will be explained in Chapters 5 and 6

Chapter 4

State Preparation

The path which starts with room temperature gases of ^{40}K and ^{87}Rb and end up with quantum simulation of the Hubbard model is long and filled with many hardships. A great deal of work is necessary to configure all of the aspects of the experiment, and coax them to cooperate in the production of a quantum degenerate Fermi gas. The thesis of Dylan Jervis [66] contains extensive details of the laser cooling and the subsequent evaporative cooling stages which we use to reduce the cloud temperature below $10\,\mu$K. Here I will describe in detail only the final stages of each experimental cycle, in which the cloud is cooled to the point of quantum degeneracy, loaded into the optical lattice, and trimmed into a single 2D slice.

4.1 Degenerate Fermi Gas

Our experiment begins by using a magneto-optical trap (MOT) to cool and trap a cloud of ^{40}K and ^{87}Rb atoms from a background vapour. After collecting roughly 3×10^9 ^{87}Rb and 1×10^8 ^{40}K atoms at a temperature of $200\,\mu$K, the cold cloud is loaded into a magnetic trap and transported into a separate chamber of the vacuum system with a lower background pressure and greater optical access. In this 'science chamber' the ^{87}Rb atoms are cooled evaporatively by using RF radiation to selectively flip the spins of the hottest atoms, leading to their ejection from the magnetic trap. ^{40}K is not directly evaporated, but is cooled sympathetically through collisions with ^{87}Rb [71, 72]. Although the centre of symmetry of the quadrupole magnetic trap is near the focus of the microscope objective, evaporation of ^{87}Rb initially takes place with the trap centre displaced by $6.6\,$mm in order to accommodate the large initial size of the cloud. The trap displacement is achieved by an imbalance in the currents running through the two quadrupole coils. After the cloud has been evaporatively cooled to $100\,\mu$K the current imbalance is reduced to zero, moving the cloud into the imaging position $800\,\mu$m from the sapphire window. Here a final stage of ^{87}Rb evaporation takes place, aided by a blue detuned 'plug' beam to keep atoms from reaching the zero-field minimum of the magnetic trap and thereby preventing Majorana loss of atoms due to spin flips [36]. After the final stage of RF evaporation the cloud temperature is $7\,\mu$K, and the previously large population imbalance of ^{87}Rb to ^{40}K has been eliminated so that about 10^6 atoms of each species remain in the trap. More detail of the RF evaporative cooling sequence can be found in the thesis of Dylan Jervis [66].

With RF evaporation finished, the cloud is transferred to a crossed optical dipole trap, which is

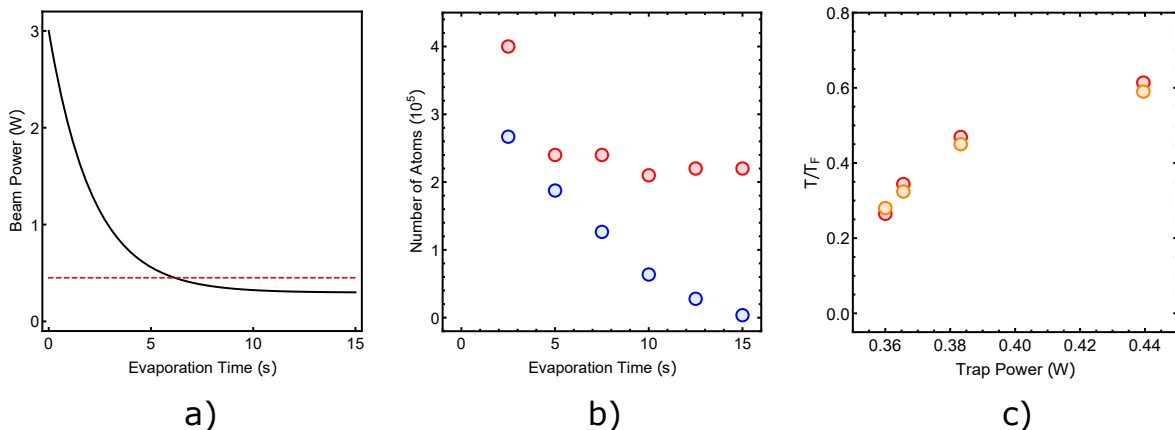a)   b)   c)

Figure 4.1: Sympathetic cooling process. (a) The power (black) of the two dipole beams ramps down exponentially with a time constant of roughly 2 s. At powers lower than 0.45 W (dotted red line), gravity causes a preferential loss of ^{87}Rb atoms from the trap, giving efficient sympathetic cooling. (b) Number of ^{87}Rb (blue) and ^{40}K (red) atoms remaining in the dipole trap during the evaporation ramp. Initially atoms of both species are lost, but when beam powers are reduced below 0.5 W we see a loss of only ^{87}Rb . The number of ^{40}K atoms is determined from absorption images with an effective cross section $\sigma_{\text{eff}} = \sigma_{\text{sc}}/3.3$ as described in Section 2.2.3 (c) Degeneracy parameter T/T_{F} during the final half of the evaporation ramp. Red points use the measured temperature of the ^{87}Rb cloud to extract T, and orange points show the temperature determined from a fit of the Fermi-Dirac momentum distribution (2.30).

formed from beams with an elliptical profile to minimize gravitational sag (as mentioned in Section 2.1.3). The change in trapping potential shape results in an adiabatic cooling of the cloud to $3.5\,\mu$K. The mean trapping frequency is $\bar{\omega}/2\pi = 250\,$Hz, so that the Fermi temperature is $T_F \approx 2.2\,\mu$K and $T/T_F \approx 1.5$. To reach lower T/T_F we lower the power of the optical trapping beams to accomplish further evaporative cooling, with the ^{87}Rb and ^{40}K clouds remaining in thermal contact. For effective sympathetic cooling of ^{40}K , we wish to selectively remove high temperature ^{87}Rb atoms from the trap. In RF evaporation ^{87}Rb can be selectively addressed since it experiences a larger linear Zeeman effect than ^{40}K , shifting its RF transitions to higher frequencies. In an optical trap evaporation, where high-temperature atoms are allowed to escape by a reduction of the trap depth, it is desirable that the trap depth for ^{87}Rb should be smaller than the trap depth for ^{40}K for efficient sympathetic cooling. Since the polarizability of ^{87}Rb is slightly larger, and because our dipole trap at 1054 nm is farther detuned from ^{40}K than from ^{87}Rb , the dipole trap potential is roughly 10% smaller for ^{40}K . Thus in the absence of any other potential hot atoms of either species may escape the trap, with the loss of ^{40}K atoms slightly more likely. This still results in cooling of the spin-polarized ^{40}K gas, as the fermions are able to collide and re-thermalize with the ^{87}Rb gas.

At the initial optical trapping powers of 3 W, atoms are most likely to be lost from the crossing region by escaping in the horizontal direction, leaving the focus of one trapping beam by moving along the optical axis of the other. In this way an atom escapes the crossing region of the trap without fully escaping into free space. On the other hand, to escape the trap along the vertical direction requires leaving the attractive potential of both beams simultaneously, giving a depth which is twice as large. However, as we reduce the beam power the additional effect of gravity on the trap depth becomes more important, tilting the trapping potential until loss along the vertical direction becomes most likely. With our trap parameters this change in loss direction happens for ^{87}Rb atoms at a trap power of 0.45 W,

and for ^{40}K at a power of 0.23 W. For powers between these two values, the trap depth for the crossing region is lower for ^{87}Rb than for ^{40}K , and so a more efficient sympathetic cooling would be expected to occur.

The details of sympathetic cooling in the dipole trap are shown in Figure 4.1. After loading from the QP trap at a power of 3 W per beam, we ramp the powers down to a final power of 350 mW. The ramp has an exponential time constant of roughly 2 s, and ends after a total evaporation time of 15 s. Over the first 5 s of evaporation we extract the temperature of both clouds from Gaussian fits to TOF absorption images, confirming that the ^{40}K and ^{87}Rb remain at the same temperature. At evaporation times larger than 5 s, where trap powers drop below 0.5 W, a Gaussian momentum distribution fails to accurately fit the shape of the expanding ^{40}K cloud. This is due the Fermi gas approaching temperatures of $T/T_F < 0.5$, where quantum statistics become relevant and the momentum distribution assumes the non-Gaussian form of Eq. (2.30). We determine the temperature through fits of Eq. (2.30) to a TOF absorption image as described in Section 2.2.3, giving a temperature of $T/T_F = 0.28$ at the end of evaporation. At the same time the ^{87}Rb momentum distribution is still Gaussian since we do not cross the critical temperature T_c for Bose-Einstein condensation. As a second check of this temperature, if we assume thermal equilibrium persists between the two species we can calculate T/T_F for the Fermi cloud by measuring the temperature of ^{87}Rb with a Gaussian fit and comparing to the known value of T_F. This calculation suggests a final temperature of $T/T_F = 0.27$, in agreement with the value obtained from the TOF shape of the ^{40}K cloud. Temperatures determined through both methods are shown for the lowest evaporation powers in Figure 4.1(c).

After the power ramp is complete, the ^{87}Rb cloud is gone, and we are left with a degenerate Fermi Gas of roughly 2×10^5 atoms. Any ^{87}Rb atoms remaining in the trap are removed with a pulse of resonant light. At this point we load the cold ^{40}K cloud into the optical lattice for quantum simulation and subsequent fluorescence imaging. When calibrating our fluorescence imaging technique, we prefer to use a cloud with low density in order to more easily identify single atoms and determine the orientation and spacing of the optical lattice. For this purpose we alter the optical evaporation sequence to perform the entire power ramp in 1 s, followed by a 14 s hold time in the trap at low depth. This process produces a smaller and hotter cloud of ^{40}K , which will fill the lattice at low density. The 14 s wait time serves to keep the overall cycle time of the experiment constant, so that thermal stability of the apparatus is not disrupted by switching between the two evaporation schemes.

4.2 ^{40}K Atoms in the Lattice

When we turn on the optical lattice beams to load our degenerate Fermi gas into the new potential, the cloud will naturally be distributed over a range of lattice sites in three dimensions. Since our fluorescence imaging can only resolve lattice sites in a single 2D lattice plane, it is necessary to reduce the dimensionality of the system at some stage before image acquisition. If we are interested in 2D physics, we may first load atoms into a deep vertical lattice, separating the cloud into a collection of disconnected 2D systems, before adding weak lattice potentials along the two horizontal axes to implement a Hubbard Hamiltonian for the gas. Before fluorescence imaging can be performed, a single 2D system must be isolated from the rest by a *plane selection* process. The plane selection can take place when atoms are confined by only one lattice axis, or after all lattice axes have been ramped on. To minimize the effect of any atom loss or heating occurring during plane selection, it is preferable to perform the selection early
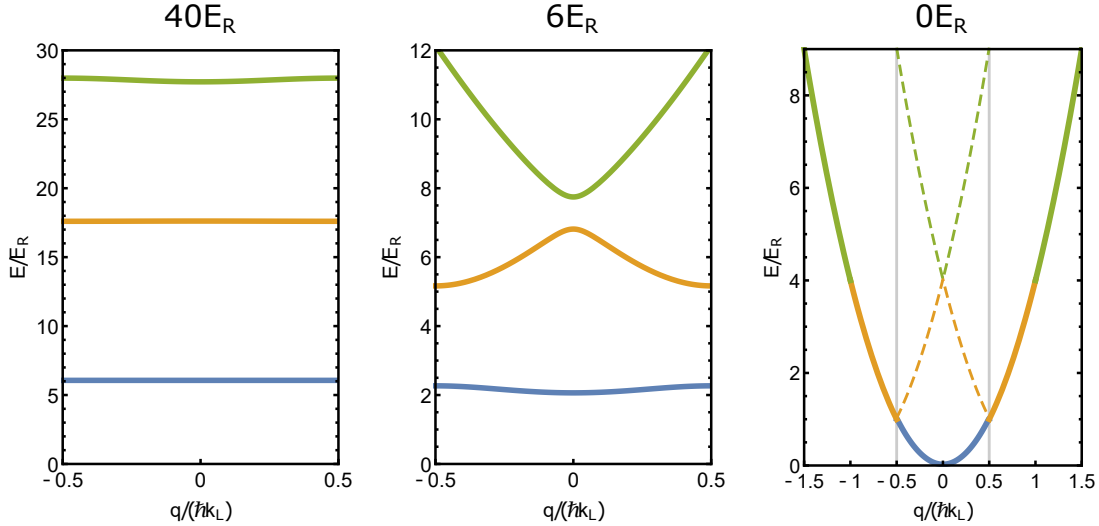
Figure 4.2: Illustration of the band mapping process. As the lattice depth is reduced the shape of the energy bands smoothly transforms back into the free-particle dispersion relation. If the lattice depth is reduced smoothly, the equally spaced quasimomentum states of the bandstructure will be transformed into the corresponding set of equally spaced plane waves with momentum $k = q$. The first two plots are shown in the reduced-zone scheme, while the third plot for free particles overlays the extended-zone (solid lines) and reduced-zone energies (dashed lines) for the excited bands. Gray vertical lines indicate the edges of the first Brilluoin zone. The ground, first-excited, and second-excited bands of the lattice (shown in blue, orange and green respectively) are mapped into distinct ranges of real momentum, and so can be counted in an absorption image.

in the cycle so that further cooling can be performed before the final state is reached. This arrangement leaves a minimal number of steps occurring between Hubbard model simulation and fluorescence state readout, maximizing the fidelity with which we determine the properties of the quantum state.

If we instead would like to study atoms in a full 3D lattice potential, we can omit the step of first loading a deep vertical lattice and instead ramp up all three lattice axes simultaneously. Before fluorescence imaging can occur, a single slice must still be isolated from the full 3D distribution by plane selection. In this case, plane selection may lead to changes of the atomic density distribution before fluorescence imaging takes place, so that the fidelity of plane selection must be much higher if accurate state determination is required. However, if we only wish to create a sparse 2D distribution of atoms to calibrate our fluorescence imaging the plane selection protocol can operate with a much greater tolerance for error. For all of the results in this thesis where we are primarily interested the fidelity of fluorescence imaging rather than properties of the many-body state, we first load a sparse 3D lattice distribution and only perform plane selection immediately before a fluorescence image is acquired.

Before we select a single plane of the 3D optical lattice, we still work with a large enough atom number to perform accurate absorption imaging using our cameras oriented along the x and y lattice axes. This allows us to perform lattice alignment as mentioned in Section 3.5.1, probe the on-site trapping frequencies by detecting excitation to higher bands, and assess the temperature limits reached by our in-lattice laser cooling techniques (see Section 5.4). The latter two measurements are most easily performed by measuring the quasimomentum distribution of the lattice gas by *band mapping* the cloud before probing the momentum distribution with an absorption image. To perform band mapping we ramp down the lattice over $1\,\mathrm{ms}$ with a time constant of $200\,\mu\mathrm{s}$, so that the potential landscape is
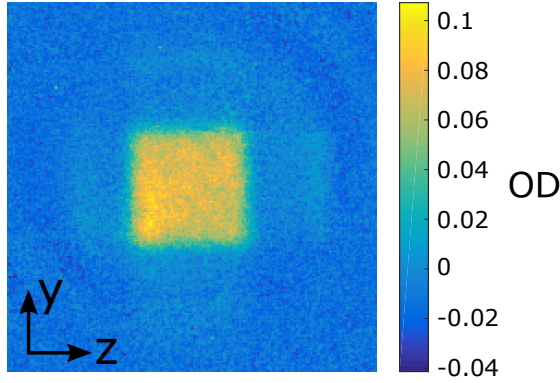
Figure 4.3: Image of ^{40}K atoms after band mapping and TOF expansion. The square momentum distribution in y and z reflects the limit $|q| \leq \hbar k_L/2$ imposed on the quasimomentum by Bragg scattering in the lattice. A small fraction of atoms are faintly visible in the second excited band of both lattice directions. This is suggestive of a excitation of atoms from the ground band by small residual amplitude modulation of the lattice.

adiabatically transformed from a lattice with energy bands into free space where atoms have the free-particle dispersion relation. Since the periodicity of the lattice is kept constant and only its amplitude is changed, quasimomentum in the lattice potential is conserved by this ramp [73]. As long as the ramp is performed slowly compared to the lattice bandgaps, the band index will be preserved as well. The result is that the Bloch waves $\phi_q^{(n)}$ of the deep lattice potential are mapped into free-particle plane waves which can be easily detected by our usual TOF absorption imaging.

An example of a band mapped absorption image of 2×10^5 ^{40}K atoms loaded into a $10\,E_R$ 3D lattice is shown in Figure 4.3. Due to the band mapping process the TOF distribution of the atoms reflects the physics of the lattice potential, with the cloud expansion lacking spherical symmetry and instead forming a cube oriented along the lattice directions. The square structure we observe when integrating the density along $\hat{x}$ is delimited by edges with momenta $\pm\hbar k_L/2$, showing that the atoms occupy the full range of unique quasimomentum values. For delocalized atoms in the centre of the lattice, all quasimomentum states will only become occupied only when atoms fill the entire energy band, which would indicate the presence of an band-insulating state with unit occupation in the centre of the lattice. We may then be tempted to conclude that this band mapped image is measuring the presence of Hubbard model physics in our lattice. However at this lattice depth the characteristic radius ξ of the harmonic confinement is only about 5 lattice sites, so that the characteristic number of particles which may occupy delocalized states in 3D is very small ($N_0 \approx 500$). This means that Figure 4.3 shows a band mapped system with $N_0/N \to 0$ where all atoms can be considered to be localized to individual lattice sites. The even distribution of all available quasimomentum states occurs because the localized atoms each occupy the lowest harmonic oscillator state inside an individual lattice well. The nth harmonic oscillator eigenstate decomposes into Bloch states $\phi_q^{(n)}$ with the same band index, and with equal population over all q [73]. Thus a collection of atoms in localized states will always appear to fill all possible quasimomenta of the ground band.

We have seen that in our optical lattice with strong harmonic confinement, only few atoms in the centre of the lattice can be considered delocalized. If we wish to use absorption imaging to readout

the results of Hubbard model simulation, this poses a severe limitation as any absorption signal will be dominated by the large population of localized atoms surrounding the lattice centre. This limitation does not apply to our site-resolved fluorescence imaging system, which can locally measure a small number of atoms in the centre of the lattice. Since we use absorption imaging for many experimental benchmarks, it is still important to be able to describe the temperature and density of the cloud as a whole. Much of the behaviour of a large cloud in a strongly confined lattice can be gained by considering only the localized atoms in the thermodynamics, ignoring entirely the delocalized atoms in the lattice centre.

If we load a large number of ^{40}K atoms into a 3D optical lattice, so that most particles occupy localized states, then we can accurately describe the resulting density distribution by working in the zero-tunneling limit. Here the density of states is determined only by the energy shift of each lattice site due to the external harmonic confinement with mean trap frequency $\bar{\omega}$. The density of states for a lattice in this limit was given in Eq. (2.65), and the same calculation for a 3D potential gives

$$g_{3\mathrm{D}}(E) = 2\pi E^{1/2} \left(\frac{2}{m\bar{\omega}^2 a_L^2} \right)^{3/2} . \tag{4.1}$$

Integrating this density distribution for energies up to the Fermi energy in the lattice $(E_F)_L$ gives the expression relating the lattice Fermi energy to the number of particles:

$$(E_F)_L = \frac{\hbar\bar{\omega}}{2} \left(\frac{4\pi}{3} \right)^{-2/3} \left(\frac{a_L}{l} \right)^2 N^{2/3} . \tag{4.2}$$

Integrating the density of states at finite T by using the Fermi-Dirac distribution, we find the expression relating $(T/T_F)_\mathrm{L}$ to the fugacity ζ_L of the gas in the lattice:

$$- \mathrm{Li}_{3/2} \left(-\zeta_L \right) = \frac{1}{\Gamma(5/2)(T/T_F)_\mathrm{L}^{3/2}} . \tag{4.3}$$

Subscripts are used for $(T/T_F)_L$, $(E_F)_L$, and ζ_L because the density of states is different for the lattice $(g_{3D}(E) \propto E^{1/2})$ and the harmonic trap $(g(E) \propto E^2)$, so that a cloud with the same atom number and entropy will have different temperature and fugacity in the two potentials.

To determine $(T/T_F)_L$ for the cloud loaded into the lattice, we can determine the entropy per particle S/N of the initial gas in the harmonic trap, and find the value of $(T/T_F)_L$ which has the same total entropy. Loading from one potential to the other is assumed to be adiabatic, so that the entropy of the gas is unchanged, but the temperature of the gas can increase or decrease due to the change in density of states. By comparing the dependence of S/N on T, it is found that for initial temperatures $T/T_F \ll 1$ heating by a factor of 2 is expected after lattice loading [55]. As shown in Figure 4.4, at temperatures greater than $T/T_F = 0.15$ the heating is even more extreme.

In order to determine what we lattice filling we might expect to see when obtaining a site-resolved fluorescence image, we can determine the density distribution of fermions in a lattice in the zero-tunneling limit. Since there is no kinetic energy term in the zero-tunneling limit, integration of the Fermi-Dirac distribution over all possible momentum states is trivial. As a result, the density distribution is specified only in terms of the energy $E_{\vec{j}} = \frac{1}{2}ma_L^2 \sum \omega_i^2 j_i^2$ of the external confinement at each site $\vec{j} = (j_x, j_y, j_z)$:

$$n_\mathrm{L}(\vec{j}) = a_L^{-3} \frac{1}{\zeta_L^{-1} e^{\beta E_{\vec{j}}} + 1} . \tag{4.4}$$
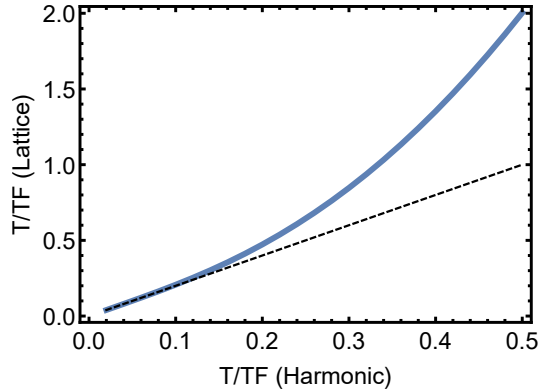
Figure 4.4: T/T_F in the lattice after loading at constant entropy from a harmonic trap. The dashed line shows the low temperature limit in which $(T/T_F)_L = 2T/T_F$. The blue solid curve shows the temperature increase determined by numerically matching S/N for the two clouds. For $(T/T_F)_{\rm SHO} > 0.2$, the heating during lattice loading becomes increasingly severe.

At the centre of the lattice the peak lattice filling $n_{\rm Peak}$ is thus determined only by ζ:

$$n_{\rm Peak} \equiv n_{\rm L}(0)a_L^3 = \frac{\zeta_L}{1 + \zeta_L} \, , \tag{4.5}$$

where the maximum value of $n_{\rm Peak} = 1$ (unity occupation of each lattice site) occurs at zero temperature when $\zeta_L \to \infty$. If we adiabatically load a cloud with initial temperature $T/T_F = 0.3$, we find that the initial fugacity $\zeta \approx 10$ reduces to $\zeta_L \approx 1.4$, and we would expect the localized particles to have $n_{\rm Peak} \approx 50\%$. However, if we speed through the dipole trap evaporation in 1 s, so that no evaporation occurs, we instead expect $n_{\rm Peak} \approx 1\%$. This low lattice filling is useful for calibration of the fluorescence imaging procedure.

The lattice filling in the centre of the lattice where atoms exist in delocalized states cannot be accurately determined using these calculations in the zero-tunneling limit. However, in order to realize the band-insulating or Mott-insulating phases of the Hubbard model with characteristic filling $\rho \approx 1$, we still expect that colder initial temperatures than $T/T_F = 0.3$ will be required before lattice loading. Evaporative cooling in the dipole trap may be improved by using two internal states of ^{40}K and tuning the interaction strength using the Feshbach resonance at 202 G. Alternatively, loading the cloud into the vertical lattice beam and performing further evaporative cooling has been shown to be effective for the production of cold lattice-gases of both ^{40}K and ^{6}Li [59,74]. This kind of evaporation can be controlled with the use of a tightly confining "dimple" trap, which we have incorporated into our experimental apparatus for the future effort to produce colder clouds.

4.3   Plane Selection

We always load our quantum gas into a 3D lattice, but our fluorescence imaging system only images the x-y plane, integrating all signal along the z-axis. We can use this total fluorescence to probe the density distribution of the 3D lattice, but we cannot infer the lattice occupancy at the level of single lattice sites. In order to read out information about site occupancy and local correlations, we image only a single plane of the z-lattice. In this case lattice sites can be uniquely associated with groups of camera pixels,
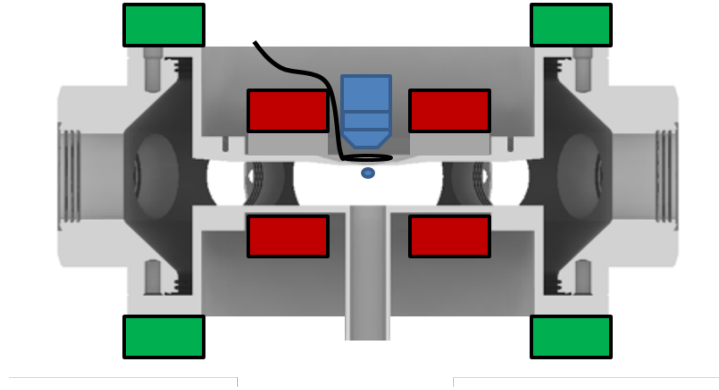
Figure 4.5: A vertical slice through the vacuum system, showing the two pairs of coils used for plane selection. The anti-Helmholtz coils used to produce the quadrupole gradient are shown in red, and the Helmholtz coils used to produce a homogeneous bias field along z are shown in green. The green coils are capable of producing fields up to $210\,\mathrm{G}$, and so can be used to access the Feshbach resonance between the two lowest energy internal states of $^{40}\mathrm{K}$. The bias coils are formed from hollow copper wire with $2\,\mathrm{mm}$ ID and $4\,\mathrm{mm}$ OD, permitting cooling water to flow through the coil. The temperature of the cooling water is stabilized using a thermoelectric chiller (Solid State Cooling Systems Thermocube 600L). The quadrupole coils are cooled with water flowing through a single loop of copper tubing (visible in Figure 3.4) stabilized by a second identical chiller.

and we can attempt to measure the lattice occupation site by site. Going beyond the imaging of only a single plane, it has also been shown that imaging and manipulation of two adjacent lattice planes can be used to provide additional information on the number of atoms in a lattice site or their internal hyperfine state [75]. The occupancy of multiple planes can be determined by taking multiple images interspersed with removal of individual planes [75], or by carefully characterizing imaging aberrations that lead a to plane-specific PSF [17]. At present we only measure the lattice occupancy using a single-plane, purely two-dimensional system.

There are many possible ways to prepare a two-dimensional sample: compressing the cloud into the node of a blue-detuned Hermite-Gauss beam [43], loading the cloud into a variable periodicity 'accordion lattice' [34], compressing the cloud with a strongly elliptical trapping beam [35], or by spectroscopically selecting a single plane out of a larger sample [20, 31, 32, 76]. We use the technique of spectroscopic selection, first flipping the internal state of a single plane of the lattice before removing all other planes with a pulse of resonant light.

In order to perform internal state manipulations on a single plane of the optical lattice we apply a large magnetic field gradient of $105\,\mathrm{G/cm}$ along the $\hat{z}$-direction, equivalent to a magnetic field difference of $5.5\,\mathrm{mG}$ from one lattice plane to the next. Since each plane of the lattice sees a different magnetic field, it is then possible to perform RF or microwave manipulations which are resonant with only one plane. After the atoms in the desired plane have been transferred to a new internal state $|F', m'_{\mathrm{F}}\rangle$, the undesired atoms in the original state $|F, m_{\mathrm{F}}\rangle$ can be removed using optical pulses or state-dependent potentials. It is important to perform this selection and removal with high fidelity since even a few atoms remaining in adjacent planes can confound the interpretation of fluorescence images.
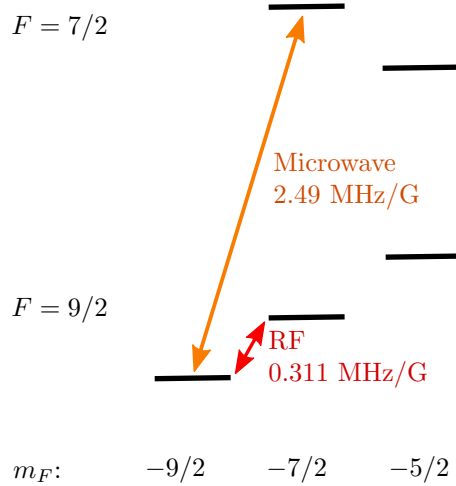
Figure 4.6: Hyperfine structure of the $4S_{1/2}$ ground state of ^{40}K relevant for plane selection. The two lowest energy states $|9/2, -9/2\rangle$ and $|9/2, -7/2\rangle$ form a closed spin-$\frac{1}{2}$ system which can be used to simulate the Hubbard model. Transfer from $F = 9/2$ to $F = 7/2$ happens at microwave frequencies near 1.285 GHz. Each pair of adjacent m_F states in the same F manifold have slightly different energy splittings determined by the quadratic Zeeman shift.

4.3.1 Magnetic Field Gradient

The magnetic field gradient for plane selection is provided by two pairs of coils, the anti-Helmholtz coil pair used for the Quadrupole Magnetic Trap (Quadrupole coils), and the Helmholtz coil pair used to apply fields up to 210 G for access to Feshbach resonances (Feshbach coils). Using both coils together results in the field

$$\vec{B}(x, y, z) = -\frac{a}{2}x\hat{x} - \frac{a}{2}y\hat{y} + (az + B_0)\hat{z}\,, \tag{4.6}$$

where a is the gradient created by the Quadrupole coils, and B_0 is the bias field generated by the Feshbach coils. The addition of B_0 displaces the Quadrupole gradient's centre of symmetry, creating a nearly linear gradient along the $\hat{z}$-direction in the position of the atoms. The remaining dependence of the field magnitude $|\vec{B}|$ on x and y is small, and for x and y much smaller than B_0/a the field magnitude can be approximated by

$$|\vec{B}|(x, y, z = 0) = \frac{a^2}{8B_0}(x^2 + y^2)\,. \tag{4.7}$$

We typically use $B_0 = 40$ G, and $a = 105$ G/cm, resulting in a vertical gradient of 105 G/cm. This splits adjacent planes by 5.5 mG, with residual field variations within each plane of less than 1 mG within a 50 μm transverse radius.

4.3.2 Microwave Spectroscopy

With each plane of the lattice experiencing a different magnetic field, it becomes possible to perform internal state manipulation on a specific plane by applying RF and microwave radiation. In an applied magnetic field, the linear Zeeman effect separates the individual m_F-states in each hyperfine manifold by aprooximately 0.311 MHz/G. In a 105 G/cm gradient the resonant RF frequency to transfer $|9/2, -9/2\rangle \rightarrow |9/2, -7/2\rangle$ will differ in adjacent lattice planes by roughly 1.7 kHz. Alternatively mi-

crowave radiation near 1285 MHz may be used to transfer $|9/2, -9/2\rangle \to |7/2, -7/2\rangle$, with a field sensitivity of 2.49 MHz/G. In this case the resonant frequency shifts by 13.8 kHz for adjacent lattice planes.

The larger field sensitivity of the $|9/2, -9/2\rangle \to |7/2, -7/2\rangle$ transition makes it appear the better choice for plane selection, since the resonant frequency of the desired plane will have a greater separation from its neighbors. However, in our experiment the accuracy with which we can address RF and microwave transitions is mainly limited by magnetic field noise. To perform state manipulation reliably we typically perform Landau-Zener frequency sweeps, with a sweep range that is larger than the spread in resonant frequencies due to the magnetic field noise. In this scenario, the field sensitivity is not a crucial parameter — for a larger sensitivity the resonant frequencies are more separated, but the frequency sweeps must also be larger, and the ratio of frequency separation to sweep width will stay unchanged. However, RF transitions with a small sweep range would be expected to be more sensitive to other additional noise sources at a smaller energy scale than the observed magnetic field noise.

After changing the internal state of a specific plane, we wish to remove all other planes in the system before a fluorescence image is taken. If we select a plane using the $|9/2, -9/2\rangle \to |7/2, -7/2\rangle$ microwave transition, then undesired planes may be removed by heating them out of the trap with an optical pulse on the $|9/2, -9/2\rangle \to |11/2, -11/2\rangle$ cycling transition of the D2 line. Such a pulse will have no significant effect on the selected atoms in $|7/2, -7/2\rangle$. On the other hand if plane selection were performed using an RF transition this optical pulse would not sufficiently discriminate between the two internal states. A subsequent microwave sweep would still be needed in order to protect the RF-selected plane from the optical pulse. For increased speed and simplicity, and for robustness against any low-energy fluctuations, we perform plane selection using a microwave transition.

Figure 4.7 shows typical microwave spectra obtained using a degenerate Fermi cloud held in the optical lattice in the presence of a vertical magnetic field gradient. In the absence of a gradient, a 20 kHz wide Landau-Zener sweep of the microwave frequency can reliably transfer the entire cloud to $|7/2, -7/2\rangle$. When a gradient is applied, the spatially varying resonant frequency leads to a broadening of the spectrum. From the known gradient we can determine the vertical density profile of the gas from the width of these spectra, indicating that the gas predominantly occupies 10 planes of the vertical lattice.

Performing high fidelity microwave transfer of a single plane while leaving all other planes unaffected requires careful choice of pulse or sweep parameters. A large microwave Rabi frequency is desired for fast sweeps and lower sensitivity to noise, but a Rabi frequency comparable to the 14 kHz spacing between planes will cause undesired off-resonant excitation if introduced suddenly. To avoid any affect on neighboring planes we slowly ramp the microwave amplitude up over 15 ms prior to performing a frequency sweep, and similarly ramp the amplitude back to zero afterwards. To achieve a high contrast between minimum and maximum power during these ramps, we control the amplitude using both the modulation input of the SRS generator and an additional voltage variable attenuator (Minicircuits ZX73-2500). Linearly ramping the control voltage of the variable attenuator results in a smooth s-shaped ramp of the microwave Rabi frequency. A Landau-Zener sweep of the detuning is performed by holding the microwave frequency fixed and sweeping the vertical magnetic field using a shim coil. The entire process is outlined in Figure 4.8, and we expect a high-fidelity single-plane transfer with negligible influence on other lattice planes.

