RELAXATION OF INTERACTING STARK-LOCALIZED STATES IN A THERMAL LATTICE BOSE GAS

BY

LAURA ROSE WADLEIGH

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2023

Urbana, Illinois

Doctoral Committee:

Associate Professor Bryce Gadway, Chair
Professor Brian DeMarco, Director of Research
Associate Professor Virginia Lorenz
Assistant Professor Lucas Wagner

Abstract

Localization in interacting quantum lattice gas systems is an active area of research. In this work, we study the dynamics of Stark-localized states in a thermal lattice Bose gas. We use a barrier beam to exclude atoms from the center of an atomic gas trapped in a harmonic potential. A three-dimensional optical lattice potential is added, creating Stark-localized states away from the center of the trap, where the gradient of the harmonic potential is large compared to the lattice bandwidth. We remove the barrier beam and observe the dynamics over hundreds of tunneling times. Using a new statistical measure, Mardia's B, we observe that the gas reaches thermal equilibrium after hundreds of tunneling times, even at high lattice depths where the gas is nearly completely Stark localized. We rule out many possible sources of relaxation, including scattering of lattice light. We conclude that physics beyond the Bose-Hubbard model may be needed to explain the long-time dynamics of interacting quantum lattice gas systems.

Acknowledgments

There are far too many people to thank to properly do them all justice, but I'd like to at least mention a few of the people who made this work possible.

First, a huge thanks to my advisor, Brian DeMarco. He has been an excellent advisor — teaching me to be a better scientist, without losing sight of being a great human.

Thank you to my entire lab group. In particular, those that worked on the rubidium apparatus with me: Carrie and Phil, who taught me the ropes, Nicholas, who spent many hours troubleshooting the apparatus with me, and Max, who is continuing on the tradition. I also really appreciate my labmates working on the potassium apparatus: Will, Pei-Wen, Nathan, and Wenchao, from whom I learned a lot.

I also want to thank my second-year housemates and friends throughout grad school, Michelle, Ashwathi, and Karmela, with whom I've shared all sorts of great discussions, both physics and otherwise.

Thank you to the many other people who made my time in Urbana-Champaign special, both inside and outside the physics department. It wouldn't have been the same without each and every one of you.

I also owe thanks to everyone who helped me on my path to a physics Ph.D. before I even started grad school. In particular, thank you to my undergraduate research advisor, Dwight Luhman, who first showed me how exciting research can be. I also want to thank the entire Carleton physics department, both the professors and my fellow students, for creating an excellent environment in which to learn and fall in love with physics.

And thank you to everyone that fostered my love of science even earlier: my teachers both in and out of school and, of course, my family, who never seemed to tire of me asking "why?" — and even managed to find ways to let me experiment and figure it out for myself sometimes.

And, lastly, I owe so many thanks to my partner, Matt, who has supported me in this pursuit, and in life, in far too many ways to list. Matthew, your encouragement, ideas, and support have been invaluable.

Table of Contents

List of Figures

Chapter 1

Introduction

Atomic gases have provided a playground for studying localization in strongly correlated systems [1–3]. In this work, we study the localization dynamics of a thermal gas of bosons trapped in an optical lattice. We engineer the system to have a large population of Stark-localized states by tuning the lattice potential depth. According to the Bose-Hubbard model usually used to describe lattice bosons, these states are prevented from moving by the presence of a potential gradient. However, we observed relaxation of these states at long times. After excluding many potential sources of thermalization, we conclude that physics beyond the Bose-Hubbard model may be needed to describe the long-time dynamics of interacting quantum lattice gas systems.[1]

We also developed tools that can be used to study the dynamics of quantum lattice gases in new ways. We developed a barrier beam to enable spatially-selective modification of an atomic gas. We also designed a new imaging system to better resolve in-trap density distributions. Lastly, we developed a statistical measure, Mardia's B, to quantify whether an atomic density distribution is in equilibrium based on *in situ* images of the gas.

1.1 Localization

Localization in many-body quantum systems is an example of a quantum system that is inconsistent with the eigenstate thermalization hypothesis and fails to thermalize [5–7], instead maintaining some memory of its initial conditions. There are multiple potential mechanisms that break ergodicity and cause localization. Work studying localization in ultracold quantum gases has primarily focused on localization arising from disorder, interactions, and the interplay between the two [1, 2, 5].

Besides interest in the fundamental physics of localization in quantum systems, there are also potential practical applications. Localization may be useful for developing quantum batteries [8, 9]. Reference [10] finds that the entanglement present in a many-body localized states may improve the work-extraction capabilities of a quantum energy storage device. Localization in many-body systems also has potential applications for quantum memories and quantum information processing [11, 12]. Development of these technologies is in very early stages. A deeper understanding of localization and its limitations is necessary to advance its practical applications. Theoretical work often focuses on infinite systems at zero-temperature, but understanding the physical limitations given real-world experimental constraints is necessary to know whether these kinds of technologies can work in practice.

[1] This experiment was also described in Reference [4].

1.1.1 Many-Body Localization

One particularly active area of research has been many-body localization (MBL), where disorder induces localization in an interacting quantum system [13]. The non-interacting analog to MBL — Anderson localization [14] — first described disorder-induced localization in 1958 [15]. Anderson localization is understood to arise from interference between the forward and time-reversed scattering path of a particle. One may expect interactions between the particles to disrupt localization (e.g. by effecting only the forward or time-reversed scattering path). However, localization has been seen to persist even in the presence of interactions. This is called many-body localization (MBL) [2, 5, 13].

Disorder-induced stopped or slowed dynamics have been observed in a variety of systems. Experimental observations of disorder-induced localization have been made in Fermi [16] and Bose [17] Hubbard systems as well as superconducting qubits [18] and many others. Systems that do not fully localize, but exhibit sub-diffusive behavior, have also been measured [19, 20]. Many studies have used quasi-periodic disorder, including Reference [21], which claims to observe localization in an interacting one-dimensional (1D) fermionic system. However, that measurement samples only 30 tunneling times. They observe a density imbalance that is approximately constant over that timescale. But, slow dynamics would be difficult to resolve on this timescale. Experimental studies are often limited to tens of tunneling times, making conclusions about the long-time dynamics of the system difficult. The difference between localization and slowly-evolving quasi-localized states may not be easily resolved on such timescales.

Numerical studies such as References [22] and [23] often calculate dynamics over much longer times scales, in some cases 1000s of tunneling times. However, these studies are limited to relatively small 1D lattices. Numerical studies on more than 20 lattice sites are uncommon. The authors of Reference [23] numerically calculated the dynamics of a 1D Heisenberg spin chain in the presence of either disorder or a quasiperiodic potential. In the disordered system, they did not observe complete localization, but rather slow power-law decay that continued for at least 1000s of tunneling times. This occurred even with a disorder strength twice that of current estimates of the critical disorder for localization. In contrast, they observed localization in the presence of a quasiperiodic potential. This suggests that a quasiperiodic potential may have a fundamentally different effect on localization compared with disorder. These results suggest that disorder may not induce MBL, but rather a quasi-localized state with very slow sub-diffusive dynamics. However, numerical studies of MBL can be strongly influenced by finite-size effects [24], thereby making it difficult to draw conclusions from numerical results.

There are many open questions regarding many-body localization, including the role of disorder or quasi-periodic disorder in localization [2, 13, 25–28]. Recently, non-ergodic behavior similar to that of disorder-induced MBL has been observed in clean interacting systems with a potential gradient [29–32]. There has been an increase in work studying localization induced by a potential gradient, known as Stark-localization [33–44]. Stark-localization exhibits many of the same properties as disordered systems and provides another tools for exploring the underlying physics of many-body localization. This type of localization is known as Stark localization (or Stark MBL) after the non-interacting Wannier-Stark effect.

1.1.2 Gradient-Induced Single-Particle Localization

In non-interacting lattice systems, potential gradients can induce localization through the Wannier-Stark effect [45]. Wannier-Stark localization and has long been understood in condensed matter systems [46]. In atomic gases, for sufficiently large gradients, atomic wave functions can be localized to a single lattice site.

Stark localization arises solely from energy shifts between sites that are comparable to the lattice bandwidth.

Wannier-Stark localization is usually described in the presence of a linear gradient, but the gradient created by a harmonic trap can have a similar effect. Near the edge of the trap, where the gradients are strong, the harmonic potential can induce Stark-localized states. These states, often referred to as single-particle localized (SPL) states in harmonic systems, have been characterized in References [47–49]. In Reference [47], Ott et al. found that tunneling and interparticle collisions were necessary to allow thermalization of SPL states in a 1D lattice. Reference [48] solved the combined harmonic – lattice potential in one dimension and calculated expansions describing the SPL eigenstates in various regimes. Reference [49] demonstrated a method to measure temperature in gases including SPL states and further characterized the SPL wave functions. The localization of single-particle states in a harmonic trap has been well-documented, but there is less work exploring how interactions affect these states.

1.1.3 Interacting Stark Localization

Single-particle Stark localization is well-understood, but there has been a recent interest in understanding the influence of many-particle quantum effects and inter-particle interactions on Stark localization. Theoretical work has been done to understand the similarities and differences of interacting Stark localized systems compared with previously studied many-body localized systems.

Recent numerical studies have probed Stark localization in interacting-spin [36] and lattice models, including the connection to many-body localization (MBL). A one-dimensional, Bose-Hubbard system [44] has been found to have energy level statistics that show a transition between ergodic behavior for small gradients to behavior consistent with MBL at higher gradients. A system of spinless fermions [37] was shown to exhibit the same behavior. Spinful fermions in a harmonic potential were demonstrated to have localized states away from the center of the harmonic potential [50].

Furthermore, logarithmic entropy growth consistent with MBL was found for an interacting fermionic system with nearest-neighbor interactions in one dimension [33]. Work in higher dimensions has been limited. One numerical study revealed Stark localization in two dimensions for hard-core bosons at higher gradients compared to the one-dimensional case [39].

There have been few experimental observations of interacting Stark localization. A chain of ions with long range spin–spin coupling has displayed a lack of thermalization and slow propagation of correlations in the presence of a linear potential gradient [30]. Interacting Stark localization has also been observed in a chain of transmon superconducting qubits [29] and in a quantum-gas tilted 1D Fermi-Hubbard model [31]. In a two-dimensional (2D) Fermi-Hubbard lattice gas, a large potential gradient along one direction exhibited sub-diffusive behavior and slow (but not fully localized) dynamics [32].

Questions remain about the cause of interacting Stark localization and whether it has the same underlying source as disorder-induced MBL. Studies of interacting Stark localization near the transition point find behavior that matches that of MBL in level statistics and entanglement entropy growth. There is also an argument that localization in the presence of gradients may be better explained by Hilbert space fragmentation as localization is dependent on the initial state, similar to quantum many-body scar states [51], and unlike disorder-induced MBL [42]. Hilbert space fragmentation occurs when local constraints divide the state space of the system into subspaces [42]. Systems with Hilbert space fragmentation also can also show sub-diffusive behavior [52], much like systems near an MBL transition. It remains an open question whether observations of interacting Stark localization are an example of MBL in a clean system, or whether some or all of the observations of interacting Stark localization arise from a fragmented Hilbert space [30, 38].

1.2 Measuring Equilibration in Sub-diffusive Systems

The same effects that can cause localization can also cause slow, sub-diffusive dynamics if the conditions are not sufficient for localization. Whether systems with slow dynamics eventually reach equilibrium is of interest. Most previous work measuring equilibration has used a measure of density imbalance as a proxy for equilibration. These measures describe a local property of the distribution that was initially out of equilibrium. An issue with using a local density imbalance to measure whether a system has equilibrated is that the entire distribution is not sampled. There are many non-equilibrium distributions that could have the same imbalance as an equilibrated system. The imbalance measure is unable to distinguish between these cases.

One type of imbalance measure is the difference between the total density in two halves of a gas. For example, Choi et al. performed a measurement of equilibration in a two-dimensional lattice [53]. They removed half of the gas and watched the imbalance return to zero. In some parameter regimes, they saw the imbalance reach zero, which is consistent with equilibrium. But, the gas could have a non-equilibrium structure that is different from its equilibrium distribution, with equal densities in each half of the gas.

Lüschen et al. also performed an equilibration measurement [54]. In a one-dimensional lattice, they measured relaxation in the presence of coupling to the environment. A charge-density wave was created with particles on every-other site. Equilibration was quantified as equal weight on even and odd sites. This measure of the imbalance between the even and odd sites is zero when the system is in equilibrium, but non-equilibrium structure of the overall distribution could be undetected by this measurement.

Hung et al. measured mass redistribution after turning on a two-dimensional lattice [55]. They measured the deviation of the density profile from equilibrium as a radial average. This measures the entire spatial distribution, but requires exact knowledge of the equilibrium distribution.

Measures used to quantify the equilibration of a quantum gas have thus far been limited in their ability to capture whether the entire density distribution has reached equilibrium. To our knowledge, the only measures sensitive to the entire distribution that have been used thus far require knowledge of the equilibrated distribution. A new measure is needed that captures the entire density distribution without requiring *a priori* knowledge of the exact equilibrated distribution.

In this work, we develop such a measure. We created a technique to measure whether the shape of the entire spatial distribution of the gas is equilibrated for a harmonically trapped gas using Mardia's B statistic. This measure quantifies the Gaussianity of the distribution independently of the size and aspect ratio of the distribution. Since equilibrated thermal distributions in a harmonic trap are Gaussian, Mardia's B is a measure of whether a distribution is in equilibrium.

1.3 Overview of this Work

In order to study the dynamics of an interacting Stark-localized gas, we create a thermal lattice gas of bosons in the presence of a barrier that excludes atoms from the center of the gas. We then remove the barrier and measure the dynamics for up to 10 seconds for a range of lattice depths. By measuring Mardia's B, we observe that the gas reaches thermal equilibrium for all lattice depths measured. We observe sub-diffusive dynamics and time scales for relaxation of hundreds of tunneling times for the largest lattice depths, where more localized states should be present. The lattice physics that determines these dynamics is described in Chapter 2. The experimental apparatus is described in Chapter 3. New tools developed to enable this

spatially resolved measure of the density distribution are described in Chapter 4. Chapter 5 describes the thermal distribution used in this work and the presence of localized states. We describe these results in more detail in Chapter 6. In Chapter 7, we conclude that the observation of relaxation of Stark-localized states suggests a continuum model may be needed to describe the long-time dynamics of interacting quantum systems. We also discuss future work enabled by the tools developed here.

Chapter 2

Lattice Systems

Understanding the observations made in this work first requires understanding the physics of atomic lattice systems. In this chapter, I will first discuss the Hamiltonian for the system and how we calculate the band energies and atomic wave functions. Then, I will describe how we use second quantization and the Bose-Hubbard Model to understand the behavior of the particles in terms of their tunneling and interaction energies. Finally, I will discuss other effects that may not be captured in the Bose-Hubbard Model.

2.1 Lattice Potentials

2.1.1 Lattice Hamiltonian

We study atoms in a three-dimensional lattice with the Hamiltonian

$$H = -\frac{\hbar^2}{2m}\boldsymbol{\nabla}^2 - s\,E_R\left[\cos^2(xk) + \cos^2(yk) + \cos^2(zk)\right], \tag{2.1}$$

where $k = 2\pi/\lambda$ is the wavenumber, λ is the wavelength of the light, and m is the mass of the atom. The strength of the lattice is $s\ E_R$, where $E_R = \frac{(\hbar k)^2}{2m}$ is the recoil energy. The Hamiltonian is separable, meaning the wave function can be treated as a product state of the independent wave functions in each direction of the lattice. We can consider the one-dimensional Hamiltonian

$$H_x = -\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2} - s\,E_R\cos^2(xk). \tag{2.2}$$

Since the potential is periodic, Bloch's Theorem [56] applies; the eigenstates $\psi_q(x)$ of this Hamiltonian can be written as a plane wave times a function with the periodicity of the lattice $u_q(x)$ such that,

$$\psi_q(x) = e^{iqx}u_q(x), \tag{2.3}$$

where q is the quasimomentum. The function $u_q(x)$ is periodic with the lattice spacing and can therefore be written in a Fourier expansion as

$$u_q(x) = \sum_n c_{n,q}e^{in2kx}, \tag{2.4}$$

where n is an integer to enforce the condition that $u_q(x)$ has the periodicity of the lattice. For an infinite lattice, n runs from negative to positive infinity. By combining Equations 2.3 and 2.4, we can write the eigenstates of the Hamiltonian in the form

$$\psi_{n,q}(x) = \sum_{n'} c_{n,n',q} e^{ix(2n'k+q)}. \tag{2.5}$$

Here we have switched the summation index to n' and added the index n to indicate that there can be multiple sets of coefficients $c_{n,n',q}$, one set for each energy level n. Each set of $c_{n,n',q}$ form one of the eigenvectors of the Hamiltonian.

We can use the identity,

$$\cos^2\theta = \frac{1}{4}\left(2 + e^{2i\theta} + e^{-2i\theta}\right), \tag{2.6}$$

to rewrite the Hamiltonian. Summing over n allows us to find the matrix elements of the Hamiltonian as

$$H_q^{n,n'} = E_R(2n + q/k)^2\delta_{n,n'} - \frac{sE_R}{2}\delta_{n,n'} - \frac{sE_R}{4}\delta_{n,n'+1} - \frac{sE_R}{4}\delta_{n,n'-1}. \tag{2.7}$$

The first term on the right side of this equation is found by taking the derivative in the kinetic energy term of the Hamiltonian, and the last three terms are found from the three terms of Equation 2.6. Multiplying Equation 2.5 by one of the exponential terms of Equation 2.6 shifts the index of the sum by one (either positive or negative, depending on the sign in the exponential). Since ψ_n are eigenstates, they are orthonormal,

$$\sum \psi^\dagger_{n,q}\psi_{n',q} = \delta_{n,n'}, \tag{2.8}$$

and all but one of the terms in the sum from Equation 2.5 are eliminated for each term of Equation 2.6, resulting in Equation 2.7 [57]. By limiting the sum to a finite n (i.e., neglecting high energy states) we can numerically calculate the band structure in the lattice.

2.1.2 Band Structure

For each particular lattice depth, the band structure can be calculated to good approximation using numerical methods. We can find the eigenenergies for each quasimomentum by explicitly writing out the matrix defined by Equation 2.7 over a finite range of n. For this calculation, we used $-10 \le n \le 10$. This includes more plane waves than are typical ($-5 \le n \le 5$ is more common). This larger plane wave basis is not computationally prohibitive and avoids small inaccuracies potentiality introduced by truncation. (These inaccuracies would primarily influence the tails of the wave functions and could have an effect on the next-nearest-neighbor tunneling and nearest-neighbor interaction terms calculated in Section 2.2.4.) The matrix can then be diagonalized to find the eigenenergies.

The four lowest energy bands of the resulting band structure can be seen in Figure 2.1. The weaker lattices are, unsurprisingly, closer to the free particle case, and there is a wider range of energies allowed in each energy band. This range of energies is called the bandwidth. For the deeper lattices, each band is flatter, and the bandwidth is smaller. The ground band bandwidth is $0.345\,E_R$ for $s = 4\,E_R$ and $0.010\,E_R$ for $s = 20\,E_R$.

The energy spacing between the ground and first excited bands is large enough relative to the temperature scales of most quantum lattice gas experiments that it is usually sufficient to only consider the ground band. In this work, we generally assume all atoms are in the ground band. In Section 6.4.3, we find that a small

fraction of the particles may be in the first excited band and consider how this population could affect the dynamics of the system.

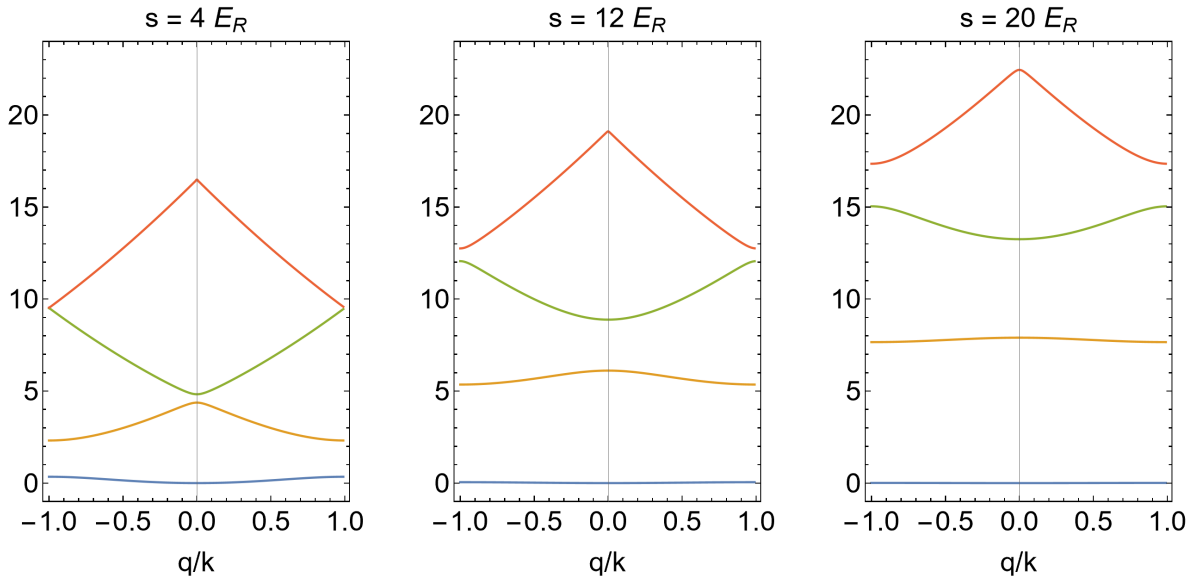

Figure 2.1: The four lowest energy bands for three different lattice depths, spanning the range of lattice depths used in this work. The energy bands have been folded into the $-q/k$ to q/k range.

2.1.3 Wannier States

Diagonalizing the matrix of Equation 2.7 not only gives us the band structure by finding the eigenenergies of the system, it also gives the particle wave functions in the form of the eigenvectors. Specifically, each eigenvector is a list of coefficients $c_{n,n',q}$ in Equation 2.5 for the nth energy levels of the system. Integrating Equation 2.5 with respect to q over a cell of the reciprocal lattice (from $-k$ to k) gives the Wannier functions of the system:

$$w_n(x - x_0) = \frac{1}{W} \int_{-k}^{k} \sum_{n'} c_{n,n',q} e^{i(x-x_0)(2n'k+q)} dq, \tag{2.9}$$

where W is a normalization constant and x_0 is the position of the lattice site at the center of the Wannier state. This convention for defining the Wannier states makes them real-valued. Each Wannier state describes a particle at lattice site x_0. For deep lattices, the tails of the wave function fall off very quickly away from x_0. The complete set of Wannier basis states consists of a Wannier function at each lattice site, that is, x_0 at each lattice position. This set of Wannier functions at each lattice site forms a set of orthogonal basis states that can be used to describe any state of the system.

The Wannier functions can be defined for each band (each n) of the lattice, but here we will focus on the Wannier states of the ground band. We calculate the Wannier function at $x_0 = 0$ with no loss of generality, since the periodic nature of the system requires that the Wannier function is the same at every lattice site given that we assume a uniform lattice.

We can find the ground-band Wannier functions numerically with the coefficients of the ground state eigenvector of the Hamiltonian. To find the Wannier functions, we must integrate over the quasimomentum.

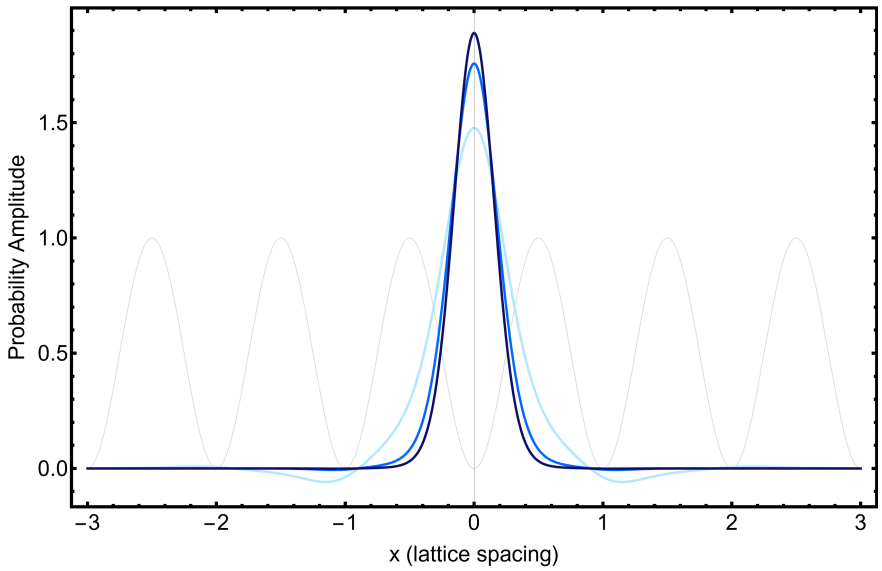

Figure 2.2: The normalized probability amplitude of the Wannier functions for a $4\,E_R$ (light blue), $12\,E_R$ (medium blue), and $20\,E_R$ (dark blue) lattice. The wave function in the lowest lattice depth is much broader than that of the highest lattice depth. The gray oscillating curve represents the lattice potential.

It is convenient to change the integration variable to q/k, which only changes the normalization constant. In this numerical calculation, the integration is approximated by a sum with small steps. A step size of $\Delta q/k = 0.01$ was used in this work. The integral runs from $q/k = -1$ to $q/k = 1$, but it is important to note that only one boundary should be included in the sum (e.g. sum from $q/k = -1$ to $q/k = 1 - \Delta q/k$). Including both boundaries introduces a non-physical oscillation into the Wannier function that persists arbitrarily far from x_0. The magnitude of this oscillation is proportional to $\Delta q/k$. The real Wannier wave function falls off away from x_0, but this non-physical oscillation persists indefinitely. The non-physical weight arbitrarily far from x_0 can introduce significant errors in calculations involving the wave function overlap between non-adjacent lattice sites.

In performing the numerical calculation, it is convenient to find the Wannier functions in terms of the unitless quantity $\mathcal{X} = x/d$. However, care must be taken to include the appropriate factor of d when integrating and taking derivatives of the functions, including correctly accounting for d in their normalization.

We find the shape of the Wannier state for each lattice depth used in the experiment[1]. Figure 2.2 shows the Wannier state for a $4\,E_R$, $12\,E_R$, and $20\,E_R$ lattice, which covers the range of lattice depths used in this work. We can see that the Wannier state for the lower lattice depth is much broader than that of the higher lattice depth. For the deepest lattice depth, particles are more completely confined to an individual lattice site and have less wave function overlap with a particle on an adjacent site.

Figure 2.3 gives another view of the Wannier function with a focus on the tails of the distribution. The majority of the wave function is close to the central lattice site, but for both the $4\,E_R$ and $12\,E_R$ lattices, non-zero weight can be seen two lattice sites away from the center.

We can use the Wannier states to calculate the energy associated with a particle tunneling to an adjacent site as well as the interaction energy of two particles occupying the same site. With these energies, we can discuss the physics of the system in the occupation number representation (i.e., second quantization) where we consider the number of particles on each lattice site.

[1]This is calculated in "Lattice_Wavefunction.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

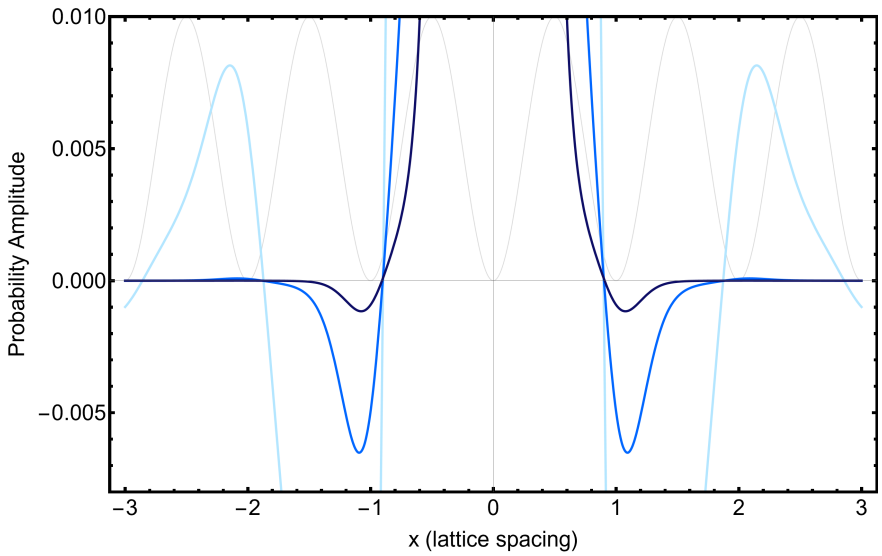

Figure 2.3: The normalized probability amplitude of the Wannier functions for a $4\,E_R$ (light blue), $12\,E_R$ (medium blue), and $20\,E_R$ (dark blue) lattice, scaled to show the tails of the wave functions. The weight more than one lattice site away from the center is small, but visibly non-zero for both the $4\,E_R$ and $12\,E_R$ Wannier functions. The gray oscillating curve represents the lattice potential.

2.2 Bose-Hubbard Model

2.2.1 Second Quantization

A useful framework to understand lattice physics is to use second quantization to rewrite the Hamiltonian. In second quantization, we write the field operator $\hat{\Psi}(\vec{r})$ as the sum of the Wannier states times the boson annihilation operator $\hat{b_i}$ on site i,

$$\hat{\Psi}(\vec{r}) = \sum_i w_i(\vec{r})\hat{b_i}. \tag{2.10}$$

Here we've changed the notation of the Wannier state to refer to each lattice site with a subscript such that $w_i(\vec{r}) = w(\vec{r} - \vec{r_0})$ where $\vec{r_0}$ is the position of lattice site i. We also assume the ground band Wannier states.

We can then write the Hamiltonian in Equation 2.1 as

$$\hat{H_s} = \int \hat{\Psi}^\dagger(\vec{r}) H \hat{\Psi}(\vec{r}) d\vec{r} = \sum_{i,j} \hat{b_j}^\dagger \hat{b_i} \int w_j^*(\vec{r}) \left\{ -\frac{\hbar^2}{2m}\boldsymbol{\nabla}^2 + s\ E_R \left[\cos^2(xk) + \cos^2(yk) + \cos^2(zk)\right] \right\} w_i(\vec{r}) d\vec{r}. \tag{2.11}$$

The Hamiltonian $\hat{H_s}$ represents the single-particle physics of the lattice system. Writing the Hamiltonian this way allows us to find the matrix elements of the Hamiltonian in the basis of particles on individual lattice sites [58].

The Wannier state is separable into each of the lattice directions,

$$w_i(\vec{r}) = w_i(x)w_i(y)w_i(z), \tag{2.12}$$

and the Hamiltonian is also separable,

$$\hat{H} = \hat{H_x} + \hat{H_y} + \hat{H_z}. \tag{2.13}$$

Each direction of the lattice is equivalent, so we can focus on $\hat{H}_x$ without loss of generality:

$$\hat{H}_x = \sum_{i,j} \hat{b_j}^\dagger \hat{b_i} \int w_j^*(x) \left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + s\ E_R \cos^2(xk) \right] w_i(x) dx \int w_j^*(y) w_i(y) dy \int w_j^*(z) w_i(z) dz. \tag{2.14}$$

The Wannier states are orthogonal so

$$w_j^* w_i = \delta_{i,j}, \tag{2.15}$$

and therefore

$$\int w_j^*(y) w_i(y) dy = 0 \tag{2.16}$$

when $i \neq j$. We can therefore eliminate every pair of i and j that have a different position along either the y or z direction. To simplify further, we must make the tight-binding approximation and assume that only tunneling between adjacent lattice sites occurs. We can then eliminate every element of the sum where i and j are not adjacent lattice sites are we are left with

$$\hat{H}_t = \sum_{\langle i,j \rangle} \hat{b_j}^\dagger \hat{b_i} \int w_j^*(x) \left(-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + s\ E_R \cos^2(xk) \right) w_i(x) dx, \tag{2.17}$$

where $\langle i, j \rangle$ represents every pair of lattice sites that is adjacent along the x direction and

$$\hat{H}_\epsilon = \sum_{i} \hat{b_i}^\dagger \hat{b_i} \int w_i^*(x) \left(-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + s\ E_R \cos^2(xk) \right) w_i(x) dx, \tag{2.18}$$

which gives the on-site energy of the terms where $i = j$. The Wannier states are real-valued in this convention, so $w_i^*(x) = w_i(x)$. Since the Wannier states are the same at every lattice site, the integrand is the same for every pair of adjacent (or same-site) lattice sites, and we can rewrite the Hamiltonian:

$$\hat{H}_t = -t \sum_{\langle i,j \rangle} (\hat{b}_i^\dagger \hat{b}_j + h.c), \tag{2.19}$$

where

$$t = -\int \left(\frac{-\hbar^2}{2m}\ w(x+d) \frac{\partial^2 w(x)}{\partial x^2} + w(x+d) s\ E_R \cos^2(xk)\ w(x) \right) dx \tag{2.20}$$

is the nearest-neighbor tunneling energy and d is the lattice spacing. Here $\langle i, j \rangle$ represents every pair of adjacent lattice sites (in any direction) where $i > j$ and the hermitian conjugate $h.c.$ includes the term where $j > i$ for the same values of i and j. Since t is constant for a given lattice depth, this representation provides a convenient way to calculate the behavior of lattice systems. Many of the properties of the lattice dynamics are known in relation to t. For example, the characteristic time of a tunneling event is $\hbar/t$. The bandwidth, the range of available energies in the ground band, is always $4t$ for a one-dimensional lattice.

The on-site energy is

$$\hat{H}_\epsilon = \epsilon \sum_i \hat{b}_i^\dagger \hat{b}_i, \tag{2.21}$$

where

$$\epsilon = \int \left(\frac{-\hbar^2}{2m}\ w(x) \frac{\partial^2 w(x)}{\partial x^2} + s\ E_R \cos^2(xk)\ w(x)^2 \right) dx. \tag{2.22}$$

2.2.2 Interactions in a Lattice

Thus far, we have discussed only single particle effects, but interaction effects can play an important role. To find the interaction energy, we must consider the product state of two particles. The interaction pseudopotential can be approximated by [59]

$$H_U = \frac{4\pi\hbar^2 a_{scat}}{m}\delta_{\vec{r_1},\vec{r_2}}, \tag{2.23}$$

where $a_{scat} = 5.82$ nm is the s-wave scattering length for ^{87}Rb at low magnetic field [60] and $\vec{r_1}$ is the position of particle 1 and $\vec{r_2}$ is the position of particle 2. The delta function imposes that the only non-zero contribution occurs when $\vec{r_1} = \vec{r_2} = \vec{r}$ and the ground-band interaction Hamiltonian is given by [61]

$$\hat{H_U} = \int\int \hat{\Psi}^\dagger(\vec{r_2})\hat{\Psi}^\dagger(\vec{r_1})H_U\hat{\Psi}(\vec{r_1})\hat{\Psi}(\vec{r_2})d\vec{r_1}d\vec{r_2} = \sum_{i,j,k,l}\hat{b_i}^\dagger\hat{b_j}^\dagger\hat{b_k}\hat{b_l}\frac{4\pi\hbar^2 a_{scat}}{m}\int w_i^*(\vec{r})w_j^*(\vec{r})w_k(\vec{r})w_l(\vec{r})d\vec{r}. \tag{2.24}$$

We take the approximation that all terms except the on-site interaction can be neglected. In this case, $i = j = k = l$. To find the interaction energy, we must integrate over all three directions of the lattice. Since the lattice potential is separable and is uniform in all three directions, this is equivalent to taking the integral in one direction to the third power. We find

$$\hat{H_U} = \sum_i \frac{U}{2}\hat{n}_i(\hat{n}_i - 1), \tag{2.25}$$

where $\hat{n}_i = \hat{b}_i^\dagger\hat{b}_i$ gives the number of atoms at site i and U is the on-site interaction energy,

$$U = \frac{4\pi\hbar^2 a_{scat}}{m}\left(\int |w(x)|^4\, dx\right)^3 \tag{2.26}$$

This is an approximation to the interaction energy, as it assume the wave functions are not changed by the interaction. In reality, the repulsive interaction will spread out the wave functions [62, 63].

Since a_{scat} cannot be tuned in our system, U depends only on the lattice depth. We can now write the second-quantization expression of the Bose-Hubbard Model. This formulation of the Hamiltonian makes the tight-binding approximation where the wave functions are assumed to be sufficiently localized that terms beyond nearest-neighbor-tunneling and can be ignored and also ignores all but on-site interaction.