By measuring the success of multiple repeated transfers, we find that 95% of atoms are successfully transferred by a single amplitude-modulated sweep. Possible improvements to this fidelity could be made
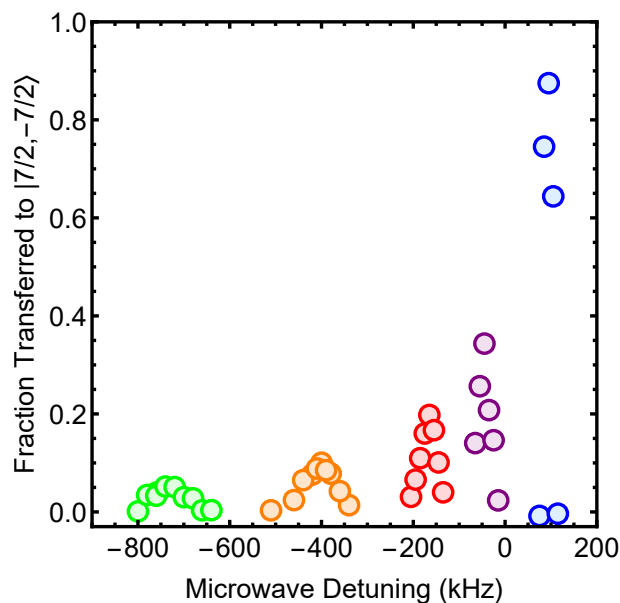
Figure 4.7: Microwave spectroscopy in linear gradients a of $0\,\mathrm{G/cm}$ (blue), $15\,\mathrm{G/cm}$ (purple), $30\,\mathrm{G/cm}$ (red), $60\,\mathrm{G/cm}$ (orange), and $105\,\mathrm{G/cm}$ (green). Landau Zener sweeps over a $20\,\mathrm{kHz}$ range are performed to transfer atoms from $|9/2, -9/2\rangle$ to $|7/2, -7/2\rangle$. As the gradient increases the finite size of the cloud causes the lineshape becomes wider and wider, with each vertical position z within the cloud corresponding to a different resonant frequency. Because the cloud does not reside at their centre of symmetry, the gradient coils also contribute a small homogeneous field, causing the lineshape to shift to lower frequencies as the gradient increases.
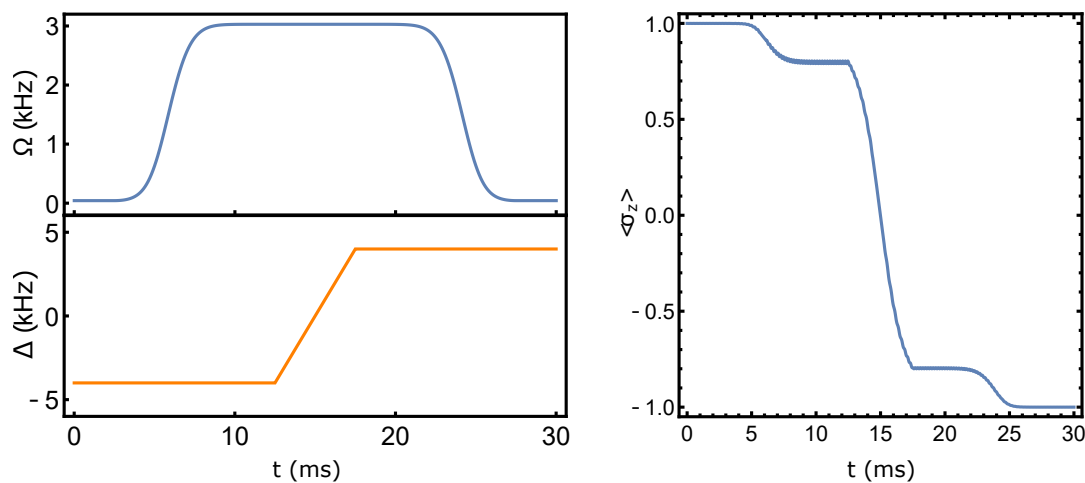
Figure 4.8: Amplitude shaped frequency sweep. (left) The Rabi frequency Ω is smoothly increased from zero to a maximum of $3\,\mathrm{kHz}$, with the detuning Δ swept by $\pm 4\,\mathrm{kHz}$ while Ω is maximal. (right) Expectation value of the σ_z (in the basis $|\uparrow\rangle = |9/2, -9/2\rangle$, $|\downarrow\rangle = |7/2, -7/2\rangle$) during the sweep. The expected transfer is over 99% efficient even if the resonant frequency is shifted (as if by a magnetic field instability) by $3.5\,\mathrm{kHz}$. Neighboring planes are predicted to be unaffected by the sweep.

by reducing noise in the microwave source or magnetic field, or by increasing the maximum attainable Rabi frequency. Increasing the Rabi frequency is possible by using a more powerful microwave amplifier, by adjusting the antenna geometry to provide better coupling through the imaging window as mentioned in Section 3.1.2, or even possibly by replacing the microwave source with a two-photon Raman coupling for state manipulation.

4.3.3 Optical Removal

With selected atoms transferred to $|7/2, -7/2\rangle$, the unwanted atoms can be heated out of the lattice using light tuned to the $|9/2, -9/2\rangle \to |11/2, -11/2\rangle$ transition of the D2 line. Ideally, even in large magnetic fields where F and m_F are no longer good quantum numbers, this is a 'cycling transition' so that atoms should not be depumped into another state before they a heated out of the trap. Improper beam polarization can still create a small depumping probability, which will allow some atoms to survive the optical removal step. In particular, atoms transferred off-resonantly to the $|9/2, -9/2\rangle$ excited state can decay to the $|7/2, -7/2\rangle$ ground state, becoming spectrally indistinguishable from atoms in the selected plane.

We typically perform our optical removal in the same 40 G vertical bias field that is used to help create the linear plane selection gradient. When using a circularly polarized beam propagating down through the imaging window, we actually saw little improvement in the depumping probability compared to the use of a linearly polarized beam. This is possibly due to birefringent effects of the sapphire imaging window, causing distortion of the incident beam polarization. When only using a linear polarization (creating σ^+/σ^- transitions), depumping effects are reduced by the 10 MHz separation between the D2 $|11/2, -11/2\rangle$ state and its adjacent m_F levels at 40 G. We find that this separation is sufficient to remove 99% of atoms without any depumping effects. Since the plane-selected atoms see the beam with a detuning of $\Delta > 200\Gamma$, their photon scattering rate is $\ll 1$ kHz and we observe no measurable loss of the selected atoms during the 500 μs pulse.

Even through the efficiency of optical removal is quite high, the few atoms remaining in the unselected planes will still be visible in a single atom sensitive fluorescence image. To ensure that this background does not appear, we typically perform the plane selection sequence a second time. This further reduces the number of remaining atoms in undesired planes, effectively reducing the confounding background in fluorescence images to an undetectable level. The second plane selection is carried out by performing a wide Landau-Zener sweep to transfer atoms in $|7/2, -7/2\rangle$ in any plane back to $|9/2, -9/2\rangle$, before transferring the desired single plane back to $|7/2, -7/2\rangle$.

4.3.4 Stability

To ensure that the same plane of the lattice is selected over many consecutive measurements, any drifts in the magnitude of the field given by Eq. (4.6) must be smaller than the 5.5 mG splitting between planes. Instability in the horizontal fields created only results in a small change in the overall field magnitude, and so the greatest effort is directed towards stabilization of the vertical field. Water cooling of the Feshbach coils with temperature regulated water provides sufficient stability in B_0, but we find that the ambient vertical magnetic field in the lab can slowly drift by 10 mG or more as subway trains travel between nearby stations. We use a 3-axis field sensor to detect these changes in the ambient field, and feed-forward the vertical field measurement to an additional coil to stabilize the field seen by the
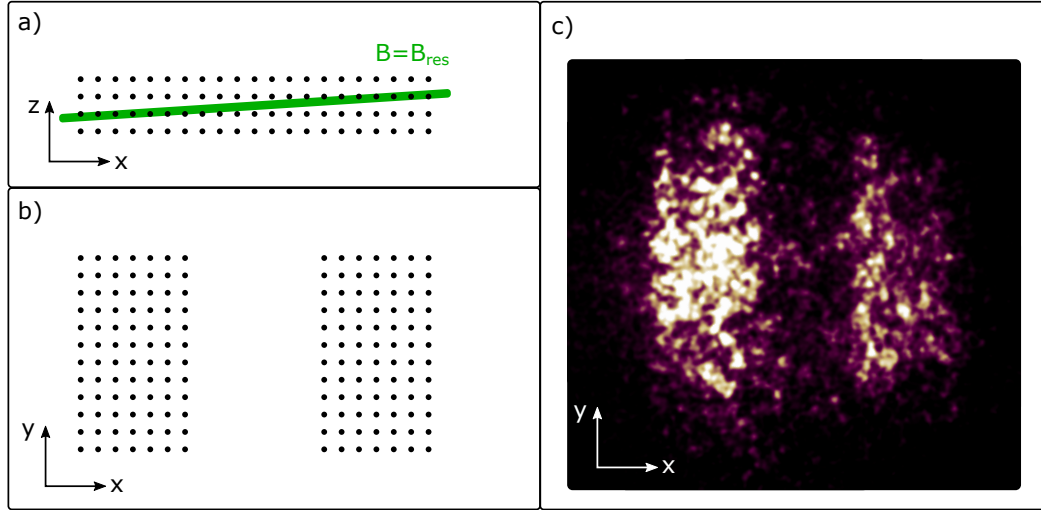
Figure 4.9: Effect of misalignment of the gradient during plane selection. a) Slight misalignments may cause the planes of constant magnetic field to cross several planes of the vertical lattice. Microwave transfer of the atoms is resonant for a specific magnetic field B_{res}, and so will transfer atoms from several planes along $\hat{z}$. b) After optical removal, the selected atoms will form a stripe pattern in the plane perpendicular to the gradient. c) An out-of-focus fluorescence image showing the stripe density distribution resulting from gradient tilt. Images of this type are used to align the gradient to the lattice axis.

atoms. We also observe a $60\,\text{Hz}$ oscillating magnetic field due to the AC power lines, with an amplitude in the vertical direction of roughly $2\,\text{mG}$. This is not stabilized, but we synchronize our spectroscopy pulses to the $60\,\text{Hz}$ phase in order to achieve stable initial conditions for the state manipulation. As a result, we find that we are able to repeatedly select the same lattice plane for over $2\,\text{hours}$, indicating that we have suppressed vertical magnetic field deviations to less than $5\,\text{mG}$ over this time span. When the field eventually drifts too much and plane selection begins to fail, we simply apply a small correction to B_0.

In order to select a single plane of the lattice with high fidelity we require that the gradient applied is not only large and stable, but also well aligned along the direction of the vertical lattice beam. If the gradient direction and lattice direction were to differ by $1°$, the projection of the $105\,\text{G/cm}$ gradient onto lattice axes would result in a $2\,\text{G/cm}$ gradient within the transverse plane. The transverse gradient would shift some of the desired lattice sites out of resonance with the plane selection pulse depending on their spatial position. With a sufficiently large transverse gradient the overall effect is that stripes of atoms will be selected, each drawn a different vertical planes as depicted in Figure 4.9. We adjust the angle between the field gradient and the lattice axes by applying additional bias fields along $\hat{x}$ and $\hat{y}$ to the field of Eq. (4.6). Applying a field $B_x\hat{x}$ tilts the direction the angle θ (given by $\tan(\theta) = B_x/B_0$) which can be adjusted by $\pm 3°$ with only a few Gauss in the transverse field. By deliberately varying the gradient direction and observing the resulting stripe patterns, we align the gradient to be collinear with $\hat{z}$. Because our quadrupole gradient is strongest along the symmetry axis of the coils, any misalignment with respect to the lattice axis also decreases the size of the gradient that we apply. However, since the overall tilt is less than $2°$, this gradient reduction is very small.

If the gradient field is purposefully misaligned with respect to the lattice axes, the stripe pattern
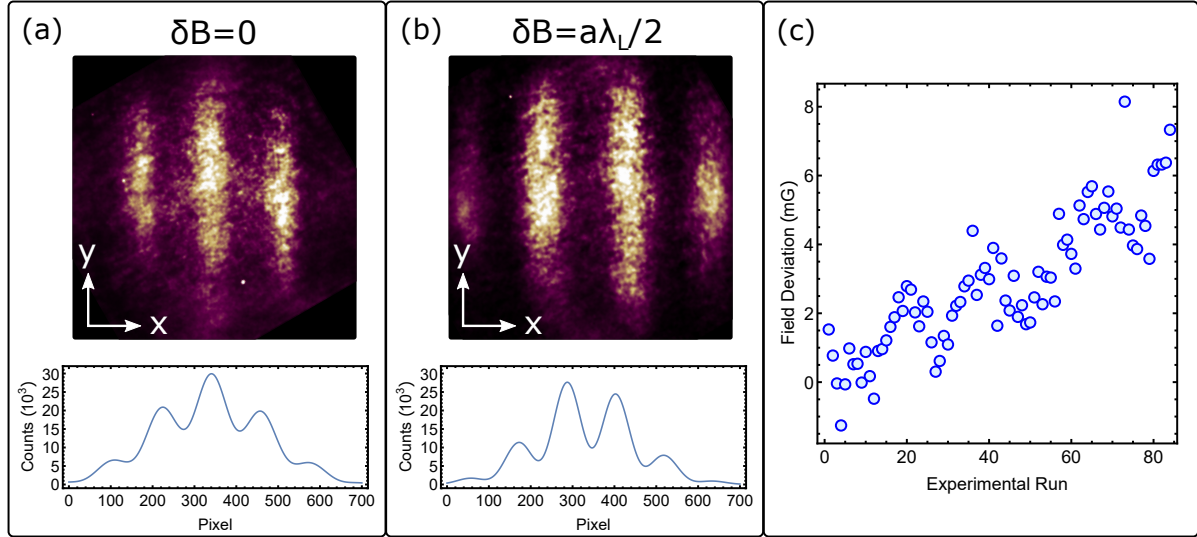
Figure 4.10: Measuring the magnetic field stability by using a deliberate tilt. The upper plots of (a) and (b) Show the selected density for two different magnetic fields during plane selection. When the field shifts by 5.5 mG during plane selection, the phase of the stripe pattern advances by 2π. By summing the fluorescence images along the y direction and then fitting a sinusoidal function with a Gaussian envelope (lower plots in (a) and (b)), we can extract the change in phase and thus sensitively determine any field drifts from one experimental run to the next. The results of tracking the field over many cycles is shown in c), with a steady field drift visible as well as a periodic oscillation.

produced can be used as a probe for magnetic field instabilities. Our use of a large gradient and a frequency sweep for plane selection ensure that small fluctuations δB in the overall field magnitude will not affect the plane selection procedure, as long as $\delta B < a\lambda_L/2$. When the gradient is tilted however, the phase of the resulting stripe pattern depends sensitively on δB, and this phase can be extracted from in-situ images of the stripe pattern in order to track the magnetic field stability over many experimental cycles. Figure 4.10 shows an example of typical magnetic field stability measured in this way. The field is stable to within 2 mG over several experimental cycles, with a slow drift of 7 mG per 100 cycles. The slow drift is likely due to slow temperature changes in the experimental apparatus, and we measure a much more severe field drift for the first hour the experiment is turned on each morning. On short timescales, the 2 mG scatter is small enough to have no effect on our plane selection. On long timescales the field drift needs to be compensated every hour or two to keep selecting the intended plane of the lattice.

In addition the steady upwards drift in field seen in Figure 4.10, an oscillation in the field is also visible with a period of about 20 experimental cycles. This is a very slow time scale, and could be correlated with cold air or humidity cycles of the HVAC system controlling the lab temperature. Alternatively, the regular experimental cycle time may periodically sample a stable noise source at a higher frequency, resulting in an aliased oscillation detected by our measurement. Regardless of the noise source, this fluctuation only has a magnitude of a few mG and so should not affect out plane selection sequence.

4.3.5   Improvements

After recent improvements to the stability of the Feshbach coil, we can perform plane selection with a 210 G/cm gradient and 120 G bias field without any additional field noise or drifts. The increased gradient provides a greater separation of resonant frequencies for different planes, making spectroscopy easier and providing more protection against the small field instabilities and drifts shown in Figure 4.10. A simultaneous increase in the bias field is necessary to minimize curvature of the increased gradient, but additionally helps to reduce depumping during the optical removal step. The $|11/2, -11/2\rangle$ excited state becomes more isolated from all other possible transitions as the field is increased, reducing the probability of off-resonant transitions. As a result, the efficiency of optical removal is improved in the 120 G bias field, and we find that it is no longer necessary to perform the plane selection twice to sufficiently remove all unwanted atoms.

4.4   Horizontal Selection

In the same way that we apply a small horizontal bias field to finely adjust the angle of the vertical field gradient, we can also apply a dominating horizontal bias field to completely rotate the magnetic gradient into the horizontal plane. This allows us to perform spatial selection of lattice sites in the field of view of the microscope. In this manner we remove atoms from certain regions of the 2D plane before imaging. Since we can only apply fields of roughly 5 G in the horizontal plane, we must ramp down the vertical bias field in order to accomplish a significant gradient rotation. With the vertical bias turned off, the resulting field is

$$\vec{B}(x, y, z) = (-\frac{a}{2}x + B_x)\hat{x} + (-\frac{a}{2}y + B_y)\hat{y} + az\hat{z}\,. \tag{4.8}$$

This produces a fairly linear gradient in the horizontal plane, which can be oriented with respect to the x and y lattice axes by adjusting the two fields B_x and B_y. The horizontal gradient is half as strong as the vertical gradient due to the symmetry of the quadrupole field, so that the effective field shift between lattice sites is 2.75 mG. We use this gradient to spectroscopically select the central 40 lattice sites of either the x or y lattice, with the same frequency sweep described in Section 4.3.2.

In order to optical remove atoms from unwanted lattice sites, we again use light resonant with the $|9/2, -9/2\rangle \rightarrow |11/2, -11/2\rangle$ transition of the D2 line. Since the overall field is small, the excited states of the D2 line are not well separated, and depumping effects are more likely. The beam used for optical removal in the horizontal magnetic field is the same resonant beam used for absorption imaging (shown in Figure 3.12). Because the beam has good circular polarization, we see that depumping effects are small and very few atoms remain in unwanted sites after the optical pulse.

By performing the horizontal selection two consecutive times with different values of B_x and B_y, we can remove all atoms from the lattice outside of the central 40×40 site region of a single vertical plane. This produces a cloud which occupies only the deepest sites of the optical lattice, so that effects of lattice inhomogeneity are minimized. At the same time these central lattice sites can be very sparsely occupied if we intentionally load a hot gas. This sparsely-filled and homogeneous lattice is the ideal system with which to calibrate our fluorescence imaging. The laser cooling technique used for fluorescence imaging, and the site-resolved fluorescence images obtained will be described in the following two chapters.

Chapter 5

Laser Cooling

With an ultracold cloud of fermionic atoms localized to the sites of a deep optical lattice, the final step required to produce high-resolution fluorescence images is the generation of a steady source of fluorescence photons. In order for atoms to stay localized during the fluorescence imaging process, it is necessary to provide a cooling mechanism that counteracts the heating arising due to repeated photon recoils. Each photon scattering event will heat an atom on average by E_R, which is equivalent to a temperature increase of $400\,\mathrm{nK}$ for $^{40}\mathrm{K}$. If hundreds of photons are to be collected from each atom, thousands of photons must be emitted due to the finite collection efficiency of the imaging system. Thus over the course of the fluorescence exposure each atom will gain energy equivalent to hundreds of $\mu\mathrm{K}$, which would lead to a high probability for loss of atoms from the confining lattice. For high-fidelity detection to be possible, this loss must be avoided and so typically in quantum gas microscope experiments laser-cooling is applied to the atoms during the entire fluorescence exposure.

Laser cooling techniques developed for free particles such as Doppler cooling or Sisyphus cooling can still be applied in a confining trap if the atomic motion can be considered classically. For instance the same beams providing doppler cooling in a MOT can be used to perform fluorescence imaging of single atoms trapped in deep ($V_0/k_B \approx 1\,\mathrm{mK}$) optical microtraps [77, 78]. When the maximum attainable trap depth is reduced, the minimum temperatures of Doppler cooling or Sisyphus cooling may not be sufficiently low to keep atoms inside the trapping potential. To build a quantum gas microscope for fermionic atoms, lighter atomic species such as the fermionic $^{40}\mathrm{K}$ and $^6\mathrm{Li}$ are used. The effective on-site trapping frequencies in a microtrap become large for these atoms, and the resulting motion is typically better described by a quantum harmonic oscillator model, with discrete bound levels. For instance, in our experiment with $^{40}\mathrm{K}$ the sites of the deep optical lattice have a depth of roughly $200\,\mu\mathrm{K}$, supporting around 10 discrete bound states. If we laser cool atoms to a temperature $T < \frac{1}{5}V_0/kB$ to prevent thermal excitation between lattice sites, we are cooling atoms near the ground state of the potential.

In previous quantum gas microscope experiments working with $^{87}\mathrm{Rb}$, atoms confined in a deep optical lattice were cooled by an applied optical molasses, so that atoms scattered many photons while staying trapped in the individual lattice sites [18, 20]. Even though the optical lattice used in this experiment achieves a similar depth, we do not find that optical molasses cooling is able to successfully cool trapped $^{40}\mathrm{K}$ for in-situ imaging. Instead, we implement a laser cooling technique used commonly in ion-trapping experiments, where discrete energy levels of the harmonic trapping potential are typically assumed and cooling is achieved by driving an atom in energy level ν into the lower level $\nu - 1$. This
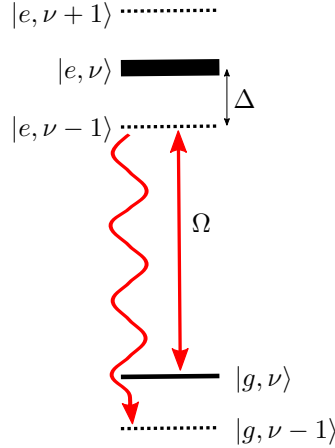
Figure 5.1: Schematic diagram of sideband cooling. Atoms in the atomic state $|g\rangle$ with vibrational level ν are coupled to a resolved motional sideband of the excited state $|e\rangle$ with Rabi frequency Ω. If $\eta^2 \ll 1$ the return to the atomic ground state through spontaneous emission will leave the atom in vibrational level $\nu - 1$. In the absence of spontaneous emission, the atom would simply undergo Rabi flopping between the levels $|g, \nu\rangle$ and $|e, \nu - 1\rangle$.

type of cooling can often cool atoms very close to the motional ground state of the trap specified by $\nu = 0$. This chapter will outline the details of the bound-state cooling technique that we use, and present details of the experimental apparatus and cooling results.

5.1 Cooling Bound Atoms

To understand how to laser cool atoms which are tightly bound in harmonic traps, we can learn a great deal from research which has been performed with trapped ions. In dynamic electric traps, such as the Paul trap, harmonic trapping frequencies can be on the order of 1 MHz or higher. Narrow linewidth intercombination lines with $\Gamma \approx 1\,\mathrm{kHz}$ can then be used for cooling, realizing the *resolved sideband regime* with $\Gamma \ll \omega_0$, as pictured in Figure 5.1. Driving the intercombination line with a detuning $\Delta = \omega_0$ can selectively excite atoms from the electronic ground state with ν vibrational quanta to the electronic excited state with $(\nu - 1)$ quanta, removing energy from the atom. This method of *resolved sideband cooling* can be used to cool a trapped ion to near its motional ground state [79].

For harmonic trapping along the $\hat{x}$ direction, a beam which propagates along $\hat{x}$ with wavenumber k can only couple the two states $|g, \nu\rangle$ and $|e, \nu - 1\rangle$ if the matrix element $\langle \nu - 1|e^{ikx}|\nu\rangle$ (describing the coupling due to the plane-wave electric field) is nonzero. Expanding the exponential to lowest order in the parameter kx, and rewriting x as a combination of raising and lowering operators for the harmonic oscillator levels, we see that

$$\langle \nu - 1|e^{ikx}|\nu\rangle \approx \frac{ikl_0}{\sqrt{2}} \langle \nu - 1|(a^\dagger + a)|\nu\rangle = i\eta\sqrt{\nu}\,, \tag{5.1}$$

where $l_0 = \sqrt{\hbar/m\omega_0}$ is the harmonic oscillator length scale, and $\eta = \frac{kl_0}{\sqrt{2}}$ is called the Lamb-Dicke parameter. The next term of the expansion will be proportional to η^2, and can commonly be ignored in experiments where $\eta \ll 1$.

For ^{40}K atoms in a deep optical lattice, the harmonic oscillator length is typically on the order of
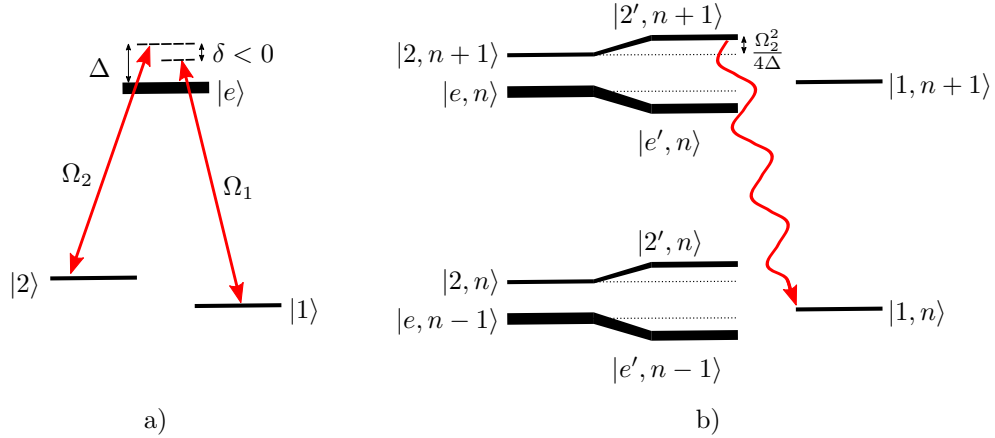
Figure 5.2: Creating a narrow excited state with a three-level system. a) Two ground states are coupled to the same excited state $|e\rangle$ by modes with frequencies ω_1 and ω_2, forming a Λ-system. b) When $\Omega_1 = 0$ the eigenstates are dressed states of the two-level system formed by $|2\rangle$ and $|e\rangle$, with $|1\rangle$ acting as an auxiliary state into which atoms accumulate through spontaneous emission. The new eigenstates of the two-level system are labelled $|2'\rangle$ and $|e'\rangle$, with $|2'\rangle$ being mostly composed of $|2\rangle$ with a small admixture of $|e\rangle$ when the detuning Δ is large and positive. In this limit, the state $|2'\rangle$ has a narrow linewidth, yet can spontaneously emit or be optically addressed due to the admixture of the excited state.

20 nm, so absorbing or emitting an optical photon with $\lambda \approx 770$ nm is associated with a Lamb-Dicke parameter of $\eta = 0.167$. Such a non-zero η allows the cooling transition shown in Figure 5.1 to occur. Rewriting η in terms of the photon recoil energy E_R gives the simple expression $\eta = \sqrt{E_R/\hbar\omega_0}$. In this form it is clear that it is the atomic recoil induced by the absorption or emission of a photon that allows coupling between motional states.

The cooling process outlined in Figure 5.1 is only completed when the atom spontaneously emits a photon to return to the ground state. As long as the atom decays to $|g, \nu - 1\rangle$, energy is removed during the entire process. The probability that the atom maintains the lower vibrational level $\nu - 1$ after decay from the excited state is proportional to η^2. Thus when implementing sideband cooling a large value of η leads to stronger coupling to atomic motional states, but a small η will tend to preserve the motional state after spontaneous emission. Many ion-trapping experiments operate in the limit $\eta \ll 1$, referred to as the Lamb-Dicke regime. When performing sideband cooling in a 3D trap with distinct motional degrees of freedom, it is not only the magnitude of the photon recoil which determines the motional state coupling, but also the recoil direction. Cooling will only occur if the recoil energy due to absorption from the cooling beam has nonzero projection along a given trap axis. Thus for a trap axis $\hat{x}$ and cooling beam wavevector $\vec{k}$, one can define an effective Lamb-Dicke parameter

$$\eta_x = \frac{l_0 \vec{k} \cdot \hat{x}}{\sqrt{2}} \, , \tag{5.2}$$

governing the coupling to motional states. If $\vec{k}$ is orthogonal to the trap axis $\hat{x}$, then photon absorption is decoupled from this motional part of the wavefunction and no cooling transitions can be driven for this axis of the trap.

5.1.1   Dressed States

In alkali atoms narrow-linewidth intercombination transitions are not available, and common laser cooling techniques like Doppler cooling and Sisyphus cooling typically operate on the D1 or D2 lines which have $\Gamma \approx 6\,\mathrm{MHz}$. Harmonic trapping frequencies for atoms trapped in a high-power optical lattice are typically limited to hundreds of kHz, making it difficult to reach the resolved sideband regime. However, transitions between different ground states of the atom have essentially zero linewidth, so that motional sidebands are easily resolved. A coupling between two ground states introduced by RF or microwave radiation cannot realize sideband cooling, because the Lamb-Dicke parameter associated with such long-wavelength radiation is essentially zero [1]. Instead a coupling between the ground states must be engineered using photons at optical frequencies.

The cooling mechanisms for bound alkali atoms can be understood using a model system with two ground states $|1\rangle$ and $|2\rangle$ coupled by two lasers to the same excited state $|e\rangle$ in a Λ-system as shown in Figure 5.2(a). The two relevant detunings in the system are the *common mode detuning* Δ of the mode Ω_2 from the excited state, and the *two-photon detuning* δ measuring the difference between the detunings of the two modes Ω_1 and Ω_2 from their respective transitions. If the lasers have frequencies ω_1, ω_2 and the energy splittings between the energy levels are denoted $\omega_{ij} = (E_i - E_j)/\hbar$, then $\Delta = \omega_2 - \omega_{e2}$ and $\delta = \omega_1 - \omega_2 - \omega_{21}$

In this system a two-photon coupling emerges between the ground states that involves large atomic recoils and so realizes a nonzero η. In the limit $\Omega_2 \gg \Omega_1$ the cooling various cooling schemes available in a Λ-system can be easily visualized by using the *dressed state picture*. For simplicity we first consider the case $\Omega_1 = 0$, and so must only rediagonalize the levels $|2\rangle$ and $|e\rangle$ due to the presence of the coupling Ω_2.

In the dressed state basis the bare states $|2\rangle$ and $|e\rangle$ are mixed by the coupling Ω_2 to form new eigenstates $|+\rangle$ and $|-\rangle$ of the form

$$|+, n\rangle = \cos(\theta)|e, n\rangle + \sin(\theta)|2, n+1\rangle \,, \tag{5.3}$$

$$|-, n\rangle = \sin(\theta)|e, n\rangle - \cos(\theta)|2, n+1\rangle \,, \tag{5.4}$$

where the notation $|2, n\rangle$ indicates the atomic state $|2\rangle$ with n photons in the mode of Ω_2. The mixing angle θ that parameterizes the superposition is determined by $\tan(\theta) = (\sqrt{\Delta^2 + \Omega_2^2} - \Delta)/\Omega_2$. When the detuning is positive (blue-detuning) and much larger than Ω_2 then $\theta \to \pi/2$, so that the state $|-, n\rangle$ is mostly composed of the atomic ground state $|2, n+1\rangle$ and the state $|+, n\rangle$ is mostly composed of $|e, n\rangle$. It is then helpful to relabel the states $|-, n\rangle \equiv |2', n+1\rangle$ and $|+, n\rangle \equiv |e', n\rangle$, where the use of such notation implicitly assumes blue detuning of the dressing mode. The lifetimes of the states $|2', n+1\rangle$ and $|e', n\rangle$ for each photon number n are determined by the admixture of $|e\rangle$, so that each can be associated with a linewidth given by

$$\Gamma_{2'} = \Gamma \sin^2(\theta) \,, \tag{5.5}$$

$$\Gamma_{e'} = \Gamma \cos^2(\theta) \,. \tag{5.6}$$

The state $|2', n+1\rangle$ can have a very small linewidth $\Gamma_{2'}/2\pi \approx 40\,\mathrm{kHz}$ for only moderately large detuning

---

[1]Although the RF or microwave photons cannot couple to the atomic motion on their own, coupling can still be created by using a spin-dependent trapping potential [80]

$\Delta = 6\Omega_2$. The small linewidth of this dressed state allows resolved sideband cooling techniques to be performed on alkali atoms.

If the narrow-linewidth dressed state $|2', n+1\rangle$ is to be used for laser cooling, the energy shifts introduced by Ω_2 must be considered. Referencing the zero-point of energy to the bare excited state $|e, n\rangle$, the shifted energy levels for a given photon number n are given by

$$E_{2'} = \frac{\hbar\Delta}{2} + \frac{\hbar\sqrt{\Omega_2^2 + \Delta^2}}{2} , \tag{5.7}$$

$$E_{e'} = \frac{\hbar\Delta}{2} - \frac{\hbar\sqrt{\Omega_2^2 + \Delta^2}}{2} . \tag{5.8}$$

The difference between these energies and the bare atomic levels is due to the same AC Stark shift discussed in Section 2.1, with the counter-rotating term already having been discarded in the current expression. Since we consider a blue-detuned beam, the Stark shift raises the energy of the 'ground' state $|2', n+1\rangle$ and lowers the energy of the 'excited' state $|e', n\rangle$. For a detuning which is large compared to the Rabi frequency, the Stark shift can be written in a simple form

$$V_{\text{Stark}} = \frac{\hbar\Omega_2^2}{4\Delta} , \tag{5.9}$$

which is reminiscent of the form given in Equation (2.2) as it is proportional to the intensity of the dressing mode ($\Omega \propto \sqrt{I}$).

If $\Omega_1 = 0$ the above description gives the correct eigenstates of the two level system composed of $|2\rangle$ and $|e\rangle$. The states $|2', n+1\rangle$ and $|e', n\rangle$ form the typical dressed state ladder, and spontaneous emission between these levels will produce the expected Mollow triplet. The third level $|1\rangle$ is still relevant in the time evolution of the system, as it can be populated via spontaneous emission. When driving the two-level subsystem with Ω_2, the atom will eventually be pumped into the state $|1\rangle$, where it will become largely decoupled from the driving field at ω_2 as long as there is a large energy difference $\hbar\omega_{21}$ between the ground states. When $\Omega_1 = 0$, the state $|1\rangle$ can be referred to as a *dark state* of the three-level system, and over time a collection of atoms would be expected to accumulate in this dark state due to multiple spontaneous emission events.

When $\Omega_1 \neq 0$, we must re-diagonalize all three levels of the Λ-system with the two applied couplings Ω_1 and Ω_2 to obtain three new dressed eigenstates. Using a notation where the state $|g, n, m\rangle$ represents the atomic state g with n photons in the mode Ω_2 and m photons in the mode Ω_1, the eigenstates for $\delta = 0$ have the form

$$|D, n, m\rangle = \frac{1}{\Omega}\left(\Omega_2|1, n, m+1\rangle - \Omega_1|2, n+1, m\rangle\right) , \tag{5.10}$$

$$|B_+, n, m\rangle = \cos(\theta)|e, n, m\rangle + \sin(\theta)|C, n, m\rangle , \tag{5.11}$$

$$|B_-, n, m\rangle = \sin(\theta)|e, n, m\rangle - \cos(\theta)|C, n, m\rangle , \tag{5.12}$$

where $\Omega^2 = \Omega_1^2 + \Omega_2^2$, θ is now defined by $\tan(\theta) = (\sqrt{\Delta^2 + \Omega^2} - \Delta)/\Omega$, and we have simplifiied the expressions using the state

$$|C, n, m\rangle = \frac{1}{\Omega}\left(\Omega_1|1, n, m+1\rangle + \Omega_2|2, n+1, m\rangle\right) , \tag{5.13}$$

which is orthogonal to $|D, n, m\rangle$ but is not an eigenstate of the system. The state $|D, n, m\rangle$ is the dark

state of the system, with no admixture of the excited state $|e\rangle$. The other two eigenstates $|B_-\rangle$ and $|B_+\rangle$ — referred to as bright states — have overlap with $|e\rangle$ and so may decay through spontaneous emission. After many emission events atoms will accumulate in the dark state, and remain there for relatively long times. In the limit that $\Delta \gg \Omega_2$ and $\Omega_2 \gg \Omega_1$, the three eigenstates $|D\rangle$, $|B_-\rangle$, and $|B_+\rangle$ are dominantly composed of $|1\rangle$, $|2\rangle$, and $|e\rangle$ respectively. It can then be helpful to relabel the states as $|D, n, m\rangle = |1', n, m+1\rangle$, $|B_-, n, m\rangle = |2', n+1, m\rangle$, and $|B_+, n, m\rangle = |e', n, m\rangle$. This notation will be used in the following discussion of laser cooling, and so the limits $\Delta \gg \Omega_2$ and $\Omega_2 \gg \Omega_1$ are implicitly assumed. These assumption of large Δ is typically required for these laser cooling mechanisms, since it creates a narrow-linewidth dressed state. The second limit is not necessary for cooling, but aligns with our experimental implementation and simplifies the conceptual picture.

To consider implementing cooling with the narrow dressed states of the Λ-system, we must consider the addition of vibrational levels to the dressed state picture introduced above. If the recoil momentum of photon absorption and emission is ignored, then each set of three dressed states $|1', \nu\rangle$, $|2', \nu\rangle$, and $|e', \nu\rangle$ with vibrational level ν (suppressing the photon numbers n and m) will be completely decoupled form all other vibrational levels of ν', forming a set of independent three-level systems. Including the effect of a small but nonzero η during spontaneous emission, a spontaneous emission event will typically keep an atom within the manifold of states with the same vibrational level ν, but will occasionally couple into nearby levels $\nu \pm 1$, and over time the vibrational level will increase in correspondence with the accumulated recoil energies of the emitted photons.

The construction of eigenstates carried out above was performed assuming a three level system only, with no coupling between the different motional states. However, as detailed in Section 5.1, the recoil momentum associated with absorption from modes Ω_1 and Ω_2 will lead to coupling between different motional states. This coupling will be small if the associated Lamb-Dicke parameter is much less than one, and we can attempt to understand the effect this coupling will have on the system by introducing it as a perturbation to the eigenstates of Equations (5.10)-(5.12) above. At zeroth order in the Lamb-Dicke parameter, the three states $|1', \nu\rangle$, $|2', \nu\rangle$, and $|e', \nu\rangle$ are orthogonal and uncoupled. At first order in the Lamb-Dicke parameter, a coupling of the form $\eta\Omega$ arises between $|1', \nu\rangle$ and $|2', \nu \pm 1\rangle$ or between $|1', \nu\rangle$ and $|e', \nu \pm 1\rangle$ (the specific form of η and Ω depends on the details of the recoil momentum and atomic couplings). If these couplings are small compared to the dissipation rate due to spontaneous emission, their primary effect will be to drive accumulated atoms out of the dark states into different vibrational levels of the dressed state manifold.

A simple cartoon for motional state cooling in this framework is presented in Figure 5.3. Atoms accumulate in $|1', n, \nu\rangle$ (suppressing the number of photons m in the mode Ω_1) after spontaneous emission, and we can consider the coupling out of this state due to vibrational-state-changing transitions. An atom in the state $|1', n, \nu\rangle$ will be resonantly coupled to the red-motional sideband of the narrow dressed state $|2', n+1, \nu-1\rangle$ when the two-photon detuning is set to $\delta = V_{\text{Stark}}/\hbar - \omega_0$ (the required detuning is not simply $\delta = -\omega_0$ because $\delta = 0$ refers to the position of the state $|2, n+1, \nu\rangle$ which differs in energy from $|2', n+1, \nu\rangle$ by the Stark shift). If an atom is transferred from $|1', n, \nu\rangle$ to $|2', n+1, \nu-1\rangle$ before finally decaying to $|1', n, \nu-1\rangle$ through spontaneous emission, then energy is removed through the entire cycle. This process of exciting to the vibrational sideband of the narrow dressed state appears analogously to the sideband cooling technique for trapped ions that is sketched in Figure 5.1. To tune the red motional sideband into resonance may require $\delta \neq 0$, so that the eigenstates presented above are not a correct description of the system and no dark state formally exists. Nevertheless, Figure 5.3
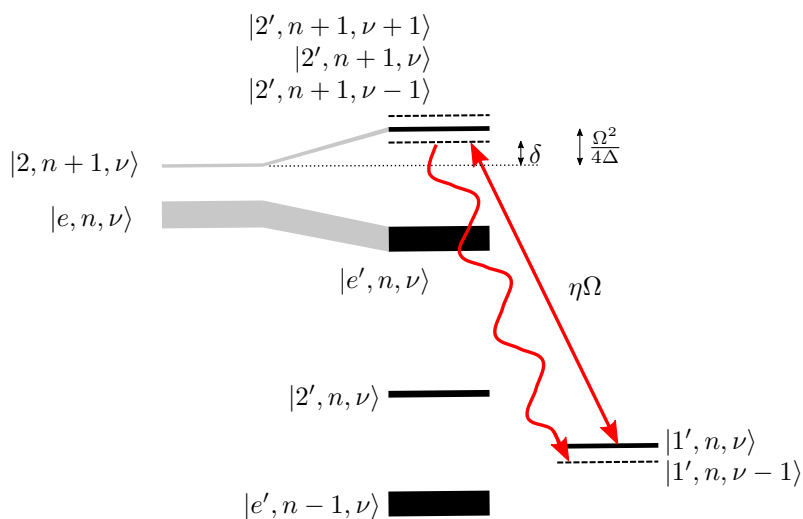
Figure 5.3: A scenario constructed to appear analogous to the sideband cooling of Figure 5.1. The states $|1'\rangle$, $|2'\rangle$, and $|e'\rangle$ are eigenstates of the three level system in the absence of any photon recoil. Atoms are taken to accumulate in the state $|1'\rangle$ over time due to spontaneous emission. From $|1', \nu\rangle$, atoms can be coupled into the motional sideband $|2', \nu - 1\rangle$ with Rabi frequency proportional to a Lamb-Dicke parameter η, and this motional sideband can be resolved since the linewidth of the state $|2'\rangle$ may be made arbitrarily narrow. Spontaneous emission events will tend to return the atom to $|1', \nu - 1\rangle$ so that the overall vibrational energy has been lowered. The two-photon detuning δ indicated on the diagram is the difference in energy between the photon energy of mode Ω_1 and the bare energy difference between the states $|2, n+1\rangle$ and $|1, n\rangle$. The use of the bare atomic states $|2\rangle$ and $|1\rangle$ in this definition is crucial for the creation of a dark state $|D\rangle$ at $\delta = 0$. The photon number m for the mode Ω_1 has been suppressed in the diagram.

(a) Bare States

(b) Dressed States of Ω_{OP}

Figure 5.4: Two diagrams for Raman sideband cooling. In (a) the two bare atomic ground states are coupled by the Raman beams Ω_1 and Ω_2. After coherent transfer from $|1\rangle$ to $|2\rangle$, the optical pumping beam $\Omega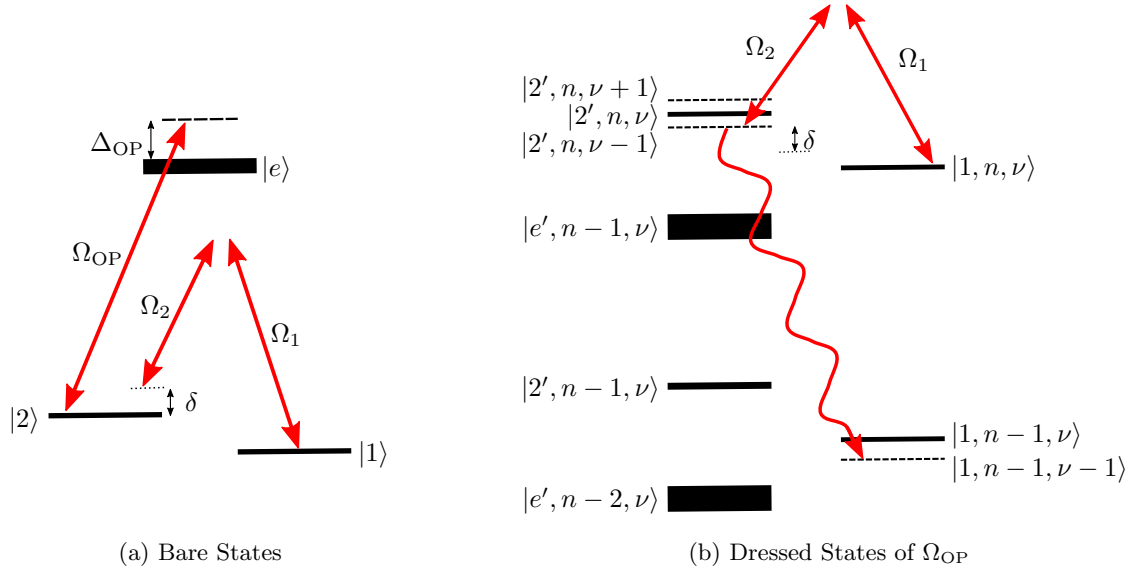_{\text{OP}}$ will incoherently transfer atoms back into state $|1\rangle$, while typically preserving the motional state of the atom. In (b) we dress the states $|2\rangle$ and $|e\rangle$ with the coupling Ω_{OP}, forming a ladder of states according to the number of photons n in the mode Ω_{OP}. The two-photon coupling from Ω_1 and Ω_2 can cause Rabi flopping of atoms from $|1, n, \nu\rangle$ to $|2', n, \nu - 1\rangle$, keeping the number of photons in the Ω_{OP} mode unchanged but reducing the vibrational energy by one. The cooling cycle is completed if the state $|2', n, \nu - 1\rangle$ spontaneously emits a photon and decays to $|1, n - 1, \nu - 1\rangle$. In both diagrams $\delta = \omega_1 - \omega_2 - \omega_{12}$ is measured from two-photon resonance between the bare atomic states, indicated by a dotted line in (b).

captures the important features of the cooling of bound atoms using narrow dressed states.

Using the basic elements of dressed states described above, it is easy to understand two common methods for cooling of bound atoms — Raman sideband cooling and EIT cooling. These two cooling schemes are very similar to one another, both sharing the general idea of cooling towards the motional ground state by coupling to a narrow dressed state, as depicted in Figure 5.3.

5.1.2   Raman Sideband Cooling

In the scenario sketched out in Fig. 5.3, there are two features of the dressed state $|2', n + 1\rangle$ which allow cooling to occur. The state $|2', n + 1\rangle$ is coherently coupled to the ground state $|1'\rangle$ in different vibrational levels, but additionally the dressed state has a finite lifetime leading to the spontaneous emission of photons necessary to complete the cooling cycle. Although in the Λ-system of Fig. 5.2 these two features arise from the same coupling Ω_2, it is also possible to generate them separately with multiple beams: two lasers with large Δ can provide the coupling between the states $|1\rangle$ and $|2\rangle$ via a Raman process, while a separate laser with smaller Δ drives the $|2\rangle$ to $|e\rangle$ transition to produce the necessary spontaneous emission. This is the essence of the *Raman sideband cooling* scheme, depicted in Figure 5.4.

Raman beams with a large detuning from the atomic excited state are desirable so that any single-photon scattering associated with these beams can be ignored and only the two-photon coupling will be important. However, for the Raman beams to provide coupling between the bare atomic states

$|1\rangle$ and $|2\rangle$ their detuning cannot be too large. Beams which are far-detuned with respect to the fine and hyperfine structure of the atom provide a coupling which is state-independent and hence does not mix different eigenstates of the bare atom. To strike a balance between the limits of far detuning for low scattering and near detuning for strong coupling between atomic states, Raman beams detuned by $10^4\Gamma$ are common. At this detuning excitation of the atom can be assumed to occur dominantly through absorption of photons from $\Omega_{\rm OP}$. The resulting spontaneous emission can be considered in the bare state basis, or by considering the dressing of state $|2\rangle$ by the beam $\Omega_{\rm OP}$. The process of forming dressed states of $\Omega_{\rm OP}$ and treating the Raman beams as a perturbation is appropriate only as long as the two-photon Rabi frequency $\Omega_R = \Omega_1\Omega_2/2\Delta$ is small compared to $\Omega_{\rm OP}$.

Since the Raman process couples states $|1\rangle$ and $|2\rangle$ via a two-photon transition, we can define a two-photon Lamb-Dicke parameter η_2 describing the strength of coupling to the motional state of the atom. This parameter η_2 takes into account the net atomic recoil $\hbar(\vec{k}_1 - \vec{k}_2) = \hbar\Delta\vec{k}$ created by absorbing a photon from Ω_1 and emitting a photon into Ω_2, and hence is sensitive to the relative alignment of the two beams. As in (5.2) only the projection of the recoil momentum onto each individual trap axis can provide coupling between motional states. Thus the two-photon Lamb-Dicke parameter coupling motional states of the x-axis is

$$\eta_{2,x} = \frac{l_0\Delta\vec{k}\cdot\hat{x}}{\sqrt{2}} \ . \tag{5.14}$$

From (5.14) it is clear that the coupling of motional states is largest for counter-propagating beams travelling along a given trap axis. Similarly co-propagating beams will have a vanishingly small $\Delta\vec{k}$ and so cannot be used for Raman sideband cooling.

Previous experiments performing Raman sideband cooling with neutral atoms such as Cs [81] and Rb [82] have generated the Raman coupling and the confining optical lattice with the same beams. In these experiments the optical lattice was moderately detuned by $\Delta \approx 10^3\Gamma$ in order to provide a nearly conservative potential while also achieving the necessary coupling between internal atomic states. If a very far-detuned lattice is used, separate Raman beams are preferable in order to achieve strong coupling between the atomic ground states [83]. A combination of a far-detuned lattice, and near detuned ($\Delta \approx 10^4\Gamma$) Raman beams has recently been used in several groups to achieve fluorescence imaging of fermions in an optical lattice [31, 34, 35]. When implementing Raman sideband cooling as a source of fluorescence photons, it is beneficial to perform the Raman coupling and optical pumping at different wavelengths [31]. This allows photons from the Raman beams to be spectrally filtered out of the fluorescence imaging system, reducing a significant source of background light.

5.1.3   EIT Cooling

For a Λ-system driven by two fields Ω_1 and Ω_2 on two-photon resonance ($\delta = 0$) there will be an eigenstate of the system $|D\rangle \propto \Omega_2|1\rangle - \Omega_1|2\rangle$ which has zero probability of excitation to the excited state $|e\rangle$. This is due to destructive interference between the two possible pathways for an atom in state $|D\rangle$ to emit a photon into the continuum. The eigenstate $|D\rangle$ is thus called the *dark state* and an atom in this state will be completely decoupled from the two light fields. An atom initially in the state $|1\rangle$ can be made transparent to a field Ω_1 through the introduction of the second field Ω_2 and the creation of a dark state, a process referred to as electromagnetically induced transparency (EIT).

In either the sideband cooling cartoon of Figure 5.3 or the Raman sideband cooling schemes depicted in Figure 5.4, transitions to a state with reduced vibrational energy are accomplished by setting the
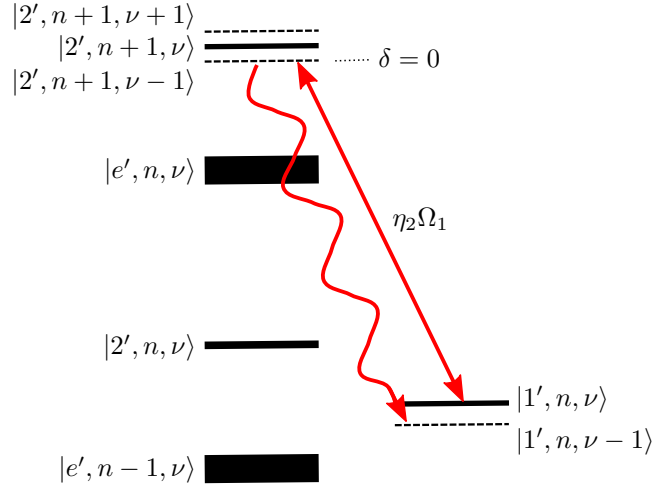
Figure 5.5: Dressed States Picture of EIT Cooling. The limit $\Omega_1 \ll \Omega_2$ is considered so that the dark ground state superposition $|D\rangle$ is mainly $|1\rangle$. Since $\delta = 0$, there is no coupling between $|1', n, \nu\rangle$ and $|2', n+1, \nu\rangle$ at zeroth order in η since these are eigenstates of the Λ-system — and this represents the cancellation of $\nu \to \nu$ transitions which appears due to EIT. At first order in the Lamb-Dicke parameter, coupling appears between different vibrational levels of the Λ-system eigenstates, and the cooling transition $|1, n, \nu\rangle \to |2, n+1, \nu-1\rangle$ can be brought into resonance by tuning the Stark shift. The cooling process is completed if the state $|2', n+1, \nu-1\rangle$ is pumped back to $|1', n, \nu-1\rangle$ by spontaneous emission. The photon number m for the mode Ω_1 has been suppressed in the diagram.

two-photon detuning to $\delta = V_{\text{Stark}}/\hbar - \omega_0$. Since $V_{\text{Stark}} > 0$ for $\Delta > 0$, it is possible that cooling is optimized with $\delta = 0$. Performing cooling while driving the Λ-system at two-photon resonance leads to the appearance of dark atomic states. The cooling process that occurs in this configuration has been previously explored theoretically [84, 85] and experimentally [86] in the context of trapped ions, and is referred to as EIT cooling.

In EIT cooling, $\delta = 0$ is fixed and cooling is tuned by adjusting the intensity of two blue-detuned fields such that $V_{\text{Stark}} = \hbar\omega_0$. This brings the red motional sideband into two-photon resonance as shown in Fig 5.5. There are two possible ways to absorb a photon from Ω_1 without changing the vibrational state of the atom: transfer to the narrow dressed state $|2', n+1\rangle$ with detuning ω_0, or transfer to the broad dressed state $|e', n\rangle$ with detuning $\Delta - \omega_0$. At $\delta = 0$ these two processes have equal and opposite amplitudes and interfere destructively, meaning that any coupling which leaves the motional state unchanged is strongly suppressed due to EIT. Transfer to other motional states such as $|2', n+1, \nu-1\rangle$ will not be suppressed at $\delta = 0$, so that cooling transitions are still possible. Thus tuning $V_{\text{Stark}} = \hbar\omega_0$ facilitates cooling transitions between $|1, n, \nu\rangle \to |2', n+1, \nu-1\rangle$ while suppressing the unwanted transitions $|1, n, \nu\rangle \to |2', n+1, \nu\rangle$.

Compared to sideband cooling and Raman sideband cooling, the suppression of the $\nu \to \nu$ transitions present in EIT cooling is unique. Suppressing these motional carrier transitions seems at first sight inconsequential, since they do not change the motional state ν. However, due to the presence of spontaneous emission and the associated photon recoil energy, any cycle between states which does not explicitly produce cooling will result in heating of the atoms after many cycles. During spontaneous emission atoms will be transferred rarely into a higher motional state ($\nu \to \nu + 1$), with the probability of these rare transfers is set by the squared Lamb-Dicke parameter $\eta^2 = E_R/\hbar\omega_0$. Averaging over many

spontaneous emission cycles, the average heating rate of these rate events is identical to the expected recoil heating. The cancellation of unwanted carrier transitions in EIT cooling removes this possible heating process from the system. The use of additional atomic levels to generate EIT suppression of a second heating process due to the blue motional sideband has also been proposed, and is referred to as double EIT cooling [87].

5.1.4   Comparing EIT Cooling and Raman Sideband Cooling

EIT cooling and Raman sideband cooling can both be implemented with alkali atoms, and expressed in the dressed state basis both cooling schemes appear quite similar. One key difference is the suppression of carrier transitions that is present in EIT cooling. To reduce the rate of off-resonant $\nu \to \nu$ excitation in a Raman Cooling scheme, it is necessary to carefully choose the strength of the fields Ω_{OP}, Ω_1, and Ω_2 [2]. Because these $\nu \to \nu$ transitions are cancelled in EIT cooling, larger Rabi frequencies and hence larger cooling rates should be possible [84]. The increased Rabi frequencies possible in EIT cooling can be expressed as a larger cooling *bandwidth*, which is relevant for ion-trapping experiments where the trapping frequencies of the 3D trap are not typically matched. In the context of in-situ fluorescence imaging there are many other factors which limit the possible Rabi frequencies and cooling rates so that this distinction between the two cooling schemes is less pronounced.

Experimentally, EIT cooling involves only two beams of one wavelength, which allows for a simplified optical setup. This is a benefit when only a cooling effect is desired [86], but actually introduces a complication when EIT cooling is used to produce fluorescence for in-situ imaging. If only one wavelength of light is present, it becomes difficult to shield the camera from background light generated by the cooling beams. In contrast, Raman sideband cooling can be performed with two separate wavelengths for the Raman coupling and optical pumping processes, allowing one beam to be spectrally filtered out of the imaging path. This means that Raman sideband cooling, which is more directly adapted to use two different wavelengths, may be easier to implement in a quantum gas microscope.

For both cooling schemes, precise tuning of the cooling mechanism is achieved in a similar way. In Raman sideband cooling, the two-photon detuning for carrier transitions is determined, before reducing δ by the trap frequency ω_0 to bring the red sideband into resonance. To stabilize the system in the cooling configuration the two-photon detuning must be well controlled, which is easily realized by using high-stability RF electronics to drive acousto-optic or electro-optic modulators. In EIT cooling, the two photon detuning is instead always maintained at $\delta = 0$ and the red sideband is brought into resonance by adjusting the Stark shift via the intensity of the beams Ω_1 and Ω_2. The cooling configuration is maintained by stabilizing the intensity of the cooling beams, a common enough task in an experiment working with ultracold atoms. In both cases, the two-photon detuning is sensitive to external magnetic fields which change the energy difference of the two ground states. Thus both cooling techniques require a stable magnetic field.

Both RSC and EIT cooling rely on the coupling between motional states that is provided by a two-photon pathway. Thus the efficiency of both cooling schemes will scale with the two-photon Lamb-Dicke parameter η_2 of Equation (5.14). The beams Ω_1 and Ω_2 must be oriented to have non-zero projection of the two-photon recoil $\Delta\vec{k}$ onto every motional axis of the trapping potential. Since both cooling cycles are completed by the spontaneous emission of a photon, it is also desirable that the single-photon Lamb

---

[2]It is also possible to implement Raman sideband cooling in a pulsed configuration, alternating the Raman coupling and optical pumping processes, in order to gain more flexibility in the choice of all Rabi frequencies

Dicke parameter η is small. This ensures that the motional state of the atom is typically preserved during spontaneous emission.

5.2   Implementation of EIT Cooling

As discussed in Section 3.5, our deep optical lattice reaches on-site trapping frequencies of $\omega_0/2\pi = 300\,\text{kHz}$, so that the single-photon Lamb-Dicke parameter for a 770 nm photon is $\eta = 0.17$. ^{40}K atoms in the sites of this deep lattice can be cooled by using Raman sideband cooling or EIT cooling, and we will resolve the motional sidebands so long as the width of the dressed state $|2'\rangle$ is smaller than ω_0. We have used EIT cooling for all results in this thesis, and in this section we detail the specific beam powers, detunings, and propagation directions which provide cooling to the three motional degrees of freedom of the optical lattice.

5.2.1   Beam Layout

To implement EIT cooling in our system, we use a laser blue-detuned from the 770 nm D1 transition between the $4\text{S}_{1/2}$ ground state and the $4\text{P}_{1/2}$ excited state, as shown in Figure 5.6(a). A strong coupling field Ω_C drives the $|4\text{S}_{1/2}, F = 7/2\rangle \rightarrow |4\text{P}_{1/2}, F = 7/2\rangle$ transition and a weak probe field Ω_P is tuned near $|4\text{S}_{1/2}, F = 9/2\rangle \rightarrow |4\text{P}_{1/2}, F = 7/2\rangle$. The beam Ω_C is generated from a homebuilt external-cavity diode laser, stabilized to the D1 transition of ^{39}K and shifted to the appropriate detuning Δ using a double-pass AOM. To create Ω_P, we send a portion of the Ω_C light though an AOM running at 321.447 MHz and aligned in the quadruple-pass configuration depicted in Fig. 5.6(b), resulting in an overall frequency difference of 1285.79 MHz between Ω_P and Ω_C. With the magnetic field tuned to $< 4\,\text{mG}$, this sets $\delta = 0$ for the two-photon transition between ground states. We typically use a much more power for Ω_C than Ω_P, and so we assume that the dark state $|D\rangle$ is predominantly composed of the $F = 9/2$ ground state. The common-mode detuning Δ is typically 40 MHz, meaning that the two ground states are mainly coupled to the $F' = 7/2$ excited state of the D1 line, forming the desired Λ-system.

In the limit that $\Omega_C \gg \Omega_P$, EIT cooling relies on transitions between the state $|1', n, \nu\rangle$ and $|2', n + 1, \nu - 1\rangle$ which involves the absorption of a photon from Ω_P and subsequent emission of a photon into Ω_C. To have a nonzero Lamb-Dicke parameter $\eta_{2,j}$ for each trap axis j (as in Equation (5.14)) want the two-photon recoil of this process $\Delta\vec{k} = (\vec{k}_P - \vec{k}_C)$ to have nonzero projection onto all three lattice axes. To achieve similar $\eta_{2,j}$ for all axes, we use multiple beam paths to provide Ω_C and Ω_P resulting in multiple possible recoils $\Delta\vec{k}$. As shown in Fig. 5.7, one Ω_C beam C_{xy} propagates in the x-y plane, 30° from the x-axis. The first Ω_P beam P_{xy} enters 120° from the x-axis, so that the recoil momentum obtained by absorbing from P_{xy} and emitting into C_{xy} has a projection $\Delta k_x = 2k_{\text{D1}}\cos(15°)$ onto the x-axis and $\Delta k_y = 2k_{\text{D1}}\cos(75°)$ onto the y-axis. The cosine factor results in smaller coupling to the trap states of the y-lattice, which reduces the cooling rate. To provide a strong recoil along the y-axis, we retro-reflect the beam P_{xy} so that it re-enters the system propagating $-60°$ from the x-axis, leading to a second possible two-photon process with a large projection of $\cos(15°)$ onto the y-axis. Thus the largest possible values of $\eta_{2,j}$ for each trap axis are $\eta_{2,x} = \eta_{2,y} = 0.47$ and $\eta_{2,z} = 0.35$. We find that cooling is still possible if the beam P_{xy} is not retroreflected, but in some cases we have noticed reduced cooling performance.
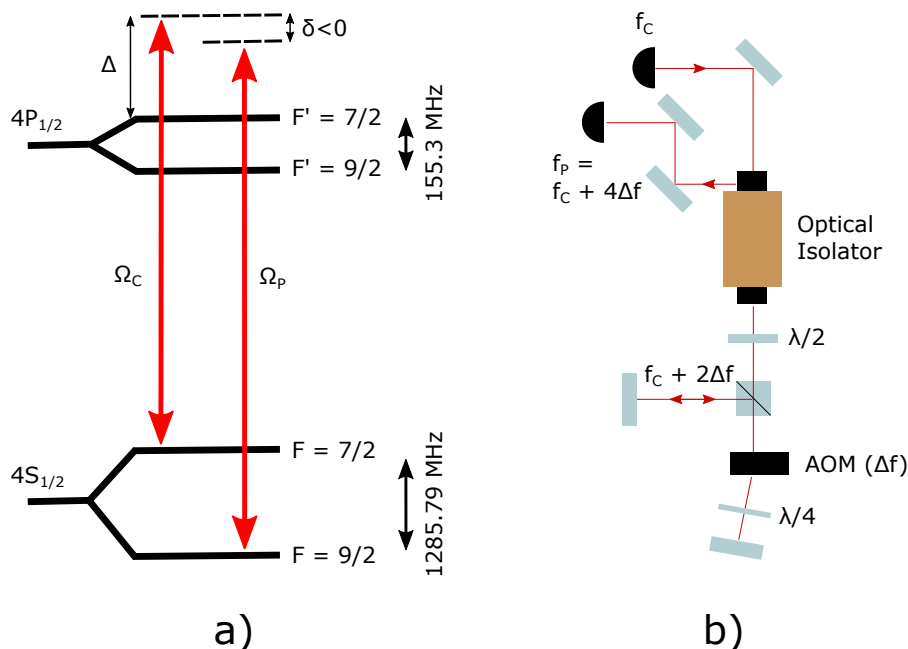
Figure 5.6: Creating the Λ-system for EIT cooling. a) Relevant atomic states for the cooling scheme. The common-mode detuning of the two beams Ω_P and Ω_C is denoted Δ, while their detuning from two-photon resonance is denoted δ. At zero field, $\delta = 0$ when $f_P - f_C = 1285.79\,\text{MHz}$. b) To create a phase-stable beam with frequency ω_P we make use of an optical isolator to send light at ω_C through an AOM four times. The AOM drive frequency Δf is provided by a custom DDS built by Alan Stummer. Using $\Delta f = 321.445\,\text{MHz}$ sets the two-photon detuning to $\delta = 0$. Ω_P can also be generated with two separate double-pass AOMs without the need for an optical isolator, or by using an EOM. This setup is easy to assemble and needs only a single low-frequency generator, but unavoidably produces a small amount of light at the output which is only shifted by $2\Delta f$. This light has no effect in our experiment.

(a)   (b)

Figure 5.7: Schematic for EIT Cooling Beams. (a) Lattices beams and cooling beams in the x-y plane. All cooling beams are linearly polarized, with P_{xy} and C_{xy} entering the system polarized perpendicular to the x-y plane. The beam P_{xy} is retro-reflected and re-enters the system linearly polarized in the horizontal plane. (b) A cooling beam C_z is 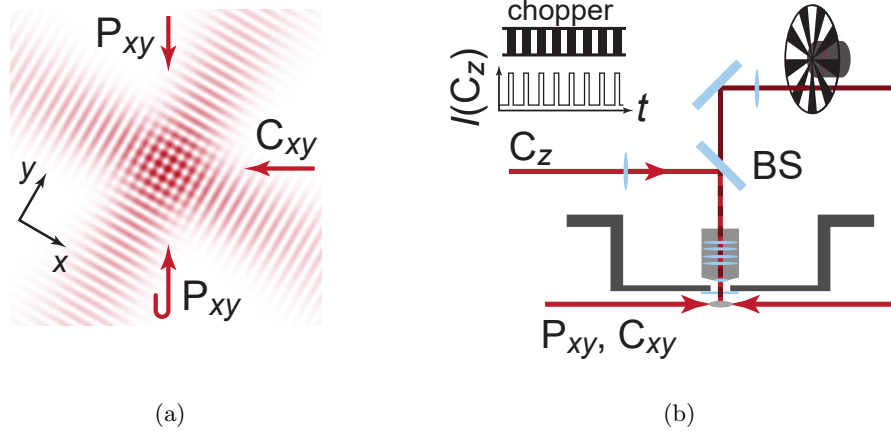sent along the vertical axis, through the microscope objective. The beam is coupled into the fluorescence imaging path using a 10% reflective plate beamsplitter (BS). Beam shaping before the beamsplitter ensures that C_z is approximately collimated after passing through the microscope objective. Reflections of C_z off of the glass surfaces of the microscope contribute a background to fluorescence images, so they are blocked by pulsing C_z synchronously with a spinning chopper wheel that shields the camera.

To cool the motion in the z-lattice, a beam propagating largely outside of the x-y plane is needed. Ideally this beam would not propagate within the acceptance angle of the microscope objective to prevent light from scattering into the fluorescence imaging path, but in our experimental geometry this is not possible. Instead we send a beam C_z directly along the imaging path as shown in Fig. 5.7. The beam is coupled into the fluorescence imaging path with a 10% reflective beamsplitter, and passes through the microscope before illuminating the optical lattice. In order to avoid strong focusing of C_z by the microscope we use a lens to focus the beam on the back focal plane of the objective, so that the lens and microscope act in combination like a reducing telescope. The position of this lens is set by overlapping C_z with a collimated 770 nm beam which propagates in the reverse direction through the microscope.

5.2.2 Background Reduction

In order to perform sensitive fluorescence imaging while we apply EIT cooling, we carefully consider the path of the EIT cooling beams. When cooling is properly tuned these beams have powers of a few μW, and so carry a flux of around 10^{13} photons per second. Any diffuse scattering of this large photon flux from the imaging window or metal chamber can reach the EMCCD camera and easily overwhelm the hundreds of fluorescence photons we collect from each atom in the ultracold cloud. To minimize any unwanted scattering of light from the horizontal cooling beams C_{xy} and P_{xy} we use $f = 750$ mm lenses to focus all of the beams in the horizontal plane to a Gaussian waist of $w_0 \approx 300\,\mu$m. Since the beam waists are smaller than the 800 μm distance from the cloud to the fluorescence imaging window, the scatter of photons off of the window is very low. Previously we found that reducing the beam size with a telescope placed near to the vacuum chamber led to a large background from scattered light, presumably due to multiple stray reflections from the glass surfaces. Switching to the 750 mm focal length for beam

size reduction allowed us to place all glass elements far from the chamber, minimizing this background. Additionally, we place irises around each cooling beam just before they enter the vacuum chamber, to block any other stray light paths originating on the lenses or polarization optics. After taking these steps, the background due to the horizontal cooling beams is reduced to the level of the camera noise floor when exposing the camera for 5 s. The statistics of the photon background on the camera is explored in more detail in Section 6.1.

The vertical cooling beam C_z can reflect from the glass surfaces of the microscope objective and contribute a large background of photons to the imaging system. One way to reduce this background is to use a polarizing element before the camera to reject the well polarized reflections while losing only 50% of the unpolarized fluorescence signal. We found this reduced the background by a factor of roughly 300. Alternatively, since the reflected light has a fundamentally different spatial mode to the fluorescence light, it can be blocked by spatial filtering. In the Fourier plane of the imaging system, the various wavevectors collected by the fluorescence imaging system are brought together to a spot size of radius 1000 μm, while the majority of the reflected cooling beam is focused to a radius of $< 50 \mu$m. Adding a circular dark spot of radius 200 μm, we can block much of the reflected light while only slightly attenuating the total fluorescence signal. Combining the polarizing and spatial filters resulted in an attenuation of the reflected light by 3 orders of magnitude, which was still insufficient for the fluorescence imaging of single atoms. Spatial filtering of the cooling light was still unable to block 100% of the beam, possibly because multiple reflected beams are generated by the many curved surfaces of the microscope objective. Each reflection would have a slightly different spatial mode, so that the sum of all scattered beams cannot necessarily be removed with a single spatial filter.

To improve the background reduction of the vertical cooling light, we implemented a pulsed cooling/imaging scheme in which a spinning chopper wheel alternately blocks and exposes the camera to the fluorescence and background light. This is synchronized with switching of the cooling beam power using an RF switch modulate an AOM. While the camera is blocked the vertical cooling beam is turned on, and when the camera is exposed the cooling beam is turned off. The spinning chopper wheel does not instantaneously block and unblock the light incident on the camera, instead taking a finite time to move across the face of a beam with finite size. This poses no issue for the exposure of the camera to the atomic fluorescence, since each possible wavevector of emitted light will still be exposed for 50% of the time on average, even though all possible wavevectors are not exposed at the same time. On the other hand, we require that the reflected cooling light be blocked entirely during the entire camera exposure, and so the cooling beam must remain off until the chopper wheel has moved into a position which completely blocks the beam. To minimize the delay required, we place the chopper wheel in the Fourier plane of the imaging system, where the reflected cooling light is focused to a small radius. This minimizes the chopper movement required to block the beam, and allows us to turn on the cooling beam for up to 35% of the total exposure while blocking all measurable background light from C_z that enters the fluorescence imaging path.

5.3   Optimizing Cooling

5.3.1   Tuning the Stark Shift

Since we desire a Stark shift equal to $\omega_0/\hbar$, we can rearrange the expression (5.9) to produce a requirement on the Rabi frequency of the beams used for cooling:

$$\Omega = 2\sqrt{\omega_0 \Delta}\,, \tag{5.15}$$

which gives $\Omega/2\pi \approx 6\,\mathrm{MHz}$ for our experimental parameters. The Rabi frequency obtained for a beam intensity I and initial and final states $|i\rangle$ and $|f\rangle$ can be calculated by evaluating

$$\hbar\Omega_R = \sqrt{\frac{2I}{c\epsilon_0}}\,\langle f|e\hat{r}|i\rangle\,, \tag{5.16}$$

where the interaction between the electric field and the atom has been expressed in the dipole approximation. The matrix element $\langle f|e\hat{r}|i\rangle$ includes all dependance on the specific atomic internal states due to angular momentum coupling. If the lambda system is formed between the $|9/2,9/2\rangle$ and $|7/2,7/2\rangle$ ground states and the $|7/2,7/2\rangle$ excited state of the D1 line, the desired Stark shift is obtained when using a beam intensity of about $20\,\mathrm{mW/cm^2}$ for the strong coupling beam. For our experimental beam waists, this corresponds to a power of about $30\,\mu\mathrm{W}$.