2.2.3 Bose-Hubbard Model

For bosons in a lattice potential, when we take the tight-binding approximation and assume only on-site interaction, we have the Bose-Hubbard model:

$$\hat{H} = -t\sum_{\langle i,j\rangle}(\hat{b}_i^\dagger\hat{b}_j + h.c) + \sum_i \frac{U}{2}\hat{n}_i(\hat{n}_i - 1), \tag{2.27}$$

where t is the (nearest-neighbor) tunneling energy, i and j are lattice indices, $\langle i,j\rangle$ represents a sum over nearest neighbors, $\hat{b}_i^\dagger$ and $\hat{b}_i$ represent the creation and annihilation operators, U is the on-site interaction energy, and $\hat{n}_i = \hat{b}_i^\dagger\hat{b}_i$ gives the number of atoms at site i. The ratio of the tunneling and interaction energies can be adjusted by tuning the lattice potential depth s, which is controlled by the optical power of the lattice

light. This Hamiltonian captures most of the physics of low-energy bosons in a lattice.

By changing the lattice depth, we can explore different regimes of the system. For weaker lattices, a quantum lattice gas will be in a superfluid state, where the particles are distributed across the lattice and can move freely. For stronger lattices, a Mott insulator state can be present. In this state, the particles are localized to individual lattice sites as the higher interaction energy creates a high energy-cost of particles occupying the same site. The phase of the system depends on the lattice depth and density. The range of densities of a gas trapped in a harmonic potential means it is possible for a single atomic gas to have components of both superfluid and Mott insulator states at different parts of the harmonic trap. More details about the phases of a quantum lattice gas can be found in References [3, 64, 65].

This work was done with a gas in a thermal state above the condensation temperature of a superfluid. While the superfluid and Mott insulator states are both ground states of the many-body systems, a thermal gas is an excited many-body state. Work studying thermal quantum gases has been relatively limited.

2.2.4 Calculating Beyond Bose-Hubbard Parameters

Although the low-energy physics of bosons in a lattice potential can be described by the Bose-Hubbard model, there are other processes that have a weaker effect on the behavior of these systems. These effects are ignored in the tight-binding approximation and are irrelevant on short time scales (hundreds of milliseconds). However, these weaker effects could influence the dynamics of the system over the long time scales that were used in this work (10 seconds). We consider the next-nearest-neighbor tunneling term and nearest-neighbor interaction term that could be relevant on long timescales.

The next-nearest-neighbor tunneling is calculated with Equation 2.14. The next-nearest-neighbor tunneling describes tunneling that skips over a lattice site and couples two non-adjacent lattice sites. In this calculation, we find the tunneling between the lattice sites at position $x - d$ and $x + d$. The next-nearest-neighbor tunneling, t_{nnn} is given by

$$t_{nnn} = -\int \left(\frac{-\hbar^2}{2m}\ w^*(x+d) \frac{\partial^2 w(x-d)}{\partial x^2} + w^*(x+d) s\ E_R \cos^2(xk)\ w(x-d) \right) dx. \tag{2.28}$$

The Wannier function $w(x)$ is real in our convention, so $w^*(x) = w(x)$, which simplifies calculation. This term is 7% of t for a $4\,E_R$ lattice and less than 1% of t for lattice depths greater than $12\,E_R$ (Table 2.1). Since this term is a fraction of t, it is likely irrelevant on the timescale of a few tunneling times. However, over hundreds of tunneling times, the effect of this term may be significant.

The other beyond-Hubbard term we considered is the nearest-neighbor interaction, U_{nn}. This can be found by taking the nearest-neighbor pairs of lattice sites in Equation 2.24,

$$U_{nn} = \frac{4\pi\hbar^2 a_{scat}}{m} \int \left(|w(x)|^2\ |w(x+d)|^2 dx \right) \left(\int |w(x)|^4\, dx \right)^2. \tag{2.29}$$

This term is less than 1% of t for all lattice depths used in this work. The contribution from this term is smaller than that of t_{nnn}, but may also be significant on long time scales.

2.2.5 Harmonic Confinement

So far, we have discussed particles in a uniform lattice. We must also have a confining potential to hold the particles in this lattice system. We use a harmonic potential $V_{\text{trap}}(r) = \frac{1}{2} m\omega^2 d^2 r^2$ to trap the particles. We

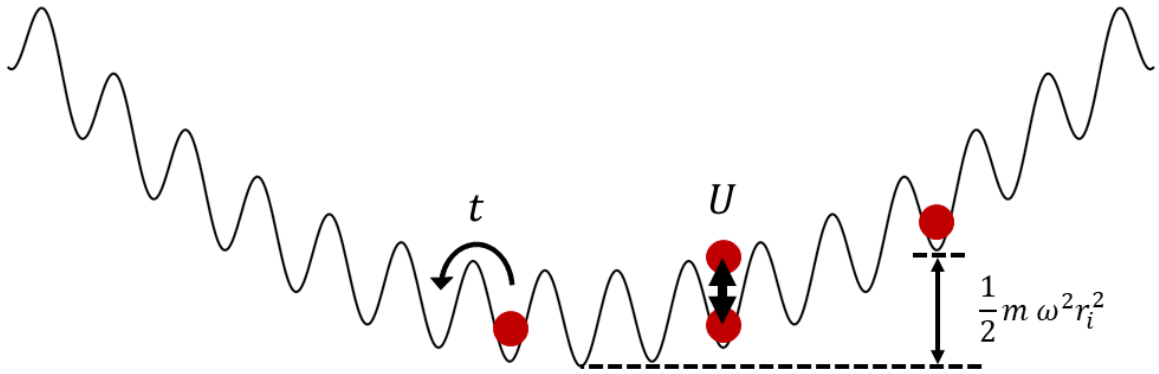

Figure 2.4: An illustration of the Bose-Hubbard Model. The tunneling energy is t, the interaction energy is U, and the harmonic trapping potential is $\frac{1}{2}m\omega^2 r_i^2$.

often give the harmonic potential in terms of Ω, such that

$$V_{\text{trap}}(r) = \frac{1}{2}m\omega^2 d^2 r^2 = \Omega r_i^2, \tag{2.30}$$

where r_i is the distance from the center in units of lattice sites and ω is the angular frequency of the parabolic trapping potential. Adding the potential to the Bose-Hubbard Hamiltonian gives

$$\hat{H} = -t\sum_{\langle i,j\rangle}(\hat{b}_i^\dagger \hat{b}_j + h.c.) + \sum_i \frac{U}{2}\hat{n}_i(\hat{n}_i - 1) + \sum_i \Omega r_i^2 \hat{n}_i. \tag{2.31}$$

This potential can also change the Wannier wave functions and make the values of t and U vary across the lattice. This work takes the approximation that $w(\vec{r})$, t, and U are not influenced by the harmonic confinement. This is a good approximation when the harmonic confinement is approximately constant across a lattice site. Specifically, the change in the confinement is small compared to the lattice depth, $\frac{dV_{\text{trap}}(r)}{dr}d \ll s$ where d is the lattice spacing and s is the lattice depth. This is true for the typical range of experimental parameters. For example, for the weakest lattice depth and trap frequency used in this work, $s = 4\,E_R$ and $\omega = 60.4 \times 2\pi\,\text{Hz}$, the gradient across the 50th lattice site from the center is $0.08\,E_R$. This is farther from the center than is typically sampled by the gas. This gradient per lattice site is approximately 2% of the lattice depth.

2.2.6 Values of Lattice Parameters

Now that we have established the effects of the lattice on the behavior of bosons in a lattice potential, we can consider how these energies vary with the depth of the lattice. Table 2.1 gives the ratio of these energies to the tunneling energy over the range of lattice depths probed in this work. The interaction energy ranges from approximately equal to the tunneling energy at $s = 4\,E_R$ to 209 times larger than the tunneling energy at $20\,E_R$. This is very far into the regime where doublon binding and unbinding is prohibited [66, 67]. The next-nearest-neighbor tunneling, t_{nnn} ranges from approximately one tenth (at $s = 4\,E_R$) to one thousandth (at $s = 20E_R$) of the tunneling energy. Although this effect is much smaller than the nearest-neighbor tunneling, over the hundreds of tunneling times probed in this experiment, the effects could become significant.

s	t	$\hbar/t$	U/t	t_{nnn}/t	U_{nn}/t	Ω/t
E_R	E_R	ms				
4	0.086	0.54	1.3	-0.0706	0.00581	0.009
6	0.051	0.91	3.3	-0.0371	0.00373	0.016
8	0.031	1.49	7.4	-0.0205	0.00234	0.030
10	0.019	2.38	15	-0.0119	0.00147	0.052
12	0.012	3.71	27	-0.0071	0.00094	0.089
14	0.008	5.65	48	-0.0044	0.00062	0.146
16	0.005	8.43	81	-0.0029	0.00041	0.236
18	0.004	12.35	131	-0.0019	0.00028	0.371
20	0.002	17.80	209	-0.0013	0.00019	0.573

Table 2.1: The energy scales of tunneling, interaction, next-nearest-neighbor tunneling, nearest-neighbor interaction, and the harmonic trapping potential for the range of lattice depths used in this work.

The nearest-neighbor interaction U_{nn} is much smaller than ($< 0.6\%$ of) the tunneling energy for every lattice depth, and its effects are expected to be smaller than those of t_{nnn}.

The accuracy of these values depends on the size of the discretization used in the numerical calculation. For this work, the step size of q/k and x/d were both taken to be 0.01. Changing this step size to 0.005 has very little effect on the results. Both t_{nnn} and U_{nn} (as well as t and U) change by less than 0.5%, suggesting this step size is sufficiently small to accurately find the values of these lattice parameters.

The ratio of the harmonic confinement, Ω, to the tunneling energy increases with lattice depth, since t decreases with increasing lattice depth. (The confinement also increase with lattice depth, as explained in Section 3.5.) The gradient of the harmonic potential is $\Omega\left((r+1)^2 - r^2\right)$ at lattice site r away from the center. When the gradient is larger than the bandwidth $4t$, the particles are localized (Section 5.4).

The beyond-Hubbard terms are very small, on the order of a hundred or thousand times smaller than the tunneling energy. The dynamics resulting from of these terms may not be discernible over typical experimental timescales on the order of tens of tunneling times. However, these terms can contribute to the long-time dynamics and influence the stability of localized states. Efforts to explore the extended Bose-Hubbard model, as well as how these effects could influence results seen in this work, are discussed in Section 6.4.4. More experiments that explore long time scales are needed to study the effects of beyond-Bose-Hubbard physics in interacting quantum lattice gases. This work measures the dynamics of localized states over hundreds of tunneling times.

Chapter 3

Experimental Apparatus

The experiment described in this thesis was performed with an ultracold rubidium apparatus that has been explained in prior theses from our group [57, 58, 61, 68, 69]. In this chapter, I give an overview of the core components of the apparatus and explain upgrades to those components. Later, in Chapter 4, I will explain the changes I made to the apparatus to enable new capabilities.

3.1 Overview

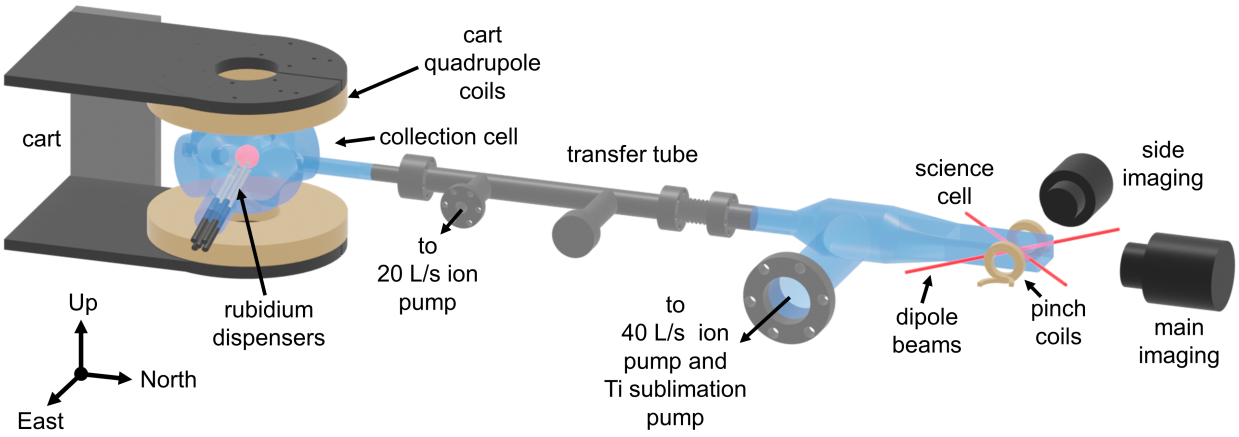

Figure 3.1: A diagram of the experimental apparatus. The experimental run starts with the atoms in the collection cell, where they are trapped and cooled in a magneto-optical trap. Next, they are trapped by the cart quadrupole coils and transported down the transfer tube to the science cell. In the science cell, they are transferred into the quadrupole trap produced by the pinch coils. After RF-forced evaporative cooling, the atoms are transferred into the optical dipole trap formed by two laser beams. At the end of the sequence, absorption images are captured on the camera with probe light that enters the apparatus on the south side of the collection cell and is magnified by a set of lenses (not pictured) after the science cell. An auxiliary side imaging system takes additional images with probe light from the East side of the science cell.

There are many steps needed to create a quantum lattice gas (see Figure 3.1 and Figure 3.2). First, a rubidium dispenser (Section 3.3) releases atoms into the collection cell. These atoms are then trapped and cooled in a magneto optical trap (MOT), as described in Section 3.2.1. From here, the atoms are trapped in

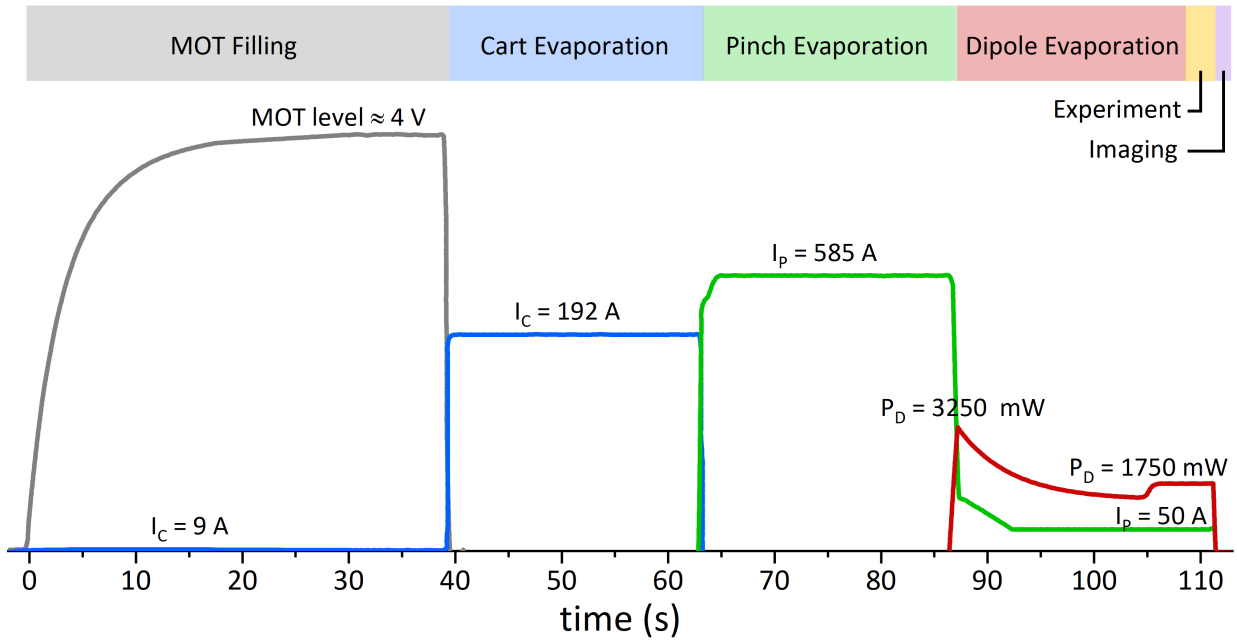

Figure 3.2: Timing diagram showing the steps of the experimental sequence used in this work. The vertical scale is arbitrary. The gray curve shows the MOT level, the fluorescence signal from the MOT, as it fills. The blue curve shows the current in the cart quadrupole coils I_C, which is on at 9 A to provide confinement during the MOT filling stage. The current in the cart coils is much higher during the transfer and cart evaporation stage. The green curve represents the current in the pinch coils I_P, which provide the sole confinement during pinch evaporation and additional confinement during dipole evaporation and the experiment stage. The red curve represents the power in the dipole beams during dipole evaporation and the experiment stage. The final dipole beam power is $P_D = 1750$ mW. All of these potentials are turned off during the final imaging stage.

a magnetic quadrupole field and transferred to the science cell (Section 3.2.2), where they are cooled with radio frequency (RF) forced evaporative cooling (Section 3.2.3).

After RF evaporation, the atoms are transferred to an optical dipole trap and undergo a final stage of evaporative cooling. This stage of cooling is capable of reaching temperatures below 100 nK and can be used to create a Bose-Einstein condensate. The optical dipole trap is the primary confining potential used to trap the atoms during the measurement sequences and is described in detail in Section 3.4.

In addition to the dipole trap, a three-dimensional optical lattice potential is created with three orthogonal counter-propagating 812 nm beams (Section 3.5). This lattice potential enables us to study the Bose-Hubbard model, which is discussed in detail in Chapter 2.

The final step of the experimental sequence is imaging, which is described in Section 3.6. After imaging, the apparatus resets, and the MOT begins filling again to collect atoms for the next run of the apparatus. The complete cycle time is about 90 to 110 seconds.

3.2 Creating an Ultracold Atomic Gas

3.2.1 Magneto-Optical Trap

The first stage of creating an ultracold atomic gas is trapping and cooling atoms in a magneto-optical trap (MOT) in the collection cell. The collection cell contains a rubidium dispenser, which creates a ^{87}Rb vapor

when sufficient current is applied. Atoms in this vapor are trapped in three mutually orthogonal sets of retro-reflected laser beams and a magnetic quadrupole field to create a magneto-optical trap (MOT).

The MOT cools atoms through scattering events with photons from the beams, which are detuned -22 MHz from the D2 cycling transition between the $F = 2$ and the $F' = 3$ states[1] (Figure 3.3). The Doppler shift brings the light closer to resonance for atoms moving towards the beam. This increases the probability of a scattering event with a photon that imparts momentum in the opposite direction of the velocity of the atom, thereby creating an overall cooling force. The light combined with the magnetic quadrupole field also provide a trapping force.

The three MOT beams are provided by three diode lasers injection-locked to a 780.24 nm optical cavity diode laser (ECDL). In 2018, after the data in this work was taken, the home-built ECDL was replaced with a Toptica DL Pro 780 nm laser controlled with a DLC pro controller. An extra laser diode was also purchased at this time. This laser is frequency-stabilized to the $F = 2$ to the $F' = 2$ to $F' = 3$ D2 crossover transition of ^{87}Rb (Figure 3.3) using a polarization spectroscopy locking scheme.

In order to prevent population from building up in a dark state, repumping light is used. Atoms can be off-resonantly excited into the $F' = 2$ state, where they can decay to the $F = 1$ state, which is a dark state to this frequency of light. An additional repumping beam is used to excite atoms from the $F = 1$ state to the $F' = 2$ state, where they are then able to decay back to the $F = 2$ state and continue to participate in the cycling transition (Figure 3.3).

At the time the data in this work was taken, the repumping light was provided by a New Focus Vortex II laser at 780 nm. In late 2021, the repump laser started to fail, requiring increasingly high current to output sufficient laser power (about 12 mW after the isolator). In February 2022, the repump laser was replaced with a distributed feedback fiber-pigtailed laser diode from Thorlabs (part number: LP780-SAD15, serial number: 180420-51) in a Thorlabs mount (LDM9LP). This laser is locked with a beat-note lock relative to the Toptica laser, as described in Reference [57]. The locking scheme was simplified for the new laser, as it only has current control. The feedback loop that had previously been used to control the piezo in the old laser is now used to control the current in the distributed feedback laser. The old current feedback loop is now unused. The new laser can output 19 mW of power, which is consistent with older benchmarks.

3.2.2 Transfer to Science Cell

With the inclusion of the repumping beam, the MOT reaches an equilibrium state with a temperature of less than 100 μK after 20 to 30 seconds. In order to reach the much lower temperatures needed for the experiments, we transfer the atoms from the collection cell to the science cell. The background vapor pressure of ^{87}Rb atoms is much lower in the science cell than in the collection cell. The low vapor pressure is necessary for cooling to the needed temperatures.

The first step in this transfer is compressing the MOT (the CMOT stage). Compressing the gas to a smaller area at the center of the trap reduces the amount of potential energy (and therefore heat) that will be introduced when the magnetic confining potential is added in the following steps. The CMOT stage involves increasing the detuning of the MOT beams which decreases in the scattering rate and therefore decreases the radiation pressure from the scattered photons, allowing the MOT to compress. Next, the gas is cooled further by turning off the quadrupole field and allowing polarization-gradient cooling in a stage of optical molasses.

Now that the gas is compressed, the repumping beam is turned off, and the atoms are optically pumped

[1] Here I follow the convention where F denotes ground hyperfine states (in this case, states in the $5S_{1/2}$ orbital) and F' denotes excited hyperfine states (in this case, states in the $5P_{3/2}$ orbital).

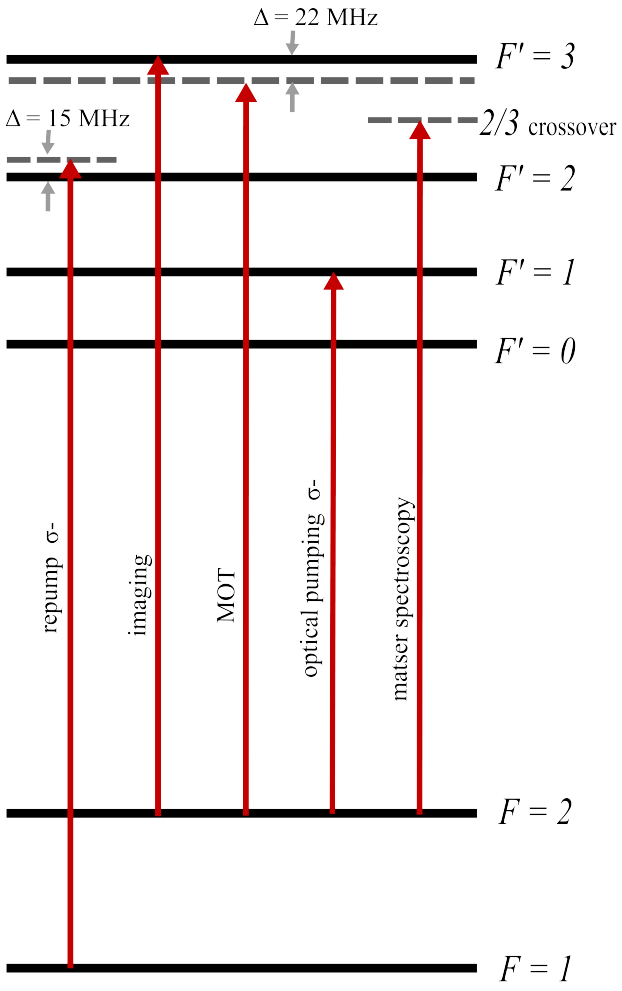

Figure 3.3: The level diagram showing the atomic transitions used in creating an atomic gas. The light used to drive these transitions is locked relative to the master spectroscopy signal.

for 1 ms with 2.5 mW of power into the magnetically trappable $F = 1$, $m_F = -1$ ground state. The atoms are transferred to this state by an optical pumping laser beam that is retro-reflected through the collection cell. Finally, the atoms are trapped in the quadrupole field created by ramping the current in the quadrupole coils up to 192 A, at which point the MOT beams are turned off. Since the quadrupole coils are attached to a cart on a precision, long-travel ball screw stage, the atoms can then be moved down the transfer tube to the science cell by simply moving the cart.

3.2.3 RF-induced Evaporative Cooling

In the science cell, the atoms experience several stages of cooling. First, the atoms are cooled with radio frequency (RF) forced evaporation in the cart quadrupole trap. An 80 MHz RF field is applied to the atoms, which allows the highest energy atoms at the edge of the magnetic trap to transition to an untrapped state. The remaining, lower energy atoms then re-equilibrate to a lower temperature through collisions. The RF frequency is then ramped to lower frequencies, selectively allowing the atoms in the high-energy tail of the

thermal distribution to escape, which progressively lowers the temperature. After an initial stage of RF evaporative cooling, the atoms are transferred to a different, more tightly confining quadrupole field that we call the "pinch trap" created by a set of coils with their axis in the East-West direction. This name is a remnant from when the pinch coils were used with their current in the same direction as each other as part of a Ioffe-Pritchard trap. In the pinch trap, the atoms undergo another stage of RF-evaporation to a final frequency of 3.5 MHz. For more details on these stages of RF evaporative cooling, see Reference [57].

After pinch evaporation, the cart is moved back part of the way towards the collection cell — just far enough to clear the laser beam paths used in the science cell. This is a change from earlier iterations of the apparatus; previously, the cart was moved all the way back to the collection cell at this point in the sequence. This change was made to facilitate the selection of planes of atoms along the direction of the transfer tube by using the magnetic field gradient generated by the cart quadrupole coils far from their center. This field gradient is nearly entirely along the direction of the transfer tube. However, the field gradient proved insufficient for useful plane selection, and this technique is not currently being used.

3.2.4 Atom Number in the Pinch Trap

The number of atoms at this stage of the experimental sequence (after evaporation in the pinch trap) is a useful measure of system performance. When the system is working very well, atom number in the pinch trap is optimally 12 million. We normally operate with atom number from 7 to 9 million, as beam alignments and other factors that affect atom number drift away from their ideal positions between optimizations. Most of the recent work on this apparatus has used relatively small gases, meaning obtaining large atom number has been unnecessary.

Atom numbers below about 6 million usually correlate with high instability in the atom number. A standard deviation of approximately 5% in the atom number can be expected when the apparatus is functioning well. When the atom number is highly variable from shot to shot (i.e., standard deviation greater than 10%), performing measurements becomes difficult. Without consistency in atom number, the variability of the conditions of the atomic gas are too high to yield useful data. We have developed multiple ways of improving atom number stability for this apparatus, which include improving laser stabilization and optimizing the evaporation sequence. One source of atom number instability that has not been previously documented for this experiment is the rubidium dispenser current.

3.3 Rubidium Dispenser

The ^{87}Rb vapor in the collection cell is produced by an alkali metal dispenser manufactured by SAES. The dispenser is a nichrome container filled with powdered rubidium chromate and a reducing agent [70] (Figure 3.4). Heating the dispenser by applying a current initiates a reduction reaction and releases rubidium vapor into the collection cell. This near room-temperature vapor serves as a source of atoms for the MOT. The magnitude of the current in the dispenser affects the vapor pressure in the vacuum system and therefore impacts the atoms number throughout the experimental sequence.

In this section, I discuss a few things we have learned about the relationship between dispenser current and atom number. First, I describe the process of switching to a new dispenser. Then, I explain the procedure for finding an ideal current for the new dispenser. This proved challenging do to the long timescales associated with changing the dispenser current. We found that surprisingly small changes in current (about 3.6 A

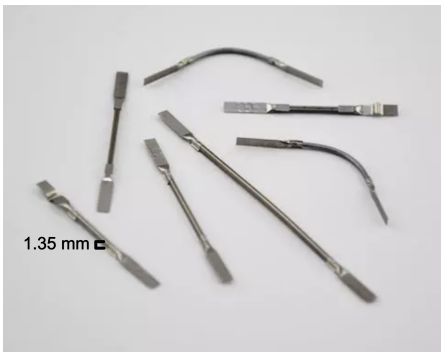

Figure 3.4: An image of alkali metal dispensers from saesgetters.com. The rubidium for this experiment is provided by a dispenser in the collection cell.

compared to the 3.3 A we typically work at) created enough vapor pressure to pollute the science cell and decrease the atom number in the pinch trap.

3.3.1 Switching Dispensers

Since building this experiment in 2003, the dispenser had been run at 3.3 A and the current was rarely adjusted. In fact, the apparatus had been using the same dispenser for approximately 20 years. To our knowledge, this is uncommon; similar experiments usually deplete dispensers in 5–7 years. So, in April 2019, when we observed that the atom number was lower than usual and had ruled out other problems, we decided it was finally time to switch to a new dispenser.

Thanks to clever planning, physically switching the dispenser was straightforward. When this apparatus was built, four dispensers were included in the collection cell, so that when a dispenser was depleted, we could switch to a new dispenser without breaking vacuum. Since the leads to the dispensers are outside the collection cell, we simply moved the connections to the power supply from the first dispenser to the second dispenser.

From there, we slowly increased the current of the power supply and monitored the pressure at the ion pumps to ensure we were not overwhelming the pumps. Using a new dispenser can temporarily cause abnormally high pressure from burning off pollutants or extra rubidium that may have adsorbed to the new dispenser in the time it was not being used. The maximum current reached was 5.2 A, at which point the dispenser was glowing a dull orange. The maximum pressure of 5×10^{-8} torr was seen slightly earlier in the process, at 4 A. More details of this process can be found on pages 84-86 of lab notebook number 21.

The following day, the pressure gauges read low pressure ($<4 \times 10^{-10}$ torr) and the dispenser was turned on to 3.3 A. At this point, we saw that the lifetime in the science cell after pinch evaporation was (8.5 ± 0.8) seconds, consistent with previous measurements. This lifetime is likely limited by Majorana losses — relatively cold gases in a magnetic quadrupole trap are susceptible to atom loss from the center of the trap, where atoms crossing through the zero-field center can make spin-changing transitions into untrapped states.

Although the lifetime was reasonable, the atom number was much lower than expected, only about 6 million atoms. After a few mostly unsuccessful weeks of trying to improve atom number by optimizing the laser beam alignments and evaporation sequence, we tried to activate the dispenser again. This time, we reached a current of 7.5 A and ran the Titanium Sublimation pump filament intermittently during and after this process. A record of this process is described on pages 106-108 of lab notebook #21. The atom number

was still less than ideal after this process, but we were able to reach sufficient atom number to proceed with the next steps for the current project. We concluded that the new dispenser had similar behavior to the old dispenser and continued to use the new dispenser.

3.3.2 Optimizing Dispenser Current

One of the reasons we had difficulty improving the atom number after changing to a new dispenser was that we had not understood that effects of changing the dispenser current appear over tens of minutes. After more investigation, we found that, although the atom number at a high dispenser current may be relatively good shortly after turning on the dispenser, for currents above 3.6 A, the atom number would quickly begin to fall.[2] This is likely because high current creates a vapor pressure that pollutes the science cell vacuum. This atom loss is difficult to measure directly, as the vacuum-limited lifetime is likely longer than the lifetime of the quadrupole trap.

Indirect Measures of Vapor Pressure

We used indirect measures of vapor pressure to optimize the current in the dispenser. Our best indicators of the vapor pressure in the collection cell comes from the MOT fluorescence. Photons scattered by the MOT are imaged onto a photodetector and a photocurrent is read out as a voltage using a transimpedance amplifier. This voltage is proportional to the number of atoms in the MOT. The equilibrium value of the MOT fluorescence, which we call the MOT level, is a measure of the equilibrium number of atoms in the MOT. In equilibrium, atoms are ejected from the MOT at the same rate they are trapped. Another useful measure of the vapor pressure is the time it takes the MOT to fill at the start of each run of the apparatus. We define the fill time τ as the time constant of an exponential fit to the MOT fluorescence signal, starting from when the MOT beams are turned on and the MOT begins filling.

In order to record these measures of the MOT fluorescence more easily, I automated the recording process. Previously, we manually measured the MOT fluorescence signal on an oscilloscope, recorded the equilibrium value, and determined the time to reach $1 - 1/e$ of the equilibrium value run-by-run. This procedure greatly limited how much data we could realistically collect. I used an Arduino to read the MOT fluorescence level into Matlab and automatically recorded the MOT filling trace for every run using the new imaging software I implemented (Section 4.3). Starting in August 2020, the MOT level and fill time for every run of the apparatus can now be easily extracted by fitting the data.

Monitoring these measures of MOT fluorescence and the atom number in the pinch trap N over time is useful to find the ideal current for the dispenser. High currents cause atom number to decays over multiple runs and low currents cause long fill times (which are undesirable because they require long cycle times of the experiment and limit the rate of taking data). Below are a few examples of how these parameters behave under different conditions of the dispenser. These examples are from late 2020, when first understood the slow effects of changing the dispenser current and optimized the current accordingly.

Measuring the Relationship Between Dispenser Current, Atom Number, and Time

In order to better understand the relationship between the dispenser behavior and atom number, we first measured the effects of turning on and off the dispenser. Figure 3.5 shows the atom number N, MOT level,

[2] This current is well-below the manufacturer's listed average current for the start of evaporation, 5.3 A. These dispensers are designed to deposit films of alkali metals.

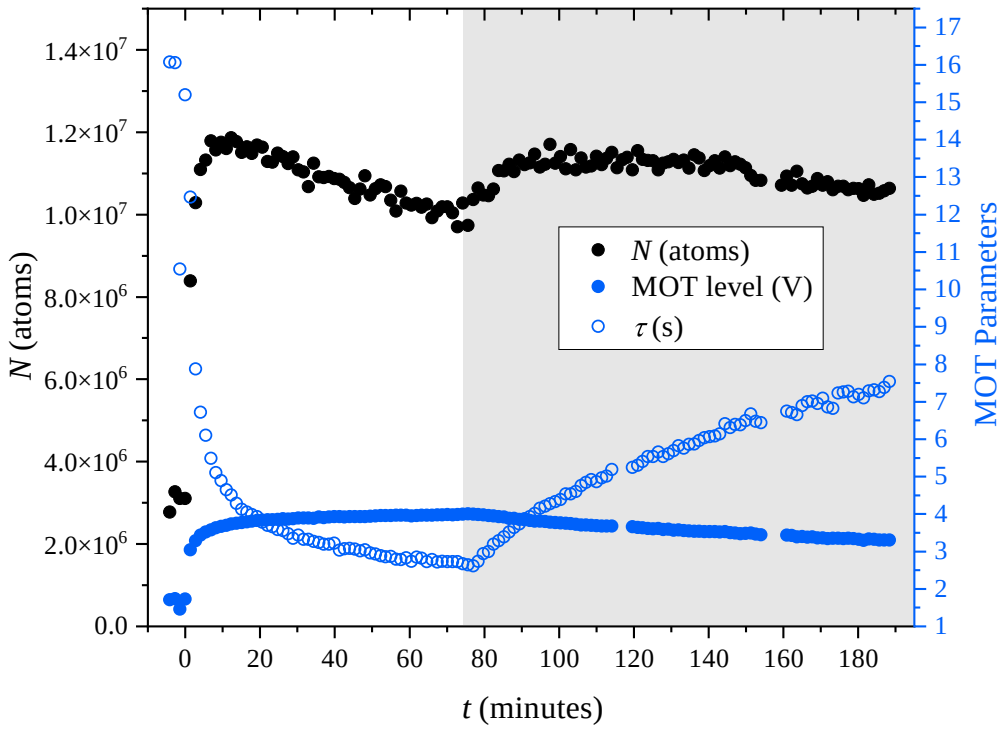

Figure 3.5: The response of atom number and MOT parameters over time to turning on and off the dispenser. The dispenser was turned on to a high current (3.65 A) at $t = 0$ and off at $t = 74$ minutes, represented by the gray region of the plot. The atom number N (black circles) increased over the first 10 minutes after turning on the dispenser, but then began to fall until the dispenser was turned off. In response, the atom number increased over about 25 minutes and then began to fall again. The MOT filling time τ (open blue circles) responded as expected; it decreased when the dispenser was on and increased when the dispenser was off. The MOT level (filled blue circles) also behaved as expected — increasing when the dispenser is on and decreasing when it is off. These data were taken on August 19, 2020.

and fill time τ for a three-hour period over which the dispenser is turned on and off. We leave the dispenser off when the apparatus is not running; at the start of this measurement, the dispenser had been off for approximately 20 hours. We turned the dispenser on (to 3.65 A) at $t = 0$ and immediately saw the atom number increase and the fill time τ drop, as expected. The rubidium vapor pressure was very low before the dispenser was turned on but quickly increased, allowing the MOT to fill. About 15 minutes after the dispenser was turned on, the atom number N began to fall, likely as the vapor pressure became high enough to pollute the science cell. The dispenser was turned off at $t = 74$ minutes, which (presumably) caused the vapor pressure to fall. At this point, the atom number began to increase and the fill time decreased. Approximately 20 minutes after turning off the dispenser, the atom number began to fall again, presumably once the rubidium vapor pressure was low enough to no longer pollute the science cell.