To confirm the Rabi frequency of the EIT cooling beams, we have measured Ω_C in our system by detecting the Stark shift it induces between the dressed states. Using a power of $80\,\mu\mathrm{W}$ for $\mathrm{C_{xy}}$ and setting $\Delta = 0$, we measure the resulting Rabi frequency by probing the absorption spectrum of the beam $\mathrm{P_{xy}}$ as we scan the two-photon detuning. Atoms are initially in the state $|F = 9/2, m_F = 9/2\rangle$, and the beam $\mathrm{C_{xy}}$ strongly couples the $|F = 7/2, m_F = 7/2\rangle$ state to the excited $|F' = 7/2, m'_F = 7/2\rangle$ state, modifying the absorption of $\mathrm{P_{xy}}$. An increased scattering of $\mathrm{P_{xy}}$ leads to loss or depumping of atoms, which reduces the number of atoms observed in a subsequent absorption image. For $\Omega_C \gg \Gamma$, we would expect the large Stark shift to split the excited state into two separate absorption resonances — the Autler-Townes effect. In the opposite limit $\Omega_C \to 0$, a minimum in the absorption of $\mathrm{P_{xy}}$ should persist at $\delta = 0$ due to EIT. For intermediate values of $\Omega_C \approx \Gamma$, both EIT and the Autler-Townes effect are important, and the absorption of $\mathrm{P_{xy}}$ is found by calculating the optical susceptibility through the optical Bloch equations [88, 89], resulting in the expression

$$R_{\mathrm{scatt}} \propto \mathrm{Im}\left(\frac{\delta + i\gamma_{21}}{\Omega_C^2/4 - \delta^2 + \gamma_{31}\gamma_{21} - i\delta(\gamma_{31} + \gamma_{21})}\right)\,, \tag{5.17}$$

where $\gamma_{31} = \Gamma/2$ is the decoherence rate of the excited state, and γ_{21} is the decoherence rate between ground states which may arise in our system from sources such as magnetic field fluctuations. The measured spectrum is shown in Fig. 5.8, and matches well to the expected form, when a fairly large value of $\gamma_{21}/2\pi = 10\,\mathrm{kHz}$ is used. This decoherence rate is not much larger than the photon scattering rates around $1\,\mathrm{kHz}$ which we detect during EIT cooling, and so is not unreasonable. The measured value of $\Omega_C/2\pi$ can be extracted from a fit of Eq. 5.17, or by determining the distance between the two maxima of absorption with peak-fitting, and is found to be $6\,\mathrm{MHz}$ for our $56\,\mathrm{mW/cm^2}$ beam intensity. This is close to the Rabi frequency of $11\,\mathrm{MHz}$ which would be expected for an intensity of $56\,\mathrm{mW/cm^2}$ driving the $|F = 7/2, m_F = 7/2\rangle \to |F' = 7/2, m'_F = 7/2\rangle$ transition, with the lower experimental value
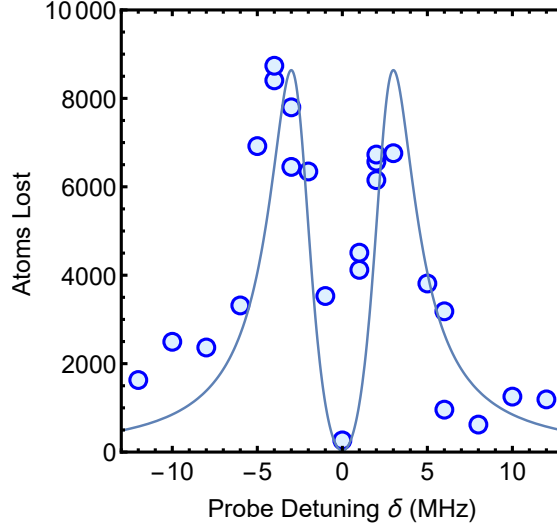
Figure 5.8: Measured excitation spectrum of the probe beam Ω_P as a function of the two-photon detuning δ, when the coupling beam Ω_C is tuned to resonance ($\Delta = 0$) with an intensity of $56\,\mathrm{mW/cm^2}$. Scattering of photons leads to heating of the atoms and loss from the trap. The excited state is split due to the Autler-Townes effect, leading to two distinct maxima for photon scattering, while EIT reduces the scattering to zero at two-photon resonance $\delta = 0$. The Rabi frequency $\Omega_C/2\pi$ is measured to be $6\,\mathrm{MHz}$, comparable to the linewidth $\Gamma/2\pi = 6\,\mathrm{MHz}$. The solid line is a plot of the theoretical absorption of Eq. (5.17) for $\Omega_C/2\pi = 6\,\mathrm{MHz}$ and $\gamma_{21}/2\pi = 10\,\mathrm{kHz}$.

possibly associated with errors in beam alignment, focussing, or polarization.

In order to optimize the Stark shift for cooling, we set the detuning of the beams to $\Delta/2\pi = 36\,\mathrm{MHz}$, hold the two-photon detuning at $\delta = 0$, and vary the intensity of the cooling beams. All beam intensities are adjusted by the same AOM, so that the ratios P_{xy}/C_{xy} and C_z/C_{xy} are held constant while the total intensity is varied. Applying the beams for $500\,\mathrm{ms}$ in a deep optical lattice leads to large heating and atom loss unless the EIT cooling mechanism can effectively cool the cloud. Thus an easy way to optimize the cooling process is to apply the EIT cooling beams for long periods of time and try to increase the survival of atoms.

When we use EIT cooling to look for site-resolved fluorescence of single atoms, we will work with atoms occupying only one plane of the vertical lattice. Such a small cloud of atoms provides little signal for a TOF absorption image, so that it difficult to initially calibrate the cooling mechanism. As a first step towards site-resolved fluorescence imaging, we instead optimize the cooling parameters while looking at a full 3D cloud of 10^5 atoms loaded into the lattice. In this case we are affected much more by lattice inhomogeneity, but the signal for absorption imaging is much stronger. Applying EIT cooling to such a large cloud, we find a maximum atom survival at an intensity of $30\,\mathrm{mW/cm^2}$ as shown in Figure 5.9, indicating that this intensity the Stark shift has moved the red-motional sideband into resonance with the beams at $\delta = 0$. Since the atoms occupy many sites with different local lattice depths, this matching between the Stark shift and the trap frequency will not occur simultaneously for all atoms.

As the power is reduced towards zero the cooling becomes less efficient, and atoms are simply heated due to photon scattering from the beams. As the intensity of all beams approaches zero, we expect the survival to increase again since the heating due to photon scattering disappears entirely. At high intensities the cooling appears less effective but the mechanism clearly remains active even with the
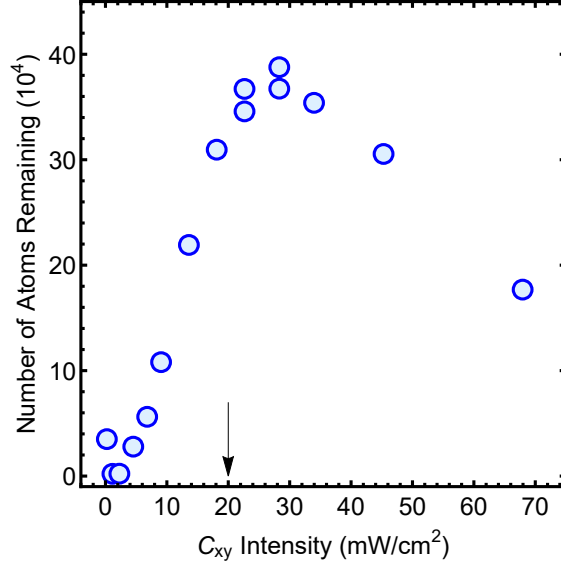
Figure 5.9: Tuning the Stark shift by measuring the number of atoms remaining in the lattice. The detuning is $\Delta/2\pi = 36\,\mathrm{MHz}$, and the probe beam is at $\delta = 0$. The power of all EIT cooling beams is increased simultaneously, preserving the intensity ratios $P_{xy}/C_{xy} = 0.04$ and $C_z/C_{xy} = 0.1$. The arrow denotes the estimated optimum intensity for the Λ-system formed by the $|4S_{1/2}, 9/2, 9/2\rangle$, $|4S_{1/2}, 7/2, 7/2\rangle$, and $|4P_{1/2}, 7/2, 7/2\rangle$ states. Depending on the particular $|F, mF\rangle$ sublevels that form the Λ-system, we infer a Stark shift of $V_{\mathrm{Stark}}/h = (100\text{-}400)\,\mathrm{kHz}$ due to C_{xy} at the optimal intensity of $30\,\mathrm{mW/cm^2}$.

intensity set to twice the optimal value. Hence EIT cooling is seen to be fairly insensitive to the precise laser intensity. This may indicate that EIT cooling is able to cool effectively over a broad range of experimental parameters. Alternatively, this may indicate that due to the many inhomogeneities in the system (arising for example from beam intensity gradients, variations in trap frequency, or the internal-state-dependence of the Stark shift) we only optimize the cooling process for some subset of atoms at any individual intensity. When imaging the fluorescence generated by EIT cooling in-situ (as discussed in Chapter 6), we will be able to distinguish many of these sources of inhomogeneity.

5.3.2 Effect of the Lattice Stark Shift

The EIT cooling mechanism relies on the existance of a long-lived dark state $|D\rangle$. Since the lifetime of this state will scale with $1/\delta$, cooling is optimized by driving the Λ-system at two-photon resonance with $\delta = 0$. This is ensured by setting the frequency difference $\Delta\nu$ between the two fields Ω_C and Ω_P to the match the precisely known Hyperfine splitting of $^{40}\mathrm{K}$

$$\Delta\nu = E_{\mathrm{HF}}/h = 1285.79\,\mathrm{MHz}\,. \tag{5.18}$$

Experimentally we find that δ must be tuned to within $50\,\mathrm{kHz}$ of two-photon resonance, or else the cooling effect is disrupted. Any effect which alters the energy difference between states $|1\rangle$ and $|2\rangle$ will change δ and therefore spoil EIT cooling. Since the magnetic field in the experiment is precisely cancelled during the cooling process, the only remaining effect which can influence δ is the Stark shift due to the confining lattice potential.

The lattice is very far detuned from all atomic transitions, so that the trapping potential which

we create is almost perfectly state-independent. This is ensured by using beams which have only linear polarization, so that only a scalar light shift is generated. A residual differential shift is produced because atom in the $F = 7/2$ ground state will see a laser detuning which is about $1\,\mathrm{GHz}$ smaller than an atom in the $F = 9/2$ ground state. This is a very small frequency difference compared with the $10^6\,\mathrm{GHz}$ detuning of the $1054\,\mathrm{nm}$ beam, and so the $F = 7/2$ state experiences only a slightly stronger trapping potential: $U_{7/2} \approx U_{9/2}(1 + 10^{-6})$. The total energy shift induced by all of our optical lattice beams is less than $\Delta E_{\mathrm{max}}/h \approx 10\,\mathrm{MHz}$, so that the difference in trapping potentials $U_{7/2} - U_{9/2}$ is less than $10\,\mathrm{Hz}$. This frequency shift between the two ground states is far to small to have any effect on the EIT cooling mechanism.

A more significant differential shift $U_{7/2} - U_{9/2}$ can result if the optical lattice beams do not have perfect linear polarization. In this case a vector light shift is generated, with the circular polarization component of the beam acting as an effective magnetic field and lifting the degeneracy between m_F states. For our far detuned lattice beams even a 100% circular polarization of the lattice beams only results in energy shifts of around $\Delta E/h \approx 100\,\mathrm{kHz}$ at large lattice depth. We have experimentally measured the presence of any vector shifts by probing the hyperfine transition frequency as a function of the optical lattice strength, and find that over the full range of lattice powers the transition shifts by less than $10\,\mathrm{kHz}$. Thus any vector shifts are sufficiently small and even if atoms in different parts of the optical lattice experience different levels of this shift they will all sufficiently satisfy the two-photon resonance condition $\delta = 0$.

Since we have confirmed that the Stark shift introduced by the lattice is almost entirely state-independent, atoms in different regions of the trap will all be coupled to Ω_C and Ω_P at two-photon resonance. This means that all atoms in the cloud can be pumped into a dark state in which photon scattering is suppressed by EIT. Due to other inhomogeneities in the system it may not be possible to optimize the *cooling* mechanism for all atoms simultaneously, as discussed below. Even with inhomogeneities in the cooling performance, we expect all atoms to experience reduced scattering, reducing the heating processes that compete with cooling.

The cooling mechanism will be spoiled for a given atom if the on-site trapping frequency and the overall Stark shift of the cooling beams are not properly matched. A mismatch between Stark shift and trap confinement can occur in many different ways. If the cooling beams are tightly focussed with a beam waist comparable to the size of the atom cloud, different atoms will experience different intensity, causing the Stark shift of the cooling beams to be non-uniform across the cloud. To minimize such a variation in the Stark shift our cooling beam waists of $w_0 \approx 300\,\mu\mathrm{m}$ are chosen to be roughly $10\times$ larger than the typical cloud size. The beams which form the the optical lattice are focussed more tightly however, so that a large and dilute cloud may sample low-intensity regions of the lattice with reduced on-site trapping strength. Since the on-site trapping frequency is proportional to the square-root of lattice laser intensity, intensity gradients in the lattice beam are actually only weakly transferred into gradients in the trap frequency. As long as we focus only on a small region in the centre of the lattice, the trap frequency can be quite uniform as shown in Section 3.5.1, so that homogeneous cooling of all atoms should be possible.

An additional inhomogeneity due to the lattice beams appears due to the strong Stark shift they apply to the D1 excited state. The Stark shift of the D1 excited state can be as large as $80\,\mathrm{MHz}$ as discussed in Section 2.3.7. Since we operate our D1 cooling beams at fixed frequency, increasing the lattice laser intensity can then Stark shift the D1 excited state significantly closer to resonance with
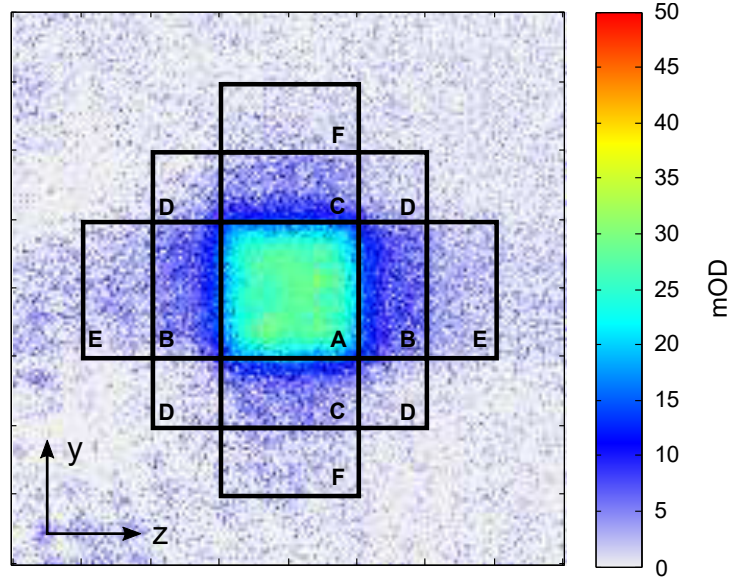
Figure 5.10: Absorption image of 40K atoms after EIT cooling has been applied for 5 s. Bandmapping has been performed by turning off the optical lattice smoothly over 2 ms. A 15 ms time-of-flight expansion gives an image of the quasimomentum distribution, showing a high occupancy of the square first Brillouin zone. The bandmapping procedure distinguishes different bands $n_{y,z}$ of the lattice, with $(n_y = 0, n_z = 0) \rightarrow$ A, $(1,0) \rightarrow$ B, $(0,1) \rightarrow$ C, $(1,1) \rightarrow$ D, $(2,0) \rightarrow$ E, and $(0,2) \rightarrow$ F. The lattice depth before bandmapping was $800\,E_R$, with negligible tunneling such that band numbers are equivalent to harmonic oscillator states for atoms isolated at each lattice site. 100 images have been averaged for improved signal to noise.

the cooling beams. The smaller resulting detuning of the EIT cooling beams generates a greater Stark shift of the dressed ground state $|2'\rangle$ used in the cooling process, spoiling the cooling process by shifting the red motional sideband out of resonance. This coupling between the lattice intensity and the cooling Stark shift is minimized if the detuning Δ of the EIT cooling beams is chosen to be large compared to the range of possible lattice Stark shifts. Using small Δ leads to poor cooling results and fluorescence images show clear spatial patterns of atom loss from either the centre or edges of the lattice. By increasing the detuning to $\Delta = 2\pi \times 40\,\text{MHz}$ we find that survival of a large cloud in the lattice is improved, and fluorescence images indicate more uniform survival and brightness of the cloud.

5.4   Cooling Results

To assess the temperature reached by EIT cooling, we measure the ground-state fraction using absorption images obtained after band mapping from the lattice. An example of the band mapped populations after several seconds of EIT cooling is shown in Figure 5.10. As discussed in Section 4.2, different harmonic oscillator levels are mapped into different regions of the time-of-flight image so that the ground state fraction along two motional axes can be calculated by finding the fraction of the total atomic density appearing in the central square region $(|p_y| < \hbar k_L, |p_z| < \hbar k_L)$. After EIT cooling we measure a ground state population of 79% in the y-lattice and 72% in the z-lattice. The difference between the two directions is likely because to the pulsed configuration of the vertical beam provides less cooling to the
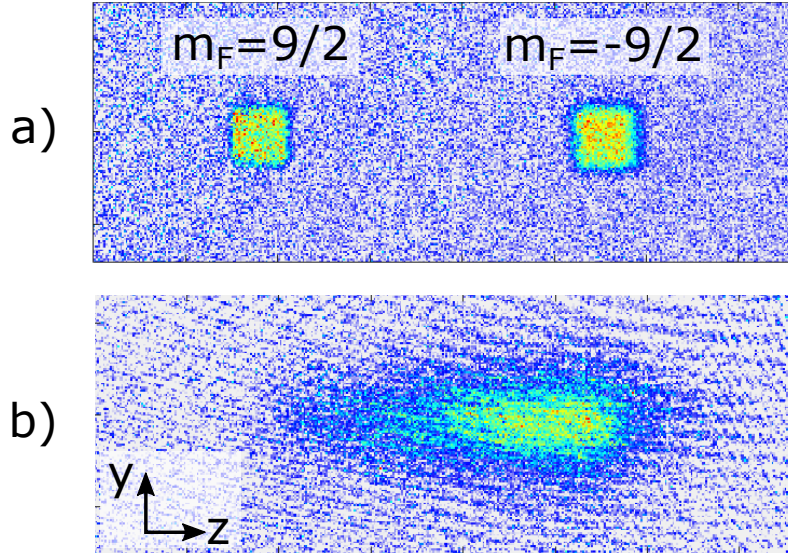
Figure 5.11: Absorption images of a cloud after band mapping from the optical lattice and applying a Stern-Gerlach pulse to spatially separate all m_F sublevels along $\hat{z}$. a) Shows the cloud when no EIT cooling is applied but with a mixture of $|9/2, 9/2\rangle$ and $|9/2, -9/2\rangle$ prepared before imaging. b) Shows the cloud after EIT cooling is applied to atoms initially in $|9/2, 9/2\rangle$, showing that the cooling process tend to distribute atoms among many m_F sublevels. In this instance, a small magnetic field of $40\,\mathrm{mG}$ is present, causing atoms to accumulate in states near $|9/2, -9/2\rangle$ and making the cooling less efficient. Precisely zeroing the ambient field to $< 5\,\mathrm{mG}$ improves cooling and distributes atoms equally among all m_F levels.

z-axis. Since the ground state population of the x-lattice is expected to be similar to that of the y-lattice, this indicates that roughly 45% of the atoms are in the 3D motional ground state. This is a lower bound for the true occupation of the ground state, since zero-temperature atoms can still appear outside the central region due to the effects of finite image resolution, nonzero initial cloud size, and non-adiabaticity in the lattice ramp [90]. For comparison, the same band mapping procedure for a degenerate Fermi gas loaded into the ground band of a $800\,E_R$ lattice and immediately band mapped (without applying EIT cooling) yields a 3D ground state fraction of 64% even through nearly 100% occupancy of the ground band is expected in this case. To better estimate the temperature increase of the previously degenerate cloud due to EIT cooling, we can take the difference between band mapped images with and without EIT cooling applied. This indicates that EIT cooling reduces the ground state fraction of the initial cloud by roughly 20%, so that we actually achieve 80% occupation in the 3D motional ground state during cooling.

Since our cooling beams form a moving polarization gradient lattice, atoms will sample all beam polarizations over time. Thus there is no net optical pumping effect for the atoms and the cloud which is initially spin-polarized will depolarize into a mixture of possible m_F after several photon scattering events. We can use a Stern-Gerlach sequence to separately image the different m_F states after EIT cooling has been applied, with adjacent states almost fully separating in TOF due to the small momentum distribution of the cloud. A Stern-Gerlach separated cloud after EIT cooling is shown in Figure 5.11, where the spread over many m_F levels is visible. Examining the m_F state distribution at various points in the EIT cooling cycle reveals that the cloud becomes depolarized after about $10\,\mathrm{ms}$ of EIT cooling. We see in this image that atoms are distributed more in the negative m_F levels near $|9/2, -9/2\rangle$, which

was due to a small bias field present when this data was taken. The field breaks the degeneracy of the m_F sublevels, and disrupts the cooling for some EIT cooling beam polarizations. After reducing the ambient magnetic field to $< 5\,\mathrm{mG}$, we found that atoms were more evenly distributed among all possible m_F states after cooling, and the lifetime of atoms was significantly improved.
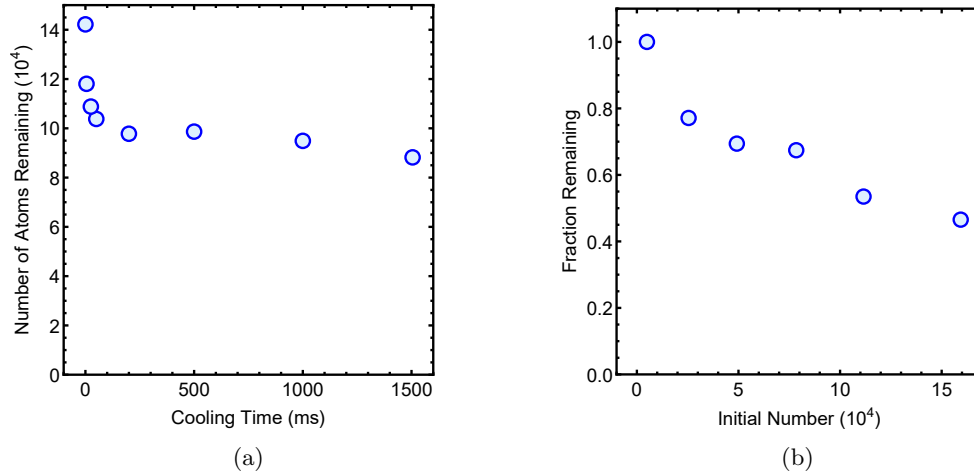

Figure 5.12: Determining the number of atoms pinned in the lattice during EIT cooling. (a) The number of atoms remaining in the lattice versus EIT cooling duration. Loss of atoms occurs in the initial $100\,\mathrm{ms}$ of cooling, after which the atom number remains nearly constant. (b) Fraction of atoms which remain pinned for $> 1000\,\mathrm{ms}$ of cooling, compared to the initial number of atoms loaded into the lattice. For small initial numbers no atoms are lost, and with increasing number the survival fraction decreases.

The survival of atoms in the lattice during application of the EIT cooling beams is a very useful metric to use when optimizing the cooling process. In order to perform fluorescence imaging, we require that the cooling beams keep atoms pinned in sites of the lattice for several seconds. In Figure 5.12(a) we see the typical survival of a large cloud of atoms during cooling. Over the first $100\,\mathrm{ms}$ of cooling roughly 30% of atoms are lost from the trap, but the remaining 70% stay in the lattice for long times. The initial loss occurs because the cloud is not small compared to the lattice beam waists, so that some atoms are loaded into lattice sites far from the central, high-intensity region. Atoms loaded into these sites are more weakly bound, and EIT cooling cannot cool them effectively. The constant fraction of surviving atoms at times $> 200\,\mathrm{ms}$ indicates that all poorly trapped atoms have been lost and only the central, well-trapped atoms remain in the trap.

To explore the initial loss of atoms during EIT cooling, we varied the number of atoms loaded into the lattice before cooling. Loading a small cloud leads to atoms occupying the deeper sites near the centre of the lattice beams, so that the cloud should be subject to less inhomogeneity of lattice depth. Figure 5.12(b) shows the fraction of atoms surviving after $1000\,\mathrm{ms}$ of cooling, for various initial cloud sizes. Larger initial clouds have greater occupation of the poorly-cooled regions at the outer region of the trap, and greater loss of atoms is observed. As we decrease the initial number of atoms, the survival approaches 100%, indicating that EIT cooling is able to keep atoms inside the optical trap for long times.
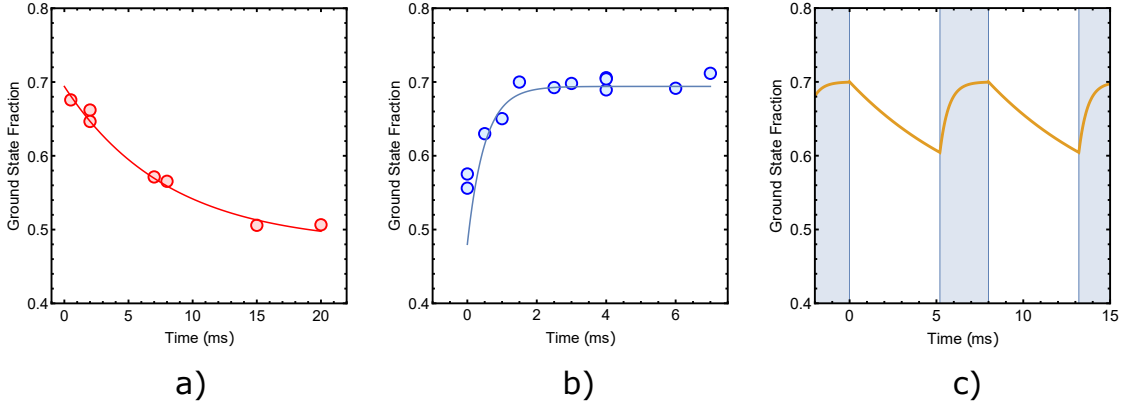
Figure 5.13: Ground state fraction in the vertical lattice during modulated vertical cooling. a) After the vertical cooling beam is switched off, the ground state fraction drops exponentially with a time constant of 8 ms. b) When the vertical cooling beam is switched on again, the cloud is cooled rapidly with a time constant of 500 μs. The solid lines in a) and b) are exponential fits to the data used to extract the time constant. c) Expected temperature evolution for a vertical cooling beam undergoing square-wave power modulation with a period of 8 ms and a duty cycle of 35%. Shaded regions indicate the periods in which cooling is applied.

5.5   Vertical Beam Parameters

Due to the background light it contributes to the imaging, the vertical cooling beam C_z requires special attention. As discussed in Section 5.2.2, we can turn this beam on 35% of the total cooling period and while shielding the camera using a spinning chopper wheel. To see the effect that this has on the atoms we can use band mapping to measure the heating of the vertical degree of freedom when C_z is turned off. As seen in Figure 5.13(a), after the vertical beam is switched off the ground state fraction in the z-lattice drops from a measured steady-state value of around 70% with a time constant of 8 ms. Figure 5.13(b) shows that the vertical degree of freedom is cooled quickly back to the steady-state value after the vertical beam is re-introduced, with a time constant of 500 μs. If the chopper modulates the vertical beam at 62.5 Hz (8 ms on / 8 ms off), we would expect the temperature evolution of the vertical degree of freedom to look like the curve in Figure 5.13(c), with a minimum ground state fraction of 60% in the vertical direction. Typically we operate the chopper at higher frequencies around 200 Hz, so that the ground state fraction stays even closer to the 70% steady state value during the entire cooling process.

Since the effectiveness of EIT cooling depends on the magnitude of the Stark shift induced by all applied cooling beams, we should also ensure that the presence or absence of the vertical cooling beam does not greatly affect cooling of the two horizontal degrees of freedom. Since we typically use only a small vertical intensity $C_z/C_{xy} = 0.1$, the Stark shift will change only by about 10% when the vertical beam is modulated. In Figure 5.9 we see that a 10% change in intensity around the optimal value has little effect on atom survival, and so we expect the cooling of the horizontal axes to persist at all times during the vertical beam modulation. The change in cooling intensity during modulation could be reduced by sending only the weaker Ω_P light through the vertical axis, or could be eliminated entirely by modulating the power of the C_{xy} beam synchronously with the vertical beam pulses.

5.6 Conclusions

We have found that EIT cooling is able to cool ^{40}K atoms in a deep optical lattice, reaching an estimated motional ground state occupancy of 80%. Even though the cooling process ideally pumps atoms into a dark state, sources of heating and decoherence present in the system will limit the lifetime of the dark state. This means that EIT cooling will steadily produce photon emission from the cloud without the atoms heating out of the sites of the lattice. Since photons emitted have the same wavelength as the cooling beams, it is necessary to shield the camera from the cooling beams. We find that EIT cooling is still effective even if the cooling of one motional degree of freedom is turned off $> 50\%$ of the time, so that we can shield our camera from the cooling beams by using a spinning chopper wheel. This allows us to exposure the camera while applying EIT cooling in order to look for the faint fluorescence signal from the individual atoms in the trapped gas.

All diagnostics presented so far have probed the cooling by using absorption images of a large cloud of atoms. The large cloud size leads to large lattice inhomogeneity, which makes an interpretation of the cooling results difficult. However we can also tune the cooling parameters while looking at fluorescence images of the cloud through the microscope objective. Looking at the lattice locally will allow us to see beyond the lattice inhomogeneity, and reducing the cloud thickness to one vertical plane should suppress additional heating processes due to the re-absorption of scattered photons [91]. For these reasons all final adjustments to the EIT cooling parameters are performed based on in-situ fluorescence measurements, as described in the next chapter.

Chapter 6

Single Atom Imaging

In Chapter 5 we have seen that our ultracold gas can be held in our deep optical lattice for several seconds, while EIT cooling beams continuously cool the atoms keeping them near the motional ground state. The laser cooling beams are close to the atomic resonance, but tend to pump atoms into dark atomic states so that photon scattering is greatly reduced compared to the atomic linewidth $\Gamma/2\pi = 6\,\mathrm{MHz}$. For the success of the EIT cooling mechanism described in the previous chapter, the scattering rate of photons should be reduced below the $250\,\mathrm{kHz}$ oscillation frequency of the confining potential. We find that while cooled near the motional ground state atoms continue to scatter photons at a rate near $1\,\mathrm{kHz}$. To produce thousands of photons from each atom for fluorescence imaging, we expose our camera to this fluorescence for several seconds.

We find that with exposure times $> 2\,\mathrm{s}$, the fluorescence signal from single atoms can clearly be distinguished from background noise. The apparent size of the individual atoms can be used to measure the PSF of our fluorescence imaging system, and we can use the locations of the atoms to determine the positions of all optical lattice sites in a given image. Once the position of the optical lattice is fixed, we can determine the most likely lattice occupation for an observed fluorescence image. The process to convert a given image into a lattice site occupation is detailed in this chapter, along with details of our optimization of EIT cooling for fluorescence imaging.

6.1   Single Atom Sensitive Images

It is crucial that all unwanted background light is suppressed in order to achieve single-atom sensitive imaging. Since the rate of photon collection from any individual atom will be very low, the number of background photons striking each pixel of the EMCCD camera must be exceptionally small. As discussed in Section 5.2.2 we have tailored the paths of the EIT cooling beams to minimize and scattering into the imaging system. This includes the implementation of a spinning chopper wheel in the Fourier plane of the imaging system which allows us to completely block the cooling beam which propagates along the z-axis. The small fraction of the vertical lattice beam which is transmitted by our imaging window is blocked by a shortpass filter with OD > 5. To reduce the occurrence of thermal dark counts the EMCCD chip is cooled to $-80°$ C. Instead of switching off all ambient lighting in the lab, we use a narrow-band interference filter in front of the camera which has OD > 4 for wavelengths in the ranges 200-750 nm and 800-1200 nm. The net result of all of these precautions is that almost no light should illuminate the
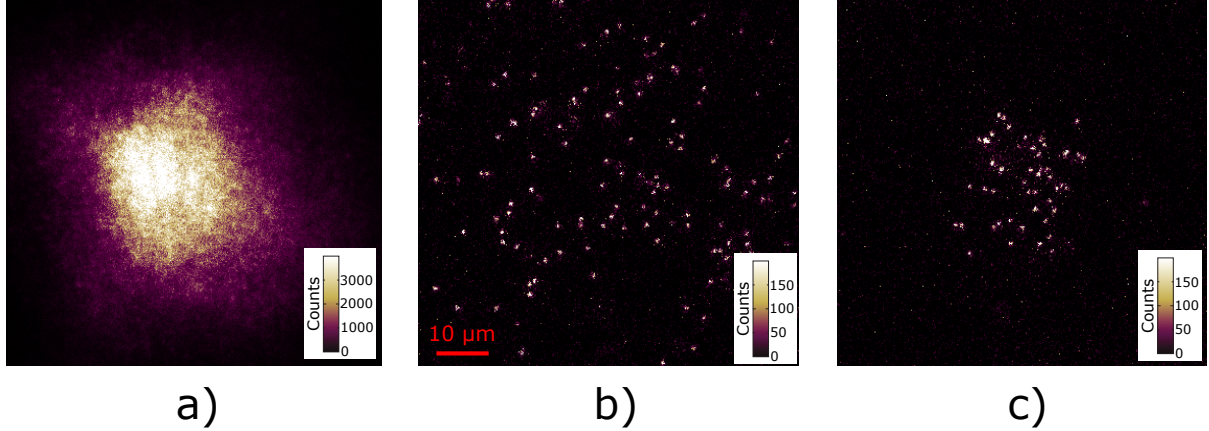
Figure 6.1: Fluorescence images obtained using EIT cooling. a) A fluorescence image of the entire cloud without plane selection. b) An image of a single selected plane, with single atoms clearly visible. c) An image of atoms after spectroscopic selection has been used to isolate a 40×40 site region in the centre of a single vertical plane. In b) and c) a hot cloud has been intentionally produced resulting in a sparse occupation of the optical lattice.

camera apart from the fluorescence of our laser-cooled cloud.

To measure the background light level that we achieve we expose the camera for $2.6\,\mathrm{s}$ with all cooling and trapping beams turned on, but without any atoms loaded into the lattice. We measure an average of 4.5 counts per pixel using EMCCD gain of 300. To determine the calibration factor between EMCCD counts and photoelectrons, we repeatedly illuminate the sensor with a weak beam while using an EMCCD gain of 300. Assuming that there are no fluctuations in the beam, the variation of the number of photoelectrons on each pixel from one image to the next should be determined by shot noise. The effect of shot noise in the creation of N_e photoelectrons should be associated with a standard deviation $\sigma_e = \sqrt{N_e}$. However, the stochastic nature of the electron multiplication process in the EMCCD leads to an effective increase in shot noise by a factor of $\sqrt{2}$ [92]. If the number of counts detected per photoelectron is $N_c = CN_e$ with standard deviation $\sigma_c = \sqrt{2}\sigma_e$, then we can determine C using the formula

$$C = \frac{\sigma_c^2}{2N_c} \,. \tag{6.1}$$

We thus determine $C = 13$ for our typical fluorescence imaging parameters. Using this value of C we find that an average of 0.35 photoelectrons is created on each pixel of the EMCCD by background noise during a typical exposure. Since the quantum efficiency of our camera is expected to be 85% at our imaging wavelength, this corresponds to 0.41 photons per pixel. The variation in background signal that we measure is bimodal as shown in Figure 6.2(a), with most pixels recording < 1 count and a few pixels receiving larger counts according to an exponentially decaying tail. This bimodal structure can be understood by modeling the noise processes in the EMCCD camera [93] showing that the Gaussian peak at low photoelectron numbers is likely due to noise in the read-out process and the exponentially decaying tail of photoelectron counts indicates noise contributed by the presence of stray light. The standard deviation of the background noise distribution is $\sigma_{\mathrm{BG}} = 0.6$ photoelectrons.

Exposing the camera for $2\,\mathrm{s}$ with a full cloud of atoms trapped in the optical lattice results in a fluorescence image as shown in Figure 6.1(a). No individual atoms are resolved in such an image because light is collected simultaneously from many vertical lattice planes, most of which are not in focus.

However, such an image can be collected in the same experimental run as a bandmapped absorption image as shown in Figure 5.10. Acquiring these two images allows us to assess the brightness and survival of a cloud during EIT cooling, and is useful to diagnose any errors in the fluorescence imaging sequence from day to day. The integrated fluorescence of the entire cloud is typically 9.2×10^7 counts at an EMCCD gain of 300, corresponding to around 8×10^6 collected photons. This signal is large enough to be distinguished above the camera noise floor even without using EMCCD gain.

The ability to obtain in-situ images of our ultracold cloud is an excellent tool in a quantum simulation experiment. However, care must be taken if one attempts to measure the cloud density profile from the shape of a fluorescence image such as Figure 6.1(a). Since there are many sources of inhomogeneity in our optical lattice, the fluorescence obtained in a given region of the camera sensor is determined not only by the cloud density but also the spatial profile of the trapping and cooling beams. For this reason, we do not use our fluorescence images to probe the density distribution of the cloud directly. Instead, we find the fluorescence images useful for characterization of inhomogeneities important for EIT cooling, such as the inhomogeneity in the lattice trapping frequencies shown in Figure 3.16.

After we use the plane selection procedure outlined in Section 4.3 to remove all planes of the vertical lattice except one, we focus the microscope on the remaining fluorescence and find that single atoms are visible as localized bright spots as shown in Figure 6.1(b). Exposure times of at least $2\,\mathrm{s}$ and an EMCCD gain of 300 allow us to easily identify these atoms above the noise floor of the camera. The difference between atom signal and noise can be seen clearly in a single image when we use further spectroscopic selection to remove atoms outside of the central 40×40 sites of the horizontal lattice as shown in Figure 6.1(c). Images of this sort are much less affected by the various inhomogeneities of the lattice, and are used for all of our image analysis. The central subset of lattice sites isolated by our horizontal selection is also close to the system size that we will eventually use for quantum simulation, so it is crucial that we optimize our imaging technique for this size of system. Equivalently, we could simply image a large cloud and analyse only the centre of the image corresponding to the lattice centre. However, we find that atoms in the outer regions of the trap are poorly cooled and often hop into the central sites, leading to errors in the analysis. Removing all poorly cooled atoms allows us to look only at the motion of the well-cooled atoms.

The signal strength and widths of each bright spot in a fluorescence image can be characterized by fitting them with a Gaussian function. The shape of atoms in the fluorescence image should be dominated by the PSF of the imaging system which does not generally have a Gaussian shape, but a Gaussian function is a good approximate form for a simple analysis. Since all atoms are identical and few spatial inhomogeneities are expected in the EIT cooling process, we would expect the number of photoelectrons created on the EMCCD at each bright spot to simply follow a Poissonian distribution with a standard deviation of $\sqrt{2N_e}$. However, if atoms are lost from a lattice site during the exposure time, they will create only a fraction of the expected photoelectron counts. We avoid the effects of atom loss in our measurements of photon statistics by taking 10 consecutive exposures of the same cloud, and collecting data only for atoms which are still pinned to their lattice sites in subsequent exposures. The distribution of collected photoelectrons from Gaussian fits to pinned single atoms is shown in Figure 6.2(b) and has $N_e = 110$ and $\sigma_e = 21$. The measured standard deviation exceeds the expected shot noise level by $\sqrt{2}$.

Since we have used the signals from several spatially separated atoms to collect the statistics in Figure 6.2(b), the larger-than-expected standard deviation might be due to spatial inhomogeneities in
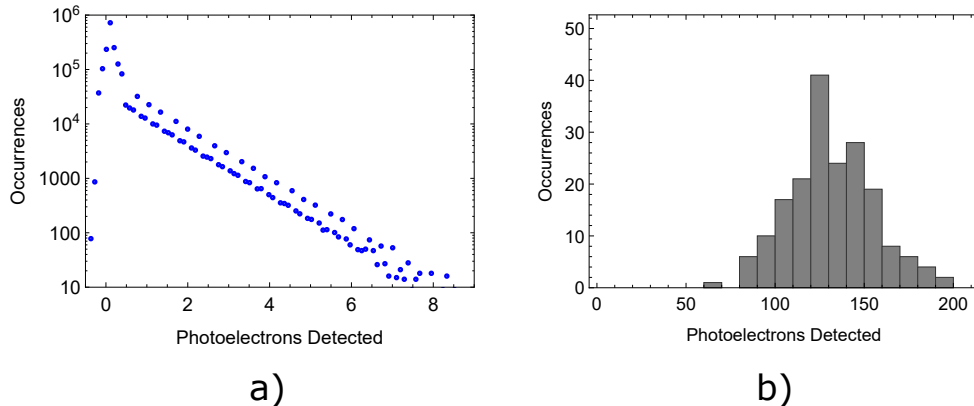
Figure 6.2: Histograms of photoelectron detections in an image. (a) Background photoelectrons detected per pixel, obtained from 2.6 s exposures which include all cooling and trapping beams but contain no atoms. (b) Photoelectrons detected from isolated atoms in several images using an exposure time of 2.6 s. Isolated atoms are fit with a 2D Gaussian function and the integral of the fit is used to determine the number of photoelectrons corresponding to a single atom.

the brightness of atoms. Such inhomogeneity could be due to spatial profile of the cooling beams, the spatially dependent trapping frequencies, or interference effects on the EMCCD chip. To test for possible spatial variations in the atom brightness, we can compare the photoelectron statistics in two scenarios: the variation in brightness for repeated images of the same atom in the same lattice site, and the variation between different atoms in different lattice sites. We find that repeated images of the same atom have a lower standard deviation of $\sigma_e = 18$, suggesting some spatial inhomogeneity. However we see that even when we measure the fluorescence of a single atom in repeated images, we still see variations which are in excess of shot noise. These large variations in signal that occur in the same region of the trap and are imaged with the same region of the EMCCD could instead be due to errors in our measurement of photon number (due perhaps to a dependence of the Gaussian fit parameters on the residual background noise) or a miscalibration of the increased noise induced by the electron multiplication process in the camera.

6.1.1   Locating Isolated Atoms

When we fit Gaussian functions to isolated atoms in a fluorescence image, the central position of the atom can be determined to much higher precision than the size of the imaging system's PSF. This is because the position x_i of each detected photoelectron can be considered as a separate sample of the atom's true position, drawn from a probability distribution specified by the PSF, camera pixel size, and background noise contributions. The mean value $\bar{x}$ of these individual measurement is an estimate of the true atom position for which the uncertainty $\Delta \bar{x}$ typically scales with $N_e^{-1/2}$. Since $\Delta \bar{x}$ for isolated emitters can be made smaller than the PSF and optical wavelength of a microscope, it becomes possible to read out certain kinds of information from an image with a much higher spatial resolution than one would naively expect based on formulas like the Rayleigh criterion. This approach to imaging is now commonly applied in the field of super-resolution microscopy. The same principles were recently applied to the optical imaging of a single trapped ion, allowing position measurement to be made with an uncertainty of only a few nanometers [94]. In our experiment locating atoms with small position
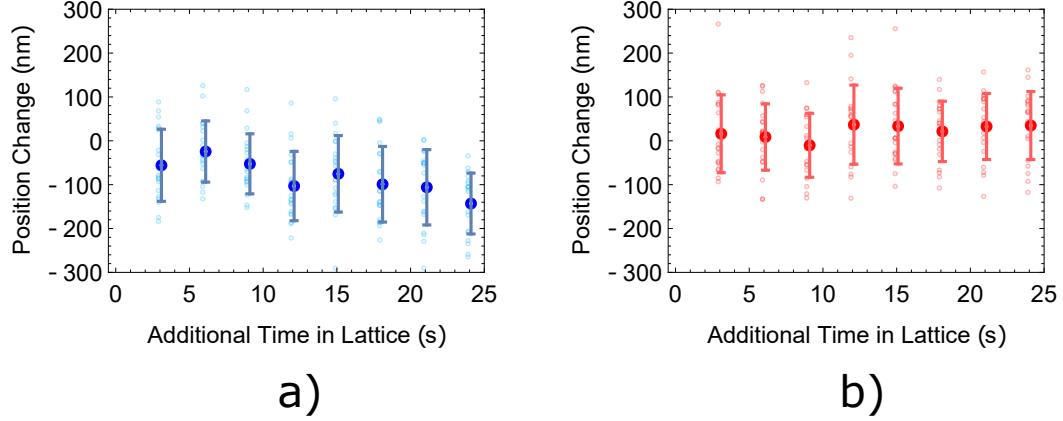
Figure 6.3: Tracking the measured position of isolated atoms during several consecutive exposures. Each exposure is 2.6 s long with 0.4 s of additional read-out time for the camera. Around 100 photons are collected from each atom, and a Gaussian fit is used to estimate the atom position. Comparing the measured position in the first exposure to each subsequent exposure allows us to look for drifts in the microscope objective or confining lattice that would be common to all atoms. (a) and (b) show the position drift over time in the horizontal and vertical dimensions of the pixel grid, respectively. Faint points are the data from individual atoms, and solid circles show the average drift of 20 atoms. Error bars show the standard deviation.

uncertainty is crucial if we wish to combine many single-atom images to examine the PSF of the imaging system, or when we attempt to locate the position of the optical lattice grid in an given image.

If we approximate our PSF by a Gaussian function with width σ_{PSF}, the combined position uncertainty from N collected photons due to the PSF, the finite pixel size d_{px}, and the standard deviation of the background noise σ_{BG} can be expressed approximately by [95]

$$\Delta \bar{x} = \left(\frac{2\sigma_{\mathrm{PSF}}^2 + d_{\mathrm{px}}^2/12}{N} + \frac{4\pi^{1/2}\sigma_{\mathrm{PSF}}^3\sigma_{\mathrm{BG}}^2}{d_{\mathrm{px}}N^2} \right)^{1/2} . \tag{6.2}$$

The first term expresses the quadrature sum of the errors associated with the Gaussian PSF distribution (with variance $\sigma_{\mathrm{PSF}}^2/N$ multiplied by a factor of 2 due to the multiplicative noise of the EMCCD) and a uniform distribution for each pixel (with variance $d_{\mathrm{px}}^2/12N$). The second term describes additional uncertainty associated with randomly occurring background counts, but our experimental background of $\sigma_{\mathrm{BG}} = 0.7$ photons is small enough that this background term contributes minimally to the overall error unless we collect $N < 10$ photons. Using experimental values of $\sigma_{\mathrm{PSF}} = 260$ nm and $d_{\mathrm{px}} = 193$ nm (as calculated in Sections 6.2 and 6.3) we use Equation (6.2) to estimate an uncertainty of 38 nm for locating an isolated atom, based on the typical value of $N = 100$ photons.

Experimentally, we can test the uncertainty with which we locate isolated atoms by tracking the measured position of a single pinned atom through several consecutive exposures. An atom might move between exposures due to a slow translation of the optical lattice or microscope objective, but in the absence of such a drift we still expect the measured position to vary from shot to shot due to the uncertainty $\Delta \bar{x}$. Using 9 consecutive exposures we find that the measured position of each atom varies with a standard deviation of 60 nm, larger than the theoretical estimate. Errors exceeding the theoretical form of Equation 6.2 are not uncommon, even in simulated images [95], and may simply indicate a
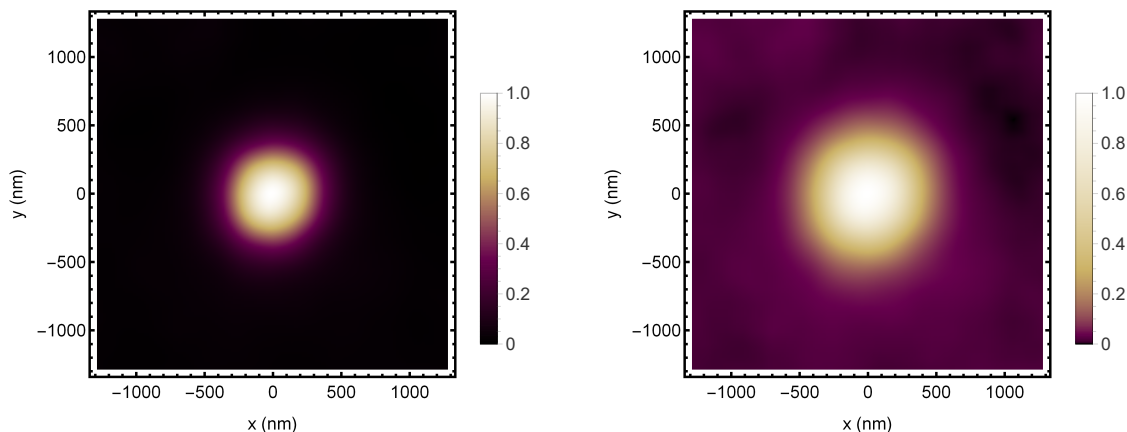
Figure 6.4: Effective PSF generated by combining the signal from > 400 isolated atoms. The colourmap on the left is used for typical fluorescence images as shown in Figure 6.1. The colour map on the right is the same as shown in Figure 3.9, adjusted to accentuate low intensity features such as the ring structure of an Airy distribution. There is no ring structure visible in our PSF, which may be due to imaging aberrations or high-NA corrections to the Airy function.

breakdown of the assumptions used in the derivation of Equation (6.2). Nevertheless, our experimental error of 60 nm is sufficiently small to accurately overlay several atoms and produce an estimate of the PSF. When we attempt to position the lattice grid on a given fluorescence image, the measurement error for a single atom's position can be further reduced by averaging the position measurements of several atoms.

Using the repeated position measurements of atoms in subsequent images, we can also test for slow drifts of the microscope or optical lattice during our long fluorescence exposures. Only slow drifts can be measured in this way, as fast mechanical vibration of the objective or jitter in the optical lattice would be averaged out over our long exposure, contributing only a blurring of sharp features in the image. Slow drifts on the time scale of seconds may occur due to thermal expansion and contraction of the microscope supports or optical lattice retro-reflection mirror, and this would likely appear as a steady change in the atom position over multiple exposures. Figure 6.3 shows the difference between an atom's measured position the first exposure and each subsequent exposure. Even though the errors with which we measure any single position are large, the mean shifts of all atoms can be used to track small displacements. In the horizontal direction of the pixel grid we observe a slight drift of $5(2)\,\mathrm{nm}\,s^{-1}$, while in the vertical direction the drift of $1(2)\,\mathrm{nm}\,s^{-1}$ is not significantly different from zero. These drifts are not large enough to produce visible differences between two consecutive 2.6 s exposures. This is useful when we compare the lattice occupancy in two subsequent images to analyze the imaging fidelity, as the same lattice grid can be used for both images.

6.2 Measured Point Spread Function

By locating many isolated atoms in a fluorescence image, we can estimate the PSF of the imaging system by overlapping the detected signals of the individual point sources. We select 400 $2.5\,\mu$m square regions each containing a single atom, and overlap them according to the atom positions determined from Gaussian fits. The resulting measured PSF is shown in Figure 6.4. Although each individual atom
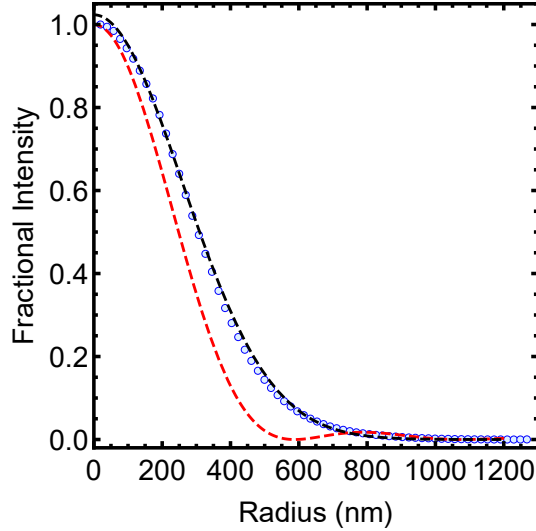
Figure 6.5: Comparison of our experimentally measured PSF (blue circles) with the theoretical Airy function for NA= 0.8 (red dashed line) and a best-fit Gaussian function with $\sigma = 260$ nm (black dashed line). A constant intensity offset has been subtracted from the experimental data.

image is detected with an effective pixel size of 200 nm, we can locate the atom centres more accurately and so we overlay the multiple atom images on a finer grid with pixel size 20 nm. This smooths out the effect of the finite pixel size on the observed PSF.

The measured PSF in Figure 6.4 can differ from the true PSF due to several effects. First, the individual atom images are only overlaid with an accuracy of $\sigma_{\text{cent}} \approx 60$ nm, and the combined random errors in atom centering will lead to broadening of the overlaid image. Similarly, although a single atom is a good approximation to a point source, the atoms we image are spatially delocalized within the confining potential of each lattice site. Since the atoms are cooled almost to the ground motional state of the harmonic potential, the probability distribution for each atom's position will be well approximated by a Gaussian distribution with $\sigma_{\text{SHO}} \approx 20$ nm. When we consider how either of these errors will effect the measured PSF, we can consider the effect of convolving the true PSF with additional Gaussian distributions of widths σ_{SHO} and σ_{cent}. If the true PSF can be approximated by a Gaussian of width σ_{PSF}, then the convolution is a Gaussian with width $\sigma_{\text{conv}} = \sqrt{\sigma_{\text{PSF}}^2 + \sigma_{\text{SHO}}^2 + \sigma_{\text{cent}}^2}$. Since σ_{cent} and σ_{SHO} are several times smaller than σ_{PSF}, they contribute only minimally to the overall width σ_{conv}.

It is also possible that the optical lattice is not stable during the fluorescence exposure. In Figure 6.3 we saw that slow lattice drifts (caused for instance by thermal drifts of the microscope supports or lattice optics) are minimal. However, mechanical vibrations of the microscope or lattice optics on timescales > 1 Hz will be averaged in our fluorescence images and will broaden the apparent size of each atom. In particular, the presence of the spinning chopper wheel in the imaging system may provide an unwanted source of mechanical vibration. If we assume that lattice and microscope instability are small, then the intensity distribution in Figure 6.4 should closely represent the true PSF of the imaging system.

In Figure 6.5 the radial average of the measured PSF is compared to an Airy distribution plotted assuming NA = 0.8. We see that the Airy distribution is significantly narrower, with a FWHM of 496 nm compared to 620 nm experimentally. The experimental data lacks the expected ring structure of the Airy function, with no local minimum of intensity visible in the radial plot. Instead we find that the experimental data is well fit by a Gaussian distribution with $\sigma = 260$ nm. Many factors may
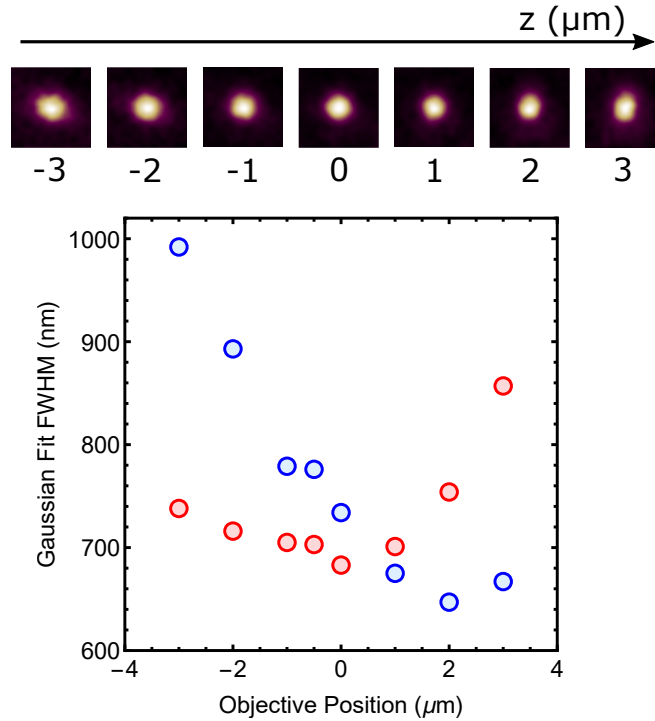
Figure 6.6: Experimental PSF as the microscope objective is scanned through the focal position. For this data a slight misalignment of lenses in the imaging path leads to astigmatism of the resulting PSF.

contribute to create an experimental PSF which deviates from an Airy function. The combined effects of uncertainty in the individual atom positions and finite wavefunction size are not large enough to account for the increase in FWHM. Lattice or microscope vibrations are not well characterized and may account for the measured PSF size. Additionally, an Airy function is only the ideal PSF for an imaging system which can be adequately described by plane wave propagation and the true theoretical PSF for our apparatus will differ due to apodization and vector diffraction effects which are relevant at large NA [96]. Above all, the theoretical PSF of the imaging system will only be attained once all imaging aberrations have been eliminated from the system.

In Figure 6.6 we show the effective PSF of the imaging system for different focal positions of the microscope objective. We see an ellipticity of the PSF which is inverted on either side of the focus, a clear sign of astigmatism in our imaging system. Data of this type can be used to eliminate alignment errors which contribute astigmatism. More generally, we can decompose our measured PSF into the basis of Zernicke polynomials in order to characterize the respective weight of all possible aberration sources. In our experiment we are particularly worried about any aberration which may be introduced to the large-NA imaging system by the birefringence of the sapphire window.

6.2.1 Comparison to Other Experiments

As shown in Figure 6.5 our microscope's measured PSF has 25% larger FWHM than would be expected from an Airy function based on a NA of 0.8 . The measured PSF likely overestimates the true PSF for several reasons outlined above, some of which may possibly be improved through optimizations of the experiment. When assessing whether we may be able to narrow our PSF with further experimental
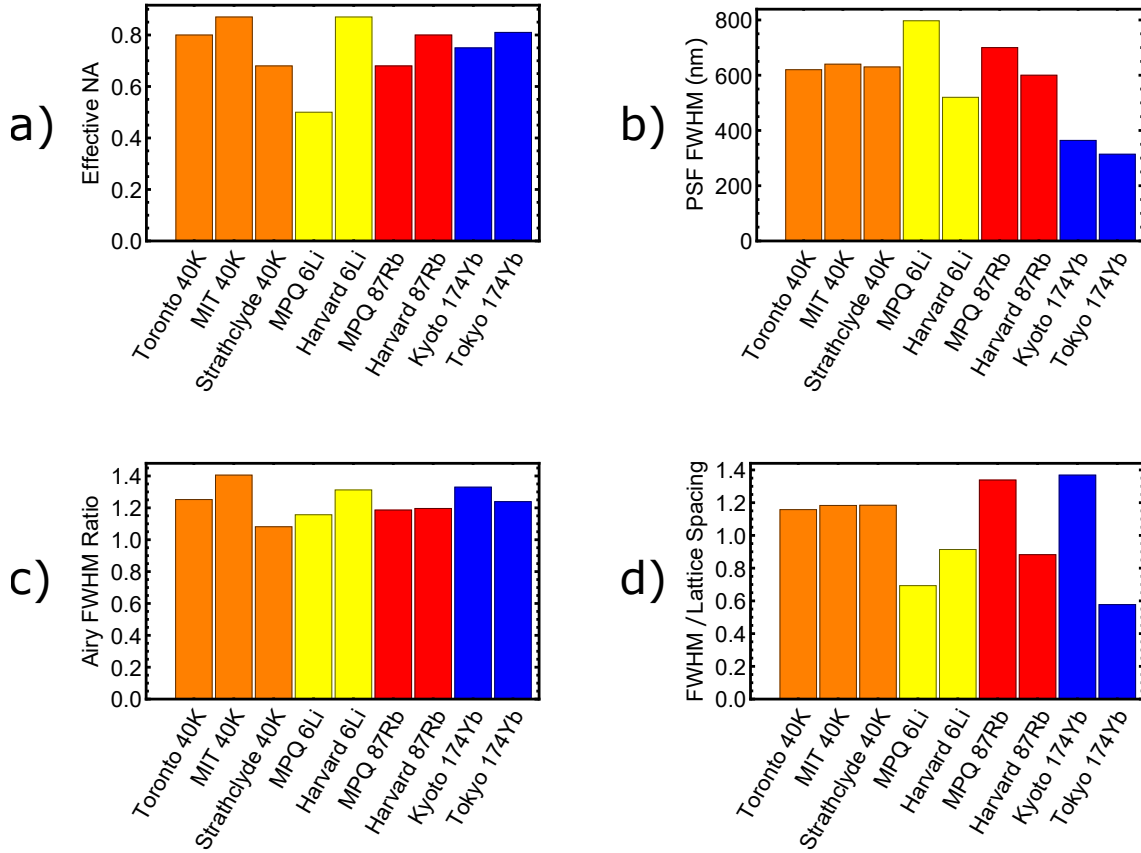
Figure 6.7: Comparison of several quantum gas microscopes. a) The NA achieved with each imaging system, including any enhancement due to solid-immersion effects. b) Reported FWHM of the point spread function of each imaging system. If two widths were reported (as in [62]), the average value is plotted. c) The ratio of the PSF FWHM to that of an Airy function based on the experimental values of λ and NA. All values are close to 1, indicating that all experiments achieve nearly diffraction-limited performance. d) The ratio of the PSF FWHM to the lattice spacing. Minimizing this ratio simplifies the process of image reconstruction. Values for these plots are taken from: Toronto [40]K [33], MIT [40]K [31], Strathclyde [40]K [32], MPQ [6]Li [35, 62], Harvard [6]Li [34], MPQ [87]Rb [20], Harvard [87]Rb [18], Kyoto [174]Yb [97], Tokyo [174]Yb [67]

improvements, it is interesting to compare our results to other quantum-gas microscope experiments around the world. Basic imaging parameters for our microscope and 8 others are compared in Figure 6.7. There is a large variation in imaging wavelengths due to the different atomic species in each experiment, and different experimental geometries are utilized to realize the large required NA for imaging. Regardless we see that all experiments realize a similar ratio of experimental FWHM to that of an Airy distribution calculated using the appropriate λ and NA. We thus consider the measured size of our experimental PSF to be similar to what is observed in other quantum gas microscope experiments.

The most important figure of merit in quantum gas microscope experiments is not the ratio of the experimental PSF size to the theoretical expectation, but rather the ratio of the experimental PSF size to the lattice constant which separates adjacent atoms. If this ratio is small, it will be easy to computationally reconstruct the lattice occupancy from an experimental image. If this ratio is large, such a reconstruction may instead by impossible. In Figure 6.7(d) we see that while all groups achieve
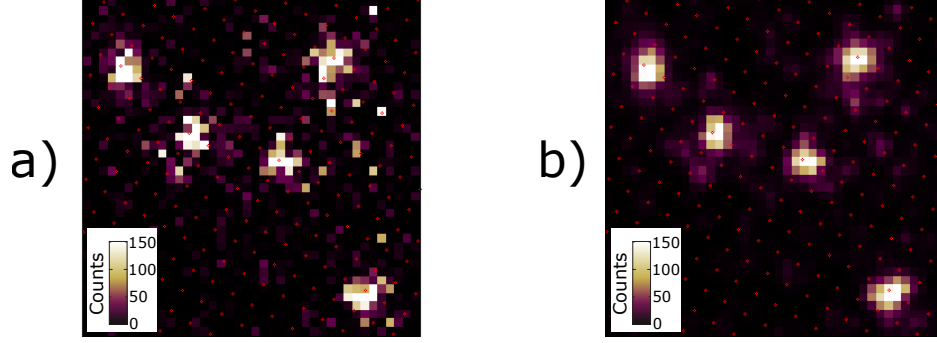
Figure 6.8: The lattice grid determined by measuring inter-atomic distances in a fluorescence image. The lattice grid is overlaid on top of a) a raw fluorescence image and b) a fluorescence image after applying a Gaussian filter with $\sigma = 1\,\mathrm{px}$.

a FWHM which is less than 140% of lattice constant, there are some strong variations between separate experiments. Some of these differences are due to smaller imaging wavelengths used with ^{6}Li and ^{174}Yb, and the larger lattice constants which can be used with a light atom like ^{6}Li while still allowing significant tunneling in the lattice. Our experiment appears quite similar to the other ^{40}K microscopes according to this metric, and so we hope to achieve similar fidelities when we attempt computational reconstruction of the lattice occupancy.

6.3   Reconstructing Lattice Occupation

6.3.1   Finding the Lattice Grid

If we can locate several isolated atoms in an image with only small errors, we can use the determined coordinates in the image to find the orientation of the optical lattice. In general we must determine 6 parameters to constrain the lattice grid: the lattice spacings a_x and a_y, the angles θ_x and θ_y between the lattice axes and the pixel coordinates, and the centre coordinates x_c and y_c for one site of the lattice. The lattice spacings and angles are considered to be constant from one image to another on a given day, and so only need to be determined once. Since the lattice may drift with respect to the microscope from one experimental cycle to the next, we must find the centre coordinates of the lattice for each individual run. Since we have seen that the lattice drift is minimal between consecutive exposures in the same experimental cycle, we only need to determine x_c and y_c once for a common set of two fluorescence images.

To determine the lattice spacings and angles, we fit the centre postions of many isolated atoms in a sparse image. With N measured atom locations, $N(N-1)/2$ different inter-atomic distances $(\Delta i, \Delta j)$ can be calculated in the pixel coordinate system $(\hat{i}, \hat{j})$. Since all atoms are located in lattice sites, these distances should have the form

$$\Delta i = na_x \cos(\theta_x) + ma_y \sin(\theta_y)\,, \tag{6.3}$$

$$\Delta j = -na_x \sin(\theta_x) + ma_y \cos(\theta_y)\,, \tag{6.4}$$

for some integers n and m. For each possible set of lattice parameters we calculate the squared error of the interatomic distances from the above form. This leads to best-fit lattice coordinates of $(a_x, a_y) =$
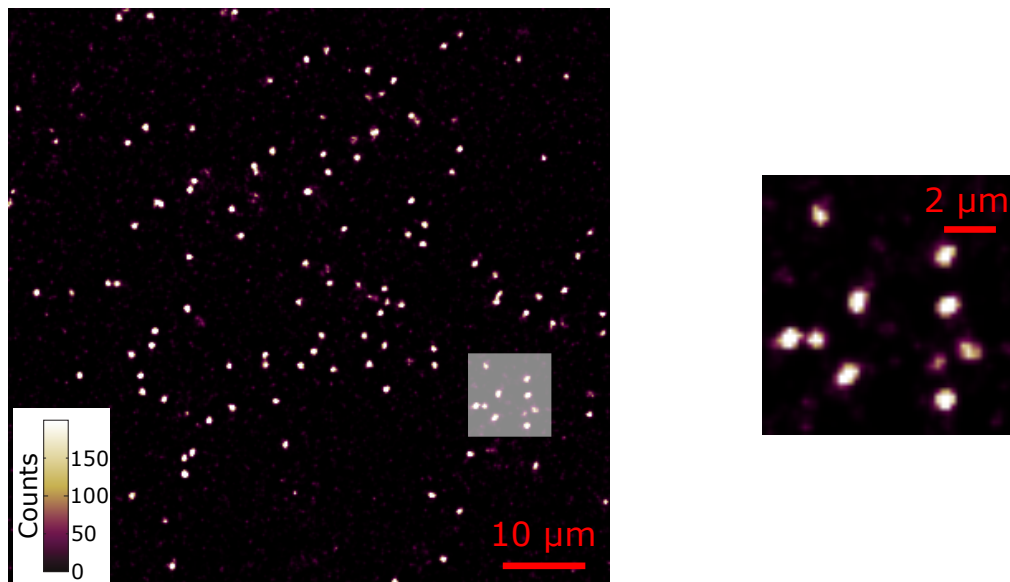
Figure 6.9: The same fluorescence image from Figure 6.1(b) after deconvolution with the measured PSF. The deconvolved image has been smoothed by a Gaussian filter with $\sigma = 1\,\text{px}$ to reduce the appearance of sharpened background noise. On the right a $10\,\mu\text{m}$ square region (indicated in gray in the main image) is enlarged to show detail.

$(2.72, 2.74)$ pixels and $(\theta_x, \theta_y) = (-30.5, -28.5)°$. The lattice spacings are precisely known to be $a_x = a_y = a_L = 527\,\text{nm}$ based on the wavelength of the lattice laser, given an effective pixel size of $193\,\text{nm}$ and constraining the magnification of our imaging system to be $(M_x, M_y) = (82.6, 83.2)$ close to the expected magnification of 80 based on the specified focal lengths.

6.3.2 Maximum Likelihood Fitting of Site Occupancy

After determining the lattice grid for a given fluorescence image, our next task is to determine the occupancy of each lattice site. If we are confident that a bright spot in the fluorescence image is a single atom, then we can use the measured centroid to determine which lattice site it occupies. If atoms are uniformly bright and sparsely distributed, this process is likely to accurately recreate the true lattice occupation. However, if atoms appear with varying brightness we must make additional decisions as to whether a bright spot in the image is even an atom at all. Additionally, the quantum states that we wish to simulate are likely to feature atoms neighboring one another rather than appearing well-separated. The bright regions in an image may then represent the combined signal of many atoms, so that finding the centre of the entire cluster is insufficient to understand the underlying lattice filling.

Rather than examining each bright spot in the image and determining where it belongs in the lattice, a more general approach examines each site of the known lattice grid and determines whether it appears to be filled by an atom. The pixels which are associated with each lattice site can be binned, to attempt to separate filled sites from empty sites based on the net fluorescence signal. Having a PSF width not much larger than the lattice spacing ensures that little fluorescence signal will spread into adjacent sites, and is necessary for the success of this site-binning approach. The size of the features in the fluorescence image can be reduced by deconvolving with the measured PSF to produce a sharpened image as shown in Figure 6.9. If we deconvolve the clear image of a point source with the PSF of the

imaging system, the result would be a delta function. However, due to the addition of finite pixel size and the presence of significant shot noise in our images, the reduction in feature size after deconvolution is limited. The amplitude of site-binned deconvolved images has been used to accurately distinguish filled and empty lattice sites in the MIT experiment, using a similar PSF width and lattice constant to our experiment [31]. However, we find that in our fluorescence images the empty and filled sites are not sufficiently discriminated by such a site-binned signal. The failure of this method for our images is likely due to the lower photon number and hence greater photon shot noise in our images. The increased shot noise creates additional variation in the site-binned signal and also limits the reduction of feature size in deconvolved images, with both effects leading to reduced separation between the binned signal for filled and empty sites.

An alternative approach to site binning is to perform a fit of a Gaussian function at each lattice site. The widths of the Gaussian function can be constrained by the known PSF, and the centre of the Gaussian is given by the calculated lattice grid. The amplitude of the function is the remaining free parameter, and the resulting best-fit value can be used to distinguish filled sites from empty ones. Simultaneously fitting the amplitudes for a cluster of sites takes into account the spread of fluorescence signal into neighboring sites, but increases the computational cost of the fit. Applying a threshold to the fitted Gaussian amplitude can be used as a criteria for site occupancy, and has been used to reconstruct the lattice occupation in two experiments with ^{6}Li [34, 35] We find that in our fluorescence images the fit amplitudes from such a procedure can distinguish filled and empty sites better than site-binned amplitudes, but the overall reconstruction of the lattice occupation still contains many errors. We expect that greater photon collection would lead to improved performance of the method, as the smooth shape of the fitted functions would be better reflected in the fluorescence images when shot noise is less significant.