The fast drop in atom number that is seen when the dispenser is on at 3.65 A, approximately 20% over an hour, is nearly impossible to work around. We need a relatively constant atom number in order to obtain data. This current is clearly higher than ideal. The atom number change that is seen after the dispenser is

turned off might be sufficiently slow to work in this regime. However, the rapidly increasing fill time would quickly become a problem, causing the cycle time between shots of the apparatus to become very long.

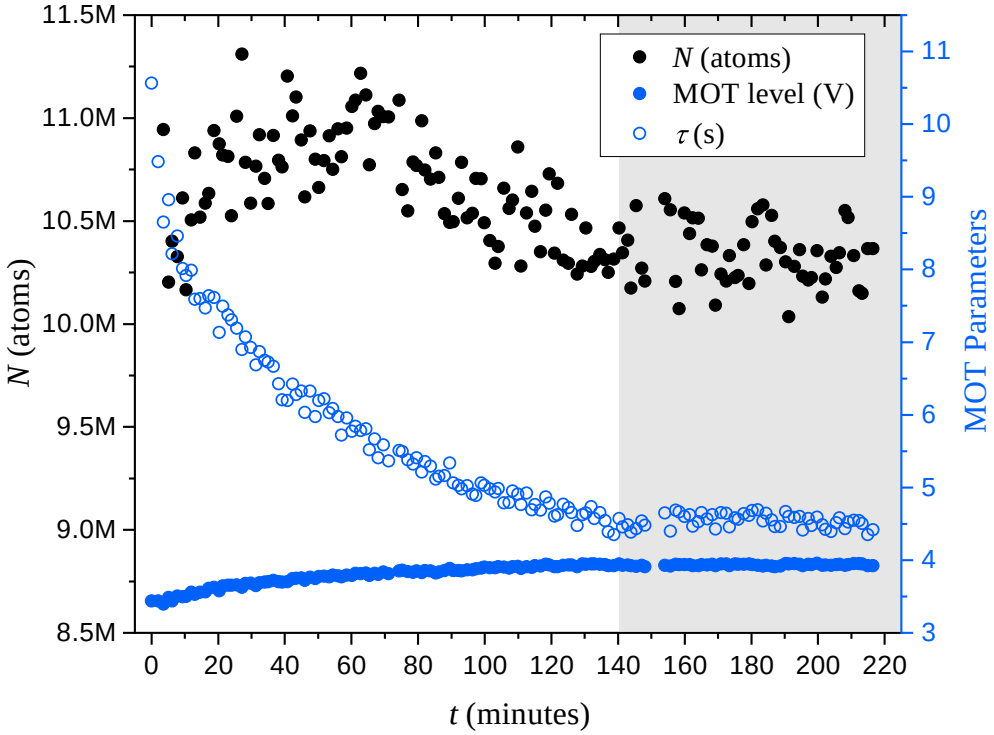

Figure 3.6: The response of atom number and MOT parameters to a change in dispenser current. The dispenser was turned on to 3.55 A at $t \approx -30$ minutes. Between $t = 0$ minutes and $t = 60$ minutes, the atom number N (black circles) was rising. After $t = 60$ minutes, the atom number began to fall. Over this time, the fill time τ (open blue circles) was decreasing and the MOT level (filled blue circles) was increasing as the vapor pressure increased. When the current was turned down to 3.43 A at $t = 140$ minutes (gray region), all measured parameters stabilized. These data were taken on September 5, 2020.

Figure 3.6 shows the process of finding an ideal dispenser current. The dispenser was turned on to 3.55 A at $t \approx -30$ minutes. At $t = 0$, the atom number was still increasing, but by $t = 60$ minutes, the atom number started to fall, indicating the dispenser current was too high. Changing the current to 3.43 A at $t = 140$ minutes (gray region of the plot) resulted in a constant atom number and fill time. This suggests 3.43 A was likely an ideal current for the apparatus at this time. The ideal current seems to drift over time, and the dispenser current should be adjusted occasionally. It may also be reasonable to use different dispenser currents for different types of work with the apparatus. Specifically, when stability is highly valued, the dispenser current could be lower, and when speed is more important, the current could be higher.

3.3.3 Possible Modes of Operation

It would be possible to use a more complicated scheme to control the dispenser current to optimize for both speed and stability. At present, our mode of operation is very simple. We turn the dispenser on when we start

the apparatus for the day and let it warm up for 30 minutes to an hour, until the atom number stabilizes. The current is not changed as part of normal operation. We periodically adjust the current to find the best balance between faster fill time and atom number stability. At the time of this writing, it is set to 3.13 A.

This method could be improved by creating an automated process to adjust the dispenser current based on recent atom number or MOT level data. Another, simpler improvement would be to turn on the dispenser to a relatively high current at the start of the day, in order to quickly increase the vapor pressure and decrease the necessary warm up time, before turning the current back down to improve stability. Although improving the number stability after evaporation in the pinch trap may be beneficial, there are other sources of number noise that limit the usefulness of improving number stability at this stage beyond a certain threshold. In particular, the stability of the next stage of the sequence, the optical dipole trap, is also a concern.

3.4 Optical Dipole Trap

The final confining potential used in the evaporation sequence, which is also the trap used to confine the atoms during the experimental sequence, is an optical dipole trap. The optical dipole trap is created with a 1064 nm IPG fiber laser (YLR-20-1064-LP) with a bandwidth of 0.5 nm. The laser beam has a waist of $(155 \pm 15)\,\mu$m and enters the science cell in the horizontal plane at an approximately 45-degree angle to the face of the rectangular science cell (Figure 3.1). This forward dipole beam can be used in conjunction with the pinch quadrupole coils to confine the atoms; the beam provides confinement in the vertical and one of the horizontal directions and the pinch coils provide confinement along the direction of the dipole beam. We call this the "hybrid dipole trap." [71]

More commonly, we use a trap called the "cross dipole trap." For this trap, an additional cross dipole beam is added that passes through the atomic cloud orthogonal to the forward beam, also in the horizontal plane. This cross beam is created by reflecting the forward beam and has a waist of $(135 \pm 15)\,\mu$m. The polarization of the forward and cross beams is orthogonal to avoid heating that can occur when the two beams have the same polarization. This pair of dipole beams can provide sufficient confinement to support the atoms against gravity without any additional magnetic fields. In December 2022, this power was measured to be 2050 mW. The power needed to support against gravity has increased in the last several years. In January 2019, 1880 mW of power was sufficient to support against gravity and approximately 1700 mW was seen in much older measurements. We do not know why the strength of the dipole trap appears to be decreasing, although degradation of the laser or damage to the optics in the beam path could be to blame.

The trap used in this work is a combination of the hybrid and the cross trap. We use both the forward and the cross dipole beams and additional weak magnetic confinement from the pinch quadrupole coils (on at 50 A). This trap was found to have improved stability over the pure cross trap. More details about this trap can be found in References [57, 69]. We can estimate the confinement of the atoms in this trap in two ways: by calculating the potentials created by the beam powers and magnetic fields used to create the trap (Sections 3.4.2 and 3.4.3), or by measuring the frequency of oscillations in the trap (Section 3.4.4).

3.4.1 Evaporation in the Dipole Trap

After the atomic cloud undergoes RF evaporation in the pinch quadrupole coils, it is transferred into the optical dipole trap by turning on the dipole beams to a high power (often near 5 W, although up to 15 W is possible) and then ramping down the current in the pinch coils from their initial value of 585 A. For this experiment, the pinch current was ramped to 50 A and left on at this low value to provide extra confinement

in addition to the confinement from the optical dipole trap. This additional magnetic confinement increases the atom number stability.

After transferring into the dipole trap, the final stage of evaporation is performed by exponentially ramping down the beam power over several seconds. Similar to RF evaporation, this allows the highest energy atoms to escape the trap, while the remaining atoms have time to thermalize to a new, lower temperature. By varying the parameters of this ramp (the total time, the time constant, the initial power, and the final power), we can control the final atom number and the temperature of the gas.

After evaporation, the dipole power is increased slightly using a sigmoid ramp, usually by at least 50 mW. This recompression of the trap increases the atom number stability. Since the thermal distribution is set by the lowest dipole power reached, the highest energy atoms are just barely trapped by the potential at that power. After increasing the confinement, the highest energy atoms have a lower energy than the edge of the trapping potential and are therefore less likely to escape the trap. Besides increasing stability, recompressing allows us to work at a consistent trap while also having some control over temperature. We always recompress to the same value for every run of a measurement in order to have a consistent trap from shot to shot. Although the final power is always the same, we may need to evaporate to different minimum powers in order to reproduce the same temperature and atom number given drifts in the conditions of the apparatus.

For this experiment, the final, recompressed value of the dipole power was 1750 mW. The specific dipole evaporation sequence used for this experiment can be seen in Figure 3.2. The dipole trap was initially turned on to approximately 3.2 W. This was varied over a range of a few hundred milliwatts to keep the atom number constant. The dipole power was then exponentially ramped down over 17 s with a time constant of 5 s. The final power of this ramp was approximately 1320 mW. This power was varied by a few tens of milliwatts to control the temperature. The dipole beam was then sigmoidally ramped up to 1750 mW over 2 seconds and held at this power for an additional 1 second to dissipate any center-of-mass motion of the gas that may have been introduced by the changing potential.

3.4.2 AC Stark Shift

The confining potential of the dipole beams comes from the AC Stark shift. The electric field of the light can induce an electric dipole moment in an atom and cause a small shift in the energy levels of an atomic transition. This effect is stronger for light with a wavelength closer to the resonance of the transition. For light red-detuned from the atomic transition, this shift lowers the energy of the ground state and atoms are attracted to the high-intensity region of the beam.

For ^{87}Rb, there are two major atomic transitions resonant with near-infrared light, the D1 (795 nm) and the D2 (780 nm) lines. The D1 line is the transition from the $5S_{1/2}$ to the $5P_{1/2}$ state, and the D2 line is the transition from the $5S_{1/2}$ to the $5P_{3/2}$ state. The potential from the AC Stark shift is determined by the linewidths Γ_{D1} and Γ_{D2} and the angular frequencies ω_{D1} and ω_{D2} of the transitions. As given in Reference [72], an approximation of the potential from both the $D1$ and $D2$ lines due to the AC Stark shift is

$$U_{\text{dip}}\left(\mathbf{r}\right) = \frac{\pi c^2}{2} I(\mathbf{r}) \left(\frac{\Gamma_{D2}}{\omega_{D2}^3} \left(\frac{2 + \mathcal{P} g_F m_F}{\omega_{D2} - \omega} + \frac{2 + \mathcal{P} g_F m_F}{\omega_{D2} + \omega} \right) + \frac{\Gamma_{D1}}{\omega_{D1}^3} \left(\frac{1 - \mathcal{P} g_F m_F}{\omega_{D1} - \omega} + \frac{1 - \mathcal{P} g_F m_F}{\omega_{D1} + \omega} \right) \right). \quad (3.1)$$

The angular frequency of the light is ω and the intensity of the light at position $\mathbf{r}$ is $I(\mathbf{r})$. The polarization is represented by $\mathcal{P} = 0$ for π light and $\mathcal{P} = \pm 1$ for $\sigma^{\pm}$ light. The magnetic quantum number is m_F, and

g_F is the Landé factor. This equation is valid when the detuning are large compared to the excited-state hyperfine splitting and the different energies of the hyperfine structure can be ignored.

Note that an approximation of the counter-rotating term is included. The counter-rotating term is resonant at $\omega = -\omega_0$ and has a denominator that is the sum of the frequencies. For small detunings, $\Delta_{\text{counter-rotating}} = \omega_0 + \omega$ results in a term with a denominator that is much larger than that of the co-rotating term and the counter rotating term is negligible. However, the 1064 nm dipole trap laser (with a frequency near 282 THz) is relatively far from resonance, which is at approximately 380 THz. At these frequencies, the counter-rotating term contributes more than 10% of the total potential.

The dipole trap used in our apparatus is created with linearly polarized light, so there are equal components of σ^+ and σ^- light. This means that there will be equal contributions with $\mathcal{P} = 1$ and $\mathcal{P} = -1$ so the polarization terms cancel, and the potential can be calculated with a simplified equation,

$$U_{\text{dip}}(\mathbf{r}) = \frac{\pi c^2}{2} I(\mathbf{r}) \left(\frac{2\Gamma_{D2}}{\omega_{D2}^3} \left(\frac{1}{\omega_{D2} - \omega} + \frac{1}{\omega_{D2} + \omega} \right) + \frac{\Gamma_{D1}}{\omega_{D1}^3} \left(\frac{1}{\omega_{D1} - \omega} + \frac{1}{\omega_{D1} + \omega} \right) \right). \tag{3.2}$$

3.4.3 The Dipole Trap Potential

The dipole trap can be modeled by considering the potential from the AC Stark shift due to both dipole beams and the additional magnetic confinement of the quadrupole field of the pinch coils. For this experiment, the dipole power was 1750 mW and the magnetic field gradient was 23 G/cm in the axial direction (created with 50 A of current in the pinch coils). Calculating the resulting the potential[3] and including the effect of gravity gives a trap depth of $1.9\,\mu\text{K} \times \text{k}_\text{B}$. Slices of the potentials at the trap minimum can be seen in Figure 3.7. In order to create a thermal gas, this trap is shallower than the typical trapping potential used for this apparatus.

3.4.4 Trap Frequency

Measuring Trap Frequency

Modeling the trapping potential gives us an approximation of the confinement the atoms feel, but uncertainty in the beam waists, beam powers, strength of the magnetic gradient, or relative positions of the component potentials can all cause inaccuracies in the calculation of the overall potential. In order to better characterize the trapping potential, we use the atoms to measure it directly.

The center of mass of a trapped gas in a harmonic potential will oscillate at the frequency of the trap after an applied impulse. For a three-dimensional potential, there are three independent trap frequencies, one along each of the principal axes. The principal axes are mutually orthogonal and serve a similar function as the axes of an ellipsoid.

We can measure the trap frequencies by creating a small Bose-Einstein condensate (BEC) of atoms and exciting motion by applying a "kick", or impulse. The gas then oscillates in the trap at the trapping frequencies. To measure this oscillation, we release the gas from the trap after various hold times and let the gas expand. This expansion projects the momentum distribution of the atoms into real space and by measuring the position of the atoms after a fixed time of flight (TOF), we can measure the oscillation of the atoms in the trap. Fitting these oscillations to combinations of oscillatory functions allows us to measure the frequencies of the trap.

[3]This potential was calculated in the Mathematica notebook "Size of trap units.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`

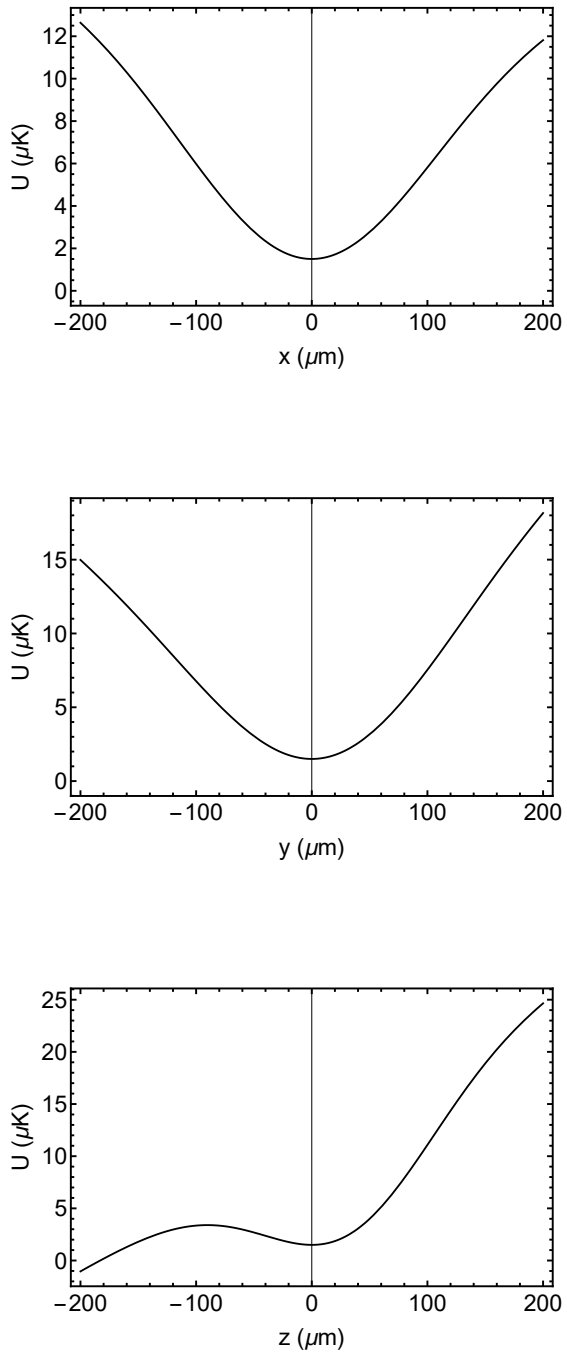

Figure 3.7: The trapping potential at the trap minimum along the x, y, and z directions. The trap is created by a combination of two crossed optical dipole beams in the horizontal plane and a magnetic quadrupole field. The gravity is also included in the potential, which creates a tilt in the potential along the z direction and limits the trap depth, as can be seen in the last plot. The center of the pinch trap is offset from the center of the dipole trap along the z direction. This creates a shallower trap, which is useful for creating thermal gas with a large spatial extent.

As the principal axes of the trap are not necessarily aligned with the imaging axes, care must be taken in resolving three separate frequencies. One tool to resolve the trap frequencies is to measure the motion in every direction simultaneously. An additional tool is to take multiple measurements, each with a kick in a different direction.

There are multiple methods of kicking the atoms, which excite motion in different directions. An ideal trap frequency measurement includes three sets of data, one for each of the principal axes where the gas is kicked along that direction. This is not usually possible, since finding a kick that aligns to each of the principal axes is not practical.

A "DDS kick" is performed by changing the frequency of the acoustic optical modulator (AOM) used to control the dipole beam by a few MHz for a few ms. Since the angle of the dipole beam exiting the AOM depends on the frequency, this changes the center of the dipole trapping potential, pulling the atoms in that direction. The change is primarily along the North-South direction, as the angle of the forward, and therefore also the cross, beams change in the horizontal plane. An "AG kick" pulses on the anti-gravity coil and kicks the atoms primarily in the vertical direction. The anti-gravity coil is a below the science cell and provides an upward force on the atoms which can be used to offset the force of gravity. The pinch coils can also be manipulated to induce motion in other directions. The exact movement of the atoms in response to a kick also depends on the other confining forces. Trial-and-error is often necessary to find a kick that moves the atoms in a particular direction.

Most commonly, we take a set of data for two or three kicks in different directions and excite motion along more than one principal axis in each set of data. We then do a global fit to a sum of two or three sine functions, fitting the data from every direction in a given set of data together to the same set of frequencies. The phase associated with each frequency is also fixed to be the same over each data set. Additionally, a given frequency may not be meaningfully present in every data set, in which case the amplitude for that frequency in that data set should be fixed to zero. This is evident when the uncertainty in the amplitude is larger than the value of the amplitude returned by the fit.

This kind of fitting is tedious and sensitive to the starting guesses. Without good starting guesses, the fitter often does not converge. Often a good strategy is to find a kick that excites motion primarily along just one of the principal axes of the trap. Then fit that data to a single sine wave. Next, take data with a kick in another direction, and use the first frequency as a fixed frequency in a two or three sine fit of the next set of data. It's best practice to get a set of data that shows the motion of the atoms in all three directions.

Trap Frequencies for this Work

The frequencies of the trap used in this work (1750 mW of dipole power and 50 A of pinch current) were found to be (73.9 ± 0.3) Hz in the near-vertical direction and (48.2 ± 0.2) Hz and (45.7 ± 0.3) Hz, both approximately in the horizontal plane. The uncertainty represents the uncertainties of the fit of a global fit to the sum of three sine functions of data taken on December 1, 2016. The uncertainties are an under-estimate. Drift in the position of the beams and other factors can cause changes in the trapping frequencies. Data taken in the same trap on December 8, 2016 had trap frequencies of (75.5 ± 1.0) Hz, (49.5 ± 0.2) Hz, and (46.3 ± 1.7) Hz. Given the drift over time and potential systematic error in the fitting method, the trap frequencies are known within 2 Hz.

The fitting equation used was

$$y(x) = y_0 + A_1 \sin(2\pi(x - x_1)f_1/1000) + A_2 \sin(2\pi(x - x_2)f_2/1000) + A_3 \sin(2\pi(x - x_3)f_3/1000), \quad (3.3)$$

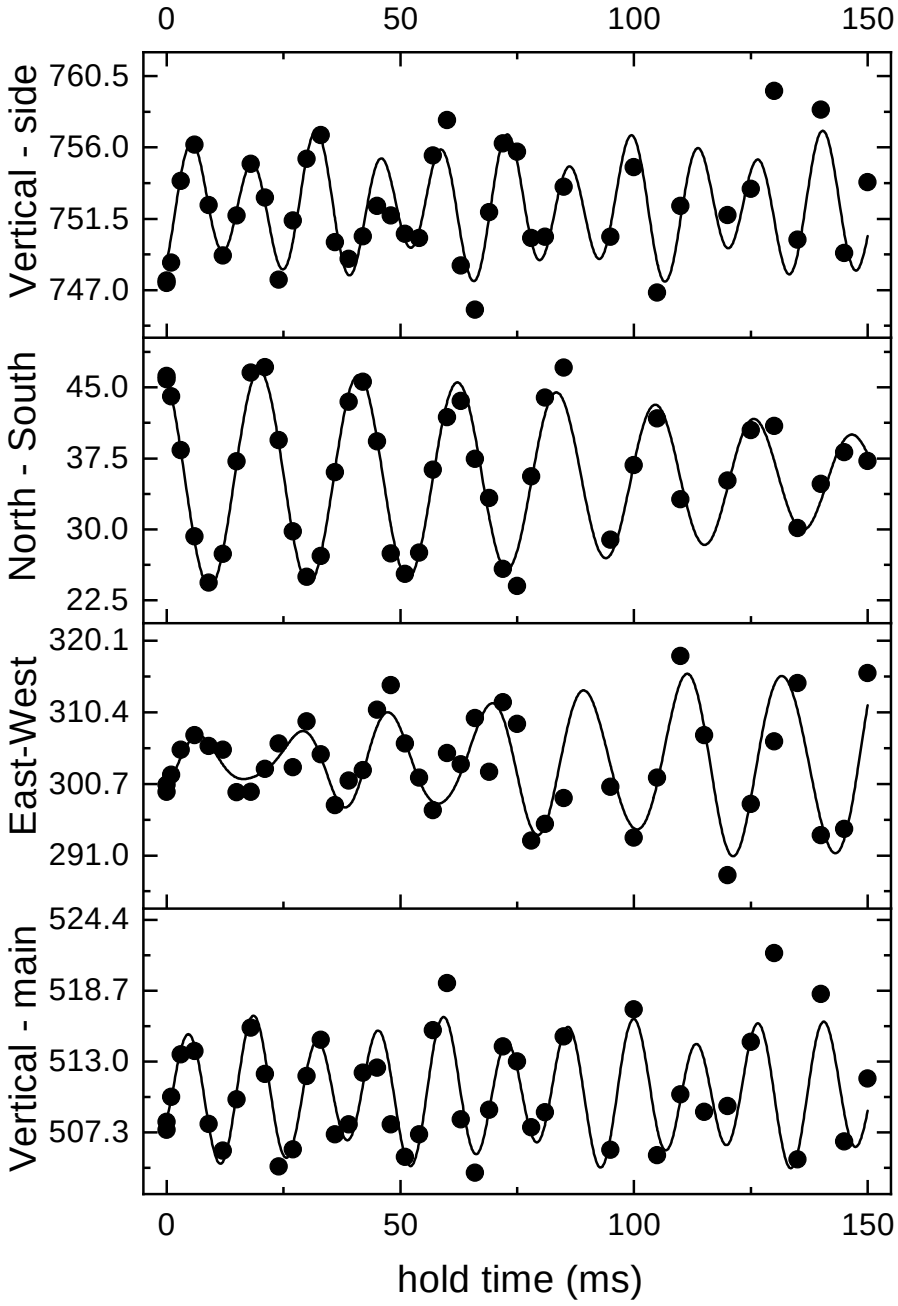

Figure 3.8: Global fit of trap frequency data to the sum of three sine functions. The vertical axis shows the position in pixels after 13 ms time of flight for varying hold times after a kick. Each pixel is 3 μm in the object plane. The East-West data were taken with the main imaging and the North-South data were taken with the side imaging.

where $y(x)$ describes the position of the gas at time x in ms. The overall offset is given by y_0. The subscripts of the remaining terms indicate which of the three sine functions they describe. The phase offsets (in units of time) are described by x_1, x_2, and x_3. The amplitudes are given by A and the frequencies (in Hz) are given by f.

The data consist of 4 data sets: the center of mass position of the gas in the x and y direction in both the main and side imaging (Figure 3.8). The main image measures the vertical and East-West directions, and the side imaging measures the vertical and North-South directions. Every data set is measuring the motion of the same atomic cloud, so the frequencies and phases of the motion are the same in every data set, and these fitting parameters are shared across all four data sets. The amplitudes and offsets are unique to each data set. A global fit was done, where each data set was fit to Equation 3.3 with this combination of shared and independent parameters. The results of the fit are given in Table 3.1 and Figure 3.8.

Three of the amplitudes were within uncertainty of zero, implying the very little motion along that principal axis was visible in the imaging direction of that data set. These parameters were fixed to zero in the fit in Table 3.1. Allowing the parameters fixed to zero to vary results in vertical side $A_1 = 0.7 \pm 0.6$, North-South $A_3 = 0.2 \pm 0.5$, and vertical main $A_2 = -0.1 \pm 0.6$. All other parameters are the same within uncertainty, and most are identical to their value when the near-zero parameters are fixed to zero.

	Vertical (side)	North-South (side)	East-West (main)	Vertical (main)
y_0	752.3 ± 0.4	35.4 ± 0.4	303.8 ± 0.4	510.6 ± 0.4
A_1	0.0 ± 0.0	-6.3 ± 0.8	8.5 ± 0.8	-1.2 ± 0.5
x_1*	3.8 ± 0.3	3.8 ± 0.3	3.8 ± 0.3	3.8 ± 0.3
f_1*	48.2 ± 0.2	48.2 ± 0.2	48.2 ± 0.2	48.2 ± 0.2
A_2	1.1 ± 0.5	-5.2 ± 0.7	-5.7 ± 0.9	0.0 ± 0.0
x_2*	4.0 ± 0.4	4.0 ± 0.4	4.0 ± 0.4	4.0 ± 0.4
f_2*	45.7 ± 0.3	45.7 ± 0.3	45.7 ± 0.3	45.7 ± 0.3
A_3	3.7 ± 0.5	0.0 ± 0.0	1.3 ± 0.5	5.0 ± 0.5
x_3*	1.6 ± 0.3	1.6 ± 0.3	1.6 ± 0.3	1.6 ± 0.3
f_3*	73.9 ± 0.3	73.9 ± 0.3	73.9 ± 0.3	73.9 ± 0.3

Table 3.1: The results of a fit of a sum of three sine functions to the trap frequency data. The parameters with "*" were set to be shared across all four sets of data. Their values were fixed to zero because they were within uncertainty of zero, as explained in the text.

3.5 Optical Lattice

3.5.1 Implementation of a Lattice Potential

After the atoms are trapped in the dipole trap, we add an optical lattice potential. The optical lattice is created with three mutually-orthogonal beams retro-reflected to create a standing wave of light at the atoms. The lattice directions are not along the directions of the cell, but rather approximately 45 degrees to one or more faces of the science cell, as illustrated in Figure 3.9. Specifically, the directions of the lattice beams, u, v, and w, are described by

$$u = \frac{-\sqrt{2}x + y - z}{2}, \; v = \frac{y + z}{\sqrt{2}}, \; w = \frac{\sqrt{2}x + y - z}{2}, \tag{3.4}$$

where x is the North-South direction, y is the East-West direction, and z is the vertical direction.

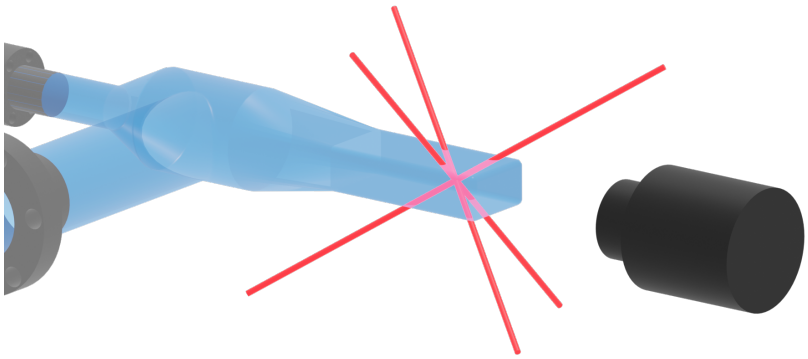

Figure 3.9: Three retro-reflected, mutually orthogonal lattice beam are used to create an optical lattice.

The lattice light is created with a Ti:Sapphire laser pumped with a 532 nm Coherent Verdi laser. The Ti:Sapphire laser is normally operated at 812 nm but is tunable over a wide range. When operated at 812 nm, the lattice spacing is $d = 406\,\text{nm}$. The waist of the lattices beams is estimated to be 120^{+30}_{-10} μm, based on the consistency between the lattice beam power, lattice potential depth, and measurements of the forces induced by the envelope of the lattice beams.

Besides providing a locally varying lattice potential, the lattice beams also create a global confining potential. The additional confining potential increases the trap frequency such that

$$\omega_{lattice} = \sqrt{\frac{4\,E_R}{M w_L^2}\left(2s - \sqrt{s}\right) + \omega_0^2} \tag{3.5}$$

Where ω_0 is the trap frequency without the lattice, s is the lattice depth (given in number of E_R), and w_L is the average waist of the lattice beams, and M is the mass of the atom [68].

3.5.2 Optical Lattice Potential

The potential created by the standing wave along each direction is

$$V(u, v, w) = s\left(\cos^2(u\ k_L) + \cos^2(v\ k_L) + \cos^2(w\ k_L)\right) \tag{3.6}$$

where $k_L = 2\pi/\lambda$ is the wavenumber and $\lambda = 812$ nm is the wavelength of the light. The strength of the lattice, s is determined by the intensity and wavelength of the light. Each lattice direction is created with a slightly different wavelength in order to avoid interference effects between lattice directions. The polarizations of the three lattice beams are mutually orthogonal.

We refer to lattice strength in units of the recoil energy $E_R = \frac{(\hbar k_L)^2}{2M}$. This experiment used lattice depths between $s = 4\,E_R$ and $s = 20\,E_R$, although higher lattice depths are possible. The strength of the lattice is calibrated by measuring band excitation.[4] The details of this process can be found in Reference [73].

[4]The code used to calculate the lattice potential given a measured frequency of band excitation oscillation is "calculate.m" in the "Code" folder on data share.

3.5.3 Ramping on the Optical Lattice

To add the lattice potential, we exponentially ramp the lattice to $4\,E_R$ over 200 ms with an exponential time constant τ of 100 ms. This process is designed to adiabatically load the lattice. The lattice is then linearly ramped to its final value, between $4\,E_R$ and $20\,E_R$, over $400\,\mu$s. This second ramp is diabatic (i.e., faster than the timescale of particle motion), so the density profile does not change from the profile at $s = 4\,E_R$. Therefore, in this experiment, every lattice depth has the same initial density profile. This process is different from previous work, where the lattice was ramped to the final lattice depth slowly. In previous work, we attempted to create an equilibrium state at the final lattice depth. However, since the timescale required for true equilibration of this thermal gas might be very long, we choose to create a starting state in equilibrium with the weaker lattice since we may not have reached the equilibrium state in the higher lattice before starting the measurement. It is necessary to start in an equilibrium state to calculate the expected starting density profile.

After adding the lattice, we can apply a variety of techniques to probe the quantum lattice gas that has been created. This apparatus has been used to perform a wide variety of experiments, a few of which are discussed in [49, 74–78]. In this work, we measure the dynamics of a nominally localized thermal lattice gas. The details of this experiment will be discussed in Chapter 6.

3.6 Imaging

In order to measure the density distribution of the atoms, we take absorption images of the gas. First, we use a repumping beam to excite the atoms to the $F = 2$ state. We then use light that is resonant with the $F = 2$ to $F' = 3$ transition on the D2 line as a probe beam. The probe beam enters the apparatus on the south side of the collection cell, travels through the transfer tube and into the science cell, where some of the light is scattered by the atoms and the resulting shadow is imaged onto a CCD camera. This camera was originally a Princeton Instruments camera, but it stopped working in 2017 and was replaced with a PIXIS camera, also from Princeton Instruments.

We have an additional imaging system called the "side imaging" that enables us to measure the position, size, and motion of atomic cloud in the North-South direction, which is along the main imaging direction and therefore not visible without additional side imaging. An additional probe be is sent into the science cell from the East and imaged with a set of optics that starts to the West of the science cell. The side imaging was originally relayed on to the same PIXIS camera that is used for main imaging, using just one part of the sensor for each imaging system. The side imaging has since been changed to use an additional camera. More details about the side imaging can be found in Section 4.2.5.

To form absorption images, we take three exposures with a 44 μs pulse of probe light, each 750 ms apart. The first image, with intensity $\mathcal{I}_{shadow}$, gives the shadow of the density distribution in the probe light. The second image, $\mathcal{I}_{light}$ is an image of just the probe beam after the atoms have left the imaging field. The last image, $\mathcal{I}_{background}$, is an image of the background light field without the probe beam. These images are combined to form an image of the optical depth OD as

$$OD = \ln\left(\frac{\mathcal{I}_{shadow} - \mathcal{I}_{background}}{\mathcal{I}_{light} - \mathcal{I}_{background}}\right). \tag{3.7}$$

This equation is applied to each pixel of the image to create an OD image. The resulting images represent the column-integrated density profile of the gas. The OD of each pixel is proportional to the number of atoms

in the column along the imaging direction that is projected onto that pixel.

In previous work, we have typically taken time of flight images where we release the atoms from the trap and allow them to expand before taking an image. This projects the momentum distribution of the atoms in to real space, and we therefore image the momentum distribution. The density of the time of flight distributions is much lower than the in-trap density, and OD saturation is not usually a problem. In this work we take in-trap images, and we must consider saturation.

Optical depth saturation can occur when the density is high enough to absorb most of the resonant probe light. The probe light is primarily resonant light with intensity I_0. However, there will be a small fraction of the probe light that is not resonant with the transition, I_β. The camera sees both parts of the light $I = I_0 + I_\beta$, but only the resonant light I_0 can be absorbed by the atoms. The fraction of this light that is absorbed α is proportional to the density. The shadow image is composed of the resonant light that is not absorbed and the off-resonant light,

$$\mathcal{I}_{shadow} = (1 - \alpha)I_0 + I_\beta, \tag{3.8}$$

and the light image is

$$\mathcal{I}_{light} = I_0 + I_\beta. \tag{3.9}$$

When α is small, I_β is negligible, but when α approaches 1, I_β becomes a significant fraction of the numerator of Equation 3.7 and the OD is no longer proportional to the density.

In order to get around this problem, we often use partial repumping for in-trap images, where the densities, and therefore, α is high. In normal imaging, we pulse the repump beam on for 300 μs, which is enough to transfer all of the atoms into the $F = 2$ state. However, to avoid OD saturation, we can pulse the repump light on for a much shorter time, usually between 10 and 30 μs. We adjust the repump time to keep $OD < 1.9$. This results in only part of the gas transferring to the $F = 2$ state, and since the repump light is uniform across the gas, the density distribution of the resulting $F = 2$ atoms is the same (up to normalization) as the density distribution of the original cloud. We use this technique to accurately measure the density distribution in-trap. However, measuring atom number from a partially-repumped image requires recalibrating at every repump time, which is typically impractical. More commonly, we measure atom number with additional runs of the apparatus to take fully-repumped, time-of-flight images of gases prepared in the same way.

Chapter 4

Tools for Spatially Resolved Measurements

Before this thesis, most of the work on this apparatus has involved measuring the momentum distribution of atomic gases using time of flight imaging. However, there is other interesting physics to probe by measuring the density profile of the gas, including measuring dynamics of far-from-equilibrium systems and the spatial distribution of localized states. In order to perform these measurements, we needed to develop new tools: a way to locally perturb the density distribution and a way to measure these changes with high-resolution imaging.

4.1 Barrier Beam

In order to measure local changes in the behavior of the system, we need some way to make spatially selective changes to the density profile. We accomplish this with a tightly-focused laser beam that propagates along the imaging axis. This barrier beam can be used to create deviations in the density profile that can be easily measured in column integrated images of the gas.

In order to add a beam along the imaging direction, we couple it into the system with a polarizing beam-splitter cube, which reflects the beam back towards the imaging lens. The imaging lens focuses the beam to the atoms. This is illustrated in Figure 4.1.

4.1.1 Gaussian Beams

We can predict the size of the beam at the atoms from the size of the beam at the first imaging lens using Gaussian beam propagation. The beams used in this work are approximately Gaussian, meaning a slice through the beam perpendicular to the direction of propagation has a two-dimensional Gaussian intensity profile. The peak intensity I_0 occurs in the center of the beam and the intensity falls off away from the center according to

$$I(r) = I_0 e^{\frac{-2r^2}{w(z)^2}}, \tag{4.1}$$

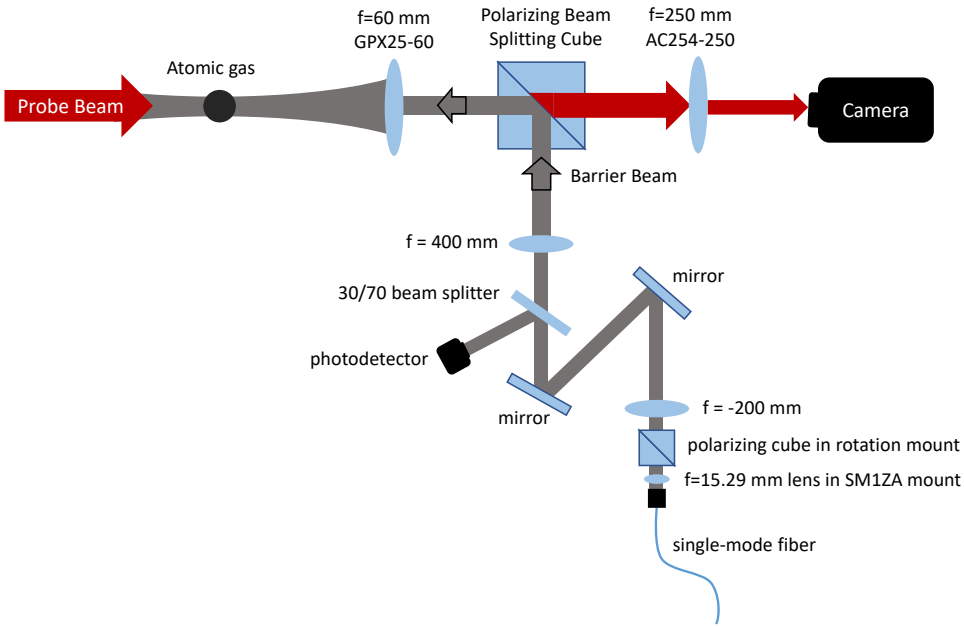

Figure 4.1: Optics used to create the barrier beam. A telescope was used to magnify the beam from a fiber, which was coupled backwards into the system along the imaging path with a polarizing cube. The beam is focused onto the atoms by the imaging lens.

where r is the distance from the center and $w(z)$ is the waist of the beam. The size of a Gaussian beam does not remain constant over the direction of propagation, but rather expands away from its focus according to

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \tag{4.2}$$

where w_0 is the waist of the beam at its focus. The distance away from the focus along the direction of propagation of the beam is given by z and λ is the wavelength of the beam. The Rayleigh range is

$$z_R = \frac{\pi w_0^2}{\lambda}. \tag{4.3}$$

When w_0 is large compared to the wavelength, z_R is large, and $w(z)$ is approximately equal to w_0 for a wide range of z around $z = 0$. However, when w_0 is on the order of λ, the beam waist changes significantly over a small range in z. We can use this relationship to calculate the necessary beam waist on the imaging lens to create a tightly-focused beam at the atoms. A 5 μm beam waist at the atoms is sufficient to make a density perturbation in a thermal gas in our experiment. In this work, we use $\lambda = 765.6$ nm and the imaging lens is 60 mm away from the focus of the beam. With these conditions, a 3 mm beam waist on the imaging lens creates a 5 μm beam waist at the atoms.

4.1.2 Barrier Beam Path

The beam is collimated in the region between the cube and the imaging lens, so sending a 3 mm waist beam to the cube creates an approximately 5 μm waist beam at the atoms. We created an approximately 3 mm

beam waist at the cube by using a telescope made of a $f = -200$ mm and a $f = 400$ mm lens to increase the size of a beam from a fiber (Figure 4.1). A negative focal length lens was used to decrease the length of the beam path, which was confined to a small area by other components of the apparatus. There is also a collimating lens on a z-axis translation stage immediately after the fiber. The beam was focused by adjusting the position of the collimating lens. A polarizing cube on a rotation mount was included in the beam path after the collimating lens. This allows us to set the polarization of the beam to maximize the fraction of the light reflected by the cube in the imaging path.

4.1.3 Barrier Beam Wavelength

This wavelength was chosen to be blue-detuned from both the D1 and D2 transitions in order to create a repulsive potential. This wavelength was also chosen to maximize the detuning that could be easily reached with the available Titanium-Sapphine laser we used to create this beam. Working at a larger detuning decreases the scattering rate, which is given by:

$$\Gamma_{\mathrm{sc}}(\mathbf{r}) = \frac{\pi c^2}{2\hbar} I(\mathbf{r}) \left(\frac{2\Gamma_{D2}^2}{\omega_{D2}^3} \frac{1}{(\omega_{D2} - \omega)^2} + \frac{\Gamma_{D1}^2}{\omega_{D1}^3} \frac{1}{(\omega_{D1} - \omega)^2} \right). \tag{4.4}$$

Increasing the detuning also decreases the dipole potential created by the beam, as given by Equation 3.2. However, the potential goes like $1/(\omega_0 - \omega)$ while the scattering rate goes like $1/(\omega_0 - \omega)^2$, so heating can be reduced by increasing the detuning and increasing the power to create the same original potential. The Ti-Sapph laser used for the barrier beam is capable of outputting powers orders of magnitude higher than necessary to create a barrier potential that completely excludes atoms from an area of the trap, even at the relatively large detuning of 765.6 nm. Heating from this beam is minimal because atoms are expelled from the high intensity region.

4.1.4 Size of the Barrier Beam

We measured the size of the beam at the atoms. We found that the beam shape was elliptical, most likely from imperfections in the beam path such as decentered or tilted optics or a beam path that is not precisely aligned to the imaging direction. As measured in the imaging plane, the axes of the elliptical shape of the beam were found to be rotated 45° from vertical with beam waists of 4.3 μm and 7.8 μm. The waists were determined by matching real column-integrated images to simulated images created using a model of the trap, including the hole beam potential. The trap model included the potential from the dipole trapping beams, the lattice confining potential, and the magnetic confining potential as described in Section 3.4.3, in addition to the barrier beam. The density was calculated with Maxwell Boltzmann statistics, neglecting the interaction energy.

A selection of the simulated images are shown next to the real images in Figure 4.2. Two real images were used to compare to the simulated images, one with $s = 4\,E_R$ and one with $s = 8\,E_R$. The first column is the real image, the second column is the images with the hole beam parameters that best matched the real images, and the third column is a set of example images with a large beam waist. The matching was done by eye.

The geometric mean of the beam waist is $(6 \pm 1)\,\mu$m. The beam power of the barrier beam, when measured before the cube in the imaging path, was 0.352 mW. The power of the barrier beam was kept fixed for the measurements in this work. The barrier has a peak potential of $V = (9000 \pm 5000)\,\mathrm{k_B} \times$nK. The

large uncertainty in the barrier potential does not introduce significant uncertainty in the initial density distribution, since this energy scale is much larger than the (210 ± 20) nK temperature of the atomic gas (Section 5.3).

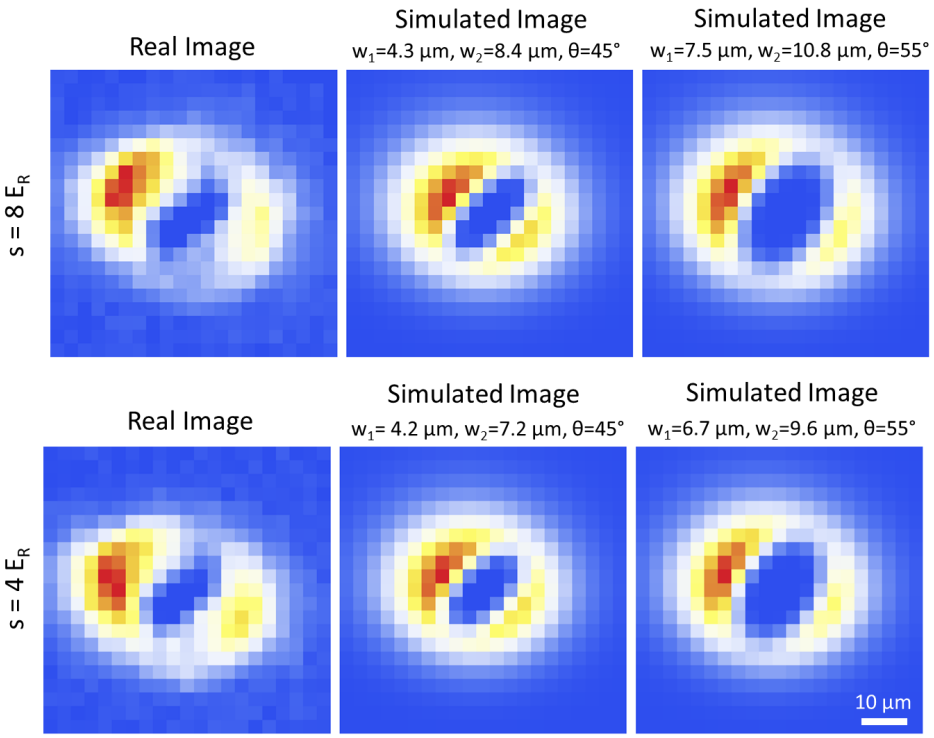

Figure 4.2: The waist of the barrier beam was found by comparing real images to images of simulated density distributions. The first column is the real image, the second column is the images with the hole beam parameters that best matched the real images, and the third column is a set of example images with a larger beam waist.

4.1.5 Barrier Beam Power

The beam power was controlled with active feedback from a home-made "noise eater" circuit. An acoustic optical modulator (AOM) before the fiber controls how much laser power is sent into the fiber. A pick-off in the beam path after the fiber directs a fraction of the light onto a photodetector. The voltage of the photodetector is compared to an analog control voltage. The power of the RF signal controlling the AOM is then adjusted using a variable gain amplifier by an integrator to minimize the difference between the control voltage and the photodetector voltage. Since the power needed for the barrier beam is small compared to the total available power of the beam, a 70-30 pick-off was used to send 70% of the light to the photodetector. Sending the majority of the light to the photodetector allows for more sensitivity in controlling the low-power beam. The polarizing cube in the cage mount containing the collimating lens and fiber mount controls the polarization so that the beam is reflected instead of transmitted through the cube in the image path (Figure 4.1). A filter to block light below 780 nm was also added between the cube and the camera to block reflected light from the barrier beam from reaching the camera.

4.1.6 Aligning the Barrier Beam

There are multiple ways to align the barrier beam, and different methods work best in different scenarios. Coarse alignment is best done with a large gas. The beam is set to a high power while staying below the power level that heats the atoms out of the trap. We have found using a very blue-detuned beam to work best for coarse alignment. We create a large thermal gas in the pinch trap, turn on the barrier beam, and then release the gas while leaving the barrier beam on until right before imaging. And image is taken after a few milliseconds time of flight. If the beam is in the region of the gas, a hole in the density profile should be visible. It is usually best to leave the anti-gravity coil on during the time of flight so that the gas stays in the area of the trap. Although, if the beam is not visible in the gas with the anti-gravity coil on, the coil can be turned off. Allowing the gas to fall will show the barrier beam if it is well-below the trap. When the beam is located, coarse alignment can be done by adjusting the last mirror before the cube in the imaging path until the hole in the density profile has the desired position in the image. Typically, the beam is aligned to the dipole trap, the position of which is found by taking separate, in-trap images of a gas.

There are other methods of finding the barrier beam that may be used if the beam is not located with a blue-detuned approach. It is possible to use a red-detuned beam to attract atoms to the beam. When this procedure is used, a small cloud of atoms should be visible at the location of the barrier beam. However, care must be taken to use a sufficiently low power as to not heat the atoms out of the trap. A high-intensity, red-detuned beam can also create a hole in the density distribution. The blue-detuned beam is usually easier to use, since a blue-detuned beam will cause a hole in the density profile at any beam power. Fluorescence could also be used to find the beam by setting the wavelength of the beam very close to resonance and imaging at a low power. We have not attempted this method.

Once the beam is coarsely aligned, it can be fine aligned. For fine alignment, the beam power and wavelength should be set at or near the values to be used in the experiment. The evaporation sequence should also be set to make a gas near the desired temperature in the desired trap. If the beam is not visible with the desired trap, a larger, hotter gas can be used for an intermediate stage of alignment. The barrier beam position should then be adjusted until the hole in the density profile is in the center (or at the desired position) of the gas. The position of the hole is very sensitive to the position of the adjustment knobs on the mirror. Half-tick adjustments may be necessary to set the beam to the center of the gas. The reproducibility of the knobs is not perfect, and the exact knob position that puts the beam in the center of the gas may not precisely match the ideal position found by a complete scan of the knob.

4.1.7 Timing of the Barrier Beam

As to avoid driving the gas out of equilibrium after evaporation, the barrier beam was turned on at the start of the dipole evaporation stage of preparing the gas. The gas was evaporated to its final temperature in the presence of the beam. After the gas had been prepared, the barrier beam was removed in order to observed dynamics. The barrier was ramped down sigmoidally over 300 microseconds (much faster than atomic motion), leaving a hole in the initial density profile. Figure 4.3 shows the timing of the barrier beam for a sequence with a 3-second hold time after ramping down the barrier. Thermodynamic calculations and measurements of the size and temperature of the starting state of the gas suggest that the gas is in equilibrium before removing the barrier beam. These calculations are discussed in detail in Section 5.3.

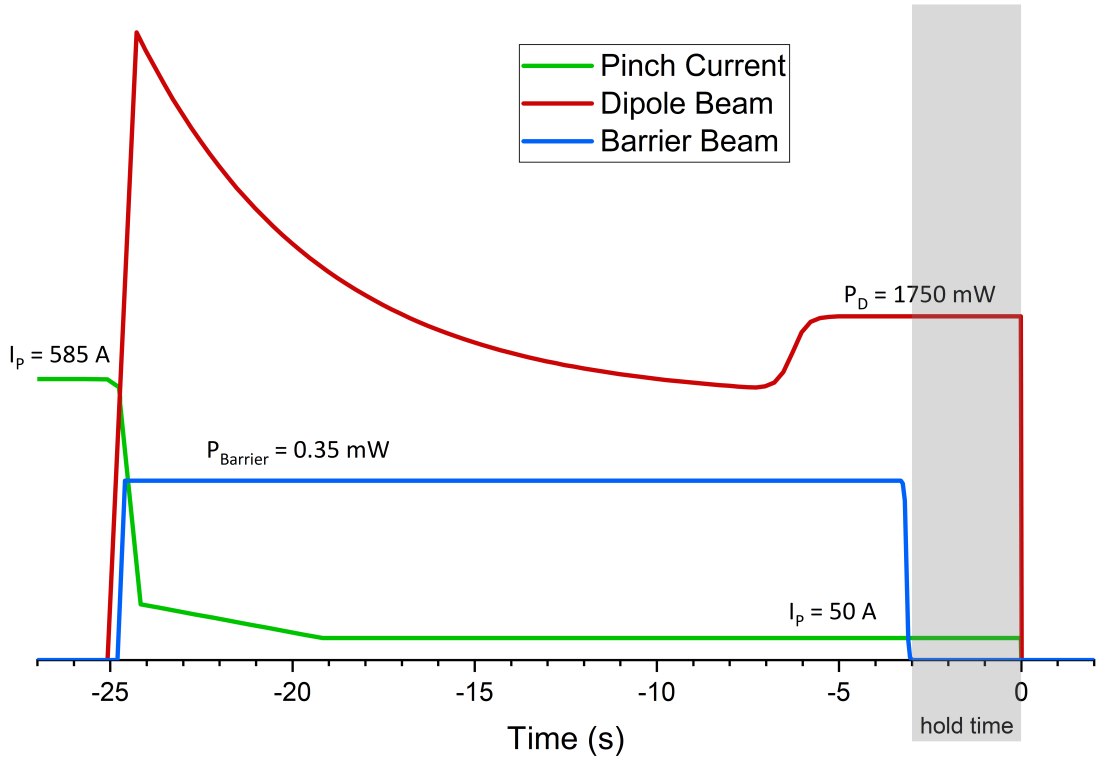

Figure 4.3: The barrier beam is turned on at the start of evaporation in the dipole trap so that the density distribution is in equilibrium with the barrier potential initially. The beam is sigmoidally ramped off over 300 μs and the atoms are held in the dipole trap for a variable hold time. A hold time of 3 seconds is shown. The vertical scale is arbitrary and not consistent between the dipole and barrier beam powers.

4.2 High-Resolution Imaging

The measurement described in this thesis was done using our old imaging system, which was designed for time-of-flight imaging. This imaging system and has a pixel-limited resolution, meaning the imaging optics are capable of reaching better resolution than the size of a pixel. Each pixel is 3.06 μm on a side in the imaging plane. This was sufficient for imaging the dynamics of a density distribution of a relatively large thermal gas. However, there is interesting physics that is not accessible at this imaging resolution. We would like to explore the density dynamics of states that are only achievable in a colder, smaller gas, such as BEC, Mott insulator, or Bose glass states. In order to do future experiments in these states, as well as to improve the detail available in measurements of thermal gases, we designed a new imaging system. Understanding how the new imaging system was designed first requires understanding how diffraction limits the resolution of an imaging system.

4.2.1 Diffraction Limit

When focusing light down to a small point or imaging light from a small area, diffraction limits the size of the features imaged or created. The problems are equivalent, the only difference between the two is the direction

the light is traveling. To explain the diffraction limit, we will discuss the problem of focusing a beam of light. Besides increasing the imaging resolution, improving the diffraction limit of the imaging system also decreases the smallest achievable barrier beam that can be created with the system, as the imaging lens is also used to focus the barrier beam.

Focusing collimated light through a lens redirects the set of plane-waves that make up the collimated light. Of course, these plane waves do not extend to infinity, but are limited by the size of the lens they pass through. Essentially, the lens performs a Fourier transform of the plane-waves to create a point of light at the focus of the lens. But just like any other Fourier transform, the size of that point must be finite when the extent of the initial waves is also finite. In media with an index of refraction near one, such as air and vacuum, the smallest achievable size of a lens focusing a given wavelength λ of light is

$$R = 1.22\lambda N, \tag{4.5}$$

where $N = f/D$ is the f-number of the lens, which describes the size of the cone of light the lens can collect by giving the ratio of the focal length f to the diameter of the lens D. (The related quantity, numerical aperture, NA $\approx 1/(2N)$ is often also used to describe lenses.) The Rayleigh criterion R gives the distance from the peak to the first minimum of the Airy disk that is created in the diffraction limit. At its best-focus, a uniformly illuminated perfect lens with a circular aperture will create a spot of light in the shape of an Airy disc.

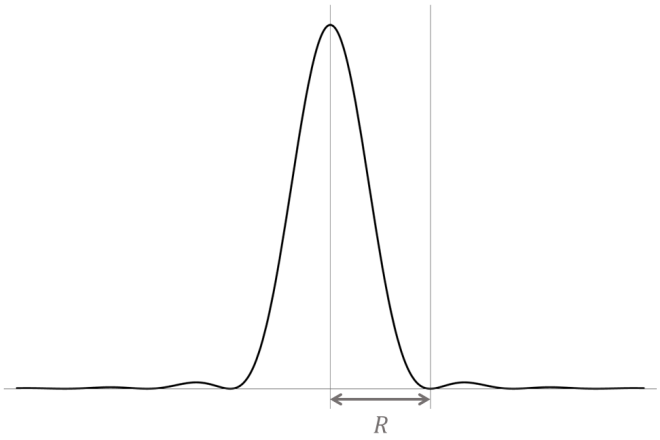

Figure 4.4: A cross-section through an Airy disc. The Rayleigh criterion is the distance from the center to the first minima.

Figure 4.4 shows a cross-section through an Airy disk. The Airy disc has a bright central spot surrounded by rings of decreasing brightness. The distance between the center and the first minima of the function is called the Rayleigh criterion. When two points are separated by more than the Rayleigh criterion, they are clearly resolvable as two individual spots (top panel of Figure 4.5). When they are separated by exactly Rayleigh criterion, they are just barely resolvable as two individual spots (middle panel of Figure 4.5). When they are within a Rayleigh criterion, the spots blend together, as seen in the bottom panel of Figure 4.5. The Rayleigh criterion is often used as a measure of the resolution of an imaging system. The distance from the peak to the minima (the Rayleigh criterion R) of the light spot created by the imaging system is a measure of the smallest feature the imaging system can resolve.

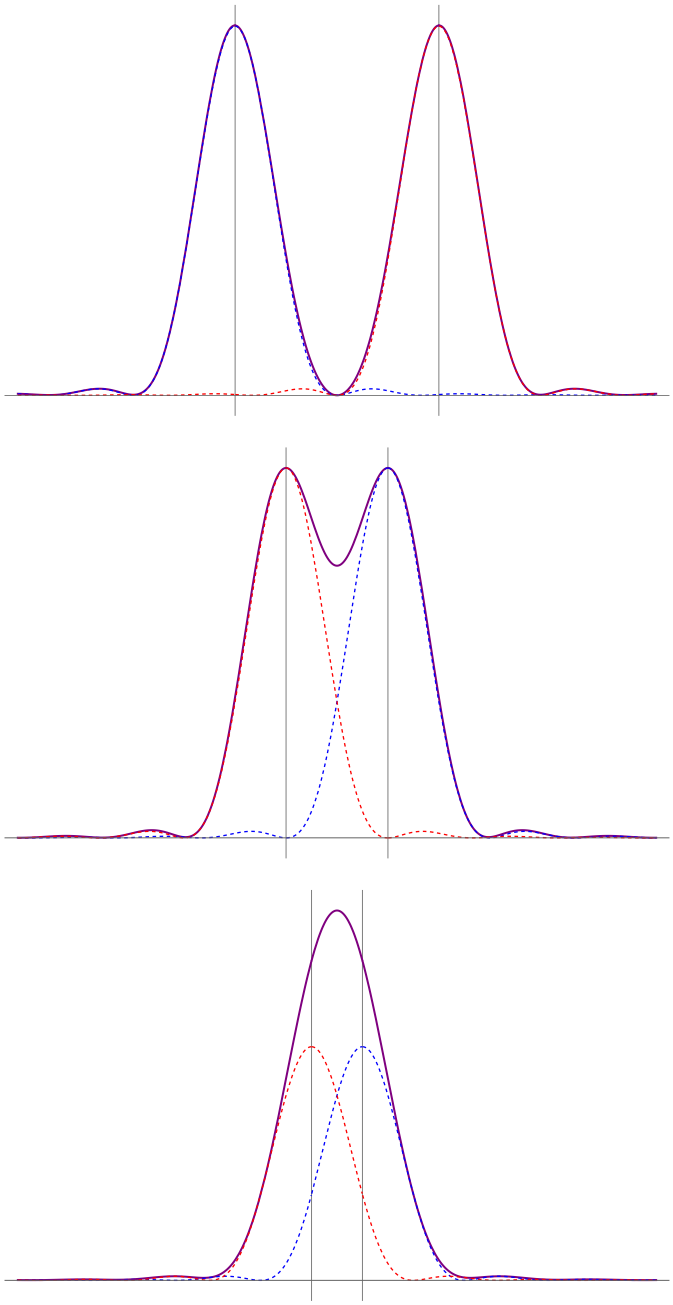

Figure 4.5: The sum (purple solid line) of two Airy functions (dashed red and blue lines) that are separated by 2 Rayleigh criterion (top), 1 Rayleigh criterion (middle) and 1/2 of a Rayleigh criterion (bottom). When the Airy discs are separated by less than a Rayleigh criterion, they appear to be a single peak.

4.2.2 Motivation for a New Imaging System

The first lens in the current imaging lens has a $f = 60$ mm focal length and a diameter of 25 mm (Figure 4.6). Given that we use 780 nm light for imaging, the current lens has a diffraction limited Rayleigh criterion of $R = 2.3\,\mu$m. We are unable to reach this resolution with the current imaging system because each pixel images an area larger than $2.3\,\mu\text{m} \times 2.3\,\mu\text{m}$.

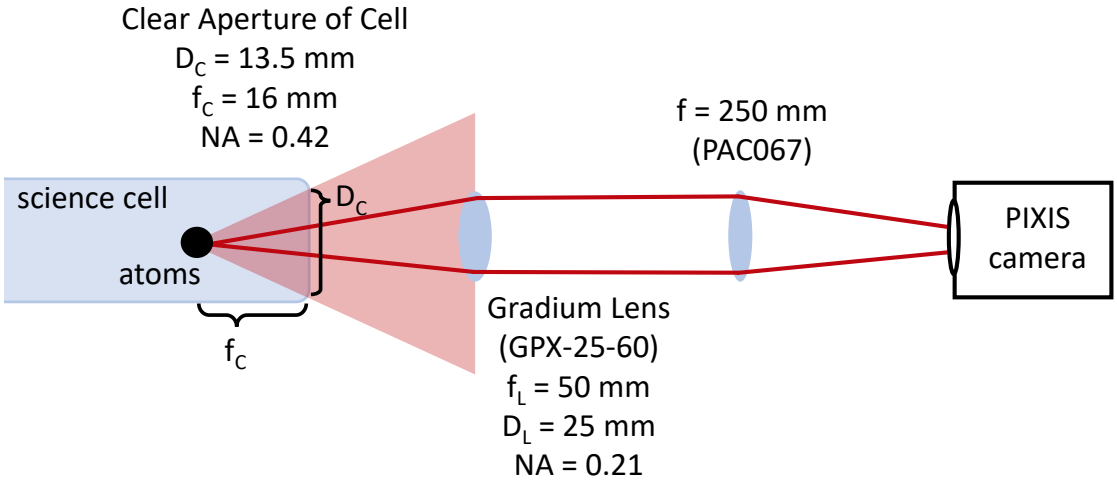

Figure 4.6: Diagram showing the Numerical Aperture (NA) for the current imaging lens and science cell. The diffraction limit of the system is limited by the geometry of the science cell. The maximum NA allowed by the cell is 0.42. The current imaging lens limits the NA to 0.21.

We could add a second stage of magnification without changing this first imaging lens. (At time of writing, the apparatus is used in this configuration. The imaging system has an addition stage of magnification, with the original lens still providing the first stage of magnification. See Section 4.2.5 for more details.) However, adding a second stage of magnification only solves one of two problems. Higher magnification creates smaller pixels in the imaging plane, so there is more data over an equivalent area. Although, the image is still limited by the diffraction limit of the first lens and spherical aberrations.

Besides the diffraction limit, we must also consider spherical aberrations. Spherical aberrations cause light from the edge of the image to focus at a slightly different point than light from the center of the image. The current first imaging lens is a Gradium lens (GPX-25-60) that is able to minimize aberrations to work near its diffraction limit by including many layers with different indices of refraction. However, the glass of the science cell introduces spherical aberrations that are not currently corrected and will limit the resolution that can be reached with the current imaging system. The Rayleigh criterion resolution of this system was measured to be 3.3 ± 0.4 μm.[1] In order to work at the best possible resolution, we need to design a new optical system to improve the diffraction limit and minimize aberrations.

We would like to work at the diffraction limit of the glass cell, which is smaller than that of the current imaging lens. The clear aperture of the science cell is 13.5 mm in its smallest direction and the atoms are 16 mm from the wall of the cell (Figure 4.6). This limits the usable f-number to 1.2. Changing the position of the atoms in the cell would require significant changes to the apparatus, since the position of the dipole

[1] This measurement is documented on page 55 of David McKay's lab notebook #2.

beams is currently tightly constrained to the current trap position. The f-number allowed by this geometry limits us to a resolution of 1.1 μm with a probe light of $\lambda = 780$ nm.

In order to maximize the resolution of the imaging system, we need to use a lens with a lower f-number than the cell, which would allow us to work at the diffraction limit of the cell. We also need to minimize aberrations in the optical system.

4.2.3 Designing a New Imaging System

To achieve high resolution, we must address spherical aberrations. Optics can introduce distortions, called aberrations, to the image by not perfectly manipulating the wavefront. For example, a perfectly spherical lens, which is the most common lens shape, actually bends the light on the edges of the lens slightly differently than the light going through the center of the lens which has the effect of focusing the light from the edges at a slightly different distance than the light through the center of the lens.

For many optical systems, the distance between optics is much larger than the waist of the beam at any point and the angle of the divergence of the beam is always small. In this limit, we can take the paraxial approximation where $\theta \approx \sin\theta$. In this case, no light passes through the outside of the lens and the lens can be treated as an ideal paraxial lens. When working at very small resolution, a high NA is necessary and the angles will not be small. In this limit, aberrations can noticeably distort the images.

Standard off-the-shelf optics have too large spherical aberration for this application. Most lenses are manufactured with spherical surfaces. Even best-form lens, where the radius of curvature of both surfaces of the lens have been chosen to minimize spherical aberrations, introduce enough spherical aberration to limit the resolution we can achieve.

Lenses are not the only possible source of aberrations. The glass window of the science cell will also introduce aberrations. Besides finding a high-quality lens with very minimal aberrations, we also need to find a way to cancel the aberrations from the glass of the science cell.

Zemax Simulation

We used the optical simulation software Zemax OpticStudio to simulate real optical elements to find an optical design that minimizes the total aberrations in the optical system. Zemax is a powerful tool to simulate optical systems. It contains a library of physical lenses that can be combined in an optical path. Zemax is then able to simulate light traveling through the optical system, which can be used to predict the system's imaging capabilities.

Each Zemax simulation consists of a list of optical surfaces and the distances between each surface. The first surface is the "OBJECT", which includes a light source. For this simulation, we used 780 nm light. The light spreads out from the source, limited by the "STOP" which is another surface in the Zemax simulation with an aperture that sets the angle of rays that are used in the simulation. Many additional surfaces are added to simulate the various lenses and other optical components of the system. The last surface is the "IMAGE" surface, where the output light field is calculated. An example of this list of surfaces can be seen in Figure 4.7.

Zemax has a built-in is a library of lens files of real lenses which can be added to the optical path, or ideal paraxial lenses can be used. The distances between the lenses can then be varied to focus the system. For a simple two-lens system, we found the ideal distance of the first lens from the object by finding the distance that produced collimated rays after the lens. The ray tracing diagram of the Zemax simulation is

Lens Data | 3: Huygens PSF Cross Section | 2: Spot Diagram | 1: Layout | 4: 3D Layout | Field Data Editor

Update: All Windows

Surface 8 Properties

	Surf:Type		Comment	Radius	Thickness		Material	Co	Clear Semi-Dia		Chi	N
0	OBJECT	Standard		Infinity	0.000				0.000		0.0..	0
1		Standard		Infinity	16.000				0.000		0.0..	0
2	STOP (aper	Standard	glass of cell	Infinity	1.587		N-BK7		6.750	U	0.0..	6
3	(aper)	Standard		Infinity	2.000				25.000	U	0.0..	2
4	(aper)	Standard	SLB-50-1500P	778.500	3.400		S-BSL7		25.000	U	0.0..	2
5	(aper)	Standard		Infinity	11.209	V			25.000	U	0.0..	2
7	(aper)	Standard		Infinity	15.500		S-LAH64		25.000	U	0.0..	2
8	(aper)	Even Asphere	AHL50-40	-31.075	-15.500	T			25.000	U	0.0..	2
10		Standard		Infinity	20.000				25.000	U	0.0..	2
11	(aper)	Standard		Infinity	50.000		N-BK7		25.000	U	0.0..	2
12	(aper)	Standard		Infinity	30.000				25.000	U	0.0..	2
13	(aper)	Standard	AC508-200-B	134.000	8.200		N-LAK22	..	25.400	U	0.0..	2
14	(aper)	Standard		-109.200	5.000		N-SF6HT		25.400	U	0.0..	2
15	(aper)	Standard		-515.200	0.000			..	25.400	U	0.0..	2
16		Standard		Infinity	160.874	V			14.232		0.0..	1
17	IMAGE	Standard		Infinity	-				5.153E-03		0.0..	5

Figure 4.7: The list of surfaces in the Zemax software. This list includes a flat glass surface and three lenses, as well as free space between the optical elements.

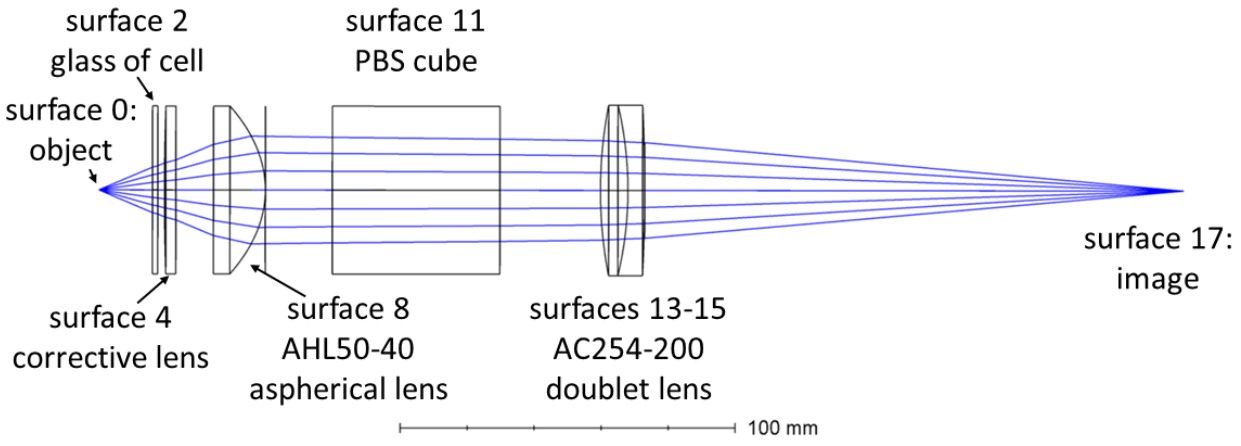

Figure 4.8: The ray tracing diagram from the Zemax software for the list of surfaces given in Figure 4.7.

shown in Figure 4.8. We found the correct distance from the second lens to the imaging plane by finding the distance that minimized the size of the point of light in the image plane. Adding additional optical elements (such as the glass representing the glass of the science cell) changed the optimal distances, and both distances were varied again to minimize the size in the imaging plane.

Three criteria were used to measure the size and quality of the point of light created by the imaging system, the Rayleigh criterion, the Strehl ratio, and the RMS spot size. Both the Rayleigh criterion and the Strehl ratio were calculated with the Zemax Huygen's point spread function calculation.

Zemax Huygen's point spread function (PSF) calculation was used to generate a plot of the PSF. The point spread function shows the intensity of light at each point in space in the image plane from a point source propagated through the imaging system. The Zemax calculation propagates rays through every optical element, and, using the path of the ray, calculates the complex amplitude of the plane wave of each ray. The complex amplitude at each point in the imaging plane is found by summing the contribution from each ray. The intensity of light at each point is found by taking the square of the complex amplitude. In this way, the diffraction and interference pattern of all the rays propagated through the system is calculated at the imaging surface [79]. To quantify the results of the Huygen's PSF calculation, we used two measures, the Rayleigh criterion and the Strehl ratio.

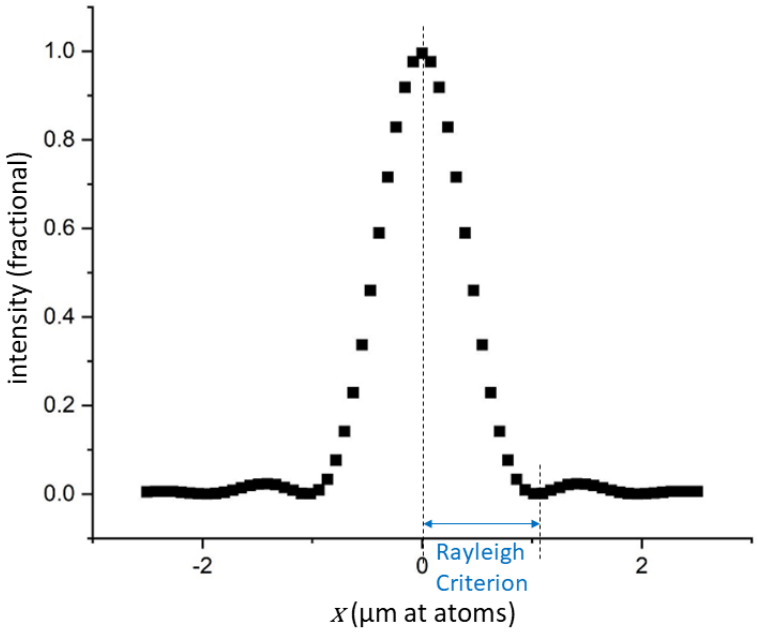

Figure 4.9: The point spread function as calculated by Zemax

The Rayleigh criterion is the distance from the peak to the first minimum of the PSF, as shown in Figure 4.9. Since the Zemax calculates the light field in the imaging plane, and this set of optics has a factor of 5 magnification, the results from Zemax are at a 5× magnification relative to the object plane. To find the Rayleigh criterion in the object plane, we divide by 5. We report all sizes in the object plane in what follows, as that is the relevant number for understanding the imaging resolution.