To reconstruct the lattice occupancy in the presence of significant shot noise, we generate simulated fluorescence images for all possible lattice site occupancies using a Gaussian approximation to the known PSF, and find the configuration which most closely agrees with the observed image. 2^n different configurations are possible for n sites since we only consider that sites may contain 1 or 0 atoms, and this generates a prohibitively large number of combinations for the full visible lattice. To minimize the computational effort we work only with subsections of the full image containing 3×3 lattice sites. For each subsection we generate all 512 possible images, and use the configuration which best agrees with the observed image to determine the occupancy of the central site of the 3×3 cluster. Fitting only the central site minimizes any possible errors due to light appearing in the subimage from atoms adjacent to it. The task of finding the underlying distribution which is most likely to generate the observed data is naturally considered using the techniques of *maximum likelihood estimation*. Since the noise in our image should follow a Poissonian distribution, the corresponding maximum likelihood estimator for goodness of fit is [98]

$$\chi^2_{\text{MLE}} = 2\sum_{i=1}^{n}(f_i - y_i) - 2\sum_{i=1}^{n} y_i \ln(\frac{f_i}{y_i})\,, \tag{6.5}$$

where y_i are the observed counts at each pixel and f_i are the counts for the distribution under consideration. The best lattice site fillings for each 3×3 cluster in our algorithm is the one which minimizes χ^2_{MLE}.

Additional speedup of the image reconstruction algorithm can be attained by restricting the number of 3×3 clusters which must be considered. For this we make use of a deconvolved and site-binned version
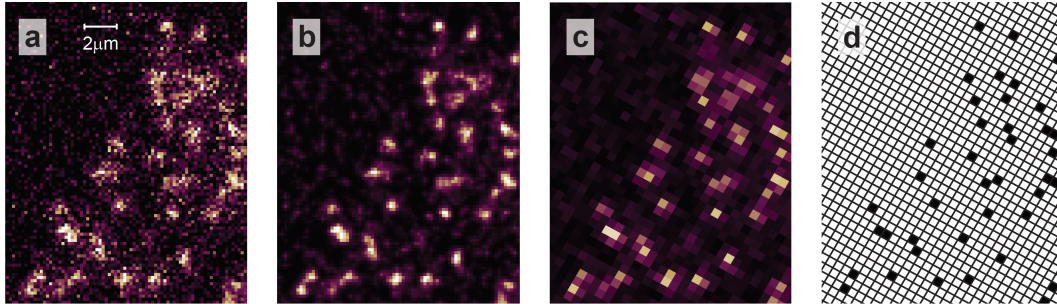
Figure 6.10: Steps for image reconstruction. a) Raw image. b) Sharpened image obtained by deconvolution with the PSF, shown here smoothed with a $1\,\text{px}$ Gaussian filter. c) Site-binned amplitudes of the sharpened image, used to preselect sites for the maximum likelihood fitting algorithm. d) Resulting lattice occupation predicted by the maximum likelihood algorithm.

of the image under consideration. Applying a threshold to the site-binned amplitudes does not allow us to identify filled sites with high fidelity as mentioned above, but can be used to isolate the subset of lattice sites that should be considered by the maximum likelihood algorithm. Pre-selecting possible filled lattice sites in this way dramatically speeds up the analysis of sparse images such as the ones used for all analysis in this thesis, but will not yield much improvement if colder and denser samples are under consideration in the future since almost all sites will exceed the threshold value.

The full sequence of steps for image reconstruction is shown in Figure 6.10. The lattice grid is located in the image based on the positions of several isolated atoms. A sharpened version of the image is generated via deconvolution with the PSF, and this sharpened image is then binned to produce a fluorescence signal for each lattice site. A threshold applied to the binned fluorescence signal identifies the potential atoms in an image for more efficient fitting. Around each site which exceeds the binned signal threshold, maximum likelihood estimation is used to determine the most probable 3×3 cluster of surrounding sites. Based on the most probably 3×3 cluster, the central site is identified as filled or empty. The result is a grid of lattice site occupancies as shown in Figure 6.10(d).

Since the presence of significant shot noise in our fluorescence images leads to large variations in the apparent brightness of atoms, we find that our maximum likelihood estimation algorithm can be improved if we consider for each 3×3 cluster not only the 2^9 possible site configurations corresponding to 0 or 1 atoms at each site, but additionally the possibility that sites are brighter or dimmer than the average brightness by 20%. This leads to 4 possibilities for each lattice site: empty, filled (dim), filled (average), or filled (bright). Thus 4^9 possible combinations now need to be considered at each cluster. This significantly increases the computational load of the algorithm, but allows us to more accurately identify all atoms in a fluorescence image.

6.4 Imaging Fidelity

Our ability to perform accurate quantum simulation relies on accurately determining the occupancy of the lattice in any given experimental run. There are many ways that errors may occur when we attempt to reconstruct a fluorescence image. If atoms cannot be perfectly confined to single lattice sites by the laser cooling mechanism, they may escape from the trap before enough photons have been collected for accurate identification. This loss process results in a site which was occupied in the true final state

appearing empty in digitization. Similarly an atom might move from one site of the lattice to another before scattering many photons, so that after digitization a site which should be empty appears occupied.

In order to probe the loss or movement of atoms during a fluorescence exposure, we take two consecutive fluorescence exposures of the same cloud, digitize both images, and compare the resulting distributions. We can then count the number of appearances N_a (for which a site is measured to be empty in the first image and occupied in the second), the number of disappearances N_d (for which a site is measured to be filled in the first image and empty in the second), and the number of occupied sites N_1 and N_2 in the first and second images. Based on these values we can determine the fraction of all N_1 atoms in the first image which were 'pinned' to a site, have been 'lost', or which have 'hopped' to a new site:

$$\text{Pinned} = \frac{N_1 - N_d}{N_1}\,, \tag{6.6}$$

$$\text{Hopped} = \frac{N_a}{N_1}\,, \tag{6.7}$$

$$\text{Lost} = \frac{N_d - N_a}{N_1}\,. \tag{6.8}$$

These three definitions sum to N_1/N_1, and we might expect them to behave like probabilities. However, it is possible for the (6.8) to become negative if the number of atoms appearing at new sites is greater than the number which disappeared from occupied sites. This can occur if we analyse only a small of the lattice, as an atom might 'hop in' to the region being considered. To minimize this effect, we use horizontal selection as describe in Section 4.4 to remove all atoms outside of the region to be digitized.

The hopping and loss processes described above are not the only possible errors which can occur during digitization. If few photons are collected from each atom, it can be difficult for the digitization algorithm to distinguish atoms from the background noise occurring due to stray light and thermal excitation of the EMCCD chip. Further, when only a small number of photons are detected and spread over several camera pixels the effects of photon shot noise can strongly influence the reconstruction algorithm. A large photon shot noise combined with the finite imaging resolution may result in atoms placed on incorrect lattice sites by the digitization procedure. These types of digitization errors cannot be measured directly for each fluorescence image since the true lattice occupation is not known *a priori*. However, with enough prior knowledge of the PSF and the noise properties of the imaging system it is possible to simulate fluorescence images for a given lattice occupation and use them to test the digitization algorithm [34]. In this way we estimate that with 100 photons collected from each atom only 1% of atoms will be incorrectly reconstructed for a sparsely filled lattice.

In the experiment we are sensitive to the combined effects of loss, hopping, and digitization errors. The optimal exposure time to use for fluorescence imaging is found by balancing these detrimental effects: hopping and loss are minimized when the shortest possible exposures are used, but long exposures result in the fewest reconstruction errors. To find an exposure time which balances these constraints we collect fluorescence images at a range of exposure times, comparing two consecutive exposures of the same lattice distribution at each exposure length. The results are shown in Figure 6.11, with a maximum pinned fraction for exposures between 2 and 5 s. In the gray region, the short exposure time restricts us to < 100 photons collected from each atom, and lattice reconstruction is poor. Errors in the digitization of the two exposures lead to an apparent hopping of atoms, and the calculated pinning is reduced. In the context of quantum simulation where high-fidelity state determination is crucial, exposure times > 2 s
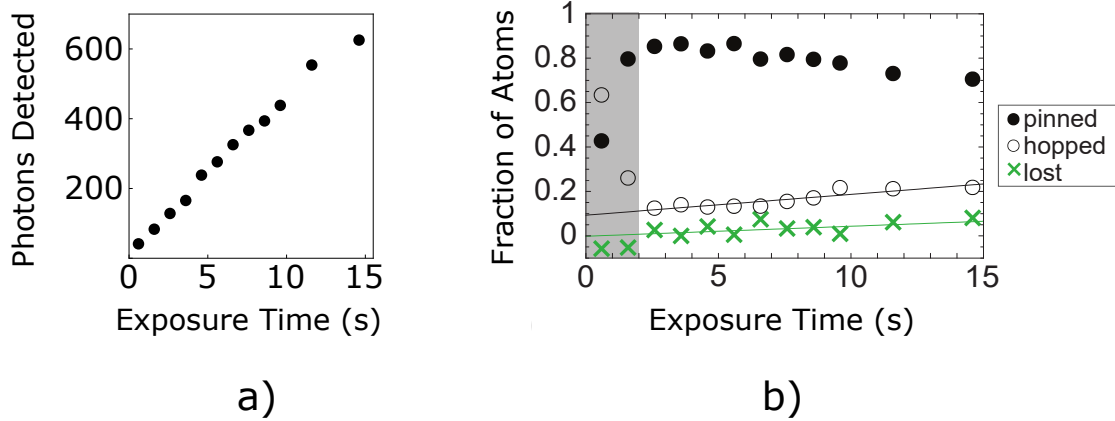
a)   b)

Figure 6.11: Finding the optimal exposure time for lattice reconstruction. a) Longer exposure times allow more photons to be collected from each atom, which makes it easier to perform lattice reconstruction. b) Taking two consecutive images of the same cloud and comparing the lattice reconstructions allows us to estimate the reconstruction fidelity. At short exposure times (gray shaded region) image reconstruction is poor due a small signal-to-noise ratio, leading to large differences in the two reconstructed lattice occupations. For long exposures reconstruction works well and any differences between the determined lattice occupations is dominated by atom movement and loss during image acquisition. We linearly fit the long-time ($> 4\,\text{s}$) loss and hopping data, finding that loss increases by $0.4(3)\%$ per second and hopping increases by $1.1(2)\%$ per second. The fit to the hopping data indicates that a finite amount of atoms hop even as the exposure time is reduced to zero

must therefore be used.

When very long exposures are taken, hundreds of photons can be collected from each atom and image reconstruction is very reliable. The hopping and loss calculated by comparing two images are then considered to be very accurate This makes long exposures ideal for optimization of cooling parameters. By considering the measured hopping and loss for exposures $> 4\,\text{s}$, we perform linear fits to estimate that loss increases by $0.4(3)\%$ per second and hopping increases by $1.1(2)\%$ per second. If we assume that all atoms are cooled and trapped identically, then we can infer from these rates that each atom escapes from its confining lattice well at a rate of $15\,\text{mHz}$ during EIT cooling. This rate is too large to be caused by tunneling in the ground band of the lattice, as t/h is many orders of magnitude smaller at lattice depths near $1000\,E_R$. The bandwidth and tunneling rate are larger if atoms are excited to higher bands of the lattice, but a tunneling rate around $10\,\text{mHz}$ is only achieved if an atom occupies the 12th excited band of the lattice well. Furthermore, tunneling would likely appear as atoms disappearing from one site and appearing in an adjacent site. In contrast, when examining sparse samples in which only a few atoms hop we can observe hopping events in which an atom appears > 10 sites away from its initial location. Thus this atom movement is unlikely to be caused by tunneling between sites in the lattice.

For a thermal distribution of particles trapped inside a potential well of depth U, the rate R at which particles high in the thermal tail of a Maxwell-Boltzmann distribution can escape the confining potential is given by the Arrhenius equation

$$R = \nu e^{-\beta U}\,, \tag{6.9}$$

where ν is the attempt frequency and is typically measured empirically for the estimation of chemical reaction rates. For temperatures small compared to the barrier height, and in the presence of minimal damping effects, the Kramers model [99] suggests an attempt frequency equal to the vibrational frequency

ν_0 of the confining well. In our $1000\,E_R$ lattice such an attempt frequency would thus be $> 250\,\mathrm{kHz}$ and escape rates of $R = 10\,\mathrm{mHz}$ are attained for a temperature of around $13\,\mu\mathrm{K}$. However, it has been suggested that since the atoms are in thermal equilibrium with the laser cooling beams and not the lattice potential, the attempt frequency will instead be determined by parameters of the cooling beams [17]. In optical molasses cooling of $^{133}\mathrm{Cs}$ and $^{87}\mathrm{Rb}$ the appropriate attempt frequency has been measured as $265\,\mathrm{Hz}$ [17] and $222\,\mathrm{Hz}$ [100], respectively. If the attempt frequency is similar for our laser cooling, then the resulting Arrhenius prediction for the hopping rate is far too small to explain the observed atom movement.

As the exposure time is reduced towards zero the fit to hopping data suggests that $11(2)\%$ of atoms still hop between one exposure and the next. This may be due to some set of lattice reconstruction errors which always occur regardless of the collected photon numbers. Such an error could occur if the true lattice grid pinning the atoms were imaged on our camera with a distortion, resulting in a poor fit with the assumed square grid. In this case certain atoms would reside at lattice sites which do not appear to correspond well to any of the assumed grid points, and the reconstruction algorithm might find several possible corresponding lattice occupancy distributions with similar goodness of fit. The nonzero intercept of the hopping data may also indicate hopping events which occur during the camera readout time rather than the exposure time used as the horizontal axis of Figure 6.11. For all exposure lengths, the same $400\,\mathrm{ms}$ readout time is used. The EIT cooling beams are kept on during this readout time, and so we can estimate the number of hopping events that occur using the measured hopping rate of $1.1(2)\%$ per second. Thus we expect $< 0.5\%$ of atoms to hop during the camera readout time, and this cannot explain the observed intercept. We thus consider reconstruction errors to be the most likely cause of the apparent offset in hopping.

Examining the linear increase in detected photons shown in Figure 6.11(a) we can also estimate the photon scattering rate during EIT cooling. Combining the photon losses due to our microscope geometry (18.5% efficient as discussed in Section 3.3.1), the 10% reflective beamsplitter used to couple in C_z, the 90% transmissive bandpass filter, and the 82% quantum efficiency of our camera, our overall photon detection efficiency is estimated to be 7%. A linear fit to the data in Figure 6.11(a) gives a photon detection rate of $38\,\mathrm{Hz}$, which corresponds to an estimated photon scattering rate of roughly $0.5\,\mathrm{kHz}$. This is a much lower scattering rate than seen in other recent quantum gas microscope experiments, and may indicate an uncharacterized loss of photons in the imaging system. Alternatively, if the EIT cooling mechanism can be alternately tuned for better cooling performance or higher scattering rates, our particular modulated vertical cooling scheme may force us to operate with parameters optimized for cooling and low scattering rates.

6.4.1   Tuning EIT Cooling Parameters

To optimize the cooling process during fluorescence imaging, we scan various EIT cooling parameters and monitor the hopping, loss, and pinning. Maximizing the pinning fraction allows us to use the maximum possible exposure time, so that more photons can be collected and digitization errors can be reduced. Scanning parameters such as the two-photon detuning δ or the power in C_{xy} does not result in large changes to the photon collection rate, so that the hopping and loss rates measured are not strongly intertwined with changes to the digitization fidelity.

Varying the two-photon detuning δ has a very strong effect on cooling performance, as shown in Figure 6.12. Since we precisely cancel the ambient magnetic field, we expect that the EIT window for
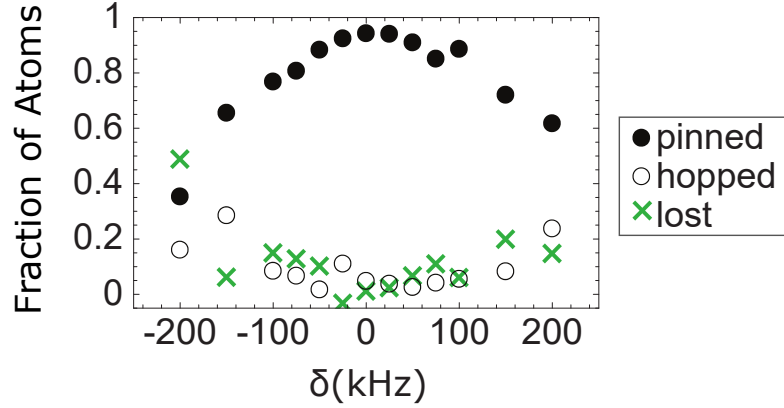
Figure 6.12: Dependance of hopping and loss on the two-photon reduning δ.
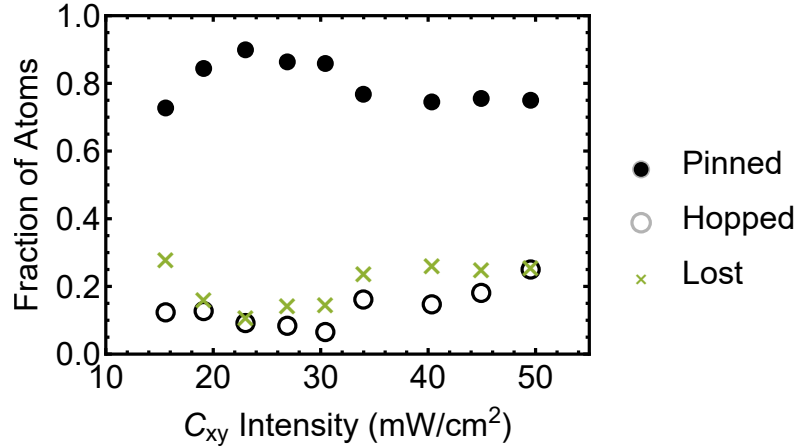
Figure 6.13: Dependance of hopping and loss on the intensity of C_{xy}. Intensities from 20-30 mW/cm^2 keep the most atoms pinned during two exposures, which is very similar to the optimal intensity range determined by observing survival of the cloud using absorption images (Figure 5.9).

all atoms occurs at $(\omega_p - \omega_c)/2\pi = 1285.79$ MHz which we take to be $\delta = 0$. Varying the frequency around this point spoils the EIT mechanism and also moves the two-photon transfer away from the red-motional sideband. Both of these consequences of nonzero δ would be expected to spoil the cooling mechanism. We find that pinning greater than 90% can only be achieved with $|\delta| < 50$ kHz. By tuning close to $\delta = -200$ kHz the two-photon transfer moves away from the red motional sideband towards the $\nu \to \nu$ carrier transition, so that few cooling transitions will result. In Figure 6.12 we see a significant increase in atom loss when the detuning is tuned towards the carrier transition, in agreement with this prediction.

Since our fluorescence images look directly at the centre of the lattice, and are not strongly affected by lattice inhomogeneity, we can perform a more direct optimization of the EIT cooling beam intensity than was mentioned in Section 5.3.1. By varying the intensity of the strongest EIT cooling beam C_{xy} we can check for an optimum intensity that minimizes atom loss and hopping. In Figure 6.13 we see that the optimal intensity is still found near 25 mW/cm^2, but the highest pinning fraction occurs at slightly lower intensities than the highest cloud survival of Figure 5.9. Even though all atoms that contribute to
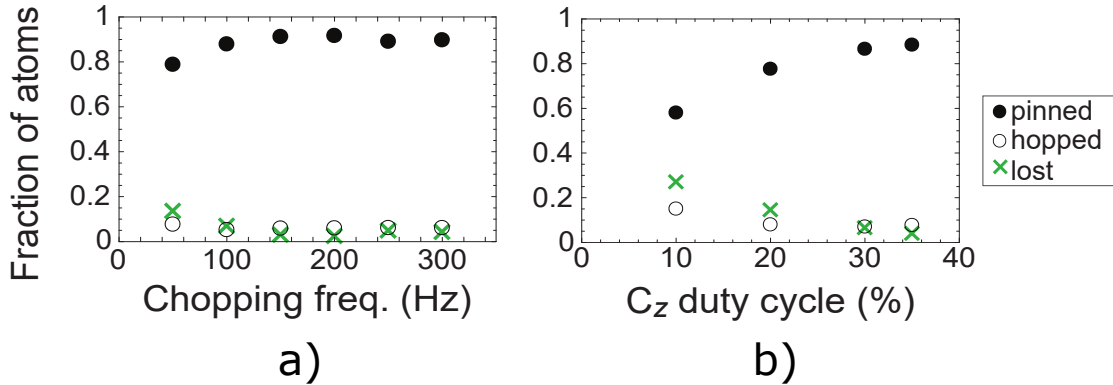
Figure 6.14: Assessing imaging performance with a spinning chopper wheel. a) Varying the frequency of the chopper wheel and the vertical cooling beam modulation, while keeping the duty cycle of the cooling beam fixed at 35%. The frequency describes the repetition rate with which the chopper blocks and unblocks the camera, not the rate of full revolutions of the wheel. Chopping frequencies greater than 300 Hz cannot be used since the vibration of the wheel leads to heating of atoms out of our optical dipole trap, hindering the production of cold gases. b) Varying the cooling duty cycle at a fixed chopping frequencies of 100 Hz.

this data should have very similar on-site trapping frequencies, we still observe a large range of intensities for which the cooling mechanism is somewhat effective. This is not surprising due to the large expected bandwidth of the EIT cooling mechanism, and a similar weak dependence on the size of the Stark shift has been reported for EIT cooling of a trapped ion [86].

6.4.2 Chopper Wheel

When studying the survival of the entire cloud in the lattice, we found that the vertical cooling beam C_z could be modulated without leading to additional loss of atoms. With single-atom images we can further test whether this modulated cooling scheme leads to hopping or loss of atoms. We cannot measure the hopping and loss fraction that would result without modulated cooling, since single atoms would not be visible above the increased amount of background light . Instead, we can gain some insight by adjusting the modulation timescale and duty cycle.

Figure 6.14 shows the results of changing the vertical beam modulation parameters. In Figure 6.14(a) the speed of the spinning chopper wheel is varied, which correspondingly adjusts the beam modulation speed. At all chopper speeds we adjust the timing of the vertical beam pulses to maintain a 35% duty cycle. At the lowest chopper frequencies of 50 Hz, corresponding to 7 ms on / 13 ms off cooling pulses, we observe an increase in hopping and loss. At all chopper speeds > 100 Hz we find similarly high fractions of atoms staying pinned to their lattice sites. Since there is little difference between the high frequency data, we conclude that a chopper speed of 150 Hz is sufficiently fast to keep all atoms strongly cooled in steady-state.

A 150 Hz chopper speed at 35% vertical duty cycle leave the vertical beam off for 4.3 ms-long periods. With the estimated photon scattering rates of 0.5 Hz, this means that only two photons are scattered on average during the time the vertical cooling is turned off. If high fidelity imaging is only possible with a few photons scattered during the chopping cycle, we would like to use the fastest possible chopping speed in order to permit similarly high photon scattering rates. Unfortunately, chopping rates > 300 Hz

led to strong vibrations on the optical table, interfering with other aspects of the experiment. Since the chopper wheel we use performs only 7 block/unblock cycles per revolution, an alternative path to increase the chopping frequency is by using a chopper wheel with an increased number of slots. This would allow us to modulate the cooling beam faster without actually spinning the motor at a higher speed. The downside to this approach is that the overall size of the opaque regions on the chopper wheel would be reduced, so that it becomes harder to properly block the background reflections which have a fixed beam size.

In Section 5.5 we measured the heating and cooling timescales for the vertical degree of freedom during modulated EIT cooling. We concluded that since the cooling rate with the vertical beam on is much larger than the heating rate with the vertical beam off, a duty cycle near 50% should still be able to keep atoms from heating out of the lattice on long timescales. In order to block the EMCCD camera from background light it is necessary that we use duty cycles $\leq 35\%$, a value set by the ratio of the chopper wheel blade size to the size of the beam which it blocks. Since we cannot increase the duty cycle without compromising the fluorescence imaging system, we instead reduce the duty cycle to observe whether atom pinning is reduced. As seen in Figure 6.14(b), reducing the duty cycle to 20% (2 ms vertical cooling pulses with 8 ms gaps) leads to significantly lower pinning. We do not observe a measurable reduction in the background light for fluorescence images when we decrease the length of the cooling pulses, so that there is no benefit to use any duty cycle less than the 35% maximum for size of chopper wheel that we use. Switching to a chopper wheel with a larger number of small blades would require the duty cycle to decrease if the same background reduction was desired, and from Figure 6.14(b) we see that this would lead to increased hopping. This constraint restricts us from using a different chopper wheel with smaller blades to increase the chopping frequency.

6.4.3 Photon Scattering Rate

The photon scattering rate of $\approx 0.5\,\mathrm{kHz}$ which we achieve during EIT cooling is an order of magnitude smaller than the $8\,\mathrm{kHz}$ scattering rate measuring in the Strathclyde experiment using EIT cooling with $^{40}\mathrm{K}$ [32] and the $5\,\mathrm{kHz}$ rate measuring in the MIT experiment using Raman sideband cooling with $^{40}\mathrm{K}$ [31]. These two experiments can thus collect sufficient photons for digitization using exposure times of 0.5-1 s, while we find we must expose for $>2\,\mathrm{s}$ as shown in Figure 6.11. Despite the longer exposure time that we use, our lower rate of hopping and loss leads to a similar overall likelihood for atoms to remain pinned during the image.

The higher scattering rate of the MIT experiment may be due to differences between EIT cooling and Raman sideband cooling. Each of the cooling mechanisms tends to pump atoms into a dark atomic state and cool them towards the motional ground state. In either cooling scheme, when an atom is cooled into the ground state photon scattering should be suppressed. Experimental imperfections such as lattice amplitude noise can cause atoms in the ground state to re-enter the cooling cycle, but these types of effects should be common to all experiments. However, when atoms are kept neat the motional ground state for a long fluorescence exposure, differences between the cooling mechanisms may lead to different steady-state photon scattering rates.

In Raman sideband cooling, transitions which preserve the atomic motional state are not forbidden since $\delta \neq 0$. These $\nu \to \nu$ transitions can couple atoms in the motional ground state out of the dark atomic state. In the absence of other experimental factors, the rate of $\nu \to \nu$ transitions may then determine the minimum scattering rate. In contrast, EIT cooling suppresses $\nu \to \nu$ transitions through
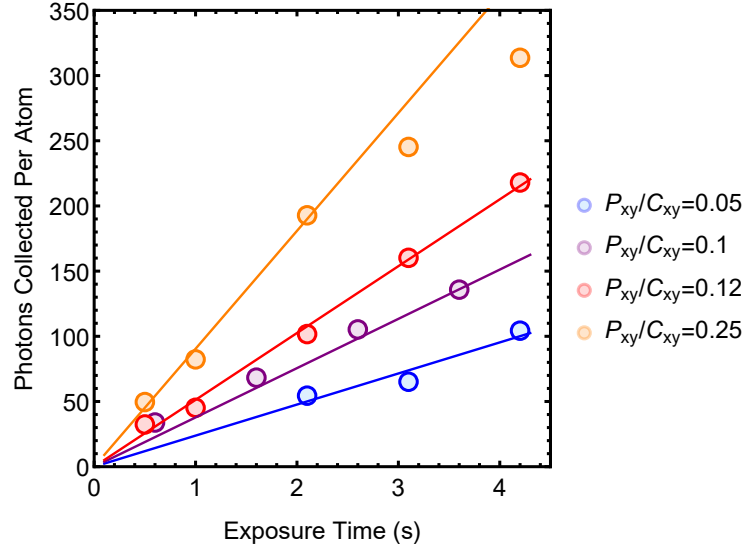
Figure 6.15: Photon collection rates with various strengths of Ω_P . The power in the beam P_{xy} is varied and the power in C_{xy} compensated slightly to keep the Stark shift similar for all trials. With more power in P_{xy} the photon collection rate increases, and linear fits to the data give photon collection rates of 24 Hz (blue), 38 Hz (purple), 51 Hz (red), and 90 Hz (orange). The implied atomic scattering rates range from 340 Hz (blue) to 1.3 kHz (orange). For the orange data, large hopping fractions are observed at exposure times >2 s, so only short exposure times are used for the linear fit.

quantum interference, so that atoms cooled into the dark motional ground state may stay there for longer times. In this way a long-lived EIT dark state in our experiment might explain the low observed photon scattering rates. The difference between the low scattering rate we observe and the higher scattering rate in the Strathclyde experiment might be explained by different decoherence rates of the EIT dark state. Since the EIT dark state is a superposition of different hyperfine ground states, it is not robust against any interaction with the atom which does not affect all internal states identically. In our experiment magnetic field fluctuations are the most likely effect leading to decoherence of the dark state. A higher photon scattering rate in the Strathclyde experiment may indicate a higher rate of decoherence, perhaps due to the additional pair of beams on the D2 transition used to cool the vertical degree of freedom. This does not necessarily suggest a benefit in our experiment, since a higher decoherence rate is not a disadvantage to a quantum gas microscope experiment unless it restricts the achievable imaging fidelity.

Although many of the EIT cooling parameters that we vary led only to changes in loss and hopping, varying the power in the beam P_{xy} is observed to be strongly correlated with the number of photons collected. Figure 6.15 shows the number of photons collected for various exposure times and P_{xy} powers. Linear fits give rates of photon collection ranging from 24 Hz to 90 Hz, which are equivalent to atomic scattering rates from 340 Hz to 1.3 kHz. The different scattering rates appear to scale directly with the power in P_{xy}, so that the scattering rate is given by $(P_{xy}/C_{xy}) \times 5.8(7)$ kHz.

The photon scattering occurring during constant EIT cooling can be broadly grouped into two categories: scattering processes involving the red or blue motional sidebands, and scattering processes which preserve the motional state of the atoms. The latter kind of scattering is suppressed by the EIT mechanism, but can still occur due to decoherence of the dark state. Both classes of scattering depend on the powers in the probe and coupling beams in the same way. In the perturbative limit $\Omega_P \ll \Omega_C$ it

has been shown that the rate of transitions $A_\pm$ to the motional sidebands scales with $\eta^2 \Omega_P^2$ [84] which is proportional to the probe beam power ($\Omega \propto \sqrt{P}$) as well as the squared Lamb-Dicke parameter as would be expected. More generally without assuming $\Omega_P \ll \Omega_C$ the rate of motional sideband transitions can be expressed as [85]

$$A_\pm \propto \eta^2 \frac{\Omega_P^2 \Omega_C^2}{\Omega_P^2 + \Omega_C^2} \; . \tag{6.10}$$

Thus the rate of red sideband scattering for atoms in excited motional states and the rate of blue sideband scattering for atoms in arbitrary motional states would both be expected to increase linearly with P_{xy} in the limit that $\Omega_P \ll \Omega_C$. This scaling behaviour of the motional sideband transitions also leads to the prediction that the fastest EIT cooling rates are achieved with $\Omega_P = \Omega_C$ [85].

To contrast the predicted scaling of the motional sideband transitions, we can also estimate the photon scattering rate for atoms which evolve out of the EIT dark state $|D\rangle$ through decoherence. Since in our experiment the dark state $|D\rangle \propto (\Omega_C |F = 9/2\rangle - \Omega_P |F = 7/2\rangle)$ is a superposition of two hyperfine states, we might expect the phase of the superposition to be altered by magnetic field fluctuations. These fluctuations would disrupt the relative phase of the two states, so that atoms evolve from the state $|D\rangle$ into a different superposition such as $|\alpha\rangle \propto (\Omega_C |F = 9/2\rangle + \Omega_P |F = 7/2\rangle)$ which is not one of the eigenstates of the Λ-system. The state $|\alpha\rangle$ is not subject to EIT, and would scatter photons at a rate

$$|\langle e|H|\alpha\rangle|^2 \propto \frac{\Omega_C^2 \Omega_P^2}{\Omega_C^2 + \Omega_P^2} \; . \tag{6.11}$$

This has the same scaling with beam powers as the motional sideband rates $A_\pm$. Thus we cannot distinguish which of the two processes leads to the increase in photon scattering rate when we increase P_{xy}/C_{xy}, but we see that such an increase is reasonable from a theoretical point of view. The apparent linear dependence of the scattering rate on P_{xy}/C_{xy} seen in Figure 6.15 is likely the limiting behaviour of the more general scaling appearing in Equations (6.10) and (6.11) when $\Omega_P^2 \ll \Omega_C^2$.

Since we can adjust the scattering rate during EIT cooling by changing the power ratio P_{xy}/C_{xy}, we might wish to tune the cooling for a higher scattering rate and then expose for a shorter time, so that fewer hopping and loss events occur during the fluorescence capture. Unfortunately, just as the scattering rate appears to be proportional to the amount of power in P_{xy}, so too do the rates of hopping and loss. This may indicate that hopping and loss events are linked to the anti-trapped excited states of the atom that may be populated during photon scattering events [31, 32]. For all of the data shown in Figure 6.15, we determine the pinned fraction of atoms and in Figure 6.16 the pinning is plotted versus the number of photons collected in each exposure. Regardless of the specific exposure time or scattering rate, we see that the optimal number of photons collected per atom is between 100 and 150. Exposing longer or increasing the scattering rate to collect more photons always results in a smaller pinning fraction. As seen before in Figure 6.11(b), collecting too few photons leads to a large number of digitization errors and a large corresponding drop in the calculated pinning fraction.

Although for any $P_{xy}/C_{xy} \leq 0.25$ we can find a suitable exposure length for high-fidelity imaging, we find that hopping and loss increase dramatically when we increase the scattering rate by using $P_{xy}/C_{xy} = 0.5$. The hopping and loss rates are so high that pinning fractions $> 50\%$ are impossible. A possible explanation for this effect is the fixed timescale imposed by the spinning chopper wheel in our experiment. As the scattering rate increases, we can decrease the exposure time in the experiment so that the total number of photons collected is fixed. However, since we cannot arbitrarily increase the
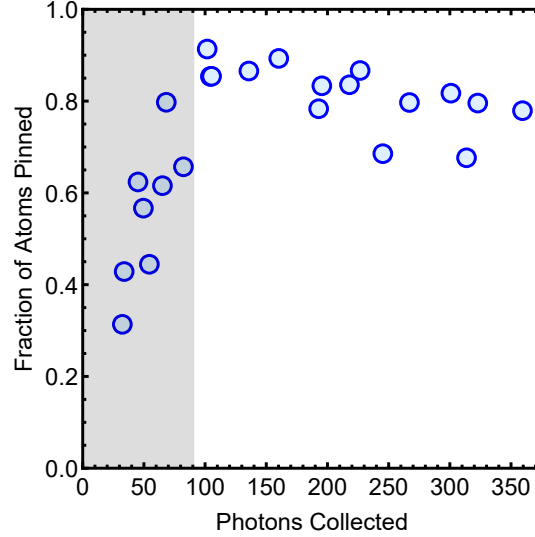
Figure 6.16: Pinning measured using all of the exposure times and Ω_P strengths from Figure 6.15. The highest pinned fraction occurs when 100 to 150 photons are collected, regardless of the length of exposure or strength of Ω_P. Collecting more than 150 photons leads to reduced pinning as atoms are able to hop between exposures. Conversely, collecting fewer than 100 photons (gray shaded region) leads to errors in lattice reconstruction so that pinning cannot be accurately measured.

speed of the chopper wheel, higher scattering rates will result in greater heating of the vertical degree of freedom during each off period of the vertical cooling beam. We have seen in Figure 6.14 that a chopper speed of 150 Hz is sufficient when the photon scattering rate is around 0.5 kHz, but reducing the speed by a factor of 3 leads to significant increase in hopping and loss. Similarly, we should expect that keeping the chopper speed fixed and increasing the scattering rate might also lead to hopping and loss as the periodic heating of the vertical degree of freedom becomes too large.

6.4.4   Reconstruction for Higher Densities

With all laser cooling parameters optimized, we have determined that 94(2)% of atoms remain pinned in a typical 2.6 s fluorescence exposure. In addition, for a sparsely filled lattice we test our digitization algorithm using simulated images and determine that only 1% of atoms are incorrectly identified. We can therefore conclude that the overall fidelity with which we determine the density distribution in the lattice exceeds 93(2)%. Although sparse distributions of atoms are ideal for calibration of the imaging process, we wish to ensure that the imaging fidelity remains high even as we approach the unity filling of the lattice which would occur in a spin-polarized band-insulator or a Mott-insulator of two spin states. Unfortunately, the process of determining the occupancy of a particular lattice site becomes more difficult when the number of occupied neighboring sites is increased.