The Strehl ratio is the ratio of the peak intensity of the point spread function to the peak of the ideal PSF, limited only by diffraction. Imaging systems that are only limited by diffraction have a Strehl ratio very close to 1. As a general rule, systems with a Strehl ratio of less than 0.8 are considered not to be diffraction

limited but rather limited by geometric aberrations.

We also measured the RMS spot size, which quantifies geometric aberrations. (In our system, the dominate source of geometric aberrations are spherical aberrations.) Unlike the Strehl ratio, which measures diffraction, the spot size is only a measure of geometric aberrations. It is the root-mean-square size of the spot size diagram created by propagating a set of rays through the optics system, accounting for geometric aberrations. A smaller spot size allows the imaging system to distinguish between more closely spaced point sources and therefore have a better imaging resolution.

Choosing New Lenses

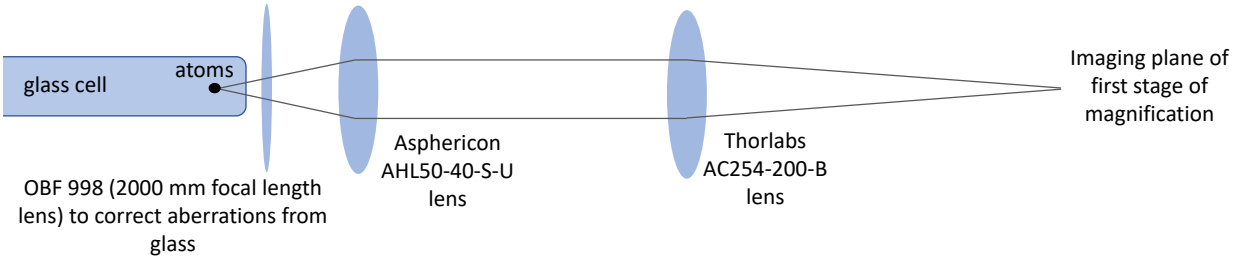

Figure 4.10: The lens diagram for the new imaging system

In searching for a new imaging lens, we primarily considered aspheric lenses from Asphericon. Reference [80] had success using one of these lenses in an imaging system, and Asphericon offers low-aberration lenses at 780 nm at a reasonable price point. We compared the lenses AHL50-40, ALL100-100, AFL50-60, and AHL25-30. All four of these lenses are designed to work at 780 nm and have a sufficiently long working distance (greater than 30 mm) to use with the current position of the trap in our science cell. We found the AHL50-40 lens to have the smallest Rayleigh criterion in our system.

The AHL50-40-S-U lens from Asphericon has a 50 mm diameter and a 40 mm focal length, making its f-number more than sufficient to allow us to reach the diffraction limit of the cell. The diffraction limit of this lens at 780 nm is 0.8 μm. This lens is an asphere; instead of the typical spherical shape, it is manufactured with a slightly modified shape that bends the light rays on each part of the lens specifically to focus at the same point. This shape is designed specifically to work with 780 nm light.

We also found that adding a $f = 2000$ mm lens (part number OBF 998) helped to correct for the aberrations introduced from the glass of the science cell (Figure 4.10). This corrective lens makes the Rayleigh criterion resolution 1.0 μm as opposed to 1.4 μm without the corrective lens. Measuring the geometric aberrations shows an even more pronounced difference. The RMS radius without the corrective lens is 3.2 μm, while the RMS radius is 0.56 μm with the corrective lens (Figure 4.11). This lens has a very long focal length and only slight modifies the image. The last element of the first stage of magnification of the new imaging system is a doublet from Thorlabs with a focal length of 200 mm and a 2-inch diameter (part number AC254-200-B).

Tolerance to Misalignment

So far, all of the Zemax calculations have been with optics perfectly aligned on the imaging axis. But in actually building the optical system, there will be some deviation in the position of the optical elements.

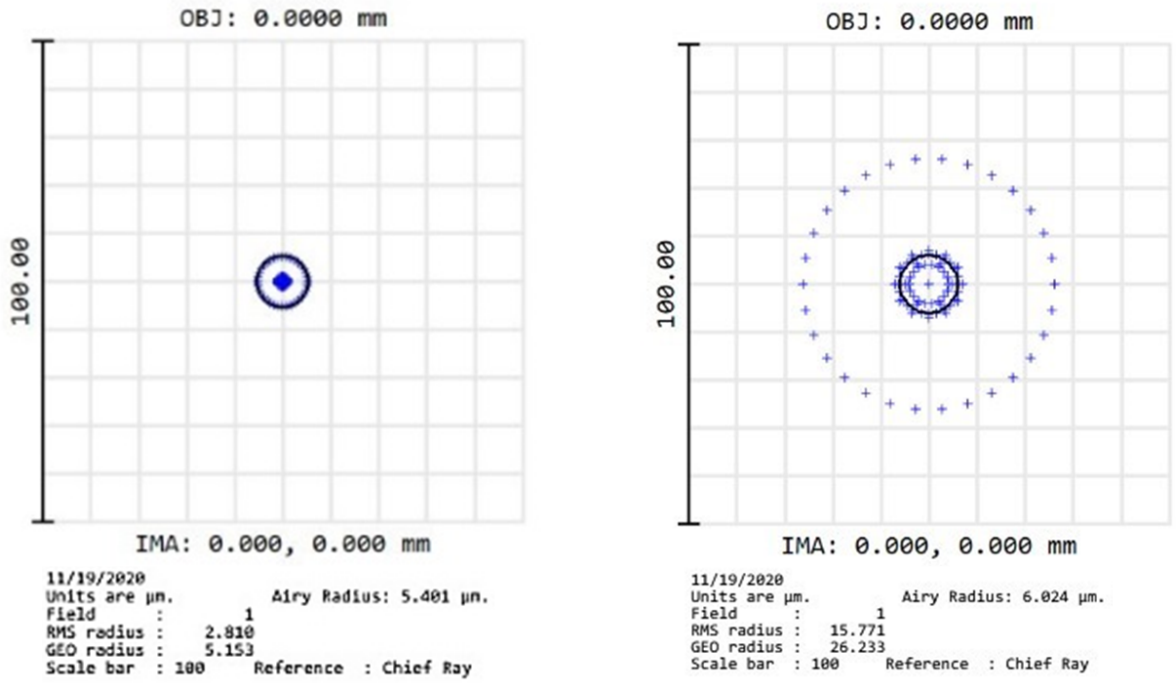

Figure 4.11: The spot size diagrams generated in Zemax for the optical system with (left) and without (right) the corrective 2000 mm lens.

In order to estimate how much tolerance there is for optics being off-axis, we simulated a shift of the main imaging lens in Zemax by an amount "decenter" and measured the resulting resolution. This simulation includes all three lenses, the glass of the cell, and an aperture representing the clear aperture of the science cell. The results are shown in Table 4.1. We found that the imaging lens can be up to 100 μm off-center and the resolution is still limited by the diffraction limit of the cell.

Decenter	RMS spot size	Rayleigh Criterion	Strehl Ratio
0 mm	0.6 μm	1.0 μm	0.996
0.1 mm	1.0 μm	0.9 μm	0.95
0.15 mm	1.4 μm		
0.2 mm	1.8 μm	0.9 μm	
1 mm	11.5 μm	0.9 μm	0.036

Table 4.1: Table of the resilience of the resolution of the imaging system to decentering the main imaging lens.

Depth of Focus

The depth of focus is another factor to consider in choosing optics. Since our images are column integrated, we need good out-of-focal-plane image clarity. Perfect image clarity at the center of the gas (i.e., the focal plane) is useless if the atoms in front of or behind the center are sufficiently out of focus as to obscure the center of the gas. To calculate the depth of field in Zemax, we shifted the simulated distance between the object plane and the first optic, effectively imaging a point away from the optimized focal plane.

With such a tightly focused imaging system, the depth of focus will be very shallow. The image will only

be in focus (defined as a Strehl ratio of the PSF of at least 0.5) over a 7 μm range along the imaging direction. In order to preserve the ability to image over a larger depth of focus (which is necessary for aligning beams and other experiments), we will add an iris to the system that we will partially close to increase the depth of focus at the expense of resolution for certain measurements. We will use a cage mountable 2 inch iris aperture (LCP50S) for this purpose. At its minimum aperture size, 2 mm, the depth of focus for our system in 2.2 mm. This should give us plenty of range to work at long depths of focus when necessary. Table 4.2 shows the depth of focus at a variety of sizes of the variable aperture.

Aperture Size (diameter)	Depth of focus
50 mm (No aperture)	7 μm
20 mm	20 μm
10 mm	100 μm
2 mm	2200 μm

Table 4.2: The depth of focus is the range over which the first minima of the PSF is still distinguishable, and the Strehl ratio is greater than 0.5 in a Zemax simulation.

Estimating the performance of the new imaging system in Zemax is only the first step in understanding the optical system. To better understand if this system will provide sufficient imaging resolution for our apparatus, we did bench-level testing. The ultimate goal in testing the imaging system is to measure the point spread function (PSF) of the system. The point spread function is the shape of the light in the imaging plane from a point source propagated through the imaging system. For a perfectly diffraction-limited system in the small-angle limit, this is an Airy function, but aberrations and large angle effects can create a PSF that is not an Airy disc. By measuring the point spread function, we can simulate how images of an atomic gas would appear with this new imaging system.

4.2.4 Testing the New Imaging System

We purchased the chosen optics and set them up for testing. We mostly focused on testing the first stage of magnification, although we did perform some tests of the whole imaging system. There were two main methods for testing the imaging system, focusing a beam with the imaging lens and imaging very small features.

For each of these ways of measuring the resolution, it is necessary to carefully set the distance between the object and the first lens and the distance between the second lens and the camera (or imaging plane, if multiple stages of magnification are used). The camera is mounted on a translation stage, so its position can be easily adjusted to the focal plane of the image. In the test set up, the object is also mounted on a translation stage to set the distance between the object and the first lens. When incorporated into the apparatus, the first lens will be on a translation stage to facilitate focusing the imaging system, instead of moving the object, which is impractical with the current state of that apparatus. We need to find the focus of each subsystem by sweeping over the range of focal distances and finding the distance that has the smallest feature or sharpest edge. An example of this is given in Figure 4.12.

Imaging a Focused Beam

We used the set-up shown in Figure 4.13 to measure the spot size created by focusing a beam with the new imaging lens. We coupled 780 nm light from a New Focus external cavity diode laser into a single-mode fiber in order to create a clean Gaussian beam. We expanded this beam to about a 1-inch diameter with

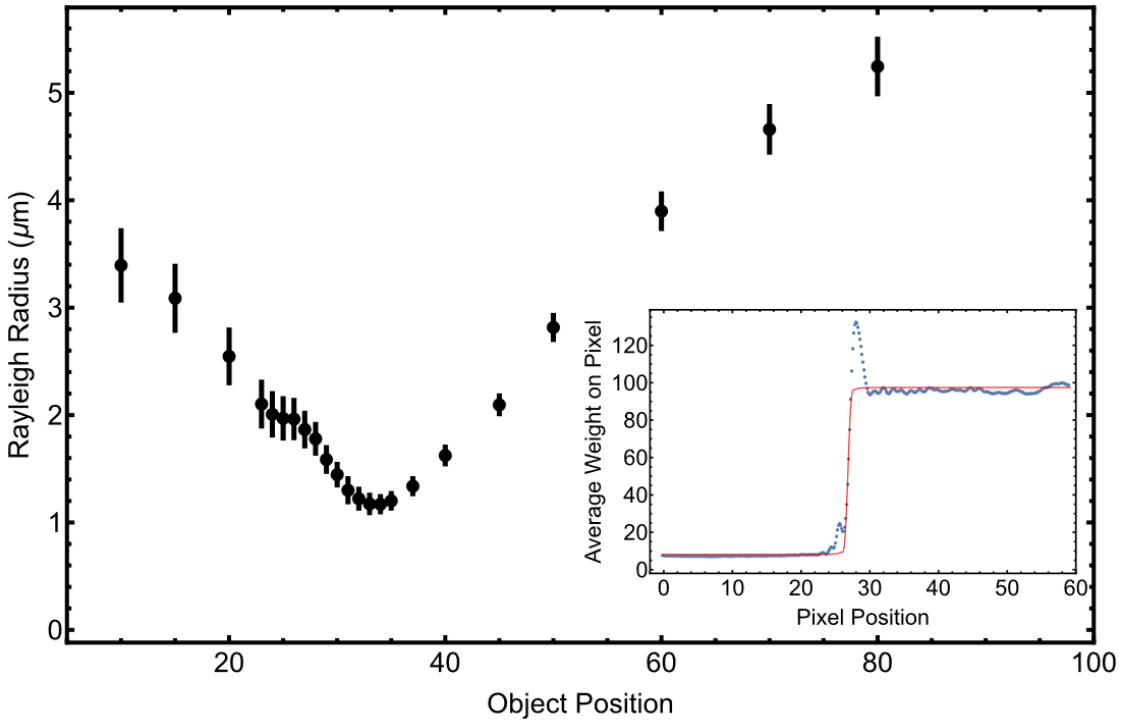

Figure 4.12: The position of the object is scanned to find the best focus. The minimum Rayleigh radius occurs at an object position of 33.5, suggesting that is the best-focus position of this optical system. The Rayleigh radius was found by fitting a cross-section of an average trace through a straight edge of the test target to the integral of an Airy function. The error bars show the fit uncertainty. The inset shows an example trace, with the blue points showing the average weight on the pixel and the red line showing the fit.

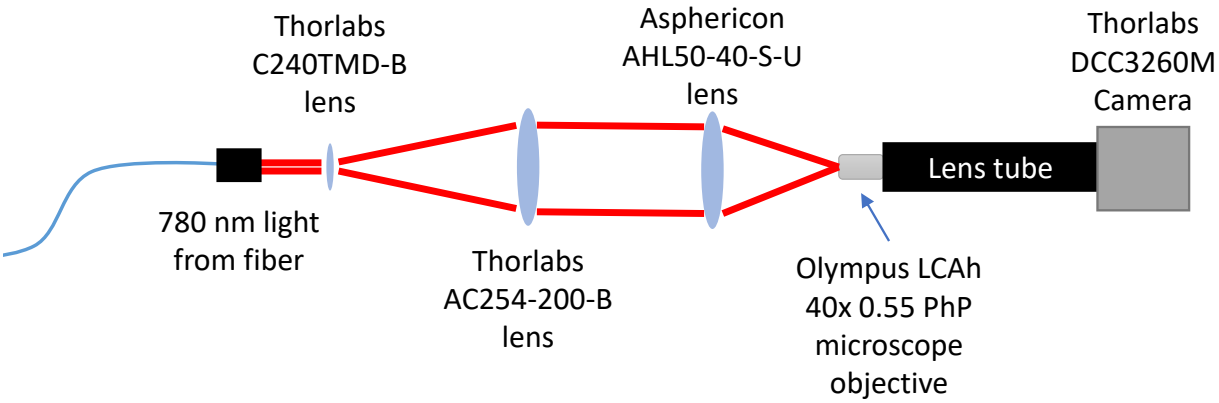

Figure 4.13: The set-up used to measure the spot size produced by focusing a beam through the new imaging lens

a C240TMD-B lens from Thorlabs and collimated it with a 2 inch, $f = 200$ mm (AC254-200-B) doublet lens from Thorlabs. We collimated the beam by adjusting the distance between the fiber and the lens while looking at the beam sent out to the wall a few meters away. The beam is collimated if the size of the beam is the same immediately after the lens and on the wall. The collimated beam was then focused with the AHL50-40-S-U imaging lens onto a 40x microscope objective (Olympus LCAh 40x 0.55 PhP). The image created by the microscope objective was imaged with a DCC3260M camera from Thorlabs.

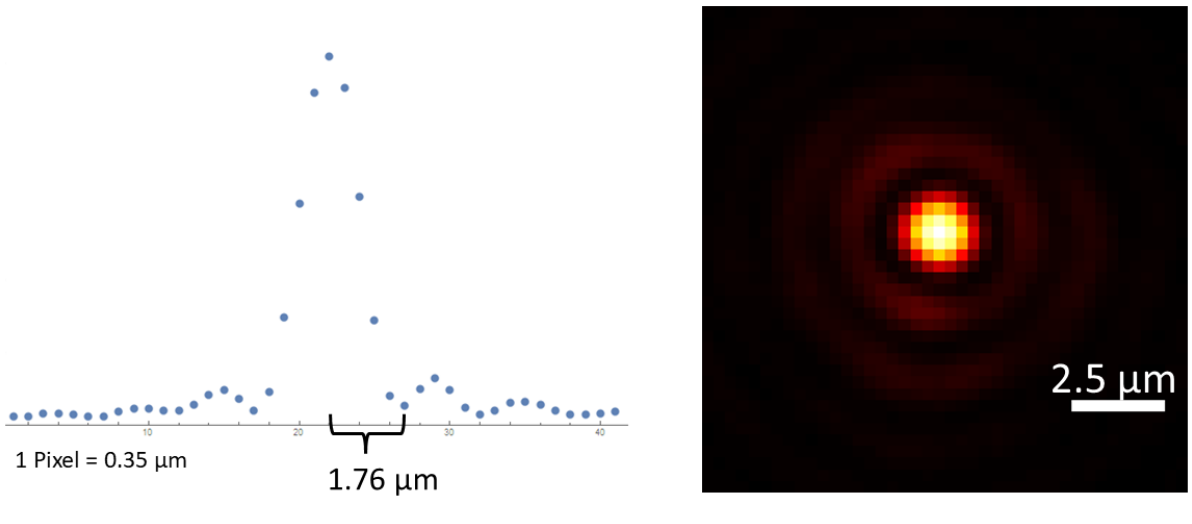

Figure 4.14: The results of measuring a large beam focused through the imaging lens

Focusing through the imaging lens produces an image that is similar to an Airy function as seen in Figure 4.14. We are able to produce a beam with a Rayleigh criterion of 1.75 μm. This result suggests we will be able to make a feature in the density distribution of the gas with a waist of less than 2 μm with the new imaging system. This is much smaller than the current barrier beam and would allow us to make smaller perturbations on smaller gases. The images of the focused beam of light were promising, but we should be able to reach even higher resolution. The resolution of the point of light could have been limited by the size of the beam on the imaging lens, the quality of the optics, or the alignment of the optics. Next, we tried to test the imaging abilities of the new optical system.

Imaging Test Target

In addition to measuring the focusing capabilities of the lens, we measured how well it can create an image. An initial test of its imaging capabilities is to image a test target. The smallest features on the test target are larger than the resolution we hope to achieve, but we can measure the resolution by measuring the sharpness of the edge in the test target.

The set-up for measuring the resolution test target is shown in Figure 4.15. The test target was back-illuminated with the incoherent light source at 780 nm (Thorlabs M780LP1). This light source has a full-width half-max bandwidth of 30 nm and emits light at a wide angle (120°). An adjustable collimator (Thorlabs SM2F32-B) was used to increase the collimation of the light. Since we are imaging a shadow instead of the light, perfect collimation is not necessary as long as a sufficient amount of light reaches through the entire imaging system.

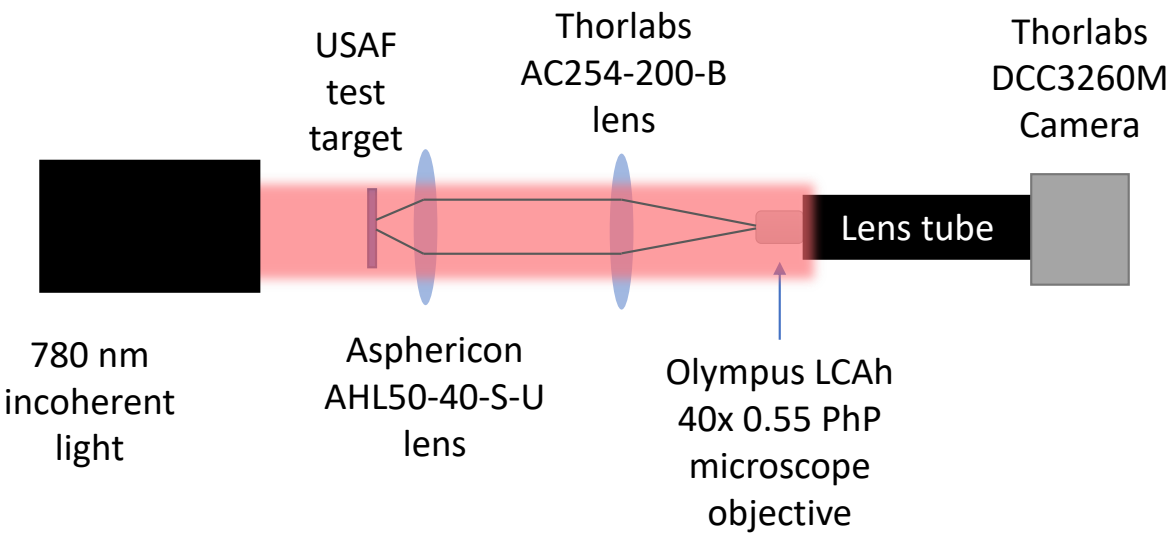

Figure 4.15: The set-up for measuring the resolution test target

The incoherent light source was chosen to decrease the effects of edge ringing, interference patterns that run next to edges in the image [81, 82]. The range of frequencies decreased (although did not eliminate) the interference pattern, but also introduced chromatic aberrations. Since the first imaging lens is designed specifically for 780 nm light, the light away from this wavelength in the incoherent light source is focused at slightly different points. The chromatic aberrations of this test set-up may decrease the contrast in the image by blurring some of the light into a plateau around each feature.

The shadow of the test target was imaged with the primary aspherical imaging lens and a $f = 200$ mm doublet. This first set of optics creates an image of the test target on the imaging plan of a microscope objective, which then magnifies the image to create a large image on the Thorlabs camera. Without the additional magnification provided by the microscope objective, we would be limited to approximately 1 μm per pixel, which is not enough resolution to measure whether the imaging system can measure 1 μm features.

The test target is a piece of glass with chrome pattered features. It is a 1951 USAF resolution test chart with features from Group 0 through Group 7. The smallest feature has a width of 2.2 μm. An image of the middle of the resolution test target can be seen in Figure 4.16. The inset shows the image captured with the full set up described above. The smallest lines imaged are 2.19 μm thick and 10.95 μm long. The bottom half of the inset image shows a set of three lines of this size aligned vertically next to a set of three lines of the same size aligned horizontally. We can see "phantom" lines between the real lines of the test target. These could potentially arise from edge ringing effects that persist with the 30 nm bandwidth of wavelengths of the incoherent light source, or possibly from the tails of the airy function overlapping.

One way to measure the resolution of the imaging system is to measure the sharpness of a straight edge. We took an image of a square in the test target that was rotated relative to the camera. The rotation helps to decrease the effect of the quantization of the pixels on the data. The image was then digitally rotated to make the edge vertical, and then horizontal slices through the image were averaged together to make an average slice across the edge of the square. We then fit this trace to the integral of an Airy function and extracted Rayleigh radius (the distance from the peak of the Airy function to the first minima) from the results of the fit. We scanned the distance between the test target and the imaging lens to find the position of best focus, where the Rayleigh radius is minimized. This is shown in Figure 4.12. This data was taken

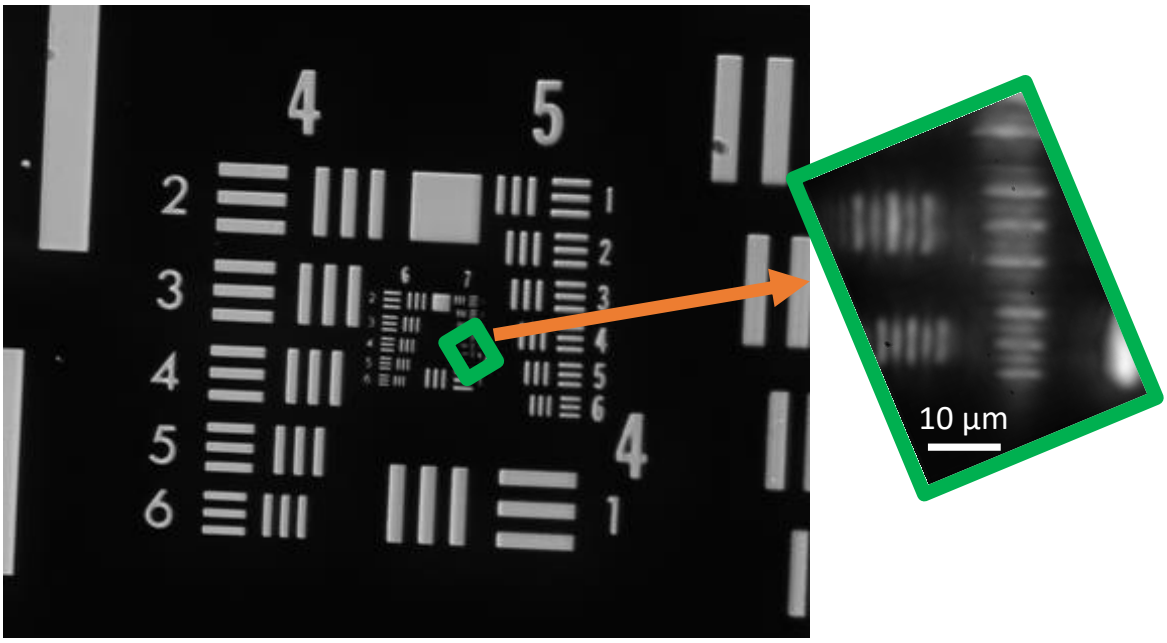

Figure 4.16: Images of the resolution test target.

with a 20× microscope objective as the second stage of magnification.

Imaging Point Sources

Directly measuring the point spread function is the ideal way to characterize the imaging system. To do this, we needed to find a point source, either a point of light or a shadow, that is much smaller than the resolution of the system. We tried a variety of different point sources, none of which gave us a definitive measure of the point spread function.

We tried to image the shadow of gold nanoparticles on a glass slide. Although we could see individual nanoparticles, we saw sections of the image that had a higher light intensity than the light background, which is nonphysical if the particles are simply blocking the light. We realized that gold nanoparticles are optically active at 780 nm [83], which distorts the image. A cross-section through the image of a 100 nm gold nanoparticle scanning the focus for a series of different focal lengths over a 75 μm range can be seen in Figure 4.17.

We also tried to measure the point spread function by imaging a point source in the form of a tapered fiber tip (MF004). The taped fiber tip, typically used for scanning near-field optical microscopes, is a fiber that has been chemically etched to a point and coated with aluminum to leave a 185 ± 35 nm aperture. This method of measuring a PSF has been previously demonstrated in Reference [80] and [84]. In our measurements, the SNOM fiber tip did not look like a point source. The reasons for this are unknown.

This project is on going. Another graduate student has been working on testing the imaging system using an array of apertures in chrome deposited on amorphous fused-silica, made by Photomask Portal. This custom-made test target has two arrays of 0.5 μm clear apertures spaced 10 μm apart and two arrays of 0.5 μm islands of chrome on a clear background.

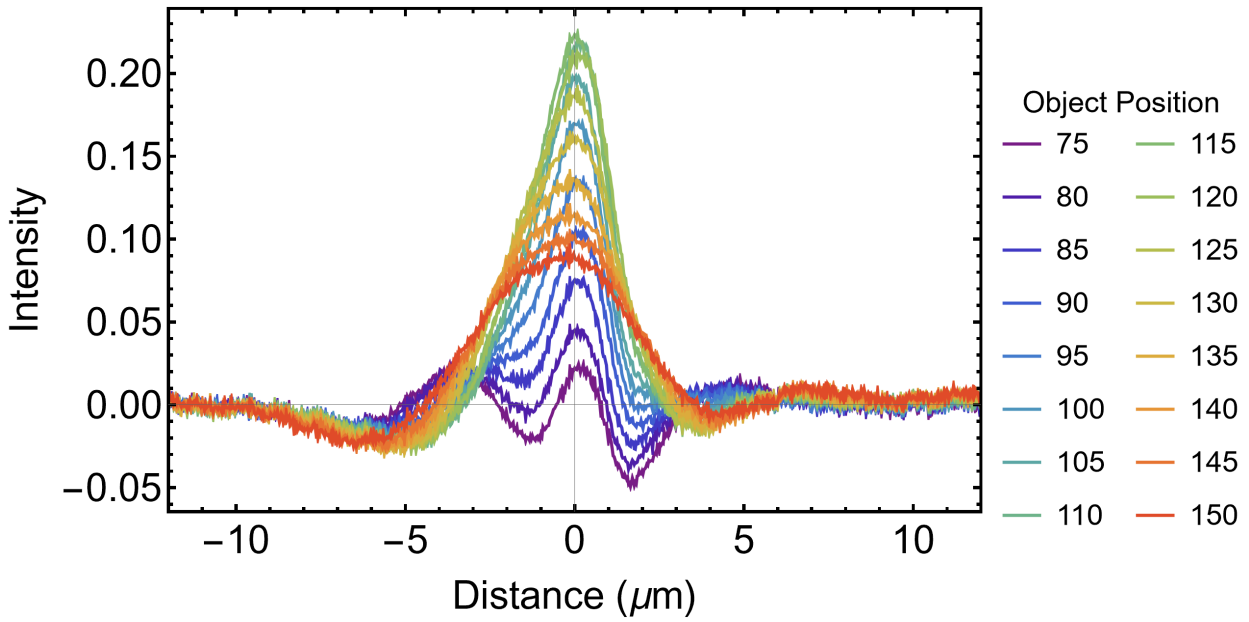

Figure 4.17: Slices through images of gold nanoparticles for a range of object positions. The intensity is the intensity of the shadow of the nanoparticles, i.e. the shadow image subtracted from the background light field.

4.2.5 Other Imaging Systems

The new first stage of imaging will allow us to reach a resolution near the diffraction limit of the system. However, to make use of that resolution, we need a second stage of magnification given the finite size of the pixels in the camera. Adding a second stage of magnification increases the usable resolution of the imaging system, but it also decreases the field of view.

Second Stage of Magnification

Our chosen set of optics for the first stage of magnification can reach a resolution near 1 μm with a factor of 5 magnification. However, with only a factor of 5 magnification, the imaging system would be pixel-limited. We need to increase the magnification even more in order to make use of this resolution.

To accomplish this, we will use a second stage of magnification. The second stage of magnification is a telescope of a $f = 100$ mm (AC508-100) and a $f = 400$ (AC508-400) lens, which adds an additional factor of 4 magnification. That brings the total magnification of the main imaging path to 20×. The imaging path will focus on a Pixis camera with 13.6μm × 13.6 μm pixels and high quantum efficiency. With the 20× magnification, each pixel will image a 0.7μm × 0.7 μm area at the atoms. This is smaller than the imaging resolution of 1.1μm we hope to achieve, meaning the system is not pixel-limited. This imaging scale corresponds to 1.7 lattice sites per pixel.

Imaging with a Large Field of View

This main imaging path will have very high resolution, but a relatively small field of view of 700 μm × 700 μm. This region is larger enough to easily fit an entire quantum gas for density measurements but limits our ability to align elements of the apparatus or use time of flight imaging. In order to maintain the ability to

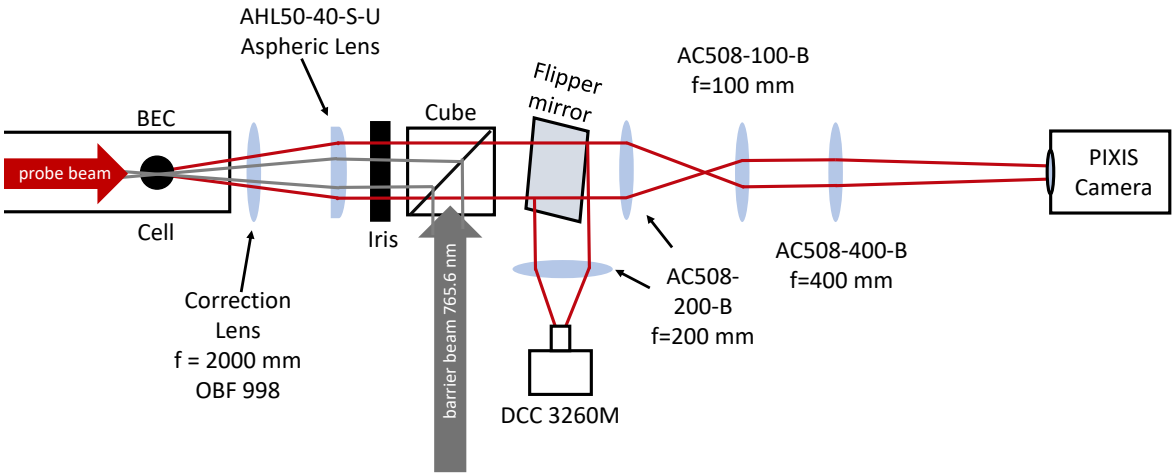

Figure 4.18: Possible scheme with two different imaging systems which preserves the ability to image at high field of view while adding the ability to image with high resolution.

image with a large field of view, we will add an auxiliary imaging path with a lower magnification onto a second camera (Figure 4.18). We can use a flipper mirror to alternatively route the image after the first 5× stage of magnification onto a DCC3260M camera that has a 1936 x 1216 pixel area with 6 μm pixels. With the 5× magnification from the first stage of imaging, we will have a 3.06 mm by 2 mm field of view at the atoms with 1.6 μm x 1.6 μm pixels. The long direction of the sensor will be orientated vertically to maximize the field of view in the vertical direction to image free-falling atoms in time of flight. This field-of-view is comparable to the field of view of our current imaging system (3.3 mm by 1.3 mm).

We chose a DCC3260M Thorlabs camera because it has precise timing capabilities and the highest quantum efficiency ($\sim 30\%$) in its price range. The quantum efficiency (the percent of the incident photons the camera detects) of this camera is much lower than the quantum efficiency (approximately 90%) of our main imaging camera (Princeton Instruments PIXIS PIX-1024BR). The lower quantum efficiency of the Thorlabs camera will lead to higher noise, specifically, the signal-to-noise ratio of the Thorlabs camera is about 40% of the signal-to-noise ratio of the Pixis camera. We will use large field-of-view imaging primarily to determine the location of the atoms, for which imaging noise is of less concern. In high magnification imaging, we will resolve changes to the density distribution. Accuracy in the counts on each pixel is necessary for an accurate density measurement, so we will use the high quantum efficiency Pixis camera for high magnification imaging and the Thorlabs camera for large field-of-view imaging.

Other setups for multiple stages of magnification are also possible. Currently, the image system is still using the original first stage of magnification, but has an additional second stage of magnification that consists of a $f = 100$ mm and a $f = 400$ mm lens that adds a factor of 4 magnification (Figure 4.19). This is the same second stage of magnification as was proposed above. However, instead of using a flipper mirror and a second camera to perform high-field-of-view imaging, this imaging system adds in an additional set of lenses to change the second stage of magnification to a 1× system (Figure 4.20). Specifically, there is a removable $f = 200$ mm lens between the primary lenses of the second stage of magnification and an additional removable $f = 100$ mm lens after the last lens of the second stage of magnification. When all four lenses are in place, the second stage of magnification has a net effect of 1× magnification, and does not appreciably change the size of the image. With the 1× stage of magnification, each pixel images a 2.96 μm square area. The old imaging

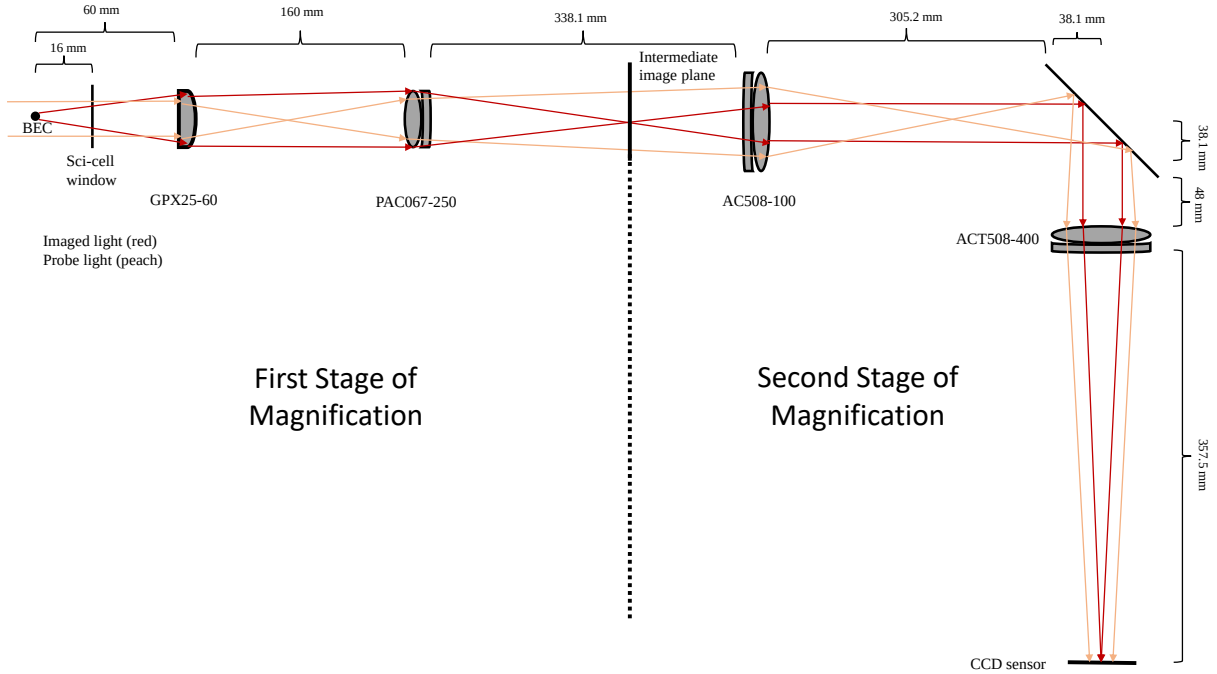

Figure 4.19: Current imaging system with second stage of magnification. Figure courtesy of lab member Max Gold.

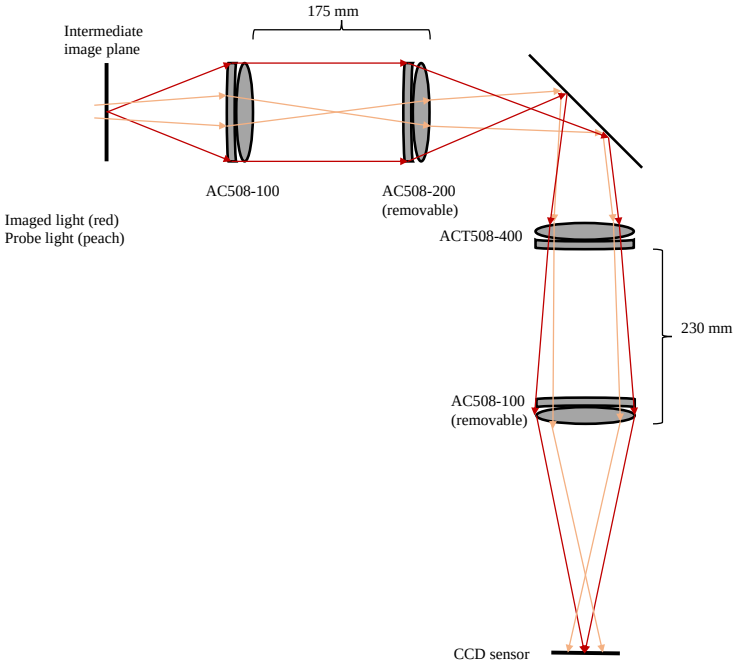

Figure 4.20: The second stage of magnification of the current imaging system with the removable lenses. The net effect of this system is a factor of 1 magnification. Figure courtesy of lab member Max Gold.

system was very similar at 3.06 μm per pixel. The additional lenses in the second stage of magnification are used when a large field-of-view is required. Otherwise, images are taken with the higher magnification. With this high-magnification system, each pixel images a 0.78 μm $\times$ 0.78 μm area at the atoms, which is a much smaller area than the old imaging system. However, the resolution is still limited to, at-best, the diffraction limit of the first imaging lens, a Rayleigh radius of 2.3 μm. This means that information on each pixel is also included in its adjacent pixels, blurring the image. This imaging system allows us to gain more information about features down to a few microns, compared to the previous imaging system — although it is not capable of resolving features smaller than about 2 μm.

Side Imaging

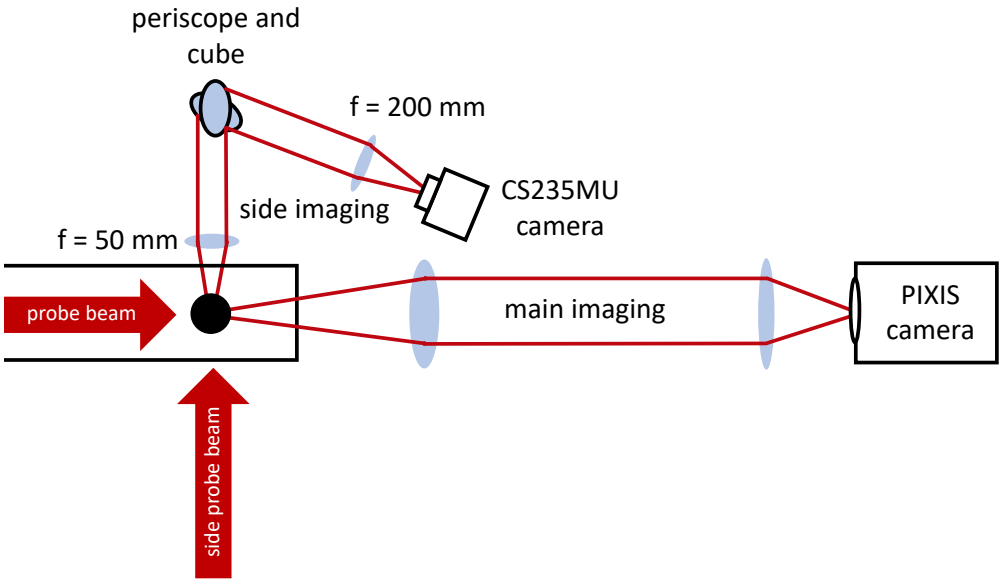

Figure 4.21: Overhead view of beam path for the side and main imaging.

Changing the main imaging also necessitates changing the side imaging. In the old imaging set up, we used the main camera to take images from two different directions. In addition to the main imaging that has been discussed, we also illuminated the atoms from the East side of the science cell and collected the image of the shadow with a lens on the West side of the cell. The light was then set through a series of lenses and mirrors to a half-mirror in front of the main camera, which sent the side image onto the CCD of the main imaging camera. More information about the old side imaging system can be found in Reference [58]. In order to use the full array of the main camera for high-resolution imaging, we decided to use switch to using an additional camera for the side imaging.

The new siding image uses part of the same optical beam path as the old side imaging beam path. The first imaging lens is the original $f = 50$ mm lens next to the science cell. Then the beam is sent down to the table with periscope. It also goes through a polarizing cube beamsplitter that can be used to couple another beam into the beam path. The light then goes through a $f = 200$mm (LA1979-B) lens. Then the light is imaged on a Thorlabs camera (CS235MU) (Figure 4.21). The magnification of the new side imaging is 4.67$\times$ and each pixel images a 1.26 μm square in the imaging plane. This beam path is much shorter than the old

side imaging beam path, which went through two more lenses and several more mirrors before reaching the main camera. Switching to using an independent camera for the side imaging required not only setting up a new beam path, but also changing our control software to accommodate more than one camera.

4.3 New Camera Software

One significant change made to the imaging system of the apparatus is the software controlling the cameras. The old system consisted of WinView, the Princeton Instruments (the camera manufacturer) software to interface with the camera and a Visual Basic program to pass information between the camera software, Matlab, where we run the fitting, and Origin, where we store and analyze the results of the fits. This system was only capable of interfacing with a single camera and would only run on a 32-bit computer. In order to use multiple cameras and switch our imaging computer to a modern 64-bit system, we developed new imaging software for our system.

We developed new camera control software that can control an arbitrary number of cameras of different types and runs on a 64-bit computer.[2] This software is mainly written in Matlab. It takes information from the sequencer, which controls the apparatus, and uses that information to control the cameras. It also processes the images in the way specified in the sequencer code. This software runs on a different computer than the sequencer computer that controls the apparatus.

4.3.1 Software Overview

Major Components Involved in the Imaging Software

The new imaging software consists of three main components: the sequencer, the Matlab program, and the cameras. First, I will discuss the general role of each of these components before describing the details of how they work together to run the imaging system for the apparatus.

The sequencer is the program that controls our apparatus. It sends instructions to the various components of the apparatus at specific times in order to orchestrate the creation and manipulation of an atomic gas. The sequencer code is the only code we change when adjusting the details of an experimental sequence. Besides manipulating the controls of the apparatus with the sequencer code, we also give commands for which images to take and what kind of image processing steps to perform. This way, only the sequencer code is changed between runs of the apparatus.

The Matlab program takes in this information from the sequencer and controls the cameras and processes the images accordingly. We use a few additional pieces of hardware and software to facilitate communication from the sequencer to the Matlab program and from the Matlab program to the cameras. These will be discussed in more detail in later sections.

The cameras are integrated into the apparatus and take images of the atomic gas. The cameras are connected to the imaging computer and are controlled by the imaging software. The timing of when the images are taken is controlled by a trigger signal that comes directly from the sequencer.

The Architecture of the New Imaging Software

In this section, I will discuss the over-arching flow through the imaging code at a high level. Later sections will include more information about how each of these steps are implemented. The order of the major steps

[2]The code for this software is available in on the DeMarco Group Gitlab.

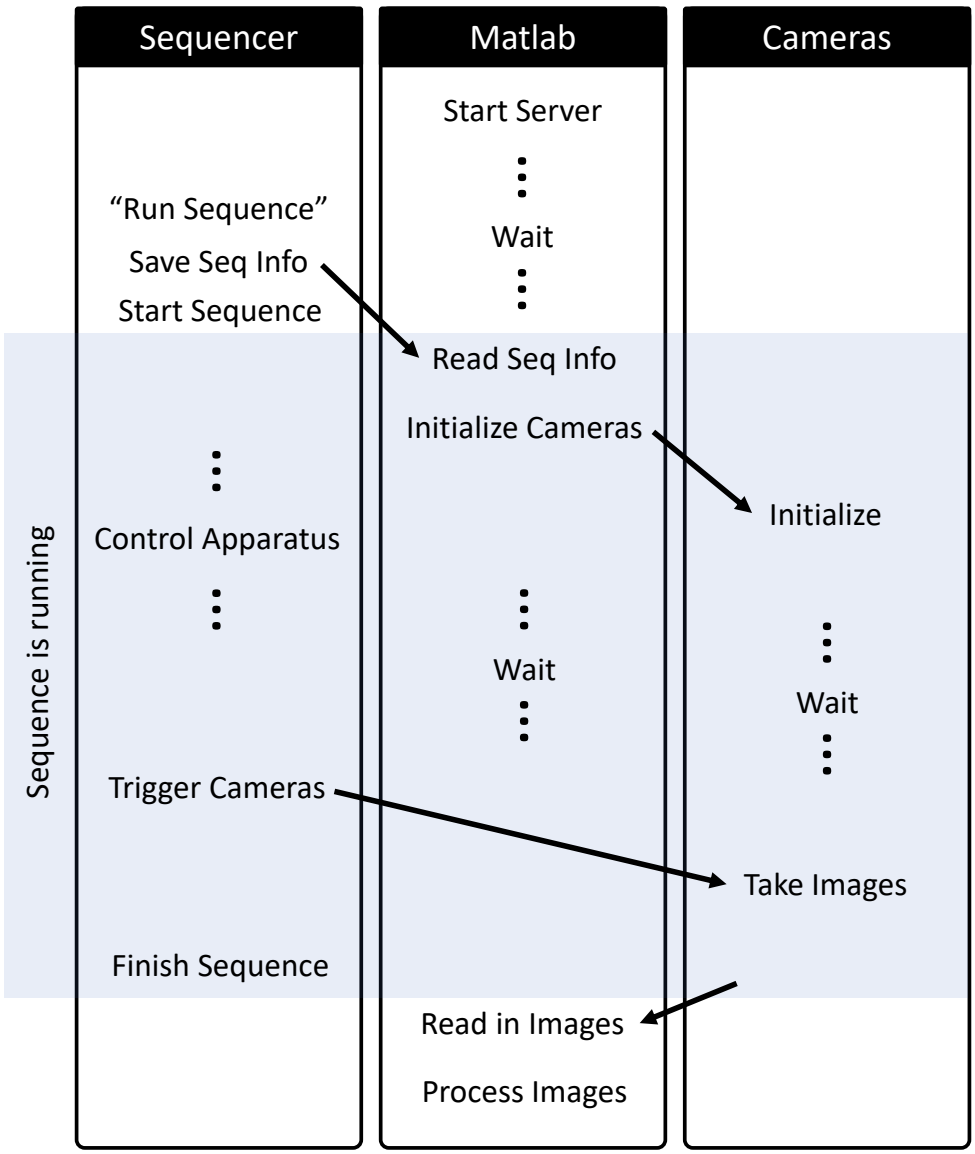

Figure 4.22: Overview of the new imaging software. Before running the apparatus, we start that Matlab program that acts as a server to receive information from the sequencer and give instructions to the cameras. The timing of the imaging software is controlled by the sequencer so that each step happens at the correct time with respect to the run of the apparatus (shaded region).

of the imaging software are shown in Figure 4.22.

First, the program is started from within the Matlab interface. This Matlab program acts like a server, interfacing with the various components of the imaging software. This step only needs to be done once per day, when turning on the apparatus. At this point, the Matlab program starts waiting for the sequence to start. Next, a sequence is started by the sequencer. This triggers a few steps. First, the sequencer saves the information about this particular sequence. Then the Matlab program sees that the sequence has started and reads in the information saved by the sequencer. The Matlab program uses this information to determine which cameras to initialize. It initializes each of these cameras, which start waiting for a trigger to take an image. At this point, the Matlab program waits until the sequence is over. In the meantime, the sequence has been performing the necessary steps to make and manipulate an atomic gas. When the sequence reaches the imaging step, it sends a trigger to the cameras. The cameras respond by taking images. Next, the sequence finishes and the Matlab program reacts by reading in the images and performing the specified image processing. At this point, the Matlab program starts over and waits for the next sequence.

4.3.2 Code Organization

The majority of the new imaging software is written in Matlab. Matlab is a computing environment and programming language that handles numerically processing data in matrix forms very well and has a wide variety of add-ons, including many to interface with various hardware components. The imaging program is run from within the Matlab interface. The code is organized into many functions that are called as needed.

This software uses Matlab's data type "structure array" to organize and pass variables. Instead of relying on global variables to access information in multiple parts of the code, the information is stored in a few structure arrays that are passed into functions as necessary. The function can then change the information on the object and pass it back out for use with the next function. There is a configuration structure array, `configObj`, that stores information global to the whole program, such as file paths and which apparatus is running the software. Each camera also has its own structure array that keeps track of all the settings for that camera and stores the images.

This software is stored on the lab's GitLab in the project "multi-camera-imaging-software" as well as on the lab server. The folder `Z:\software\Multi-Camera_Imaging_Software_Folder` contains the code stored on GitLab, ".dll" driver files, and the folder `control_PI_camera` that contains the code for controlling the Pixis camera (this folder is rarely changed and contains files not allowed on the GitLab server). Using Git version control allows us to make a new branch when making a significant change or adding a new feature to the software and not affect the branch of the software that is used on the apparatus until it is ready for integration. Version control also provides a record of the changes that were made to the software. Storing the software on GitLab allows it to be easily kept in sync between both apparatuses in the lab. When a change is make to the software, it can be pushed to GitLab and then pulled to each apparatus. Since the software works nearly identically on both apparatuses, we wanted to keep the software versions in sync so both apparatuses benefit from improvements made to the software on either apparatus. The differences to settings between the apparatuses are accommodated by setting the values of `configObj` based on the text file stored on the computer that contains either "Rb" or "K" to indicate which apparatus is running the software.

4.3.3 Initialization

The highest level function of this software is `start_imaging.m`. This function performs a few simple steps to initialize the software. One of the most important initialization steps is connecting to the Arduino which facilities communication from the sequencer to the Matlab program (Section 4.3.4). This function also configures the software to work with the correct apparatus and creates a folder for the current day to save data. Additionally, the `start_imaging.m` function starts an instance of Origin, the program we use to view and process the data (or connects to the open instance if one already exists). Besides these crucial tasks, this function also does a few house-keeping steps that make the software more robust against different starting conditions of the Matlab software. The `start_imaging.m` function closes any preexisting connections to cameras. This is useful if the program was previously exited unexpectedly; normally, the program closes the connections to the cameras after it reads in the images. The `start_imaging.m` function also adds the path of the file to the search path for Matlab so that Matlab can find all the functions of the program. This step can also be done manually through the Matlab interface. Lastly, this function declares the window size of the images, which helps facilitate running statistics on the images in a way that persists over multiple runs of the apparatus.

After these initialization steps, the software calls `take_and_process_images()` in a while loop. This function is the workhorse of this software and handles an entire run of the apparatus. For as long as `start_imaging.m` is running, it repeatedly calls `take_and_process_images()` which handles every repeated action of the software.

The `take_and_process_images()` handles timing the function based on when the apparatus is running, logging the MOT signal, reading in commands for the sequencer, initializing the cameras, reading in the images, displaying the images in a Matlab graphical user interface, and running any functions from the sequencer, including fitting and saving the images and writing the results to Origin. The details of these steps are discussed below.

The `take_and_process_images()` function includes a flag that states if the apparatus is running. This is simply for testing. If the flag is zero, the function can be run without waiting for the Arduino triggers and it initializes the software at the start of this function, so that it can be run independently of `start_imaging.m`. Whenever this software is being used in the apparatus, this flag should be set to one.

4.3.4 Communication from the Sequencer to Matlab

Communicating Timing Information from the Sequencer

The sequencer computer controls the apparatus and determines the timing of the experimental sequence. In order to handle the timing of the cameras correctly, the Matlab program needs to know the current stage of the experimental sequence. This information is relayed to the Matlab program from the sequencer through an Arduino.

The timing of the new imaging software is controlled with a digital signal from the sequencer that is read into the imaging computer by an Arduino Uno as seen in Figure 4.23. The sequencer is programmed to set a digital line to "high" when it starts a run (when the apparatus captures from the MOT and moves the atoms to the science cell) and set the same line back to "low" at the end of the run after the apparatus takes images, moves the cart back to the collection cell, and resets to fill the MOT again.

This signal is connected to one of the input pins of the Arduino. The Matlab program creates a connection to the Arduino when the software is initialized and repeatedly checks the input pin and reads in the value

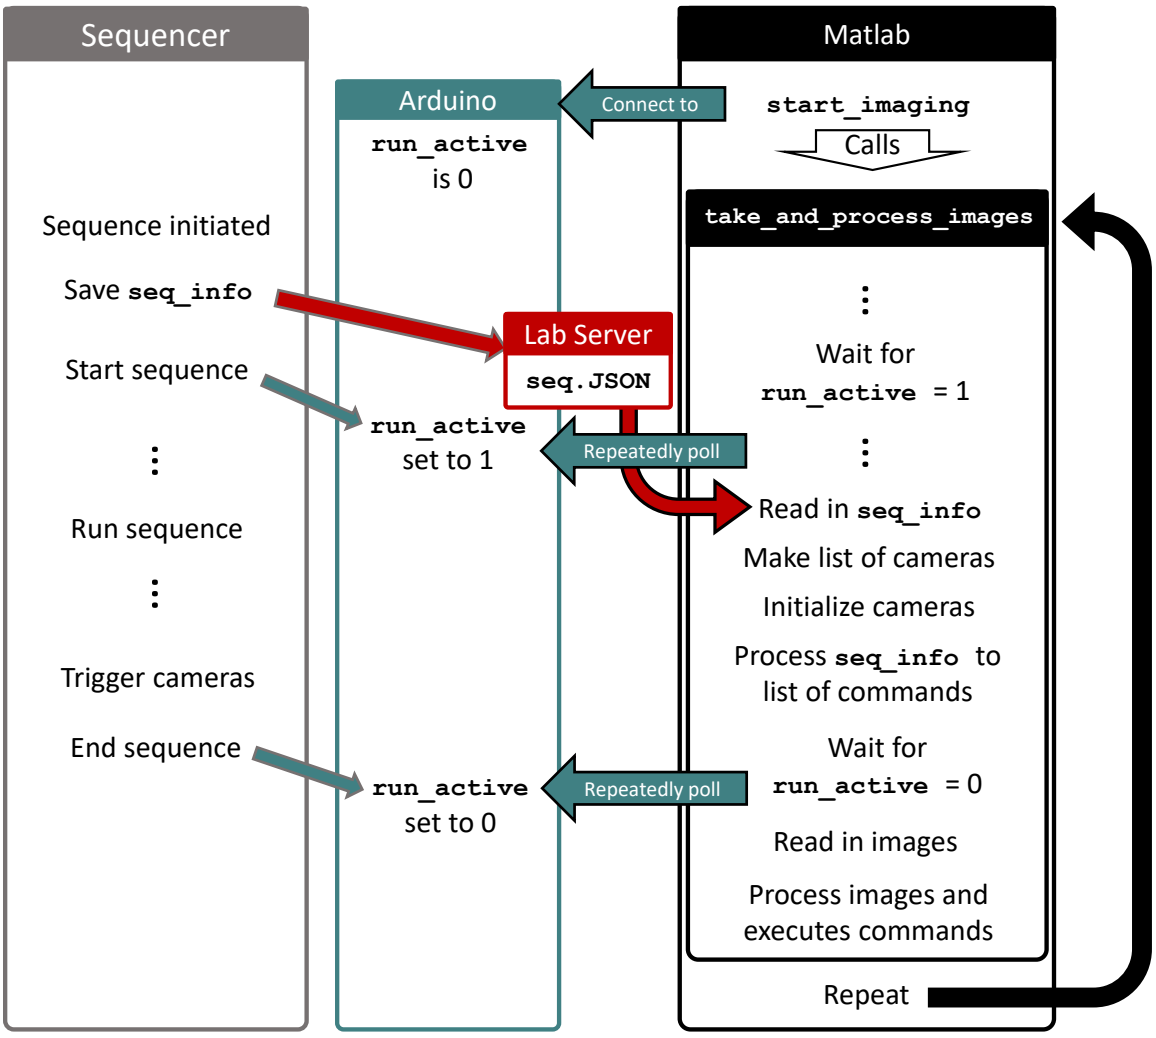

Figure 4.23: Communication from the sequencer to Matlab is facilitated by an Arduino and the Lab server.

to the variable `run_active`. The Arduino is repeatedly polled (waiting 1 ms between checks) whenever the software is at a step where it needs to wait for the apparatus to change states before moving on. The Matlab software uses the Matlab add-on "MATLAB Support Package for Arduino Hardware" to communicate with the Arduino.

It is important to note that the Arduino signal is not telling the camera when to take the image, which needs to be controlled more precisely than is possible through the imaging computer. The timing of the images is handled by a digital trigger directly to the camera. The Arduino signal just tells the program when the sequence starts and stops running. When the sequence starts, the Matlab program uses information from the sequencer to determine which cameras to initialize so that the cameras are ready to take an image. The Matlab program waits until after the sequence is finished to read in the images from the cameras, which guarantees the images are ready.

Importing Sequence Information from the Sequencer

Besides timing information from the sequencer, Matlab also needs the information about the conditions of each run of the sequence. This information includes the values of parameters of the experimental sequence, such as the hold time or time of flight for a given run. These values are used to label the images produced by that run. This information also includes the commands given in the sequencer. The post-processing steps of the imaging software are determined by commands typed directly into the sequencer code. The most common command is `runfit()`, which tells Matlab information about what fit to run on the images. These commands are discussed in more detail in Section 4.3.7.

Since the sequencer and the imaging software are running on different computers, we needed to find a way to pass information between these computers. We do this by writing the sequencer information to a JSON file on the lab server. Both computers have a drive mapped to the lab server, so the sequencer information can be written to a specified file path on the server and the Matlab program can simply read the file from the server at the start of each run of the apparatus.

The JSON file transfer of information specifically works with the new sequencer, which we started using in 2021. We started using the new imaging software before we switched to the new sequencer. The first version of the new imaging software ran an executable created by the old sequencer as a way of passing information. This executable was also stored on the lab server.

The new imaging software brings in the information from the sequencer in the function `import_seq_info()`. This function finds and reads the JSON file saved to the server by the sequencer. It then parses the file into a list of variables `seqexevars` which contains information about time of flight, repump time, and any other parameters passed from the sequencer. It also saves the list of commands from the sequencer to the variable `commands`. These two variables are returned from this function back into `take_and_process_images()` where they are used for later steps of the process.

4.3.5 Logging the MOT Signal

Besides controlling the cameras of the apparatus, the imaging software performs the added function of logging the MOT signal. This is new functionality over the old imaging software. While the software is waiting for the run to start (that is, while `run_active` is zero), the software reads in the voltage from the MOT level signal and checks if the run has started every 50 ms. After the run starts, the software saves this list of MOT level voltages and the corresponding times to a CSV file with the current time in the file name. These files

can be processed later to find the fill time and equilibrium MOT level for each run of the apparatus as a way of monitoring the MOT over time. The usefulness of this information is explained in Section 3.3.

4.3.6 Controlling the Cameras

The new imaging software is designed to control an arbitrary number of cameras. The names and settings for each of the available cameras are given in the function `init_camObject`. New cameras can be added by adding a new case in this function. The imaging software determines which cameras to use for a given run based on which camera names were included in the commands passed in from the sequencer. Each command from the sequencer has an associated camera name. After reading in the list of commands for the sequence, the software finds the unique camera names in the list of commands. It creates a structure array for each camera and stores all of these structure arrays in a list. Anytime the imaging software interacts with the cameras, it loops over this list to act on every camera. The software was originally written with support for two kinds of cameras: a Thorlabs DCC3260M camera and a Pixis camera from Princeton Instruments. Figure 4.24 shows the steps involved in controlling these two cameras.

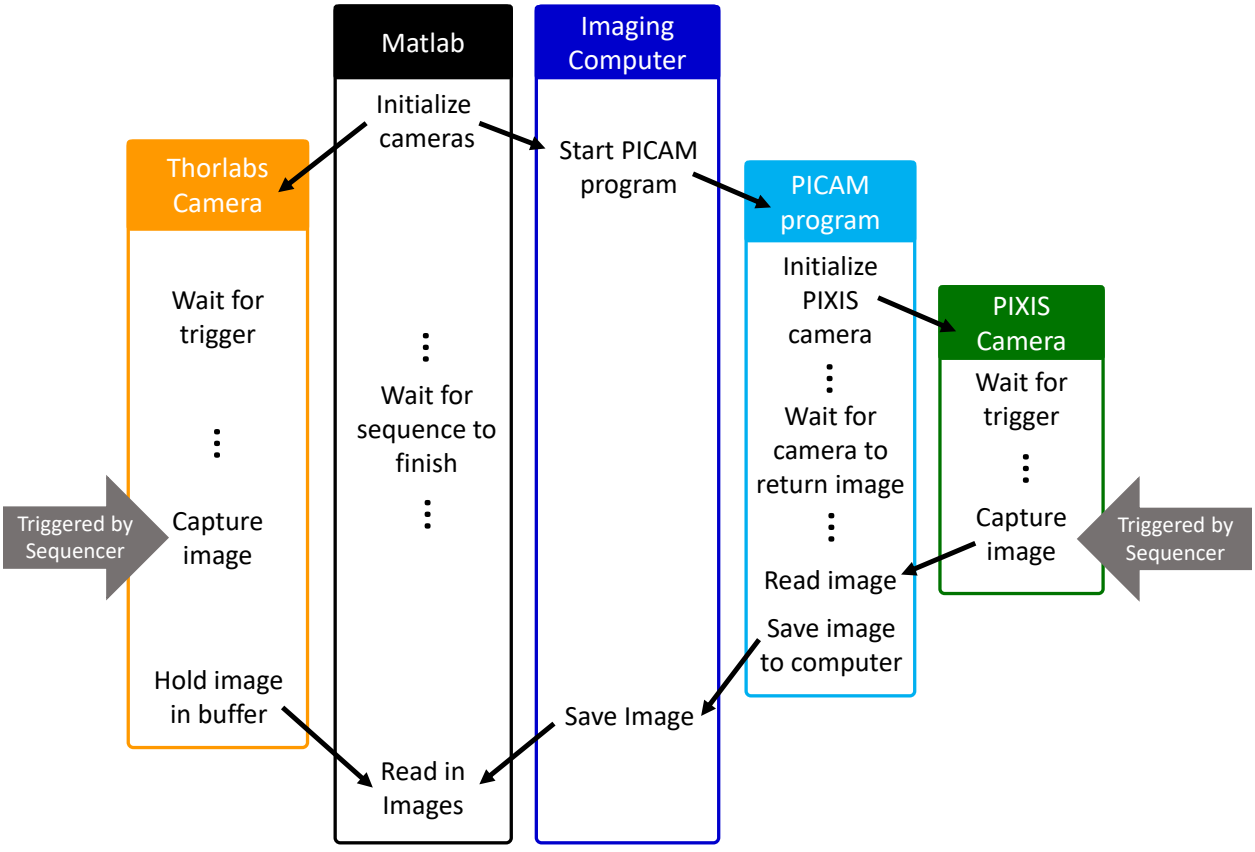

Figure 4.24: Controlling the cameras. Matlab controls the Thorlabs camera directly with the ".NET" programming interface. The PIXIS camera is controlled by launching a separate executable function created with the PICAM software development kit.

Control of the Thorlabs camera is done directly in Matlab through the ".NET" programming interface. Drivers and documentation about this interface are available with Thorlab's Windows Software Development

Kit (SDK) for DCx series cameras. The interface provides commands for controlling the camera that can be used in Matlab code. Figure 4.24 shows the program controlling one Thorlabs camera, but any number of additional Matlab-controlled cameras should be possible.

The Pixis camera does not have direct Matlab integration, and controlling it required a more complicated scheme. The Pixis camera is controlled by executing a PICAM program in the command line of the computer from within Matlab using Matlab's `dos()` command. PICAM is the driver for the Princton Instrument cameras, including the Pixis camera, and can be accessed with the PICAM SDK.

The PICAM program to run the Pixis camera is an executable called `control_PI_camera` in a folder by the same name. The settings of the camera (exposure time, number of frames, etc.) are passed in as arguments of the function this executable runs. Matlab launches a command prompt window and runs the PICAM program, which initializes the camera. After the camera is triggered by a digital signal from the sequencer, the executable saves the image with the requested number of frames to the computer. Matlab later reads in that image.

As long as certain parameters for the camera stay the same, this executable will not need to be changed. Some of the parameters are inputs to this function and can be changed without generating a new executable. Specifically, the inputs to the executable are: the folder name to save the data (with a slash at the end), the file name to save the data, the pixel coordinates of the region of interest (start x, start y, stop x, and stop y), exposure time in ms, and number of frames. For example, running the application from the command line may look like:

```
path\control_PI_camera.exe path\control_PI_camera\ active_PIXIS_image 0 0 1024 1024 10 3
```

Note the space between the folder name and the file name and that the region of interest has to start at zero (start x = start y = 0). It may also be necessary for the stop x and stop y values to be the full range of the sensor, 1024 pixels.

If changes other than these parameters are needed to the control of the Pixis camera, it will be necessary to generate a new PICAM executable. The C++ source code for this executable and notes about how to run it can be found in the Cad Share folder `Z:\software\Multi-Camera_Imaging_Software_Folder\control_PI_camera`. The executable was made with the PICAM SDK, Princeton Instruments' software development kit. The Installer and Manual for PICAM is saved in Cad Share folder, `Z:\software\PICAM Princeton Instruments SDK`. This source code is a modified version of "MATLAB-For-Picam" by Joe Lowney [85] which is also saved in `Z:\software\PICAM Princeton Instruments SDK`.

All of the control of the cameras is called from within the `take_and_process_images()` function. After finding the list of cameras from the list of commands from the sequencer, the function initializes all the cameras. It then processes the list of commands into a new format. This new format saves the results of each command to the structure array associated with the camera of that command. Then the function waits for the apparatus to finish running.

After the run of the apparatus is finished, `take_and_process_images()` reads in and processes the images. For the Pixis camera, the Matlab program reads in the image from the file on the computer where the PICAM program saved the image. For the Thorlabs camera, the program reads the image out of the buffer of the camera. For each camera, the program creates a window to display the images with the function `display_raw_data()`.

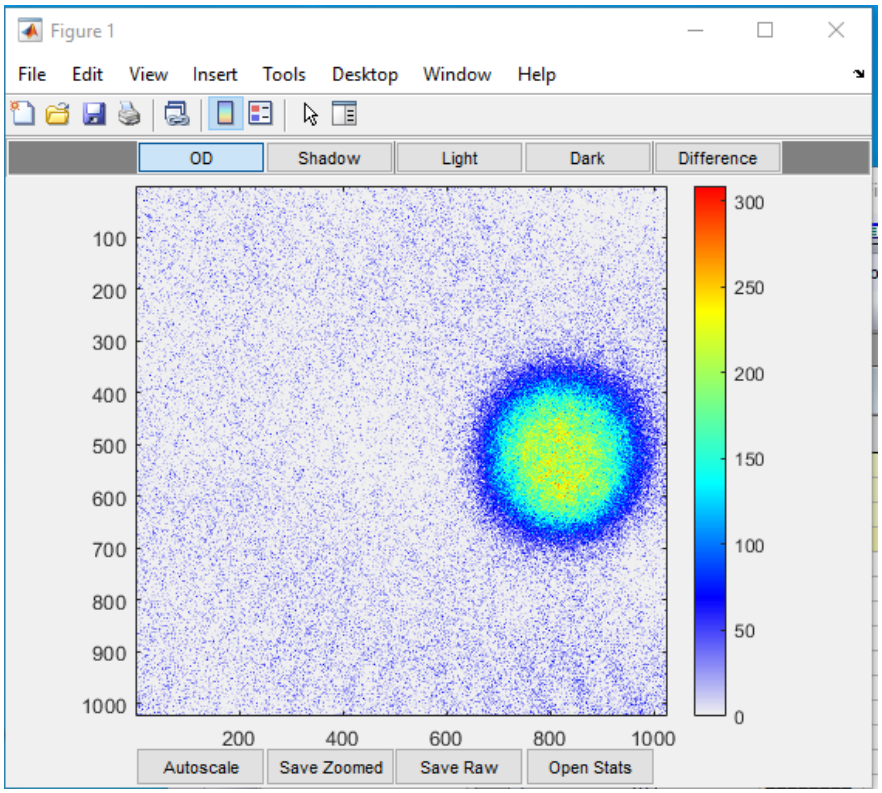

Figure 4.25: Interface for viewing images in the new imaging software

4.3.7 Processing and Analyzing the Images

Displaying the Images

For each camera, the imaging software launches a window to view the image collected with that camera. An example of this window with an in-trap image of a large thermal gas is shown in Figure 4.25. Each optical depth image consists of three constituent images (Section 3.6). This window has a tab for each image: the optical depth (OD) image that represents the density distribution of the atoms, and the shadow, light, and dark images that make up the OD image. It also has a tab for the difference between the light and shadow images. Matlab's built-in image tools such as zoom and pan can be used to look at these images. The data cursor feature can also be used to find the coordinates of a selected pixel in the image.

Besides the built-in Matlab tools for examining an image, there are also a few custom options in this image viewer. These buttons are below the image. The "Autoscale" feature changes the color scale to range from 5% to 95% of the range of pixel values within the currently displayed window.[3] This feature is used after zooming in on a region of the image. The "Save Zoomed" feature saves the currently displayed image. The "Save Raw" feature saves all three frames of the raw image.

The last feature, "Open Stats", opens a new menu within the image viewer that is used to find statistics on the image. The image viewer with the statistics menu open can be seen in Figure 4.26. This feature calculates the mean and standard deviation of the counts on each pixel within a specified region. It is most commonly used to find the average counts in a region of the light or background image for adjusting the

[3]At time of writing, the numbers of the color scale in the image viewer are arbitrary and do not match the true OD values of the image. The fits return the correct OD values.