Using simulated images we have estimated that with our current photon collection, the fidelity with which we assign site occupancies in a dense lattice can be as low as 50%. The simulations show that increasing the number of detected photons leads to a recovery of high reconstruction fidelity for dense lattice fillings, and to achieve $> 90\%$ fidelity we must increase the number of collected photons to 400. As we have seen in Section 6.4.3 increasing the photon scattering rate by tuning the laser cooling parameters can lead to an increase in atom hopping, so that increasing photon collection in this manner

is not optimal. Keeping the laser cooling parameters unchanged, we can increase the photon collection simply by using a longer exposure time. An increase in the exposure time to 7.2 s would result in sufficient photon collection. Based on our measured hopping and loss rates, this would increase the fraction of atoms hopping or escaping the lattice during the exposure to around 10%, so that an overall fidelity greater than 80% could still be attained for the density measurement.

If we wish to improve the photon collection in our microscope without increasing the exposure time, it may be possible to adjust the photon scattering rate by using a different laser cooling technique. Another experiment using EIT cooling with ^{40}K has implemented a set of Raman beams to couple different motional states of the confining lattice, so that no EIT cooling beam must propagate along the microscope optical axis [32]. Implementing such a technique would allow us to remove the chopper wheel, doubling the rate of photon collection. Alternatively, an experiment using Raman sideband cooling for ^{40}K has reported photon scattering rates of 5 kHz [31], an order of magnitude higher than the scattering rate obtained in our experiment with EIT cooling. Thus a fruitful path to increasing the fidelity of fluorescence imaging in our experiment is the addition of Raman beams for sideband cooling or motional state coupling.

6.5 Conclusions

We have found that EIT cooling can be used for single-atom sensitive fluorescence imaging, when a spinning chopper wheel is used to significantly reduce the background light level. Due to the presence of an atomic dark state, the scattering rate of photons in steady state during EIT cooling is strongly reduced from the level expected from the power and detuning of the cooling beams. We find that scattering rates are typically ≈ 1 kHz, so that several-second exposures are required to collect hundreds of photons at our experimental collection efficiency. Hopping and loss events occur during the cooling and imaging process, and the number of these events appears to scale with the number of photons scattered.

We have optimized the cooling process by minimizing the occurrence of hopping and loss, and we find that in optimized conditions we can collect > 100 photons per atom while keeping 94(2)% of atoms pinned to their lattice sites. This allows us to accurately reconstruct the lattice occupation for our dilute samples. Image reconstruction becomes more difficult when the density of atoms in the lattice is increased, and a larger number of reconstruction errors would be expected if the density were increased to half-filling of the lattice. From simulations we have calculated that our reconstruction fidelity would fall to a value near 50%, greatly hindering our ability to read out the result of a simulation. We also estimate that adjusting the laser cooling and imaging parameters to achieve 300% more collected photons would be sufficient to ensure high fidelity state reconstruction with our algorithm. Thus, although our fluorescence imaging sequence has single-atom sensitivity, we should still seek to improve this experimental step before quantum simulation of arbitrary many-body states will be possible.

Chapter 7

Conclusions and Outlook

We have successfully constructed a quantum gas microscope experiment for the study of fermionic atoms in a lattice potential. We prepare ultracold quantum gases of ^{40}K atoms and load them into a three-dimensional cubic lattice with a periodicity of $527\,\mathrm{nm}$. The experimental apparatus is designed for quantum simulation of the Hubbard model with full control over all relevant parameters. The correlated many-particle states which are produced in such a simulation are driven by Fermi statistics, as well as a competition between kinetic energy and interparticle interaction. These types of highly correlated many-body states are of great theoretical interest, and our high resolution single-atom sensitive fluorescence imaging allows us to read out the density of these states at a site by site level. High resolution density profiles can be used to probe the equation of state of a lattice gas through a local pressure [64], can detect the suppression of number fluctuations in insulating phases [35, 59, 74] and can be used to detect spin correlations related to magnetic orderings [60–62, 101].

We realize quantum gas microscopy using EIT cooling to keep the atoms in our quantum gas cooled near the motional ground state of the confining lattice while they scatter photons at a rate of around $500\,\mathrm{Hz}$. We collect the scattered photons with a high-NA microscope to produce high resolution fluorescence images. After collecting roughly 100 photons from each atom, we use a maximum likelihood algorithm to reconstruct the occupation of each lattice site. Comparing multiple exposures of the same cloud we investigate possible atomic motion during out fluorescence images, and we find that in optimal cooling conditions fewer than 2% of atoms move per second. In a typical fluorescence exposure we see that 94(2)% of atoms remaining pinned to their lattice sites by the EIT cooling mechanism. This is comparable to the atomic pinning observed in other quantum gas microscopes developed in parallel with our experiment.

When we move towards simulation of the Hubbard model, we have estimated based on simulated images that the accuracy of our image reconstruction will decrease. Increased photon collection will be necessary for high fidelity lattice reconstruction of a half-filled lattice, and we have explored the possibilities of increasing the photon scattering rate of our current cooling method. We have seen evidence that the spinning chopper wheel that we use for background suppression may limit the photon scattering rates that are achievable, and so a change in our imaging paradigm may be required. During the preparation of this thesis, alternative cooling methods have been implemented in the lab which allow for higher numbers of photons to be collected without the use of a chopper wheel. With these improvements, this experiment is now fully capable of performing high-fidelity readout of a quantum

simulation of the Hubbard model.

To explore the most interesting phases of the Hubbard model, it will be necessary to prepare deeply degenerate quantum gases at lower temperatures than we currently reach using evaporative cooling. To optimize evaporative cooling it may be useful to tune the interparticle interaction using a Feshbach resonance [14]. In this work, when implementing spectroscopic plane selection, we have made many improvements to the magnetic field stability of the experiment, which will allow for precise control over the atomic interactions at the Feshbach resonance. An second useful tool for the improvement of evaporative cooling is the use of a tightly confining dimple trap, which can allow greater control over the final number and temperature of ultracold clouds. We have constructed a broadband 850 nm laser system to be used for this purpose.

The presence of a large-NA microscope objective in our experiment also invites the possibility of projecting optical potentials onto the atoms with high spatial resolution. Projected potentials have been used to divide a lattice gas into subsystems for measurements of entanglement entropy [29], number-resolved detection of lattice site occupancy without the common parity projection [102], and the preparation of localized impurities in a lattice gas for the study of quantum walks [28] or spin impurity dynamics [30]. Sophisticated optical potentials are also at the heart of proposals for the production of low entropy states of a lattice gas [103].

Appendices

Appendix A

Imaging on the $4S_{1/2} \to 5P_{3/2}$ Transition

When building a quantum-gas microscope, a lot of work goes towards reducing the ratio of the imaging resolution to the lattice spacing. This ratio can easily be made quite small by using an optical lattice with a large periodicity [17], but the tunnel coupling decreases exponentially with increasing lattice period so that quantum simulation of many-body Hamiltonians is no longer possible on realistic timescales. A notable exception can be made for quantum gas microscopes which use Lithium, where the very small mass allows strong tunnel coupling even with lattice spacings larger than 1 μm [35]. For a heavier atom like ^{40}K , effort must instead be directed towards improving the imaging resolution, as discussed for our experiment in Section 3. When powerful microscopes of NA ≈ 0.6 are used, the imaging resolution can be reduced to the size of the wavelength λ of the light collected. To further improve the imaging resolution, it is attractive to consider imaging using a different atomic transition for which λ is small. For instance, recent demonstrations of high-resolution imaging of Ytterbium [67,97] use light scattered on a transition with $\lambda = 399$ nm, achieving a PSF with FWHM near 300 nm.

To try to achieve a small imaging PSF with ^{40}K , we built our experiment to allow imaging on the $4S_{1/2} \to 5P_{3/2}$ transition at $\lambda = 404.5$ nm. This $J = 3/2$ transition, the hyperfine levels of which are shown in Figure A.2 has many similarities to the $4S_{1/2} \to 4P_{3/2}$ D2 line typically used to collect ^{40}K in a MOT and perform optical molasses. Our group has demonstrated that a MOT can be formed with this $n \to n+1$ transition [104], and similar results have been reported for ^{6}Li [105]. Two ^{87}Rb quantum gas microscopes have realized high-fidelity imaging using the D2 transition, generating fluorescence and cooling atoms using an optical molasses. We attempted to realise similar optical molasses cooling of ^{40}K at 404.5 nm to create a high-resolution quantum-gas microscope.

On the D2 line, optical molasses cooling of ^{40}K is complicated by the small and inverted hyperfine structure. The $|F' = 11/2\rangle$ excited state typically used for Doppler cooling is also the lowest energy state in the hyperfine structure, so that cooling beams which are red-detuned from the $|F' = 11/2\rangle$ state are even farther red-detuned from all other excited states in the D2 line. Although in this configuration one naively expects the optical molasses mechanism to be hindered less by interference from the other excited states, calculations have shown that the capture velocity of molasses at large detunings is actually worse than for a non-inverted hyperfine structure [106]. However, in the context of cooling tightly bound atoms, the requirements for effective cooling are quite different and limitations such as a weak capture
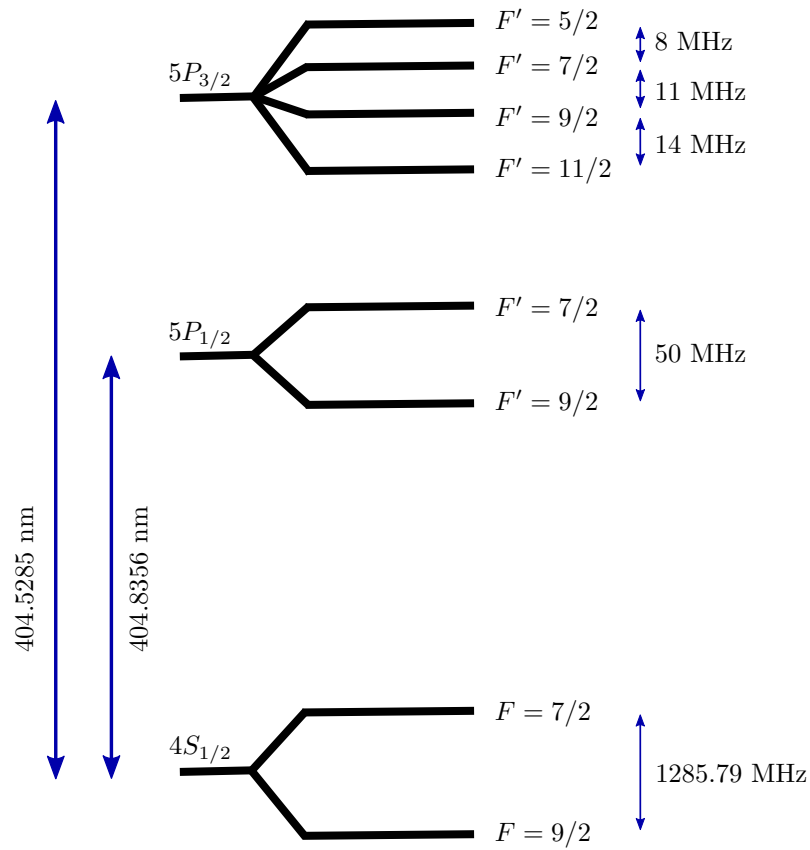
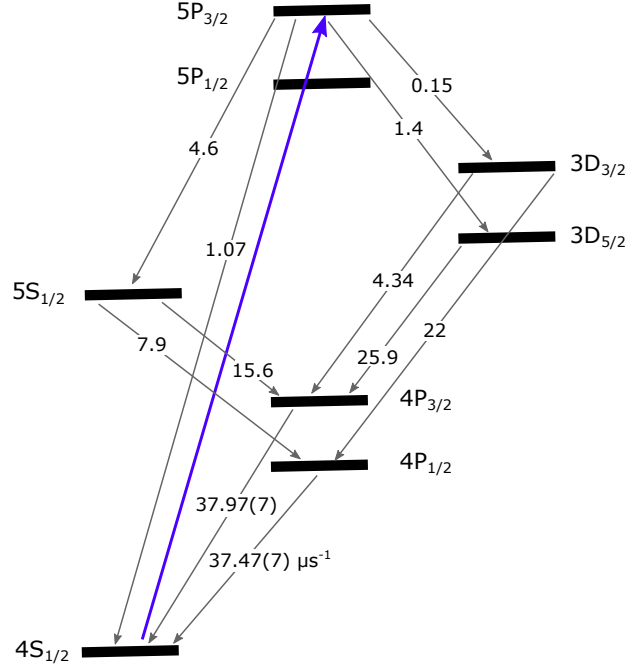Figure A.1: The 4S to 5P transitions of ^{40}K .

Figure A.2: Possible decay pathways from the $5P_{3/2}$ excited state. Transition probabilities are given as $A = 1/\tau$ in units of $\mu\mathrm{s}^{-1}$. Values are taken from [107] and [108]

velocity may not pose any trouble.

In our demonstration of a MOT at 404.5 nm, we did not observe a temperature reduction when attempting an optical molasses on the $5P_{3/2}$ transition. This does not preclude effective cooling of atoms bound in an optical lattice, and so further experiments were needed to assess the outlook for operating a quantum-gas microscope with 404.5 nm light. When investigating cooling with this more highly excited state, we have paid close attention to the possible spontaneous emission pathways, the possibility of photoionization, and the influence of the lattice wavelength on the trapping of atoms in the $5P_{3/2}$ excited state. Details of these are summarized below.

A.1 Radiative Cascade

The $5P_{3/2}$ transition is much weaker than the D2 transition in ^{40}K , with a transition probability of $2\pi \times 185$ kHz. An atom excited to the $5P_{3/2}$ state has several possible three-photon decay paths, as shown in Figure A.2, leading to a broadening of the state to $\Gamma = 2\pi \times 1.19$ MHz. The overall linewidth is smaller than on the D2 line, leading to a lower Doppler temperature [104]. This linewidth is about ten times smaller than the hyperfine splitting Δ_{HF} in the $5P_{3/2}$ excited state, leading to a better resolved hyperfine structure than on the D2 line. In the context of sub-Doppler cooling, a well resolved hyperfine structure $\Gamma \ll \Delta_{HF}$ is desirable as it allows the optical pumping rates and light shifts of individual hyperfine states to be controlled [106].

Since Sisyphus cooling relies on atoms being pumped into specific hyperfine ground states, the presence of the radiative cascade would be expected to reduce its efficiency. Cooling occurs in a polarization gradient when the presence of a given light polarization tends to optically pump atoms into states with
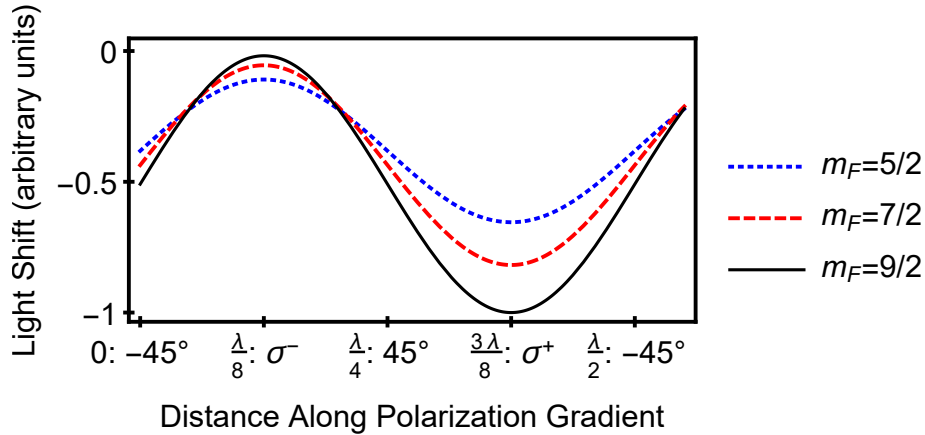
Figure A.3: Light shifts for three states in the $F = 9/2$ ground state, in the presence of a polarization gradient. In the presence of σ^+-polarized light atoms will tend to be optically pumped to higher m_F states, which also lowers their overall energy as shown. Similarly, in the presence of σ^--polarized light atoms are pumped into states with lower m_F, which is also seen to lower their overall energy although for the states plotted here the difference is small.

lower energy due to the light shift. This leads to the Sisyphus cooling condition that an atom moving with respect to the polarization gradient tends to travel 'uphill' into regions with a higher energy (converting kinetic energy into potential energy) before 'rolling to the bottom of the hill' by being pumped into a lower energy state (reducing its potential energy through photon emission). On $J = 3/2$ transitions this condition is satisfied when beams are red-detuned, so that absorption of $sigma^+$-polarized light tends to pump atoms to states with larger m_F and correspondingly lower energy (and vice-versa for $sigma^-$-polarization). Light shifts in a polarization gradient for a few relevant energy levels of the $F = 9/2$ ground state are schematically plotted in Figure A.3, to help illustrate the effect.

When attempting to realise optical molasses on the $5P_{3/2}$ line, the various decay channels through which an atom can return to the ground state can reduce the overall cooling efficiency given by the Sisyphus mechanism. Consider an atom initially in the $|9/2, 7/2\rangle$ ground state excited in the presence of σ^+-polarized light to the $|11/2, 9/2\rangle$ state of $5P_{3/2}$. This atom has a 16% chance to emit a 404.5 nm photon and decay to $|9/2, 9/2\rangle$, which would reduce the potential energy as required for Sisyphus cooling. However, it is mre likely that the atom will decay first to the $5S_{1/2}$ state, returning to the ground state via one of $4P_{1/2}$ or $4P_{3/2}$. These decay channels may cause the atom to end up in the $|9/2, 5/2\rangle$ ground state with *lower* m_F and *higher* energy, and so can cause heating rather than cooling to occur due to the Sisyphus mechanism.

Overall there are still many favourable decay channels, and so Sisyphus cooling would not necessarily be prohibited on the $5P_{3/2}$ line, but the lowest attainable temperatures may be expected to be higher than for a molasses on the $D2$ line. Even if this transition could be used for effective laser cooling of atoms trapped in an optical lattice, the radiative decay paths also present challenges to any attempt to produce fluorescence images using the light scattered during cooling. The benefit to working on the $5P_{3/2}$ transition is the possibility to collect photons at 404.5 nm for improved imaging resolution. However, as discussed above, there is a high probability that an atom excited to the $5P_{3/2}$ state will not emit at 404.5 nm, instead following one of the other decay paths in Figure A.2 and emitting photons of a larger wavelength. Thus this transition is not only likely to be less efficient in terms of cooling, it is certain to
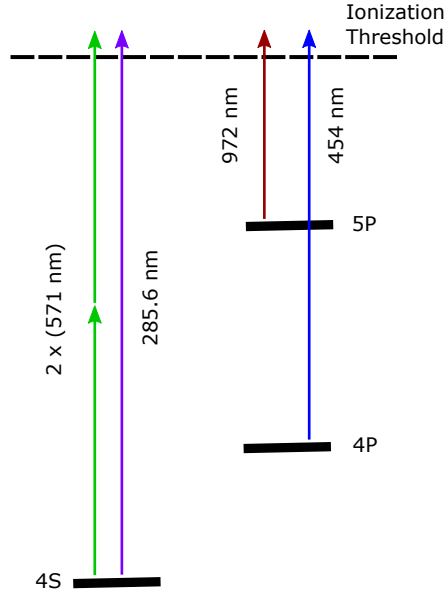
Figure A.4: Some possible ionization processes for a ^{40}K atom. Indicated wavelengths are the longest wavelengths which can ionize an atom in a particular state. Both one-photon and two-photon ionization are shown for the $4S$ ground state. Hyperfine structure has been omitted.

be less efficient at producing the desired fluorescence signal. When attempting to resolve the locations of atoms in a closely spaced optical lattice, an imaging system with a smaller PSF will require fewer detected photons in order to reconstruct the distribution of atoms in lattice sites. Thus, the reduced fluorescence brightness expected on the $5P_{3/2}$ transition is not necessarily a hindrance to quantum-gas microscopy.

A.2 Photoionization

In typical experiments with ultracold ^{40}K , laser wavelengths used for manipulation and trapping will be in the near-infrared range, and ionization is not energetically possible except through rare multi-photon processes. In contrast, when working working with 404.5 nm light, several ionization processes become energetically favourable as shown in Figure A.4. It is then important to consider the ionization rates of all processes to determine whether they will impose any bound on high-fidelity fluorescence detection.

The rate γ_1 for one-photon ionization of a ground state atom is given by

$$\gamma_{1,g} = \sigma_g^{(1)} \frac{I\lambda}{hc} \,, \tag{A.1}$$

where I is the light intensity, λ is the wavelength, and $\sigma^{(1)g}$ is the one-photon ionization cross section of the ground state. $\sigma^{(1)g}$ may be wavelength-dependent, scaling with the excess energy of the photon beyond the ionization energy, but this is neglected here for simplicity. The quantity $I\lambda/hc$ gives the photon flux incident on the atom. In the case of ground state ^{40}K atoms exposed to 404.5 nm light, the photon energy is not high enough to cause direct ionization, so ionization is only possible through the

absorption of two photons with a rate given by

$$\gamma_{2,g} = \sigma_g^{(2)} \left(\frac{I\lambda}{hc} \right)^2 , \tag{A.2}$$

where $\sigma^{(2)g}$ is the two-photon ionization cross section of the ground state. For ^{87}Rb atoms in their ground state, $\sigma^{(2)g} \approx 10^{-48}$ cm^4 s [109], so that $\gamma_{2,g} \ll 1$ Hz except in the case of very large intensities ($I > 100$ kW cm^{-2}). Two-photon ionization would similarly be expected to be negligible for ^{40}K at low intensities, so that the dominant ionization process occurring during 404.5 nm laser cooling will be single-photon ionization of atoms in an excited electronic state. In a simple two-level picture, this will occur with a rate proportional to the excited state fraction ρ_{ee}:

$$\gamma_{1,e} = \sigma_e^{(1)} \rho_{ee} \frac{I\lambda}{hc} . \tag{A.3}$$

The cross section $\sigma_e^{(1)}$ for ionization from the excited state may differ from $\sigma_g^{(1)}$, and in ^{87}Rb it is larger by about one order of magnitude. We have measured trap loss processes in a 404.5 nm MOT and estimate that $\sigma_e^{(1)} < 8 \times 10^{-18}$ cm^2 for ionization from the $5P_{3/2}$ excited state.

If we wish to image the fluorescence generated on the $5P_{3/2}$ transition, we would need to generate scattering of 404.5 nm photons at a rate of at least 1 kHz in order to collect a sufficient amount of light during a < 5 s exposure. This means achieving $\rho_{ee} > 5 \times 10^{-3}$ for the $5P_{3/2}$ state. If effective optical molasses cooling requires several beams with intensities of approximately 30 mW cm^{-2} (as used for the 404.5 nm MOT) then we can estimate an ionization rate $\gamma_{1,e} = 10$ mHz during fluorescence image capture. During a 1 s long exposure, we would then expect only 1% of atoms to be ionized, so that the atomic distribution could still be resolved with high fidelity. Concersely, if we attempted to perform an optical molasses using a detuning large compared to the excited state hyperfine structure, as recently shown to provide effective sub-Doppler cooling for ^{6}Li [110], intensities as high as 5 W/cm^2 might be required which would lead to unacceptably high ionization rates on the order of 1 Hz.

An alternative route to imaging at 404.5 nm would be laser cooling atoms with beams on the D1 or D2 transitions while adding a weak 404.5 nm beam to scatter photons. In this case, atoms would occupy both the $4P$ and $5P$ excited states. When in the $5P$ state, an atom can be ionized by either a 404.5 nm photon or a 767 nm photon. In the $4P$ state, only 404.5 nm photons will cause ionization. To collect enough photons at 404.5 nm, an excited state fraction greater than 5×10^{-3} will still be required, placing a bound on the allowable intensity on the D1 or D2 lines. Cooling methods such as optical molasses or EIT cooling which operate with near detuned beams and low intensities could still function in this case without excessive losses. On the other hand, Raman sideband cooling techniques used for quantum gas microscopy [31, 34, 35] typically require large detunings of $> 1000\Gamma$ and high intensities of several $W\,cm^{-2}$, and would lead to unacceptable ionization rates if operated simultaneously with a 404.5 nm beam. These kinds of multi-wavelength ionization problems could be mitigated by alternately pulsing the high-intensity cooling and short-wavelength imaging beams.

A.3   Magic Wavelength Optical Lattice

When attempting to laser cool atoms trapped in a deep optical lattice, an additional complication arises due to the effect of the trapping light on the atomic energy levels. A red-detuned optical lattice traps
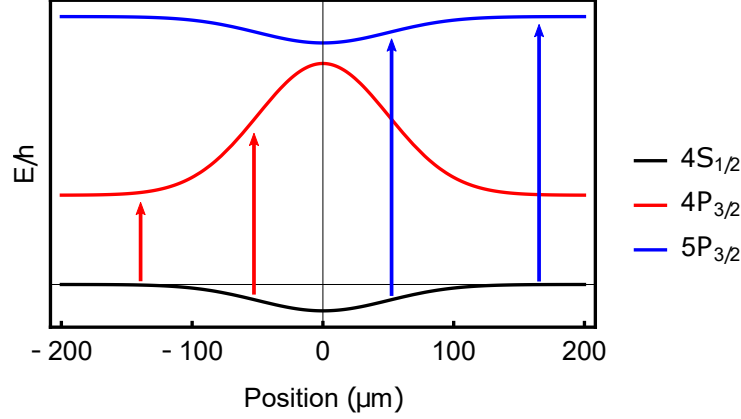
Figure A.5: Effect of light shifts on atomic transitions. The $4P_{3/2}$ state is strongly anti-trapped by a trap at $\lambda = 1054\,\text{nm}$, meaning that the resonant frequency of the transition becomes dependent on trap strength. At a magic wavelength, the potential of the $5P_{3/2}$ state can be matched to the $4S_{1/2}$ ground state, so that the resonant frequency is independent of the strength of the trap potential.

atoms by lowering the energy of the ground electronic state in regions of high lattice intensity. The strength of the trap can be estimated by finding the energy eigenvalues of a two-level atom perturbed by the dipole interaction $H_I = -\hat{d}E$, giving the energy shifts

$$\Delta E = \pm \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I \,, \tag{A.4}$$

where Δ is the detuning of the trapping beam, ω_0 is the resonant frequency, Γ is the linewidth of the excited state, and I is the beam intensity. The negative (positive) sign corresponds to the ground (excited) state of the two-level system. It is seen in this picture that the excited state is shifted to higher energy, meaning that an atom in the electronic excited state is repelled by regions of high laser intensity. This has two obvious consequences for laser cooling. First, the detunings of all laser cooling beams need to be adjusted to correspond to a new resonance location in the optical trap (and correspondingly, trap inhomogeneity will lead to inhomogeneous detuning of laser cooling beams). Second, if the excited state population becomes large, the trapping potential at each site of the optical lattice will be reduced.

For an atom with many excited levels interacting with an far-detuned laser beam, the energy shift of a particular state i is calculated by summing over the transitions to all possible final states j:

$$\Delta E_i = \frac{3\pi c^2 \Gamma}{2\omega_0^3} \times \sum_j \frac{c_{ij}^2}{\Delta_{ij}} \,, \tag{A.5}$$

taking into account the different detunings Δ_{ij} and line strengths c_{ij}^2. Now depending on the frequency of the laser beam, it is possible that multiple transitions have comparable contributions to the sum, and so E_i may be positive, negative, or zero for an excited state i. For effective laser cooling of atoms in an optical trap, it is desirable that the energy shift of the excited state of the cooling transition is as close as possible to the shift of the ground state energy. The trapping wavelength which satisfies $\Delta E_g = \Delta E_i$ for some excited level i is referred to here as a 'magic wavelength'. This choice decouples the cooling laser detuning from the trap intensity, and the trap potential from the excited state population. A similar term 'magic-zero wavelength' is commonly used to refer to wavelengths for which $\Delta_g = 0$ so that atoms
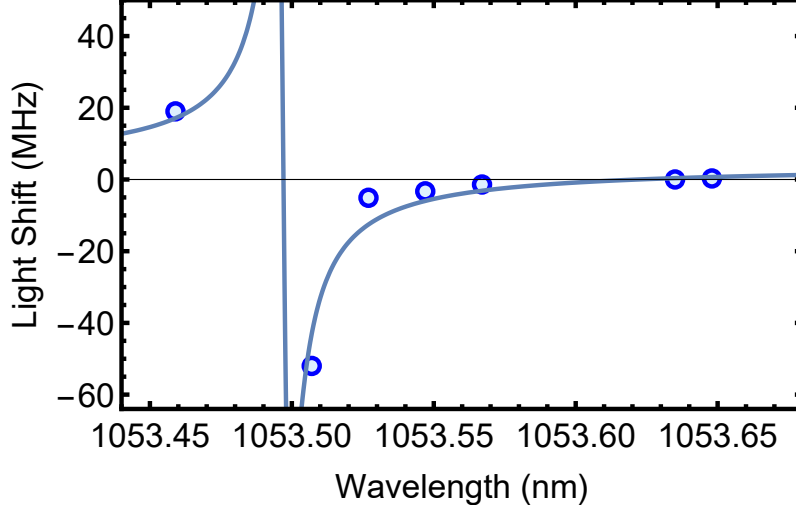
Figure A.6: Locating the magic wavelength near the $12D_{3/2}$ transition. The measured differential light shift between the $|4S_{1/2}, F = 9/2, m_F = 9/2\rangle$ and $|5P_{3/2}, F = 11/2, m_F = 11/2\rangle$ states is plotted as a function of the wavelength of a $400\,E_R$-deep optical lattice which confines the cloud. A dispersion curve centred near the predicted $5P_{3/2} \to 12D_{3/2}$ resonance at $1053.49\,\text{nm}$ is plotted as a guide to the eye. Within the tuning range of the laser, the differential shift can be tuned to be strongly positive, strongly negative, or zero.

feel no potential of any kind in the presence of an intense laser [111].

Theoretical calculations have shown that the $4S \to 5P$ transitions of ^{40}K have several magic wavelengths near $1064\,\text{nm}$ for which high-power laser sources are available [112]. We chose to create our optical lattice at a wavelength of $1053.6\,\text{nm}$ to take advantage of the magic trapping of $5P_{3/2}$ due to a strong contribution of the $5P_{3/2} \to 12D_{3/2}$ transition in the sum (A.5). Similarly, a previous demonstration of Doppler cooling on the $2S_{1/2} \to 3P_{3/2}$ transition of ^{6}Li used a dipole trap at $\lambda = 1070\,\text{nm}$ to achieve nearly-magic conditions, allowing Doppler cooling to continue inside the deep trap [105].

Since the $5P_{3/2} \to 12D_{3/2}$ is fairly weak, the magic wavelength occurs very close to the resonant wavelength $\lambda_0 = 1053.49\,\text{nm}$. This small detuning means that the wavelength of the optical trap must be set very precisely. Our master laser has a tuning range of $30\,\text{GHz}$ which can be slightly extended with manual thermal tuning, giving access to a wavelength range of $1053.45 - 1053.65\,\text{nm}$. Scanning the wavelength of the master laser through this range, and probing the $4S \to 5P_{3/2}$ transition for atoms trapped in a deep optical lattice, we map out the differential light shift $\Delta E_{5P_{3/2}} - \Delta E_{4S}$ as a function of λ. The results are shown in Figure A.6. We find that the differential light shift vanishes for $|F = 9/2, m_F = 9/2\rangle \to |F' = 11/2, m'_F = 11/2\rangle$ transitions at a wavelength $\lambda_{\text{magic}} = 1053.64(2)\,\text{nm}$. The tuning range of our laser also allows us to reach both sides of the transition to $12D_{3/2}$, realising either large trapping or anti-trapping for the $5P_{3/2}$ excited states. This technique to create large shifts of an electronic excited state has been used in Yb to produce deep trapping during single atom imaging [67], and in Rb to image different regions of a trapped cloud in-situ [113]

Since the $5P_{3/2}$ excited state has $J = 3/2$, it is subject to both scalar and tensor light shifts in the presence of a perturbing electric field $\vec{E}$ such as our far-detuned optical trap. The energy shifts for the different internal states of the $5P_{3/2}$ manifold are best expressed in the $|I, J, m_I, m_J\rangle$ basis, and are given

by

$$E_{m_J} = -\frac{1}{4}\left(\alpha^S + \alpha^T \frac{3m_J^2 - J(J+1)}{J(2J-1)}\right)\vec{E}_0^2, \tag{A.6}$$

where α^S and α^T are the scalar and tensor polarizabilities of the $5P_{3/2}$ state. The second term depending on m_J leads to different overall light shifts for states with $|m_J| = 1/2$ and $|m_J| = 3/2$, which was neglected in writing (A.5) previously. Thus there cannot be a single magic wavelength for the entire $5P_{3/2}$ manifold. A wavelength can be chosen such that only the scalar light shift is matched between ground and excited states, or alternately the light shift for states with either $|m_J| = 1/2$ or $|m_J| = 3/2$ can be tuned to the magic condition. This complication does not arise for states with $J = 1/2$, as there is no tensor light shift for these states.

For the data in Figure A.6 we probed the light shifts using the $|F = 9/2, m_F = 9/2\rangle \rightarrow |F' = 11/2, m'_F = 11/2\rangle$. This final state has $m_J = 3/2$, and so by using this transition we exclusively probe the $|m_J| = 3/2$ light shifts, finding $\lambda_{\mathrm{magic},|m_J|=3/2} = 1053.64(2)\,\mathrm{nm}$. This is close to the theoretical prediction of $1053.59(2)\,\mathrm{nm}$ [112]. Our wavelengths were measured using a Burleigh WA-1000 Wavemeter, and calibration errors may influence our measurement of λ_{magic}. Since we can determine the location of the $12D_{3/2}$ resonance, we can also report $\lambda_{\mathrm{magic},|m_J|=3/2} - \lambda_{5P_{3/2}\rightarrow 12D_{3/2}} = 0.14(3)\,\mathrm{nm}$, free of calibration errors.

The second magic wavelength $\lambda_{\mathrm{magic},|m_J|=1/2}$ is expected to be $0.07\,\mathrm{nm}$ larger [112], outside of the tuning range of our laser. With the optical trap operating at $\lambda_{\mathrm{magic},|m_J|=3/2}$, we can probe transitions to other excited states $|F = 9/2, m_F = 9/2\rangle \rightarrow |F', m'_F\rangle$ in order to determine the size of the tensor light shifts. We can detect a weak transition to a state roughly $20\,\mathrm{MHz}$ lower in frequency than $|F'11/2, m'_F = 11/2\rangle$ state. This frequency shift is not due to the hyperfine interaction, but is comparable to the hyperfine splittings of the excited state. An energy shift of this size and direction would be expected for states with $|m_J = 1/2$ when at $\lambda_{\mathrm{magic},|m_J|=3/2}$, due to the size of the tensor polarizability. This large tensor shift indicates that the dipole interaction between the atoms and the optical trap is not simply a perturbation to the hyperfine Hamiltonian, but the two contributions are instead comparable to the overall atomic Hamiltonian. This indicates that neither $|I, J, F, m_F\rangle$ nor $|I, J, m_I, m_J\rangle$ will form a set of good quantum numbers for the $5P_{3/2}$ excited state in this regime, similar to the intermediate-field regime associated with the Zeeman interaction between atoms and static magnetic fields. This breakdown of hyperfine structure due to a strong optical potential has also been observed in ^{87}Rb [114].