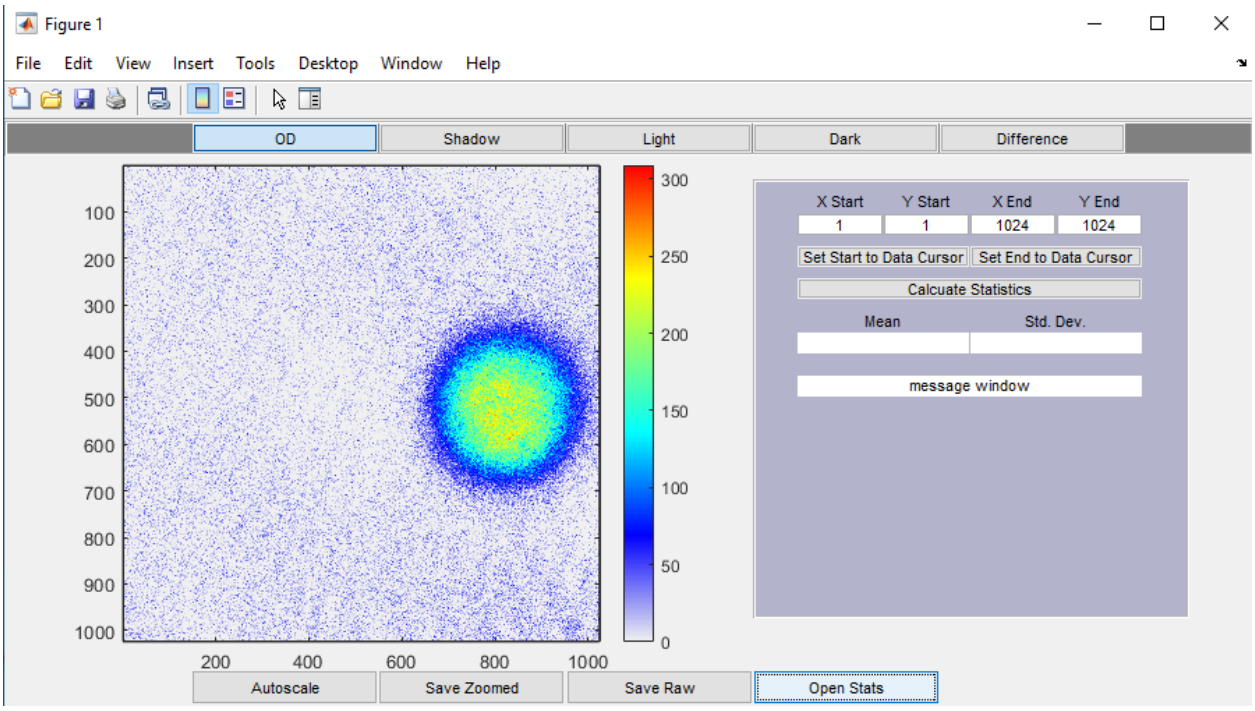

Figure 4.26: Interface for viewing images in the new imaging software with the "statistics" menu open

probe beam or determining if background light levels are too high. The region of interest to average over defaults to the whole image. It can be changed by typing in new values or setting the upper left or lower right corner of the region to the position of the data cursor. This region of interest persists over multiple runs of the apparatus.

Fitting and Processing the Images

Lastly, the new imaging software runs the commands from the sequencer. The available commands are `runfit()`, `showres()`, `writetoOrigin()`, and `writetoFile()`.

The `runfit()` command is almost always called for every image collected. The results of the fit are used by most of the rest of the commands. This function runs one of a library of fitting functions on the image and returns the results. These functions are the same fitting functions as were used in the old imaging code. An explain of how `runfit()` might be called is: `runfit(PixisRb, 1,''gauss'', ''fix'', ''offset'', ''ROI'', [130 1; 1020 450])`.

The first argument to any of these commands is always the name of the camera it applies to. In this case, the camera name is `PixisRb`. These camera names are defined in `init_camObject.m` and contain information about the settings and type of each camera.

The second argument is either a number or the name of the fit to run. If the second argument is a number, more than one fit is being run for this camera, and the number indicates which fit this command refers to. Calling multiple fits for one camera is most commonly used to fit both parts of a gas separated into different spin components.

The next argument (or the second if only one fit per camera is used) is the name of the fitting function to use. The one given in this example is `''gauss''`. There are also functions to fit BECs, partially thermal gases, rotated Gaussians, and a number of other distributions. The available fitting functions are in the `fits`

folder of the imaging processing software. These functions define not just the function to fit to, but also the Jacobian which defines the derivatives used by the fitter.

The remaining arguments are all optional and are given in pairs. The first argument of a pair defines the option, and the following argument gives the values for that option. In this example, ‘‘fix’’ is telling the fitter to fix (hold constant) the parameter ‘‘offset’’. The next set of options is defining a range of interest (ROI) so that the fitter only fits the part of the image within that range. The ROI is given in the format [start_x start_y; end_x end_y]. Another commonly used optional argument is ‘‘AutoROI’’, which automatically picks a ROI of the size specified by the following argument in the form [x_size, y_size].

The showres() command displays plots showing the residuals of the fit generated in runfit(). The writetoOrigin() command writes the results of the fit to Origin. It generates (or adds to, if it already exists) an Origin worksheet for each camera, labeled with the camera name. Typically, the program writes to the same worksheet over multiple runs of the apparatus. Changing the name of the worksheet will cause the Matlab program to create a new worksheet with the old name. The writetoFile() command saves the image to a file. The file path for these files is created in init_configObject.m based on the apparatus and the current date.

After take_and_process_images() finishes running the last command from the sequencer, it is finished running this instance of the function. The program returns out to start_image.m where a new instance of the take_and_process_images() function is immediately called and begins reading in MOT data and waiting for the digital signal that indicates the apparatus has started a new run.

4.3.8 Current Use

We switched to using this imaging software in 2019. We generally stop the start_imaging.m program at the end of each day and start it again when we turn on the apparatus the next day. Running this program indefinitely causes issues with a time-out of the connection to the Arduino. The software was originally tested with a Pixis camera as the main imaging camera and a Thorlabs DCC3260M camera as an additional camera. Since then, the Thorlabs DCC3260M camera has been discontinued. The CS235MU Thorlabs camera is the successor of the DCC3260M camera and was installed as the new side imaging camera. The new imaging software was designed to be camera agnositic, and able to handle this kind of switch. Some small modifications were made to the software to accommodate this new camera.

Chapter 5

Understanding the Thermal Atomic Gas

In this chapter, we will describe initial density distribution and the presence of localized states. We will give the parameters of the initial gas and describe how we model the gas to understand the density distribution. We will also discuss the temperature of the gas and how we measure it, as well as how we expect the temperature to change over time. Lastly, we will describe the fraction of the particles that are in localized states and the multiple ways for calculating this fraction.

5.1 Creating the Initial State

The starting state of the experiment described in this work is a thermal gas trapped in a hybrid optical dipole and magnetic quadrupole trap (Section 3.4.3). This gas is initially in equilibrium with the barrier beam potential (Section 4.1) and a $s = 4\,E_R$ lattice potential (Section 3.5). The creation of the initial density profile is illustrated in Figure 5.1.

The thermal gas is composed of $(61\,000 \pm 6000)$ atoms confined in an approximately harmonic trap with a (54.6 ± 0.4) Hz geometric mean trap frequency. To create a gas that is initially in equilibrium, the gas is evaporatively cooled in the presence of the optical barrier, which excludes atoms from a central region. For the measurements discussed here, the optical power is kept fixed, resulting in a barrier with a peak potential of $(9000 \pm 5000)\,k_B$ nK. This energy scale is much larger than the (115 ± 10) nK temperature of the gas before the lattice is turned on (Section 5.3). Given this condition, a hard-wall potential is formed, and the atoms are completely excluded from a cylindrical region with a 14 d radius (where $d = 406$ nm is the lattice spacing) that penetrates through the gas along the imaging direction (the z direction). The size of the region of exclusion is calculated in Section 5.2. After creating a thermal gas, the lattice potential is smoothly ramped on over 200 ms to $s = 4\,E_R$. The temperature of the gas in the lattice is (210 ± 40) nK, which is determined by fitting the tails of the density distribution (Section 5.3).

To study relaxation at higher lattice depths, the lattice potential is quickly increased over 0.4 ms, which is slow enough to avoid band excitation but too fast to allow the density profile to adjust. Therefore, for all data in this paper, the initial density profile is approximately fixed to the distribution realized at $s = 4\,E_R$.

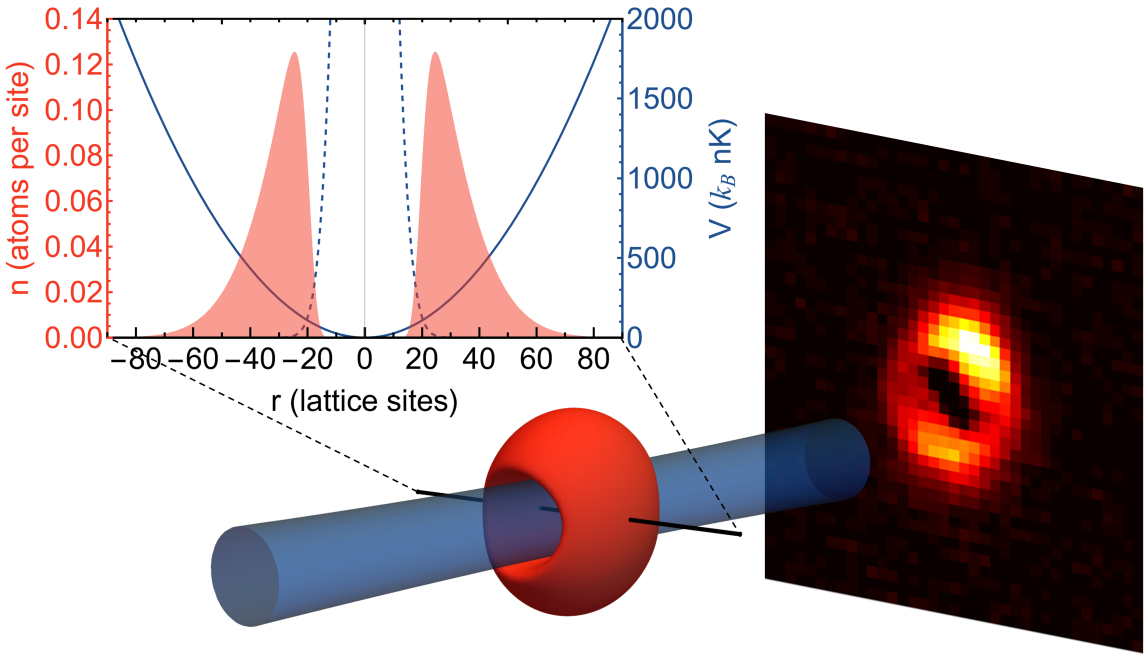

Figure 5.1: Creating the initial far-from-equilibrium density profile. After trapping and cooling a thermal gas (red) in a harmonic potential with a repulsive optical potential (blue) present, a cubic optical lattice is slowly turned on. The plot shows a slice through the predicted density profile (filled red curve) along with the parabolic trapping (solid blue) and barrier (dashed blue) potentials. A sample column-integrated image taken of the initial state at $s = 10\,E_R$ is displayed. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

5.2 Modeling the Density Distribution

5.2.1 Atomic Limit Calculation

The density distribution of the initial state was modeled using an interacting atomic limit calculation. The distribution was discretized into lattice sites. The potential energy, including the harmonic confining potential, the barrier beam potential, and the additional confinement from the $4\,E_R$ lattice beams, was determined for each site. The potential energy at each lattice site and interaction energy for each occupation (0-3 particles) were used to calculate the Maxwell-Boltzmann probability of each state at the experimentally measured temperature.

In the atomic limit, we assume the particles are localized to individual lattice sites. In this limit, tunneling energy is neglected, and we can define the energy for each atom at each individual lattice site. The energy per atom is described by

$$\frac{E(x,y,z)}{n_{x,y,z}} = \left[U\,\frac{n_{x,y,z}-1}{2} + \frac{1}{2}m\,d^2\left(\omega_x^2\,x^2 + \omega_y^2\,y^2 + \omega_z^2\,z^2\right) - V_{\text{barrier}}(x,y)\right] \tag{5.1}$$

where x, y, z are quantized to the lattice sites and describe the position of the lattice site in the East-West, Vertical, and North-South directions, respectively. The occupation in the lattice site at position x, y, z is given by $n_{x,y,z}$. The interaction energy is U, m is the mass of the atom, and d is the lattice spacing.

The trap frequencies (including the confinement from the lattice given by Equation 3.5) along direction i are given by ω_i. In this calculation, without a lattice potential, the vertical trap frequency is $\omega_y = 73.6 \times 2\pi$ Hz and the horizontal trap frequencies are $\omega_x = \omega_z = 47.0 \times 2\pi$ Hz. This model is axis-aligned and cylindrically symmetric around the vertical axis.

This model is a good approximation of the trap based on the measured trap frequencies (Section 3.4.4). The measured principal axes are approximately, but not exactly, aligned to the vertical and horizontal directions. The measured trap frequency of 73.9 Hz contributed 80% of the amplitude of the observed vertical oscillation of the trap frequency measurement and very little (0% and 8%) of the amplitude of the horizontal oscillations, meaning this principal axis is primarily along the vertical direction. In this calculation, we take the approximation that one of the principal axis is exactly vertical, with a trap frequency of 73.6 Hz. The slight discrepancy between the measured trap frequency (73.9 Hz) and the trap frequency used in the model (73.6 Hz) exists because an updated fitting procedure was used to find the trap frequencies after the density calculation was performed.

The two remaining principal axes are approximately in the horizontal plane, although not axis-aligned. The motion of the trap frequency measurement in the x and z horizontal directions had approximately equal contributions to their amplitudes from the 45.7 Hz and 48.2 Hz oscillations, so they were averaged to find one trap frequency for both horizontal directions. Taking the horizontal trap frequencies to be equal creates a symmetry around the vertical axis and simplifies calculations with little loss in accuracy since the measured trap frequencies are similar.

The additional confinement from the 120 μm waist lattice beams (Equation 3.5) was also included in calculating the density profile. With the confinement from the 4 E_R lattice, the horizontal trap frequency is $\omega_x = \omega_z = 53.6 \times 2\pi$ Hz and the vertical trap frequency is $\omega_y = 78.0 \times 2\pi$ Hz.

The potential from the barrier beam $V_{\text{barrier}}(x, y)$ is given by Equation 3.2 where the intensity is

$$I(x, y) = \frac{2P}{\pi \; w_x \; w_y} \exp\left[-2\left(\frac{x^2}{w_x^2} + \frac{y^2}{w_y^2}\right)\right]. \tag{5.2}$$

Here P is the power of the beam and w_i is the beam waist in the i direction. In this calculation, we assume the Rayleigh range of the barrier beam is long with respect to the size of the gas. As given by Equation 4.2, the Rayleigh range of the barrier beam is 135 μm, and the waist changes by less than 4% over 10 μm.

To find the density distribution, we found the energy for each lattice site at each allowed occupation. In this calculation, we allowed up to three particles on a lattice site and used a cubic lattice that was 200 sites on each side. Since the distribution is symmetric, we performed these calculations on just one corner of the gas in order to improve calculation time. We centered the harmonic trap at one corner of the lattice, so that the lattice in the calculation was an octant of the real distribution. We kept track of lattice sites at positions greater than zero in every direction, but the density on those sites also represents the density on the equivalent sites that are reflected across each axis.

Since this is a thermal gas, the density can be found with the Maxwell-Boltzmann distribution. The probability of finding lattice site (x, y, z) with occupation n is

$$p(x, y, z, n) = \exp\left(-n\left(E(x, y, z, n) - \mu\right)/(k_B \; T)\right), \tag{5.3}$$

where μ is the chemical potential. The chemical potential determines the atom number N, which can be

found by

$$N = \sum n\ p(x, y, z, n). \tag{5.4}$$

In this calculation, the chemical potential was varied to find the value of μ for which N matched the N of our system. The ideal N for this calculation is $N = 8500$. This is one eighth of the measured atom number for the whole distribution, since this calculation was performed on one octant of the distribution. The temperature used in this calculation was $T = 210$ nK (Section 5.3).

5.2.2 Simulated Initial Density Distribution

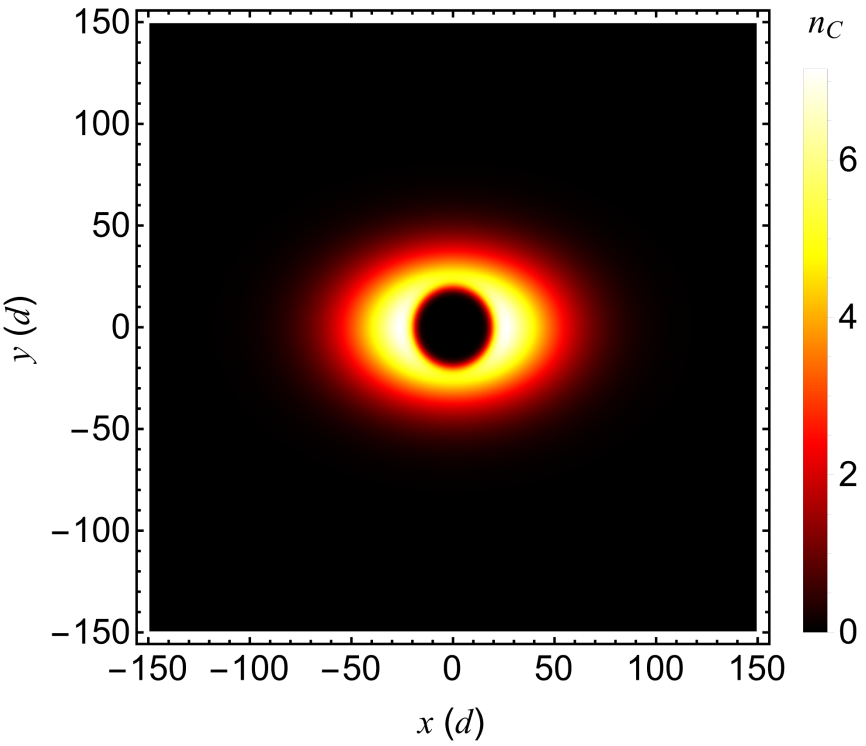

Figure 5.2: A simulated column-integrated image of a realistic density distribution. The color scale is the column integrated density n_C in units of particles per column of lattice sites.

We find the density distribution[1] shown in Figure 5.2. The maximum density is 0.18 atoms per lattice site and the density-weighted density is 0.06 atoms per lattice site. The density-weighted density is the average density and is calculated by $n_{dwd} = \sum \frac{n(x,y,z)}{N} n(x, y, z)$. This density is low compared with typical densities for atomic gas experiments. Most experiments use colder, denser gases. Simulations of Stark localization often use a density near one particle per lattice site. There are more empty lattice sites in this gas than is typical. Most of the atoms are in singly occupied sites, with 9% of particles on a doubly occupied site and 1% on a triply occupied site. The barrier completely excludes atoms from the central region, the density is < 0.01 atoms per lattice site within 16 lattice sites from the center. The peak density occurs just outside the barrier potential, 26 lattice sites from the center.

[1]This is calculated in "Atomic Limit Density.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

This simulation is an approximation of the real density profile and makes some assumptions. This simulation used a uniform barrier beam, but the real barrier beam is elliptical in nature. The hole in the simulated distribution is round, unlike the real distribution. The principal axes of the real distribution are also rotated from the x and y axes by approximately 10°, unlike in the simulation. Although this simulation does not perfectly match the real density distribution, this approximation captures the general trends of the density profile, in particular scaling with lattice depth in the same way as the real density profile.

The size of the gas can be described by fitting the tails of the gas to a Gaussian distribution. In understanding the density profile of a thermal gas, it is useful to note that the density profile of a thermal gas in a harmonic trap has the shape of a Gaussian distribution, $e^{\frac{-r^2}{2\sigma^2}}$. A thermal gas has a density distribution that goes like $e^{-E(r)/(k_B T)}$ where $E(r)$ is the energy at position r. For a perfectly harmonic trap, $E(r)$ goes like r^2 and the density distribution is perfectly Gaussian.

The initial density profile in this work is obviously not Gaussian, as the barrier potential in the center of the gas creates a hole in the density distribution. However, the tails of the gas, where the barrier potential is not present, will be approximately Gaussian. We find that the tails of the simulated density distribution fit to a Gaussian distribution with $\sigma_x = 32.0$ lattice sites in the horizontal direction (Figure 5.3) and $\sigma_y = 22.6$ lattice sites in the vertical direction (Figure 5.4).

The sizes of the simulated density profile are similar to the sizes measured from the real initial density distribution. (The size along the near-horizontal axis is $\sigma_{x'} = 33.2 \pm 1.4$ lattice sites and the size along the near-vertical direction is $\sigma_{y'} = 21.4 \pm 1.0$ lattice sites.) This agreement is unsurprising, as the size of the initial density profile was used to determine the temperature of the distribution (Section 5.3.1). The slight discrepancy between the sizes of the simulated and real density profiles may result from the different shape of the barrier beam.

This simulation of the density profile assumes the atoms are in a harmonic trap. The statistic we use to measure equilibrium, Mardia's B (Section 6.2), also requires that the trapping potential is harmonic. The

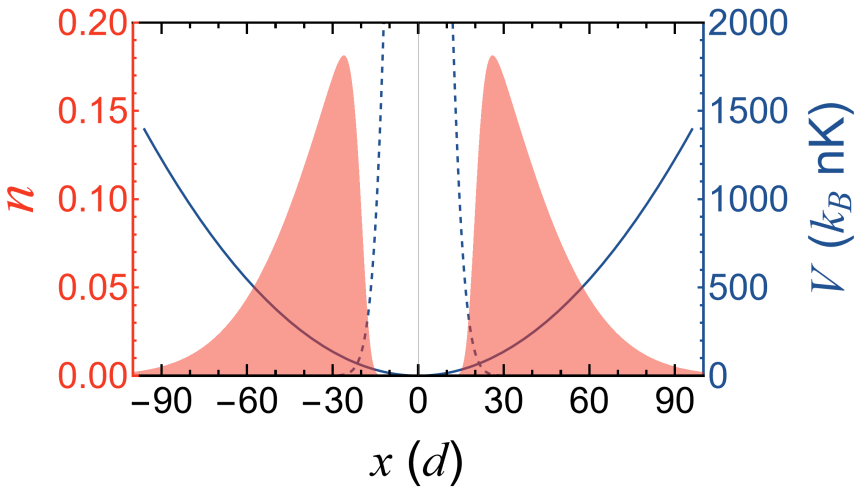

Figure 5.3: A slice though the center of the simulated density distribution along the horizontal direction.

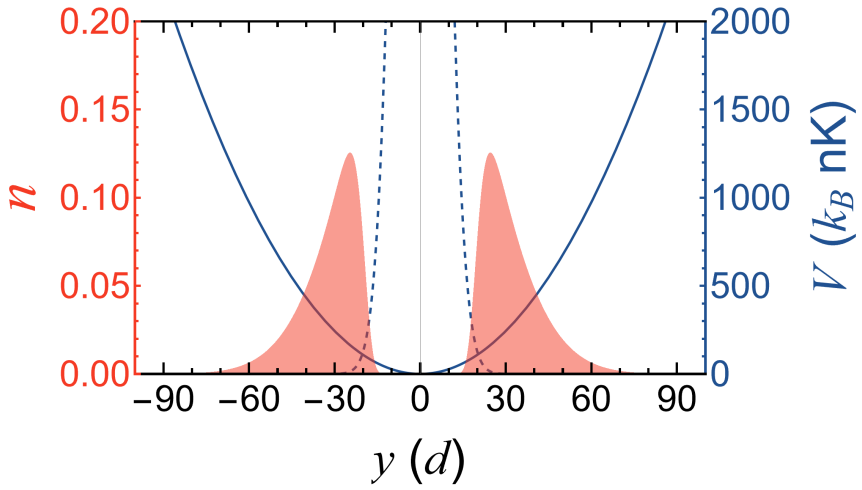

Figure 5.4: A slice through the center of the simulated density distribution along the vertical direction.

real trapping potential used in this work is well-approximated as a harmonic potential near the center of the trap, but the potential deviates from harmonic far from the center of the gas. Near the edges of the trap, the laser intensity becomes small, and the potential becomes flat. We can calculate how much the trapping potential deviates from a harmonic potential using the trap model defined in Section 3.4 with the additional confining potential from the optical lattice beams. For a 20 E_R lattice, in the horizontal direction, the potential is 2% and 4% lower than the harmonic potential at 64 lattice sites ($2\sigma_x$) on either side of the center. The trap is slightly asymmetric due to the alignment between the magnetic and optical traps. In vertical direction, the potential is 4% higher than the harmonic potential at 45.2 lattice sites ($2\sigma_y$) above the center of the trap and 8% lower than the harmonic potential 45.2 lattice sites below the center of the trap. In the vertical direction, gravity is a source of asymmetry in the potential at the atoms. The deviation of the potential from a harmonic potential is small in the area of the trap that contains the atoms, and treating the trap as harmonic is likely a reasonable approximation. This was further confirmed with measurements of gases in equilibrium (Section 6.2.2).

5.2.3 Other Approximations of the Initial Density Distribution

An interacting atomic limit calculation is not the only way to estimate the density profile of the atoms. In some situations, it makes sense to neglect certain effects in order to simplify the calculation. This section discusses a few other ways of calculating the density distribution and how the results differ from the full atomic limit calculation above.

One possible approximation is to neglect the interaction term. Since the density is low in our system, one might expect the effect of interactions to be small. Calculating the density in the same way as above, except setting the interaction U to zero, yields a very similar density profile. The black dashed line in Figure 5.5 shows a slice through this density profile in the x direction. The peak density is 2% higher than the peak

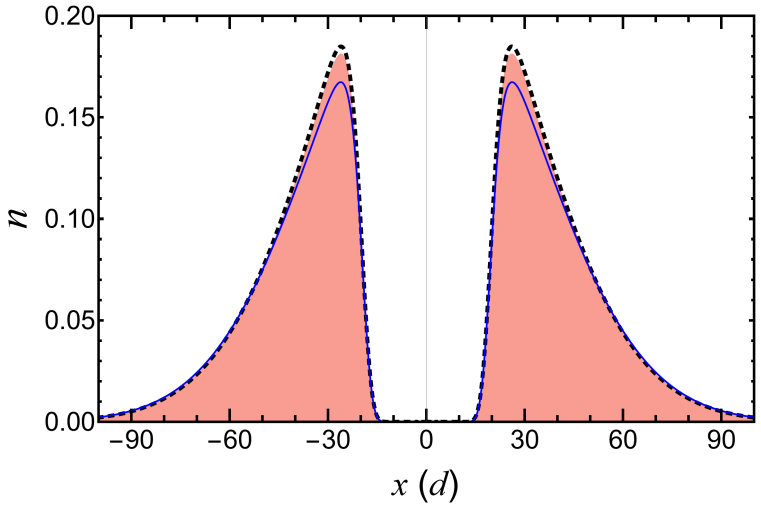

Figure 5.5: A slice through the simulated density distribution along the horizontal direction with multiple types of calculations. The red filled curve is the atomic limit calculation with the interactions. The black dashed line is the atomic limit calculation without interactions. The blue solid line represents a calculation of the density distribution, neglecting the quantization of the lattice.

density with interactions (red filled curve) and the tails of the two distributions are indistinguishable. The interaction slightly decreases the density in the high-density regions of the distribution, but otherwise has little effect. This is as expected; the repulsive interaction causes the atoms to spread out, and the effect is strongest where the density is highest. We would also expect the effect of interaction to be weak, since the density is low ($n_{dwd} = 0.06$ atoms per site).

An even simpler approximation is to completely ignore the quantization of the lattice and take

$$n(x, y, z) = \frac{e^{-V(x,y,z)/(k_B T)}}{\int e^{-V(x,y,z)/(k_B T)}}, \tag{5.5}$$

where $V(x, y, z)$ is the total potential and x, y, and z are now continuous variables. The x direction slice through this density distribution is shown as the blue solid line in Figure 5.5. The peak of this density distribution is 9% lower than the peak of the non-interacting quantized density distribution, and the tails of both distributions are similar. Although the non-quantized calculation would misrepresent the entropy of the system, we would expect it to give a reasonable approximation of the density profile. This non-quantized estimation of the density profile is used in Section 5.4.2 to find the fraction of particles in localized states.

5.3 Temperature

Understanding the temperature of the gas is necessary for understanding its density distribution. In this section, I will first discuss two different ways of measuring the temperature of a thermal gas in a lattice. Then I will explain how the temperature is expected to change over the 10-second hold time of this work and our measurements of the change in temperature.

5.3.1 Temperature from the Size of the Gas in Trap

The temperature can be estimated from *in situ* images of a gas in the lattice. If we take the approximation that the density is determined by Equation 5.5 and that the confining potential is harmonic, then the initial temperature is related to the size of the gas by

$$\frac{1}{2}M\omega^2\sigma^2 = \frac{1}{2}k_B T, \tag{5.6}$$

where ω is the angular trap frequency and σ is the RMS (root mean square) size of the gas. Since the trap frequencies and sizes can be different in each direction, the trap frequency and size along the same direction should be used. The center of the gas is effected by the barrier potential and is not in a harmonic trap, but the tails of the gas only experience the harmonic confining potential and are not affected by the barrier beam. Since the gas was prepared adiabatically, the tails of the distribution are in thermal equilibrium with the entire gas and the shape of the tails of the distribution can be used to find the temperature of the whole gas.

To find the temperature, we masked the center of the image of the gas and fit the tails of the distribution to a two-dimensional Gaussian with the rotation angle as a free parameter. We find that the rotation angle of the fit is about 11 ± 4 degrees. Since this is close to zero, the principal axes of the gas are nearly aligned with the imaging directions, which matches the results we found for the principal axis from the trap frequency measurement. With the rotated fit, we found the size along the near-vertical axis to be $\sigma_y = 21.4 \pm 1.0$ lattice sites and the size along the near-horizontal axis to be $\sigma_x = 33.2 \pm 1.4$ lattice sites.

For this calculation we used a vertical trap frequency of (74 ± 2) Hz and a horizontal frequency of (47 ± 2) Hz and then added the confinement from a $4\,E_R$ lattice. We found the temperature to be $T = (190 \pm 20)$ nK from the calculation in the vertical direction and $T = (220 \pm 20)$ nK from the horizontal direction. Taking the geometric mean trap frequency and size (assuming the horizontal values represent both horizontal directions) gives $T = (210 \pm 20)$ nK.[2]

5.3.2 Temperature from Expansion Velocity

To measure the temperature of the gas independently from the in-trap images, we do an expansion velocity measurement. When an atomic gas is released from the trap, it expands based on its velocity according to

$$\sigma(t)^2 = \sigma_0^2 + (vt)^2, \tag{5.7}$$

where $v = \sqrt{k_B T/M}$ is the average velocity, t is time after releasing the gas from the trap, $\sigma(t)$ is the RMS size of the gas at time t, $\sigma_0 = v/\omega$ is the size at $t = 0$, and ω is the trap frequency. The temperature of an atomic cloud can be found by taking images at several different expansion times and fitting the resulting sizes. For longer time of flight, when $vt \gg \sigma_0$, we can approximate Equation 5.7 as $\sigma(t) = b + vt$ and find v, and therefore T, with a linear fit to the data. For the experiments discussed in this work, the gas was evaporated to between (108 ± 3) nK and (120 ± 30) nK, which was measured by snapping off the barrier beam and taking an expansion velocity measurement before ramping on the lattice potential.

[2] This calculation was done in "Temperature from Size in Lattice.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

Entropy Matching

To estimate the temperature after the lattice was turned on, we did an entropy-matching calculation. Since the lattice is ramped on slowly relative to the tunneling time, we assume that the process is adiabatic and entropy is conserved. We calculate the entropy of the gas at the measured temperature before adding the lattice and then find the temperature that gives the same entropy when the lattice potential is present.

This calculation is an approximation in that it treats the gas as spherically symmetric and ignores the barrier potential. Using the distribution with the barrier beam present may be possible with more advanced methods. Without the barrier potential, the density of the distribution of the gas is not accurately represented. This calculation is only accurate in so far as the change in entropy from adding the lattice is the same for the density distributions with and without the barrier. Adding the lattice potential primarily changes the momentum space of the gas. The change in the entropy due to the new dispersion relation should not be dependent on the density distribution of the gas, and this aspect of the change in the entropy should be captured by this calculation. However, the lattice also adds additional confinement and therefore changes the density distribution as well. Since the density distribution of this calculation is only a rough approximation of the real density distribution, the results may not capture the change in entropy due to the changing density distribution, and the results are only a rough approximation of the temperature in the lattice.

To calculate the entropy, we assume Maxwell Boltzmann statistics. The entropy is given by

$$S = \frac{1}{T}\left[E(T,\mu) - \mu N + N k_B T\right], \tag{5.8}$$

where $E(T,\mu)$ is the total energy, μ is the chemical potential, N is the total atom number, and T is the temperature. For a spherically symmetric harmonic trap without the lattice, the energy per particle is

$$\epsilon(p,r) = \frac{p^2}{2m} + \frac{1}{2}m\omega^2 r^2. \tag{5.9}$$

The phase space density is

$$\rho(p,r) = \frac{1}{h^3} e^{(\mu-\epsilon(p,r))/(k_B T)}, \tag{5.10}$$

and integrating over the phase space density gives the atom number,

$$N = \int_{r,p} \rho(p,r) 4\pi r^2 4\pi p^2 \,\mathrm{d}r\,\mathrm{d}p. \tag{5.11}$$

By integrating over the momentum p and position r we can solve for μ in terms of the total number of atoms, finding

$$\mu = k_B T \ln\left[N\left(\frac{\hbar\omega}{k_B T}\right)^3\right]. \tag{5.12}$$

We can then find an explicit expression for the entropy per particle in a gas of free particles in a harmonic trap,

$$\frac{S}{N\ k_B} = 4 - \ln\left[N\left(\frac{\hbar\omega}{k_B T}\right)^3\right]. \tag{5.13}$$

The N in the right side of the equation is a result of finding μ in terms of N. Now that we know the entropy in the gas without the lattice, we need to find the entropy in the lattice in terms of temperature in order to find the starting temperature of the gas after the lattice has been added.

In the lattice, the energy per particle depends on the quasi-momentum q such that

$$\epsilon(r, q_x, q_y, q_z) = \frac{1}{2} m\ \omega_L^2 r^2 + 2t\left[1 - \cos\left(\pi q_x/q_b\right)\right] + 2t\left[1 - \cos\left(\pi q_y/q_b\right)\right] + 2t\left[1 - \cos\left(\pi q_z/q_b\right)\right], \tag{5.14}$$

where q_x, q_y, and q_z are the quasi-momentum in each of the three lattice directions, $q_b = \hbar\pi/d$ where d is the lattice spacing, and t is the tunneling energy. The trap frequency ω_L now includes the additional confinement from the lattice potential.

We can now use the same method as above to find an expression for the entropy per particle in the lattice. Although the expression is now much more complicated, it is still analytically solvable, and we find

$$\frac{S}{N\ k_B} = \frac{5}{2} - \frac{6t\ I_1\left(\frac{2t}{k_B T}\right)}{k_B T\ I_0\left(\frac{2t}{k_B T}\right)} + \ln\left[\frac{1}{N}\left(\frac{2k_B T}{m\pi}\right)^{3/2}\left(\frac{q_b}{\hbar\omega_L}\right)^3 I_0\left(\frac{2t}{k_B T}\right)^3\right] \tag{5.15}$$

where $I_n(z)$ is the modified Bessel function of the first kind.[3]

Since the lattice ramp-on is adiabatic, the entropy of the gas in the lattice (Equation 5.15) is equal to the entropy of the gas before adding the lattice (Equation 5.13). The temperature of the gas before adding the lattice was measured with a time of flight measurement (115 ± 10 nK), so the entropy in the trap can be found. We can set that equal to the expression for entropy in the lattice and (numerically) solve for T.

We find that the temperature of the initial gas in the $4\,E_R$ lattice is $T = 210 \pm 40$ nK.[4] This estimate of temperature is only a crude approximation of the actual temperature, since the barrier potential was neglected in this calculation. This approximation is only accurate to the extent that the change in entropy due to the lattice potential is the same both with and without the barrier present. However, we do find that this estimate agrees with the temperature found from the in-trap size of the gas in the lattice ($T = 210 \pm 20$ nK). The agreement between the temperature measured from the size of the gas and the temperature calculated by entropy matching suggests that the ramping on of the lattice potential was, in fact, adiabatic and that the initial distribution of gas is in equilibrium with the trapping potential, the barrier potential, and the $s = 4\,E_R$ lattice potential.