It is difficult to determine the effect that the large tensor shifts in the $5P_{3/2}$ state would have on the possibility of performing optical molasses in a deep optical potential at $\lambda_{\mathrm{magic},|m_J|=3/2}$. Optical molasses typically performs best with zero external magnetic field, so that all ground state hyperfine levels are degenerate in energy. However, if the ground state sublevels are off-resonantly coupled to the $5P_{3/2}$ manifold (which experiences large state-dependent shift) then this degeneracy of ground state levels may be broken in unfavourable ways. We previously estimated that an excited state population of $\rho_{ee} > 5 \times 10^{-3}$ would be needed in a $5P_{3/2}$ optical molasses to generate the necessary photon scattering rate for fluorescence imaging. Since we estimate the tensor shifts to be as large as $20\,\mathrm{MHz}$ in our deep optical trap, we might expect various ground state levels to experience shifts as large as $\Delta E/h = \rho_{ee} \times 20\,\mathrm{MHz} = 100\,\mathrm{kHz}$. This state dependant shift would then be comparable to the shift arising from a magnetic field of a few hundreds of mG, and might destroy the Sisyphus cooling mechanism. If large tensor shifts thus preclude an effective optical molasses on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition, it may instead be possible to simply scatter photons at $404.5\,\mathrm{nm}$, while relying on other
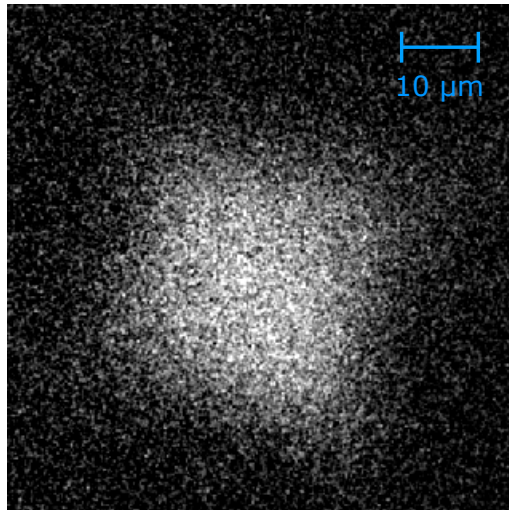
Figure A.7: Fluorescence image of ^{40}K in an optical lattice taken at 404.5 nm. A full three-dimensional cloud loaded from the crossed dipole trap into the optical lattice is imaged, without any plane-selection. Cooling is provided by EIT cooling on the D1 line, and scattered photons from a separate 404.5 nm beam are imaged using a bandpass filter on the iXon EMCCD camera. Exposure time is 500 ms, with 200× EMCCD gain. Neither signal from single atoms nor density modulation due to the optical lattice are visible in such an image, and all fine structure arises from noise sources.

methods such as EIT cooling to keep atoms pinned to lattice sites. In the case of EIT cooling, it would be crucial to alternately pulse the cooling and scattering beams, so that any large ground state shifts associated with the scattering beam are kept from spoiling the cooling mechanism.

A.4   Fluorescence Imaging with 404.5 nm Light

Attempts to perform an optical molasses using 404.5 nm beams to cool atoms trapped in an optical lattice were not successful. We typically observed that atoms were quickly heated out of the lattice when a 404.5 nm polarization gradient was applied. The time for atoms to heat out of the lattice was consistent with the expected recoil heating due to the scattering rate of 404.5 nm photons. Since we had some evidence that EIT cooling on the D1 transition was able to cool atoms trapped in the optical lattice, we thus investigated a combination of D1 EIT cooling and 404.5 nm scattering for fluorescence imaging. We used the same imaging setup described in Chapter 3, with a NA = 0.6 microscope for $\lambda = 405$ nm, and a 405 ± 5 nm bandpass filter in front of the iXon camera to remove stray light from the cooling beams.

We found that it was possible to combine EIT cooling and scattering of 404.5 nm photons, but the intensity of the 404.5 nm beams had to be kept near $0.1\,mW\,cm^{-2}$. Increasing the intensity beyond this point lead to a loss of atoms from the trap, indicating that the cooling method had either been compromised or was unable to remove the additional heat generated by photon scattering. Lowering the intensity would preserve the EIT cooling, but lead to fewer scattered photons at 404.5 nm. We could use this dual-wavelength scheme to produce fluorescence images as shown in Figure A.7. The number of photons collected during a 1.5 s exposure was very low, leading the fluorescence images to be dominated by photon shot noise and the read noise of the EMCCD camera. We were able to discern the fluorescence
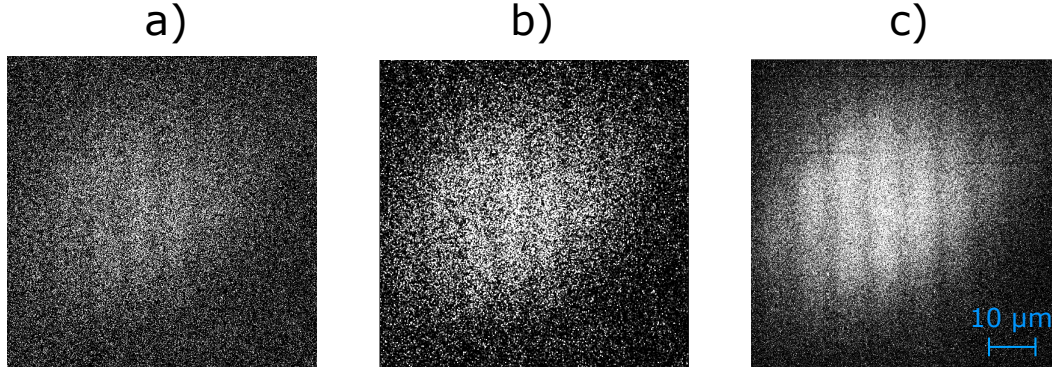
Figure A.8: Fluorescence images of a two-dimensional sample of ^{40}K in an optical lattice taken at 404.5 nm, with similar parameters to the image shown in Figure A.7. Plane-selection is used to reduce the system size to 2-3 planes. Two 404.5 nm beams are used for illumination, crossing at a small angle to create an intensity modulation in the fluorescence. This modulation is used to help focus the microscope. The density pattern is shown with a) raw data from the EMCCD, b) Gaussian filtering with $\sigma = 250$ nm to suppress noise, c) an average of ten consecutive exposures to simulate a tenfold increase in brightness.

of a large three-dimensional cloud of roughly 5×10^4 atoms, but when we selected a single plane of the lattice to look for single-atom fluorescence the signal to noise ratio was very poor as shown in Figure A.8. The brightness of these 404.5 nm fluorescence images was about 350 times less bright than the single atom images obtained using EIT cooling.

Since we worried that the 404.5 nm beam was disturbing the EIT cooling mechanism when its intensity was increased to produce more fluorescence, we attempted to alternately pulse the cooling and scattering beams during image capture. We used AOMs to pulse the beams with a period of around 1 ms, with the 404.5 nm beam turned on for 1/6 of each period and the EIT cooling on the for the remaining 5/6 of each period. While this should have prevented the scattering beam from interfering directly with the cooling mechanism, we were still unable to increase the intensity of the scattering beam without observing loss of atoms from the trap. Thus we suspect that the EIT cooling mechanism was simply unable to cool the atoms quickly enough to overcome the heating due to scattering of 404.5 nm photons.

A.5 The $4S_{1/2} \to 5P_{1/2}$ Transition

Another way to approach cooling and imaging with a short wavelength is to use the $5P_{1/2}$ transition instead of the $5P_{3/2}$ transition. The D1 transition with $J = 1/2$ has been used in many groups to achieve better sub-Doppler cooling in atoms with poorly-resolved hyperfine structure such as ^{39}K [115, 116], ^{40}K [117, 118], ^{6}Li [118, 119], ^{7}Li [120], and ^{23}Na [121]. The simpler hyperfine structure and larger hyperfine splitting of the $J = 1/2$ states make it easier to satisfy the conditions of Sisyphus cooling, with additional coherent effects as described in [120] leading to especially low temperatures. Transitions with $J = 1/2$ are also suitable for implementing EIT cooling, as described in this thesis for the D1 transition. Separately, the implementation of magic wavelength optical traps is also simpler for $J = 1/2$ states, since the lack of a tensor light shift allows a state-independent trapping potential to be created.

To investigate the possibility of performing fluorescence imaging using EIT cooling on the $4S_{1/2} \to 5P_{1/2}$ transition, a laser system was constructed to provide a few mW of power at 404.8 nm. The system is built around a home-built external cavity diode laser assembled by Will Caincross [122], with

an AR-coated laser diode purchased from Sacher Lasertechnik and a narrow band interference filter purchased from Iridian Spectral Technologies. The laser is locked using saturation spectroscopy to the $4S_{1/2} \to 5P_{1/2}$ transition of ^{39}K and shifted into resonance with ^{40}K using a double-pass AOM, as described in the thesis of Daniel Nino [123]. An EOM purchased from Qubig generates sidebands at 1.285 GHz as required for EIT cooling in zero magnetic field. At this point, no laser cooling has been attempted at 404.8 nm.

Bibliography

[1] Ryan Olf, Fang Fang, G Edward Marti, Andrew MacRae, and Dan M Stamper-Kurn. Thermometry and cooling of a Bose gas to 0.02 times the condensation temperature. *Nature Physics*, 11(9):720–723, 2015.

[2] Richard P Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.

[3] Cheng Chin, Rudolf Grimm, Paul Julienne, and Eite Tiesinga. Feshbach resonances in ultracold gases. *Rev. Mod. Phys.*, 82:1225–1286, Apr 2010.

[4] Thomas Uehlinger, Gregor Jotzu, Michael Messer, Daniel Greif, Walter Hofstetter, Ulf Bissbort, and Tilman Esslinger. Artificial graphene with tunable interactions. *Phys. Rev. Lett.*, 111:185307, Oct 2013.

[5] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch. Realization of the Hofstadter hamiltonian with ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 111:185301, Oct 2013.

[6] BK Stuhl, H-I Lu, LM Aycock, D Genkina, and IB Spielman. Visualizing edge states with an atomic Bose gas in the quantum Hall regime. *Science*, 349(6255):1514–1518, 2015.

[7] D. Schlippert, J. Hartwig, H. Albers, L. L. Richardson, C. Schubert, A. Roura, W. P. Schleich, W. Ertmer, and E. M. Rasel. Quantum test of the universality of free fall. *Phys. Rev. Lett.*, 112:203002, May 2014.

[8] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. *Rev. Mod. Phys.*, 87:637–701, Jun 2015.

[9] P. Hamilton, M. Jaffe, P. Haslinger, Q. Simmons, H. Müller, and J. Khoury. Atom-interferometry constraints on dark energy. *Science*, 349(6250):849–851, 2015.

[10] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, 1995.

[11] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. *Science*, 285(5434):1703–1706, 1999.

[12] C. Ospelkaus, S. Ospelkaus, K. Sengstock, and K. Bongs. Interaction-driven dynamics of ^{40}K-^{87}Rb Fermion-Boson gas mixtures in the large-particle-number limit. *Phys. Rev. Lett.*, 96:020401, Jan 2006.

[13] S Aubin, S Myrskog, MHT Extavour, LJ LeBlanc, D McKay, A Stummer, and JH Thywissen. Rapid sympathetic cooling to Fermi degeneracy on a chip. *Nature Physics*, 2(6):384–387, 2006.

[14] R Jördens, N. Strohmaier, K Günter, H Moritz, and T Esslinger. A Mott insulator of fermionic atoms in an optical lattice. *Nature*, 455:204–207, April 2008.

[15] Ruth S. Bloom, Ming-Guang Hu, Tyler D. Cumby, and Deborah S. Jin. Tests of universal three-body physics in an ultracold Bose-Fermi mixture. *Phys. Rev. Lett.*, 111:105301, Sep 2013.

[16] S. Ospelkaus, K.-K. Ni, G. Quéméner, B. Neyenhuis, D. Wang, M. H. G. de Miranda, J. L. Bohn, J. Ye, and D. S. Jin. Controlling the hyperfine state of rovibronic ground-state polar molecules. *Phys. Rev. Lett.*, 104:030402, Jan 2010.

[17] Karl D Nelson, Xiao Li, and David S Weiss. Imaging single atoms in a three-dimensional array. *Nature Phys.*, 3:556–560, 2007.

[18] Waseem S Bakr, Jonathon I Gillen, Amy Peng, Simon Foelling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature*, 462(7269):74–U80, 2009.

[19] Nathan Gemelke, Xibo Zhang, Chen-Lung Hung, and Cheng Chin. In situ observation of incompressible Mott-insulating domains in ultracold atomic gases. *Nature*, 460(7258):995–U75, 2009.

[20] Jacob F Sherson, Christof Weitenberg, Manuel Endres, Marc Cheneau, Immanuel Bloch, and Stefan Kuhr. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature*, 467(7311):68–U97, 2010.

[21] Waseem S Bakr, A Peng, M E Tai, R Ma, J Simon, J I Gillen, S Foelling, L Pollet, and M Greiner. Probing the Superfluid-to-Mott Insulator Transition at the Single-Atom Level. *Science*, 329:547–550, 2010.

[22] Ruichao Ma, M Tai, Philipp Preiss, Waseem Bakr, Jonathan Simon, and Markus Greiner. Photon-Assisted Tunneling in a Biased Strongly Correlated Bose Gas. *Phys. Rev. Lett.*, 107:095301, 2011.

[23] Jonathan Simon, Waseem S Bakr, Ruichao Ma, M Eric Tai, Philipp M Preiss, and Markus Greiner. Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472(7343):307–312, 2011.

[24] Waseem S Bakr, Philipp M Preiss, M Eric Tai, Ruichao Ma, Jonathan Simon, and Markus Greiner. Orbital excitation blockade and algorithmic cooling in quantum gases. *Nature*, 480(7378):500–503, 2011.

[25] Manuel Endres, Takeshi Fukuhara, David Pekker, Marc Cheneau, Peter Schauß, Christian Gross, Eugene Demler, Stefan Kuhr, and Immanuel Bloch. The 'Higgs' amplitude mode at the two-dimensional superfluid/Mott insulator transition. *Nature*, 487(7408):454–458, 2012.

[26] Peter Schauß, Marc Cheneau, Manuel Endres, Takeshi Fukuhara, Sebastian Hild, Ahmed Omran, Thomas Pohl, Christian Gross, Stefan Kuhr, and Immanuel Bloch. Observation of spatially ordered structures in a two-dimensional Rydberg gas. *Nature*, 491(7422):87–91, 2012.

[27] Takeshi Fukuhara, Peter Schauß, Manuel Endres, Sebastian Hild, Marc Cheneau, Immanuel Bloch, and Christian Gross. Microscopic observation of magnon bound states and their dynamics. *Nature*, 502(7469):76–79, 2013.

[28] Philipp M Preiss, Ruichao Ma, M Eric Tai, Alexander Lukin, Matthew Rispoli, Philip Zupancic, Yoav Lahini, Rajibul Islam, and Markus Greiner. Strongly correlated quantum walks in optical lattices. *Science*, 347(6227):1229–1233, 2015.

[29] Rajibul Islam, Ruichao Ma, Philipp M Preiss, M Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83, 2015.

[30] Takeshi Fukuhara, Sebastian Hild, Johannes Zeiher, Peter Schauß, Immanuel Bloch, Manuel Endres, and Christian Gross. Spatially resolved detection of a spin-entanglement wave in a Bose-Hubbard chain. *Phys. Rev. Lett.*, 115:035302, Jul 2015.

[31] Lawrence W Cheuk, Matthew A Nichols, Melih Okan, Thomas Gersdorf, Vinay V Ramasesh, Waseem S Bakr, Thomas Lompe, and Martin W Zwierlein. Quantum-Gas Microscope for Fermionic Atoms. *Phys. Rev. Lett.*, 114:193001, 2015.

[32] E Haller, J Hudson, A Kelly, D A Cotta, B Peaudecerf, G D Bruce, and S Kuhr. Single-atom imaging of fermions in a quantum-gas microscope. *Nature Phys.*, 11:738–742, 2015.

[33] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen. Imaging and addressing of individual fermionic atoms in an optical lattice. *Phys. Rev. A*, 92:063406, Dec 2015.

[34] Maxwell F Parsons, Florian Huber, Anton Mazurenko, Christie S Chiu, Widagdo Setiawan, Katherine Wooley-Brown, Sebastian Blatt, and Markus Greiner. Site-Resolved Imaging of Fermionic Li6 in an Optical Lattice. *Phys. Rev. Lett.*, 114:213002, 2015.

[35] Ahmed Omran, Martin Boll, Timon A. Hilker, Katharina Kleinlein, Guillaume Salomon, Immanuel Bloch, and Christian Gross. Microscopic observation of Pauli blocking in degenerate fermionic lattice gases. *Phys. Rev. Lett.*, 115:263001, Dec 2015.

[36] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969–3973, Nov 1995.

[37] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical Dipole Traps for Neutral Atoms. *Advances in Atomic Molecular and Optical Physics*, 42:95–170, 2000.

[38] Hidetoshi Katori, Tetsuya Ido, and Makoto Kuwata-Gonokami. Optimal design of dipole potentials for efficient loading of Sr atoms. *Journal of the Physical Society of Japan*, 68(8):2479–2482, 1999.

[39] V. Boyer, R. M. Godun, G. Smirne, D. Cassettari, C. M. Chandrashekar, A. B. Deb, Z. J. Laczik, and C. J. Foot. Dynamic manipulation of Bose-Einstein condensates with a spatial light modulator. *Phys. Rev. A*, 73:031402, Mar 2006.

[40] A. L. Gaunt and Z. Hadzibabic. Robust digital holography for ultracold atom trapping. *Scientific Reports*, 2, 2012.

[41] N. Friedman, L. Khaykovich, R. Ozeri, and N. Davidson. Compression of cold atoms to very high densities in a rotating-beam blue-detuned optical trap. *Phys. Rev. A*, 61:031403, Feb 2000.

[42] K Henderson, C Ryu, C MacCormick, and M G Boshier. Experimental demonstration of painting arbitrary and dynamic potentials for BoseEinstein condensates. *New Journal of Physics*, 11(4):043030, 2009.

[43] Steffen P. Rath, Tarik Yefsah, Kenneth J. Günter, Marc Cheneau, Rémi Desbuquois, Markus Holzmann, Werner Krauth, and Jean Dalibard. Equilibrium state of a trapped two-dimensional Bose gas. *Phys. Rev. A*, 82:013609, Jul 2010.

[44] Alexander L. Gaunt, Tobias F. Schmidutz, Igor Gotlibovych, Robert P. Smith, and Zoran Hadzibabic. Bose-Einstein condensation of atoms in a uniform potential. *Phys. Rev. Lett.*, 110:200406, May 2013.

[45] A. Ramanathan, K. C. Wright, S. R. Muniz, M. Zelan, W. T. Hill, C. J. Lobb, K. Helmerson, W. D. Phillips, and G. K. Campbell. Superflow in a toroidal Bose-Einstein condensate: An atom circuit with a tunable weak link. *Phys. Rev. Lett.*, 106:130401, Mar 2011.

[46] Stuart Moulder, Scott Beattie, Robert P. Smith, Naaman Tammuz, and Zoran Hadzibabic. Quantized supercurrent decay in an annular Bose-Einstein condensate. *Phys. Rev. A*, 86:013629, Jul 2012.

[47] Russell A Hart, Pedro M Duarte, Tsung-Lin Yang, Xinxing Liu, Thereza Paiva, Ehsan Khatami, Richard T Scalettar, Nandini Trivedi, David A Huse, and Randall G Hulet. Observation of antiferromagnetic correlations in the Hubbard model with ultracold atoms. *Nature*, 519:211–214, February 2015.

[48] Georg M. Bruun and Charles W. Clark. Ideal gases in time-dependent traps. *Phys. Rev. A*, 61:061601, May 2000.

[49] Marcius Extavour. *Fermions and Bosons on an Atom Chip*. PhD thesis, University of Toronto, 8 2009.

[50] Charles Kittel. *Elementary Solid State Physics: A Short Course*. Wiley, 1962.

[51] J. C. Slater. A soluble problem in energy bands. *Phys. Rev.*, 87:807–835, Sep 1952.

[52] U. Schneider, L. Hackermüller, S. Will, Th. Best, I. Bloch, T. A. Costi, R. W. Helmes, D. Rasch, and A. Rosch. Metallic and insulating phases of repulsively interacting fermions in a 3d optical lattice. *Science*, 322(5907):1520–1525, 2008.

[53] Ana Maria Rey, Guido Pupillo, Charles W. Clark, and Carl J. Williams. Ultracold atoms confined in an optical lattice plus parabolic potential: A closed-form approach. *Phys. Rev. A*, 72:033616, Sep 2005.

[54] H. Ott, E. de Mirandes, F. Ferlaino, G. Roati, V. Türck, G. Modugno, and M. Inguscio. Radio frequency selective addressing of localized atoms in a periodic potential. *Phys. Rev. Lett.*, 93:120407, Sep 2004.

[55] Michael Köhl. Thermometry of fermionic atoms in an optical lattice. *Phys. Rev. A*, 73:031601, Mar 2006.

[56] Dieter Jaksch, Ch Bruder, Juan Ignacio Cirac, Crispin W Gardiner, and Peter Zoller. Cold bosonic atoms in optical lattices. *Physical Review Letters*, 81(15):3108, 1998.

[57] Lorenzo De Leo, Jean-Sébastien Bernier, Corinna Kollath, Antoine Georges, and Vito W. Scarola. Thermodynamics of the three-dimensional Hubbard model: Implications for cooling cold atomic gases in optical lattices. *Phys. Rev. A*, 83:023606, Feb 2011.

[58] D. C. McKay and B DeMarco. Cooling in strongly correlated optical lattices: prospects and challenges. *Rep. Prog. Phys.*, 74:054401, 2011.

[59] Daniel Greif, Maxwell F. Parsons, Anton Mazurenko, Christie S. Chiu, Sebastian Blatt, Florian Huber, Geoffrey Ji, and Markus Greiner. Site-resolved imaging of a fermionic Mott insulator. *Science*, 351(6276):953–957, 2016.

[60] Maxwell F. Parsons, Anton Mazurenko, Christie S. Chiu, Geoffrey Ji, Daniel Greif, and Markus Greiner. Site-resolved measurement of the spin-correlation function in the Fermi-Hubbard model. *Science*, 353(6305):1253–1256, 2016.

[61] Lawrence W. Cheuk, Matthew A. Nichols, Katherine R. Lawrence, Melih Okan, Hao Zhang, Ehsan Khatami, Nandini Trivedi, Thereza Paiva, Marcos Rigol, and Martin W. Zwierlein. Observation of spatial charge and spin correlations in the 2d Fermi-Hubbard model, 2016.

[62] Martin Boll, Timon A. Hilker, Guillaume Salomon, Ahmed Omran, Immanuel Bloch, and Christian Gross. Spin and charge resolved quantum gas microscopy of antiferromagnetic order in Hubbard chains, 2016.

[63] B Zimmermann, T Mller, J Meineke, T Esslinger, and H Moritz. High-resolution imaging of ultracold fermions in microscopically tailored optical potentials. *New Journal of Physics*, 13(4):043007, 2011.

[64] Eugenio Cocchi, Luke A. Miller, Jan H. Drewes, Marco Koschorreck, Daniel Pertot, Ferdinand Brennecke, and Michael Köhl. Equation of state of the two-dimensional Hubbard model. *Phys. Rev. Lett.*, 116:175301, Apr 2016.

[65] Dave McKay. *Quantum Simulation in Strongly Correlated Optical Lattices*. PhD thesis, University of Illinois at Urbana-Champaign, 7 2012.

[66] Dylan Jervis. *A Fermi Gas Microscope Apparatus*. PhD thesis, University of Toronto, 8 2014.

[67] Martin Miranda, Ryotaro Inoue, Yuki Okuyama, Akimasa Nakamoto, and Mikio Kozuma. Site-resolved imaging of ytterbium atoms in a two-dimensional optical lattice. *Phys. Rev. A*, 91(6):063414, 2015.

[68] Aaron Reinhard, Jean-Flix Riou, Laura A. Zundel, and David S. Weiss. Dark-ground imaging of high optical thickness atom clouds. *Optics Communications*, 324:30 – 33, 2014.

[69] S. Blatt, A. Mazurenko, M. F. Parsons, C. S. Chiu, F. Huber, and M. Greiner. Low-noise optical lattices for ultracold ^{6}Li. *Phys. Rev. A*, 92:021402, Aug 2015.

[70] Oliver Morsch and Markus Oberthaler. Dynamics of Bose-Einstein condensates in optical lattices. *Rev. Mod. Phys.*, 78:179–215, Feb 2006.

[71] G. Roati, F. Riboli, G. Modugno, and M. Inguscio. Fermi-Bose quantum degenerate ^{40}K-^{87}Rb mixture with attractive interaction. *Phys. Rev. Lett.*, 89:150403, Sep 2002.

[72] J. Goldwin, S. Inouye, M. L. Olsen, B. Newman, B. D. DePaola, and D. S. Jin. Measurement of the interaction strength in a Bose-Fermi mixture with ^{87}Rb and ^{40}K. *Phys. Rev. A*, 70:021601, Aug 2004.

[73] A. Kastberg, W. D. Phillips, S. L. Rolston, R. J. C. Spreeuw, and P. S. Jessen. Adiabatic cooling of cesium to 700 nK in an optical lattice. *Phys. Rev. Lett.*, 74:1542–1545, Feb 1995.

[74] Lawrence W. Cheuk, Matthew A. Nichols, Katherine R. Lawrence, Melih Okan, Hao Zhang, and Martin W. Zwierlein. Observation of 2d fermionic Mott insulators of ^{40}K with single-site resolution. *Phys. Rev. Lett.*, 116:235301, Jun 2016.

[75] Philipp M. Preiss, Ruichao Ma, M. Eric Tai, Jonathan Simon, and Markus Greiner. Quantum gas microscopy with spin, atom-number, and multilayer readout. *Phys. Rev. A*, 91:041602, Apr 2015.

[76] Michał Karski, Leonid Förster, Jai-Min Choi, Andreas Steffen, Noomen Belmechri, Wolfgang Alt, Dieter Meschede, and Artur Widera. Imprinting patterns of neutral atoms in an optical lattice using magnetic resonance techniques. *New J Phys*, 12:065027, 2010.

[77] Nicolas Schlosser, Georges Reymond, Igor Protsenko, and Philippe Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027, 2001.

[78] T Grünzweig, A Hilliard, M McGovern, and MF Andersen. Near-deterministic preparation of a single atom in an optical microtrap. *Nature Physics*, 6(12):951–954, 2010.

[79] F. Diedrich, J. C. Bergquist, Wayne M. Itano, and D. J. Wineland. Laser cooling to the zero-point energy of motion. *Phys. Rev. Lett.*, 62:403–406, Jan 1989.

[80] Leonid Förster, Michał Karski, Jai-Min Choi, Andreas Steffen, Wolfgang Alt, Dieter Meschede, Artur Widera, Enrique Montano, Jae Hoon Lee, Worawarong Rakreungdet, and Poul S. Jessen. Microwave control of atomic motion in optical lattices. *Phys. Rev. Lett.*, 103:233001, Dec 2009.

[81] Andrew J. Kerman, Vladan Vuletić, Cheng Chin, and Steven Chu. Beyond optical molasses: 3d Raman sideband cooling of atomic cesium to high phase-space density. *Phys. Rev. Lett.*, 84:439–442, Jan 2000.

[82] Y. S. Patil, S. Chakram, L. M. Aycock, and M. Vengalattore. Nondestructive imaging of an ultracold lattice gas. *Phys. Rev. A*, 90:033422, Sep 2014.

[83] Helene Perrin, Axel Kuhn, Isabelle Bouchoule, and Christophe Salomon. Sideband cooling of neutral atoms in a far-detuned optical lattice. *EPL (Europhysics Letters)*, 42(4):395, 1998.

[84] G Morigi, J Eschner, and C H Keitel. Ground state laser cooling using electromagnetically induced transparency. *Phys. Rev. Lett.*, 85(21):4458–4461, 2000.

[85] G Morigi. Cooling atomic motion with quantum interference. *Phys. Rev. A*, 67:033402, 2003.

[86] C F Roos, D Leibfried, A Mundt, F Schmidt-Kaler, J Eschner, and R Blatt. Experimental demonstration of ground state laser cooling with electromagnetically induced transparency. *Phys. Rev. Lett.*, 85:5547, 2000.

[87] J. Evers and C. H. Keitel. Double-EIT ground-state laser cooling without blue-sideband heating. *Europhys. Lett.*, 68(3):370–376, 2004.

[88] Michael Fleischhauer, Atac Imamoglu, and Jonathan P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Rev. Mod. Phys.*, 77:633–673, Jul 2005.

[89] Petr M. Anisimov, Jonathan P. Dowling, and Barry C. Sanders. Objectively discerning Autler-Townes splitting from electromagnetically induced transparency. *Phys. Rev. Lett.*, 107:163604, Oct 2011.

[90] Stefan S. Natu, David C. McKay, Brian DeMarco, and Erich J. Mueller. Evolution of condensate fraction during rapid lattice ramps. *Phys. Rev. A*, 85:061601, Jun 2012.

[91] S Wolf, S J Oliver, and D S Weiss. Suppression of recoil heating by an optical lattice. *Phys. Rev. Lett.*, 85(20):4249–4252, 2000.

[92] Mark Stanford Robbins. *Electron-Multiplying Charge Coupled Devices – EMCCDs*, pages 103–121. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[93] Andrea Alberti, Carsten Robens, Wolfgang Alt, Stefan Brakhane, Michał Karski, René Reimann, Artur Widera, and Dieter Meschede. Super-resolution microscopy of single atoms in optical lattices. *New Journal of Physics*, 18(5):053010, 2016.

[94] JD Wong-Campos, KG Johnson, B Neyenhuis, J Mizrahi, and C Monroe. High resolution adaptive imaging of a single atom. *Nature Photonics*, 10:606–610, 2016.

[95] Tingwei Quan, Shaoqun Zeng, and Zhen-Li Huang. Localization capability and limitation of electron-multiplying charge-coupled, scientific complementary metal-oxide semiconductor, and charge-coupled devices for superresolution imaging. *Journal of Biomedical Optics*, 15(6):066005–066005, 2010.

[96] Min Gu. *Diffraction Theory*, pages 7–35. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[97] Ryuta Yamamoto, Jun Kobayashi, Takuma Kuno, Kohei Kato, and Yoshiro Takahashi. An ytterbium quantum gas microscope with narrow-line laser cooling. *New Journal of Physics*, 18(2):023016, 2016.

[98] Ted A Laurence and Brett A Chromy. Efficient maximum likelihood estimator fitting of histograms. *Nature Methods*, 7(5):338–339, 2010.

[99] Peter Hänggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after Kramers. *Rev. Mod. Phys.*, 62:251–341, Apr 1990.

[100] Christof Weitenberg. *Single-atom resolved imaging and manipulation of an atomic Mott insulator.* PhD thesis, Ludwig-Maximilians-Universität, 2011.

[101] JH Drewes, LA Miller, E Cocchi, CF Chan, D Pertot, F Brennecke, and M Köhl. Antiferromagnetic correlations in two-dimensional fermionic Mott-insulating and metallic phases. *arXiv preprint arXiv:1607.00392*, 2016.

[102] Philip Zupancic, Philipp M. Preiss, Ruichao Ma, Alexander Lukin, M. Eric Tai, Matthew Rispoli, Rajibul Islam, and Markus Greiner. Ultra-precise holographic beam shaping for microscopic quantum control. *Opt. Express*, 24(13):13881–13893, Jun 2016.

[103] Jean-Sébastien Bernier, Corinna Kollath, Antoine Georges, Lorenzo De Leo, Fabrice Gerbier, Christophe Salomon, and Michael Köhl. Cooling fermionic atoms in optical lattices by shaping the confinement. *Phys. Rev. A*, 79:061601, Jun 2009.

[104] D. C. McKay, D. Jervis, D. J. Fine, J. W. Simpson-Porco, G. J. A. Edge, and J. H. Thywissen. Low-temperature high-density magneto-optical trapping of potassium using the open $4s \rightarrow 5p$ transition at 405 nm. *Phys. Rev. A*, 84:063420, Dec 2011.

[105] P. M. Duarte, R. A. Hart, J. M. Hitchcock, T. A. Corcovilos, T.-L. Yang, A. Reed, and R. G. Hulet. All-optical production of a lithium quantum gas using narrow-line laser cooling. *Phys. Rev. A*, 84:061406, Dec 2011.

[106] M. Landini, S. Roy, L. Carcagní, D. Trypogeorgos, M. Fattori, M. Inguscio, and G. Modugno. Sub-doppler laser cooling of potassium atoms. *Phys. Rev. A*, 84:043432, Oct 2011.

[107] J. E. Sansonetti. Wavelengths, transition probabilities, and energy levels for the spectra of potassium (KIKI through KXIXKXIX). *J. Phys. Chem. Ref. Data*, 37, 2008.

[108] H. Wang, J. Li, X. T. Wang, C. J. Williams, P. L. Gould, and W. C. Stwalley. Precise determination of the dipole matrix element and radiative lifetime of the ^{39}K 4p state by photoassociative spectroscopy. *Phys. Rev. A*, 55:R1569–R1572, Mar 1997.

[109] D. Ciampini, M. Anderlini, J. H. Müller, F. Fuso, O. Morsch, J. W. Thomsen, and E. Arimondo. Photoionization of ultracold and Bose-Einstein-condensed Rb atoms. *Phys. Rev. A*, 66:043409, Oct 2002.

[110] Alexandra Behrle, Marco Koschorreck, and Michael Köhl. Isotope shift and hyperfine splitting of the $4s \rightarrow 5p$ transition in potassium. *Phys. Rev. A*, 83:052507, May 2011.

[111] William F. Holmgren, Raisa Trubko, Ivan Hromada, and Alexander D. Cronin. Measurement of a wavelength of light for which the energy shift for an atom vanishes. *Phys. Rev. Lett.*, 109:243004, Dec 2012.

[112] M. S. Safronova, U. I. Safronova, and Charles W. Clark. Magic wavelengths for optical cooling and trapping of potassium. *Phys. Rev. A*, 87:052504, May 2013.

[113] J. P. Brantut, J. F. Clément, M. Robert de Saint Vincent, G. Varoquaux, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer. Light-shift tomography in an optical-dipole trap for neutral atoms. *Phys. Rev. A*, 78:031401, Sep 2008.

[114] Andreas Neuzner, Matthias Körber, Stephan Dürr, Gerhard Rempe, and Stephan Ritter. Breakdown of atomic hyperfine coupling in a deep optical-dipole trap. *Phys. Rev. A*, 92:053842, Nov 2015.

[115] G Salomon, L Fouché, P Wang, A Aspect, P Bouyer, and T Bourdel. Gray-molasses cooling of 39K to a high phase-space density. *Europhys. Lett.*, 104:63002, 2013.

[116] Dipankar Nath, R Kollengode Easwaran, G Rajalakshmi, and C S Unnikrishnan. Quantum-interference-enhanced deep sub-Doppler cooling of K-39 atoms in gray molasses. *Phys. Rev. A*, 88:053407, 2013.

[117] D Rio Fernandes, F Sievers, N Kretzschmar, S Wu, Christophe Salomon, and F Chevy. Sub-Doppler laser cooling of fermionic K-40 atoms in three-dimensional gray optical molasses. *Europhys. Lett.*, 100:63001, 2012.

[118] Franz Sievers, Norman Kretzschmar, Diogo Rio Fernandes, Daniel Suchet, Michael Rabinovic, Saijun Wu, Colin V. Parker, Lev Khaykovich, Christophe Salomon, and Frédéric Chevy. Simultaneous sub-doppler laser cooling of fermionic ^{6}Li and ^{40}K on the D_1 line: Theory and experiment. *Phys. Rev. A*, 91:023426, Feb 2015.

[119] A. Burchianti, G. Valtolina, J. A. Seman, E. Pace, M. DePas, M. Inguscio, M. Zaccanti, and G. Roati. Efficient all-optical production of large ^{6}Li quantum gases using D_1 gray-molasses cooling. *Phys. Rev. A*, 90:043408, 2014.

[120] A T Grier, I Ferrier-Barbut, B S Rem, M Delehaye, L Khaykovich, F Chevy, and C Salomon. Λ-enhanced sub-Doppler cooling of lithium atoms in D 1 gray molasses. *Phys. Rev. A*, 87:063411, 2013.

[121] Giacomo Colzi, Gianmaria Durastante, Eleonora Fava, Simone Serafini, Giacomo Lamporesi, and Gabriele Ferrari. Sub-Doppler cooling of sodium atoms in gray molasses. *Phys. Rev. A*, 93:023421, Feb 2016.

[122] Will Cairncross and Ian Kivlichan. Interference-filter stabilized ECDLs. Technical report, University of Toronto, 8 2012. NSERC USRA Project.

[123] Daniel Nino. Master's thesis, University of Toronto, 8 2015.