5.3.3 Temperature Dynamics

We have demonstrated an understanding of the initial temperature of the gas. However, the temperature does not remain constant over the course of a measurement, but rather increases over time. This temperature increase is explained by two main sources: thermalization of the potential energy that was present before removing the barrier potential and heating from scattering of the lattice light.

Heating from Thermalization

We can calculate the expected change in the temperature due to thermalization with energy conservation. The barrier is removed diabatically and the entropy per particle is not conserved. Rather, since the gas is isolated, the energy of the system is conserved. We can calculate the energy of the initial density distribution immediately after the barrier is removed. To find the expected temperature after thermalization, we use this initial energy in the system and find the equilibrated state with the same energy.[5]

[3] This is Mathematica function `BesselI[n,z]`.
[4] This calculation is done in "Thermodynamics of Trap.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.
[5] This calculation is done in "Atomic Limit Density.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

We calculate the initial density distribution and energy after removing the barrier using the same atomic limit calculation as described in Section 5.2. Since the barrier is turned off quickly relative to the atomic timescales, the initial density distribution is the same as the distribution in equilibrium with the barrier potential.

We took the initial density distribution to be the equilibrated state with the barrier potential still on and in a $4\,E_R$ lattice present. Then we calculated the energy per particle of that distribution after removing the barrier potential and increasing the lattice to its final value. For this calculation, we looked at the final expected temperature for the weakest ($4\,E_R$) and strongest ($20\,E_R$) lattice depths used in the experiment. We then found the new equilibrium distribution that has the same energy as the initial distribution. In the $4\,E_R$ lattice, the heating we would expect only from thermalization resulted in $T = 240$ nK and for $20\,E_R$ we found $T = 480$ nK to be the thermalized temperature.

Heating from Lattice Scattering

However, this is not the only expected source of heating. Occasionally, the 812 nm lattice light will off-resonantly scatter with an atom, transferring energy to the atom. We calculated the energy added by this scattering process. An expression for calculating the lattice heating rate is given in Reference [61], Appendix B and C. Here we modify this expression to avoid taking the approximation of being far from resonance, where the D1 and D2 lines can be treated identically. For a retro-reflected lattice with light near resonance, the heating rate per lattice depth is

$$\frac{\dot{E}}{V_{lat}} = E_R \frac{2\Gamma_{D2} + \Gamma_{D1}}{3} \frac{H_{sc}}{V_{sc}}, \tag{5.16}$$

where V_{lat} is the lattice potential,

$$H_{sc} = \frac{\pi c^2 \omega^3}{2\hbar} \left(2A_{D2} + \frac{\Gamma_{D1}}{\Gamma_{D2}} A_{D1} \right), \tag{5.17}$$

$$V_{sc} = \frac{-1}{2} \pi\ c^2 \left(2A_{D2} + A_{D1} \right), \tag{5.18}$$

and

$$A_{Di} = \frac{\Gamma_{Di}}{\omega_{Di}^6} \left(\frac{1}{\omega_{Di} - \omega} + \frac{1}{\omega_{Di} + \omega} \right)^2. \tag{5.19}$$

Here Di is either $D1$ or $D2$.[6] This expression describes the heating rate from a single retro-reflected lattice beam. Multiplying by three gives the heating from all three beams. We find the total added energy over 10 seconds (the longest hold time used in this work) is 220 $k_B\times$nK for the $s = 4\,E_R$ case and 1110 $k_B\times$nK for the $s = 20\,E_R$ case. We can find the new expected equilibrium by adding this energy from lattice scattering to the initial energy of the system and finding the new equilibrium temperature. We find an expected temperature of $T = 390 \pm 40$ nK for $s = 4\,E_R$ and $T = 1200 \pm 100$ nK for $s = 20\,E_R$. The dominant source of uncertainty in these estimates is the uncertainty in the lattice waists and therefore in the trap frequencies used to define the potentials. The expected heating from the 1064 nm light of the optical dipole trap is 10% of the expected heating from the lattice light for a $s = 20\,E_R$ lattice and was ignored in this calculation.

We can also calculate the observed temperature after 10 seconds from the size of the gas in the long-time

[6] This is essentially a restatement of Equation C. 70 in Reference [61] with both the D1 and D2 lines included and corrected for the missing $\hbar$ in the denominator.

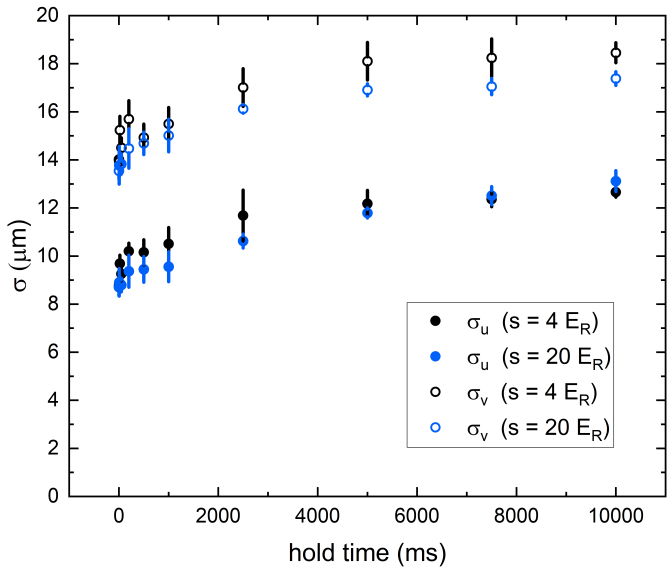

Figure 5.6: The size over time of a gas in an $s = 4\,E_R$ (black) and $s = 20\,E_R$ (blue) lattice. The size of the gas is the σ from the rotated Gaussian fit to the tails of the density distribution. The near-vertical size is σ_u (closed circles) and the near-horizontal size is σ_v (open circles).

images (Figure 5.6). We find that, for $s = 4\,E_R$, $T = 390 \pm 40$ nK based on the images. This is consistent with the expected temperature given the known heating effects. For $s = 20\,E_R$, we find $T = 700 \pm 200$ nK. This is colder than the expected temperature. This colder temperature could potentially be explained by energy lost from atom loss. Approximately 40% of the atoms are lost over the 10 seconds of the experiment, and it is possible that high-energy atoms were selectively lost, resulting in a cooling effect on the gas.

5.3.4 Conclusions on Temperature

We measured the temperature of the initial density distribution to be $T = 210 \pm 20$ nK based on the size of the images of the initial density distribution. This agreed with an entropy matching calculation estimating the increase in temperature that results from turning on the lattice potential. We conclude that the initial density distribution was in equilibrium in the presence of the barrier beam and initial $s = 4\,E_R$ lattice potential. We saw that, after the barrier was removed, the temperature of the gas increased over time. This is generally consistent with the thermalization of the potential energy left behind when the barrier was removed and heating from lattice scattering. The measured temperature increase at the highest lattice depths was less than expected. The discrepancy may be explained by atom loss. Our observations of the temperature of the atomic distribution are generally consistent with expectations.

5.4 Presence of Localized States

The initial density distribution (which is influenced by the temperature of the gas) is of interest because it is used to determine the initial population of localized states. In this section, we discuss the physics of Stark

localization and then calculate the presence of localized states in two different ways.

5.4.1 Stark Localization

A particle in a lattice potential can become localized if the change in energy needed to move to a different lattice site is larger than the bandwidth in the lattice. When a particle is localized its wave function only extends over a fraction of the system, that is, its motion is constrained. The primary source of localization considered in this work is Stark localization.

Stark localization occurs when there is a potential gradient present. For a linear gradient in a non-interacting system, this effect is known as Wannier-Stark localization [46]. Particles in the ground band of the lattice have a limited range of allowable energies. This range, the bandwidth, is $4t$ (Section 2.1.2). When the energy gradient F is on the order of the bandwidth, the wave function of the particle will only extend over the range of positions (and range of potential energies) within the range of energies allowed in the band. Specifically, the localization length ℓ_L of the wave function is

$$\ell_L \approx \frac{4t}{F}. \tag{5.20}$$

When F approaches $4t/d$ (where d is the lattice spacing), then ℓ_L approaches d and the wave function is concentrated mainly on one lattice site. [86]

The gradient need not be linear to induce localization. A harmonic trapping potential can cause a similar effect, particularly away from the center of the trap, where the gradient is stronger. Localized atoms are trapped near the edge of the gas between a classical turning point and Bragg reflection from the lattice. This is illustrated in Figure 5.7, where the particle can move over the range of energies ΔE within the bandwidth $4t$. One can imagine the particle falling towards the center of the trap until it acquires as much kinetic energy as can be accommodated in the band. At this point, it Bragg reflects and moves towards the outside of the trap until all of its kinetic energy has been converted to potential energy, and it once again turns around to fall back towards the center of the trap.

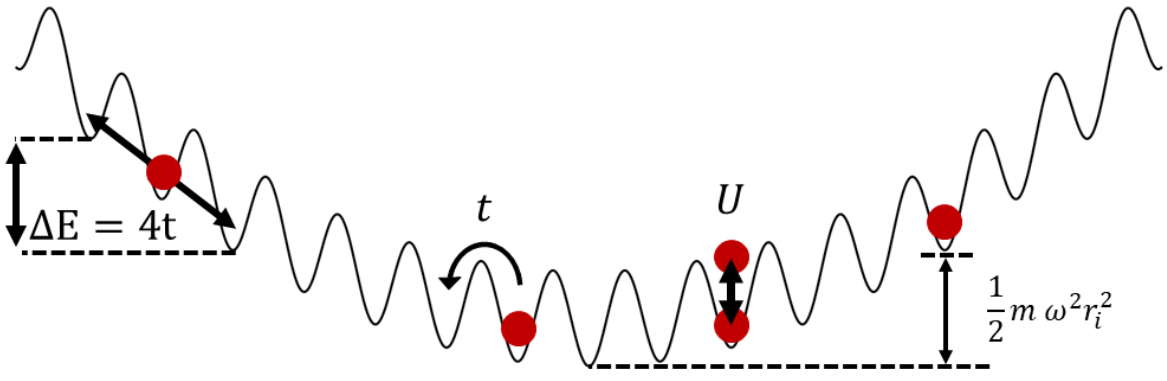

.

Figure 5.7: An illustration of Stark localization in the Bose-Hubbard model. Particles towards the outside of the harmonic trap oscillate over the range of lattice sites accessible within the bandwidth $4t$.

This kind of localization can occur sufficiently far from the center for any harmonic trap, but experiments

generally are done in a regime where all particles are sufficiently close to the center to avoid this localization. This work uses a thermal gas and a barrier beam to populate the states far from the center of the trap and create localized states.

In this work, we measured thermalization dynamics in the presence of localized states. In order to understand how the localization affected the dynamics, we must first understand how much of the gas was initially localized. We calculated the localization fraction in two different ways to capture the two opposite extremes of the definition of localization. In the first calculation, we used energy constraints to determine how many particles are constrained to a single lattice site. This provides a lower bound on the fraction localized, where the lack of any tunneling of the particle would clearly prohibit thermalization. However, more relaxed constraints on the particle position can also influence dynamics. In the second calculation, we exactly solved a one-dimensional Hamiltonian that approximates our system and found the fraction of particles that are in states that are excluded from the central (lowest energy) lattice site. This provides an estimate of the other extreme of localization, where localized states are any state that cannot sample the entire trap.

5.4.2 Localization to a Single Lattice Site Based on Energy Constraints

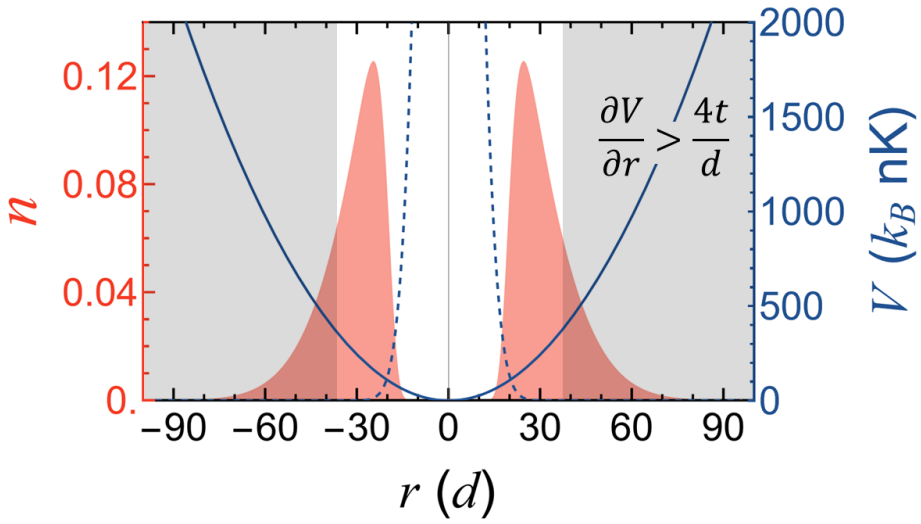

Figure 5.8: An illustration showing how localization to a single site is calculated. The condition for localization is shown along a slice through the density profile. Particles in the gray shaded region are localized.

The lower bound on localized states was calculated by finding where the potential gradient of the trapping potential is greater than the bandwidth, $4t$. When this is the case, tunneling to the adjacent lattice site is prohibited as the change in potential energy cannot be accommodated in the band. Figure 5.8 illustrates this condition. We then find the fraction of particles along each lattice direction where this condition is true.

The first step in finding the fraction of particles in a localized state is to find how far away from the center of the harmonic trap the gradient becomes greater than $4t$. We approximate the potential as a harmonic trap,

$$V_{\text{harmonic}}(x, y, z, s) = \frac{1}{2} m\, d^2 \left[\omega_x(s)^2\, x^2 + \omega_y(s)^2\, y^2 + \omega_z(s)^2\, z^2\right], \tag{5.21}$$

where the trap frequencies depend on the lattice depth and x, y, and z are in units of lattice sites d and describe the position in the East-West, Vertical, and North-South directions, respectively. This is the same confining potential used in the atomic limit calculation of the density profile in Section 5.2.1. The barrier is not present for this calculation because the localization is determined by the potential after the barrier is removed. The trap frequencies depend on the lattice depth, as the lattice light adds additional confinement (Equation 3.5). With the confinement from a $4\,E_R$ lattice, the horizontal trap frequencies are $\omega_x(4) = \omega_z(4) = 53.6 \times 2\pi$ Hz and the vertical trap frequency is $\omega_y(4) = 78.0 \times 2\pi$ Hz. For $s = 20\,E_R$, the horizontal trap frequencies are $\omega_x(20) = \omega_z(20) = 78.5 \times 2\pi$ Hz and the vertical trap frequency is $\omega_y(20) = 96.8 \times 2\pi$ Hz.

Localization is determined along each lattice direction independently, as the particles only tunnel along one lattice direction at a time. The lattice directions are not along the x, y, and z directions. The directions of the lattice beams, u, v, and w, are described by

$$u = \frac{-\sqrt{2}x + y - z}{2},\ v = \frac{y+z}{\sqrt{2}},\ w = \frac{\sqrt{2}x + y - z}{2}. \tag{5.22}$$

To find the localization fraction along each lattice direction, we convert to the u, v, and w coordinate system. We then take the derivative along each lattice direction to find the condition for localization along that direction. For example, a particle is localized along the u lattice direction when

$$\left|\frac{\mathrm{d}}{\mathrm{d}u}V(u,v,w,s)\right| \geq 4t(s), \tag{5.23}$$

where t is the tunneling energy, which depends on s. Since u is in units of lattice sites, the derivative of V with respect to u describes the change in the potential over a distance of one lattice site. Finding the fraction of particles that are localized is simply a matter of finding what fraction of the density distribution is at a position where this condition is true.[7]

The density distribution was calculated assuming Maxwell Boltzmann statistics and neglecting the interaction energy and quantization of the lattice (Section 5.2.3). We include the barrier potential and use the trap frequencies for $s = 4\,E_R$ to find the initial density distribution. The density distribution is

$$\frac{n(x,y,z)}{N_{\text{total}}} = \frac{1}{\mathcal{V}}\, e^{-V_{\text{inital}}(x,y,z)/(k_B T)}, \tag{5.24}$$

where $\mathcal{V}$ is the volume of the unit cell of calculation (d^3 for a single lattice site) and N_{total} is the total number and

$$V_{\text{inital}}(x,y,z) = V_{\text{harmonic}}(x,y,z,4) - V_{\text{barrier}}(x,y) \tag{5.25}$$

and $V_{\text{barrier}}(x,y)$ is given by Equations 3.2 and 5.2.

To find the localization along the lattice directions, we converted $n(x,y,z)$ to $n(u,v,w)$ using the relationships described in Equation 5.22. Then we pre-calculated $n(u,v,w)$ for every even-numbered lattice site on a grid extending to 150 lattice sites to either side of the trap center and stored these values in a list with their coordinates. Calculating $n(u,v,w)$ (in particular, $V_{\text{barrier}}(u,v,w)$) is computationally expensive, so precalculating is important to reduce the computation time. The sum of every element in the list N_{total} was found and $\frac{n(u,v,w)}{N_{\text{total}}}$ is taken to be the fraction of the total distribution at position (u,v,w). Specifically, each element in this list represents the fraction of the distribution in the $2 \times 2 \times 2$ lattice site cube ($\mathcal{V} = 8d^3$) centered at position (u,v,w) as a result of calculating $n(u,v,w)$ at only even lattice sites (which speeds up

[7] This calculation was done in "Single Site Localization.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

the calculation).

To find the localization along a particular lattice direction, we loop over this list and check if the condition described in Equation 5.23 holds at that lattice site i. If it does, that is, particles are localized at lattice site i, then $n(u_i, v_i, w_i)$ is added to the total localization number N_{loc}. The fraction of particles localized along that direction is $N_{\text{loc}}/N_{\text{total}}$.

The calculated fraction localized along the u and v directions can be seen in Figure 5.9. The w direction is the same as the u direction. We see that no particles are localized to a single lattice site at the lowest lattice depth, and the localization fraction starts to increase at $s = 7\,E_R$ until nearly all of the particles are localized in at least one direction by $s = 20\,E_R$. We also calculated the fraction of particles that are localized in all three lattice directions. This calculation is the same as for the single lattice direction case, except the condition of localization is for the gradient to be greater than $4t$ in all three lattice directions, that is,

$$\left|\frac{\mathrm{d}}{\mathrm{d}u}V(u,v,w,s)\right| \text{ and } \left|\frac{\mathrm{d}}{\mathrm{d}v}V(u,v,w,s)\right| \text{ and } \left|\frac{\mathrm{d}}{\mathrm{d}w}V(u,v,w,s)\right| \geq 4t(s). \tag{5.26}$$

We see that by $s = 20\,E_R$, approximately 80% of the atoms are localized in all three lattice directions.

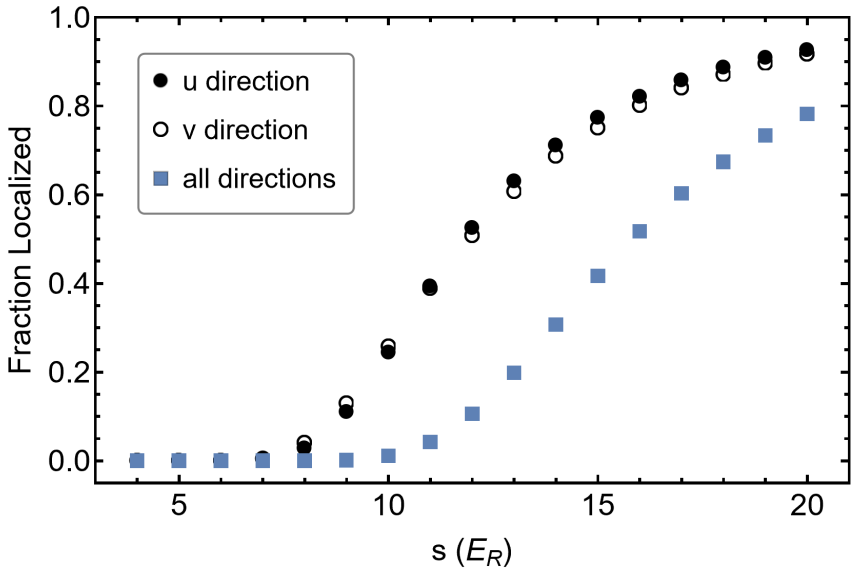

Figure 5.9: A lower bound of localization. Here localization is defined constrained to a single site. Blue squares represent the fraction localized in all three directions. Black circles show the localization along one lattice direction. Open and closed circles represent different lattice directions; slight differences arise from the geometry of the lattice relative to the nonuniform principal axes of the trap. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

The initial state is very strongly localized at high lattice depths, with nearly all the atoms localized to a single lattice site along at least one lattice direction. Lower lattice depths experience less localization. This is as expected: the gradient is weaker at lower lattice depths and, more importantly, the bandwidth is larger. At lower lattice depths, there are no particles in a region where the gradient is larger than the bandwidth, and

the particles can always extend over more than one lattice site. However, that does not mean the particles do not experience any localization at lower lattice depths. Even if the gradient is not strong enough to fully constraint a particle to one lattice site, it may limit the extent of the wave function to a subset of the system that includes several lattice sites.

5.4.3 Localization as Exclusion from the Central Lattice Site

We calculated the fraction of particles localized when we define localization as constrained to a single lattice site. But less extreme forms of localization also exist. When the gradient is half as strong as is necessary to localize a particle to a single lattice site, that is, when $\frac{dV}{dr} = 2t/d$, the particle will be localized to a state that oscillates between two lattice sites. In the extreme limit, localization will occur whenever the range of lattice sites the particles can occupy does not include the lowest-energy, central lattice site. When the localized state consists of a particle oscillating across several lattice sites, describing the exact position of an individual particle is no longer a useful framework. Instead, we consider the eigenstates of the system. Each eigenstates describes the probability amplitude of finding a particle with that particular eigenenergy at each position in the lattice.

These localized eigenstates are often described as single-particle localized states (Section 1.1.2) and are discussed in details for a harmonic trap in Reference [48] and [49]. One characteristic property of the localized states is that they do not extend into the center of the harmonic trap. We use this definition and framework to calculate the other extreme of localization. The fraction of particles that are in a state that does not extend to the central lattice site is an approximation for an upper bound on localization. Since the definition of localization is less extreme than the single-site case, the fraction of particles localized by this definition will be higher.

Since the lattice Hamiltonian is separable and localization occurs along each lattice direction independently, we use the one-dimensional lattice Hamiltonian to find the condition of localization. Only after we have determined the condition of localization in one dimension do we use the three-dimensional density distribution to find the fraction of the particles that are localized according to this condition.

In the above calculation of single-site localization, we used the gradients along the true lattice directions, which are rotated from the imaging directions. This allowed us to accurately represent the density distribution relative to the lattice directions. However, in this calculation we did not rotate the lattice but rather used a spherically symmetric approximation of the trapping potential to find the density distribution. In the spherically symmetric case, the rotation of the lattice is irrelevant. Of course, the barrier beam along the imaging direction breaks spherical symmetry. Instead of trying to rotate the lattice to the correct angle relative to the barrier beam, we calculated the localization of lattice directions along and perpendicular to the barrier beam to capture the two extremes and found the results to be similar.

To find the condition of localization, we use a single-particle tight-binding Hubbard Hamiltonian with a harmonic potential

$$\hat{H} = -t \sum_{\langle i,j \rangle} (\hat{b}_i^\dagger \hat{b}_j + h.c) + \sum_i \Omega r_i^2 \hat{n}_i, \tag{5.27}$$

where $\Omega = \frac{1}{2} m\omega^2 d^2$. For this calculation we use the measured geometric mean trap frequency for the bare trap, $\omega = 54 \times 2\pi$ Hz and modify ω for each lattice depth to include the confinement from the lattice as given by Equation 3.5.

We create a matrix representing the Hamiltonian in a one-dimensional lattice of 300 sites centered at the

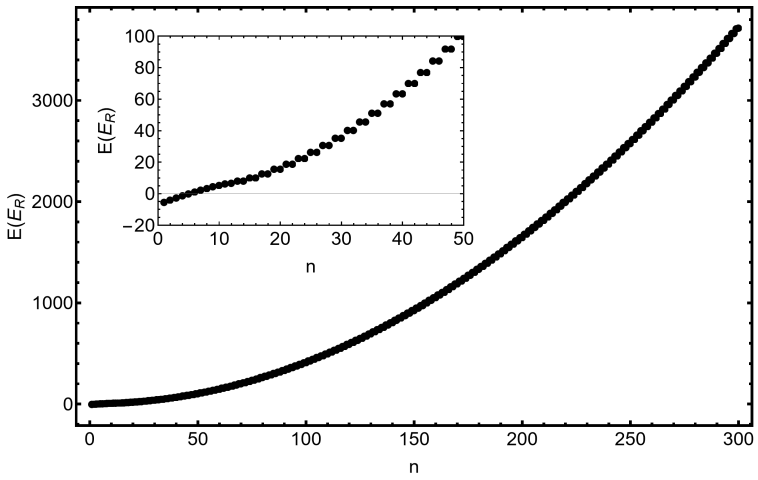

Figure 5.10: Energy spectrum for $s = 10\,E_R$. The inset shows the lowest energy eigenstates. The double degeneracy above $n \approx 14$ is characteristic of localization.

center of the harmonic potential.[8] The basis states are the set of single-particle Fock states on each lattice site $|i\rangle$ such that:

$$|1\rangle = \begin{pmatrix} 1 \\ 0 \\ ... \\ 0 \\ 0 \end{pmatrix} \qquad |2\rangle = \begin{pmatrix} 0 \\ 1 \\ ... \\ 0 \\ 0 \end{pmatrix} \qquad |i\rangle = \begin{pmatrix} 0 \\ ... \\ 1 \\ ... \\ 0 \end{pmatrix} \qquad |300\rangle = \begin{pmatrix} 0 \\ 0 \\ ... \\ 0 \\ 1 \end{pmatrix} \tag{5.28}$$

where 1 in the ith row represents that the particle is fully on the ith lattice site. We diagonalize the Hamiltonian to find the eigenstates and eigenenergies. The spectrum of eigenenergies can be seen in Figure 5.10. The lowest energy states are shown in the inset. Above about $n = 14$, the eigenenergies have two-fold degeneracy. This is a signature of localized states as described in Reference [48]. This method for finding the eigenenergies of the system gives an eigenenergy spectrum that is within 0.05% of the analytical expressions for the eigenenergies of the single-particle localized states give in Reference [49], when the calculation in done with a lattice site at the center of the harmonic potential, as is the case in Reference [49].

The eigenstates $|E_n\rangle$ are also found from the diagonalization. A few example low-energy eigenstates are shown in Figure 5.11 for $s = 10\,E_R$. The first two plots ($n = 1$ and $n = 11$) are not localized, and the last plot ($n = 16$) is localized. We check for localization by measuring whether the probability on the lattice site nearest to the center of the harmonic trap, p_c is less than 0.001. Specifically, the localization l_{En} of eigenstate E_i is given by

$$l_{En} = \begin{cases} 1, & p_c < 0.001 \\ 0, & p_c \geq 0.001 \end{cases} \tag{5.29}$$

The exact choice of threshold is somewhat arbitrary, but must be sufficiently low as to not exclude states that have a node in the center, but are not localized, such as the $n = 11$ state. A threshold of 0.01 is too high, as it labels states like the $n = 11$ state as localized. A threshold of 0.00001 gives localization fractions that are approximately 10% different from the localization fractions found with a 0.001 threshold. Below 0.001, the choice of threshold does not significantly modify the localization fraction.

[8] This calculation was done in "Exact Diagonalization one particle" in `Z:\Pubs\Laura Wadleigh\thesis code`.

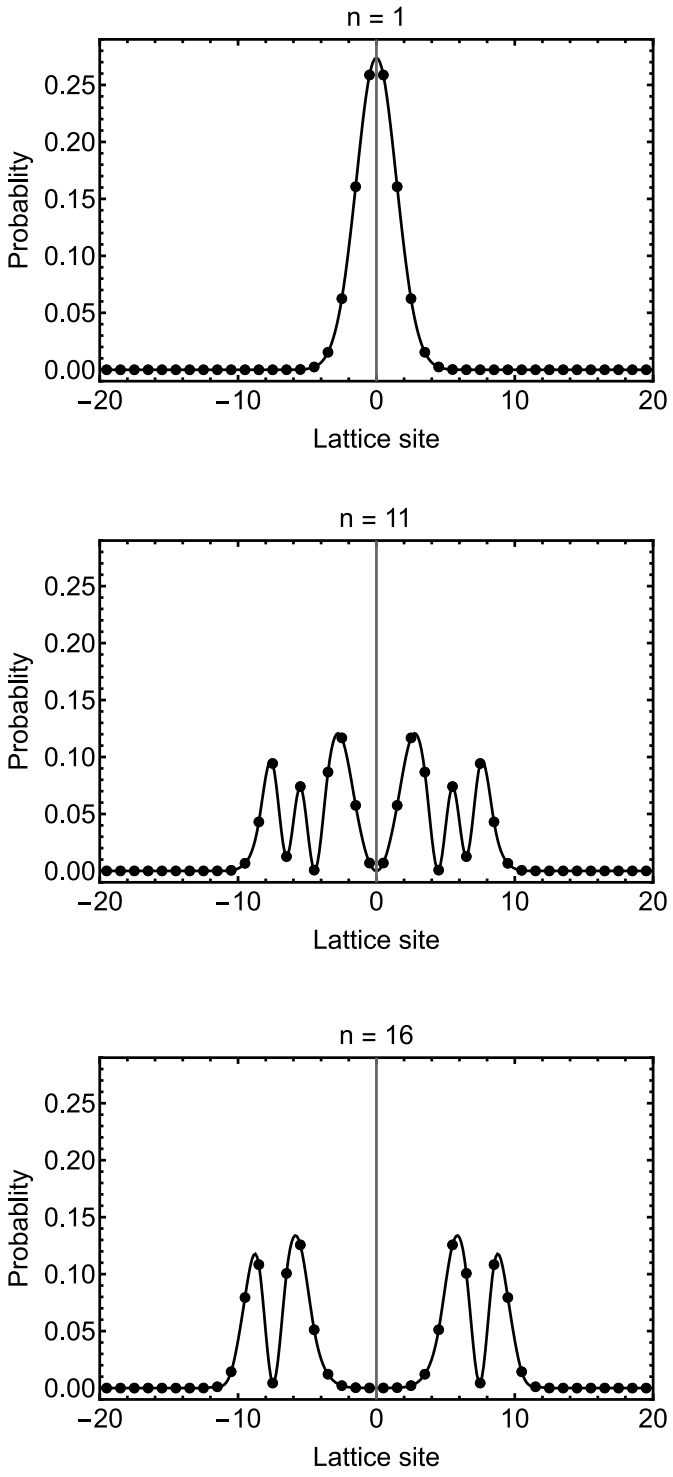

Figure 5.11: Probability distribution for a few example eigenstates for $s = 10\,E_R$. The first distribution is the ground state, $n = 1$. The second distribution, $n = 11$, is an example eigenstate that has a node at the center, but is not localized. The last distribution, $n = 16$, is a localized eigenstate. The points represent the calculated probability at each lattice site and the line is included to guide the eye.

We use an even number of lattice sites (and therefore do not have a lattice site exactly at the center of the harmonic trap) so that we can take the weight on the lattice site $d/2$ from the center of the harmonic trap. This avoids the condition where a delocalized state that has a node at the center of the trap appears to be localized.

This procedure for determining if a state is localized shows general consistency with the definition of localization given in Reference [48]. In Reference [48], eigenstates where $n \geq n_c - 1$ are considered to be localized eigenstates, where $n_c = 2||\sqrt{4t/(2\Omega)}||$ and $||a||$ represent the nearest integer to a. The calculation done here finds similar results; the energy level n of the lowest-energy localized state varies from $n = n_c - 1$ at $s = 4\,E_R$ to $n = n_c + 3$ at $s = 20\,E_R$. Some of the discrepancy may be explained by the difference in how the lattice is aligned to the center of the harmonic potential. This calculation may slightly underestimate localization compared to the definition in Reference [48], but follows the same general trend.

Once we have characterized the eigenstates as localized or delocalized, we need to determine how the observed density distribution populates these eigenstates. To find the localization of our real-space density profile, we use a basis state conversion to project the single-particle Fock state of an atom on each lattice site into the eigenstate basis. The set of single particle Fock states $|i\rangle$ are the basis states of this calculation.

To change the basis, we form a matrix of the eigenstates

$$\mathcal{E} = \Big(|E_1\rangle \quad |E_2\rangle \quad ... \quad |E_n\rangle \quad ... \quad |E_{300}\rangle\Big) \tag{5.30}$$

We use the inverse of this matrix to find the representation of each Fock state in the eigenstate basis

$$\sum_n c_n \,|E_n\rangle = \mathcal{E}^{-1}\,|i\rangle\,, \tag{5.31}$$

which gives the probability amplitude c_n of each eigenstate that makes up the Fock state $|i\rangle$. We can then find the fraction of each Fock state that is in a localized eigenstate as

$$l_i = \sum_n c_n^2 l_{En}. \tag{5.32}$$

Each Fock state corresponds to a position in the lattice, so l_i can be understood as the localization fraction of a particle at lattice site i. Since the potential is separable, this localization fraction along a direction depends only on the lattice site position along that direction. To find the localization along a single direction, we weight the localization l_i at each lattice site by the fraction of the density distribution at that lattice site ρ_i. The one-dimensional normalized density distribution is found by integrating the density in the other two directions and normalizing so that $\sum \rho_i = 1$. We can then find the average fraction localized along a particular direction as

$$l_{1\mathrm{D}} = \sum_i \rho_i l_i. \tag{5.33}$$

We can find the fraction of particles that are localized in every direction by multiplying together the fraction localized in each direction for each lattice site and weighting that by the normalized three-dimensional (3D) density distribution,

$$l_{3\mathrm{D}} = \sum_{(x,y,z)} \rho_{(x,y,z)} l_x l_y l_z. \tag{5.34}$$

The 3D localization $l_{3\mathrm{D}}$ is one only if all the directions are localized and zero if any direction is delocalized.

For this calculation, we took the approximation that the harmonic trap is spherically symmetric, but

the barrier potential breaks the spherical symmetry. In the real system, the barrier beam is neither along nor perpendicular to any of the lattice directions. Calculating the localization with the beam at a realistic alignment to the lattice would be challenging. Instead, we calculated the localization fraction for the two most extreme lattice alignments, perpendicular to the barrier beam, and along the barrier beam and find the fraction to be similar in both directions. The actual alignment of the lattice beams is between these two extremes for every lattice direction.

In this upper bound of localization, where localized states are defined to be excluded from the central lattice site, we see approximately 70% of the particles are localized in at least one direction even at the lowest lattice depths. The localization fraction increases with lattice depth to a plateau near 95% for the most localized direction (Figure 5.12). The localization fraction does not reach one because the barrier only excludes atoms that are central in both the x (East-West) and y (vertical) directions. Although a cross-section through the center of the distribution, that is, at $y = 0$, shows that there are no atoms at $x = 0$, a cross-section through the edge of the gas, at $y = 25$ lattice sites, for example, would show atoms are present at $x = 0$. Atoms at $x = 0$ cannot be localized along the x direction.

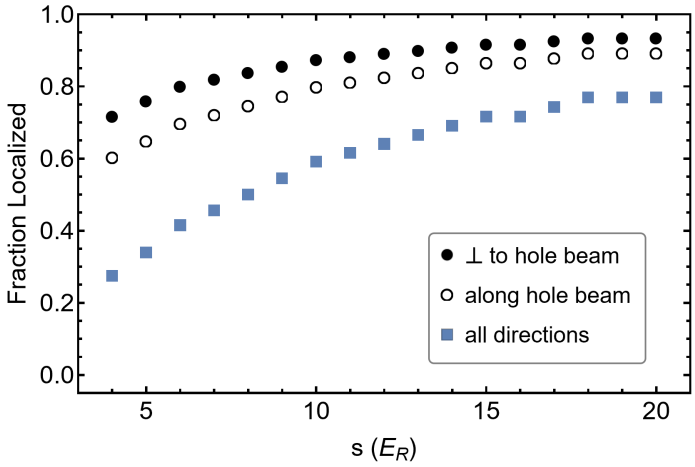

Figure 5.12: An upper bound of localization with localized states defined as states without weight on the central lattice site. Black circles represent the fraction localized along one direction: perpendicular (filled) and along (open) the barrier beam. Blue square represent the fraction localized in all three lattice directions. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

When we define localization as exclusion of the wave function of a particle from the central lattice site, we see that a significant fraction of the particles are localized even at the lowest lattice depths used in this work. We would expect (and observe in Section 6.3) Stark localization to influence the dynamics — not only at high lattice depths, but also in the lowest lattice-depth regime. However, this calculation only accounted for the single-particle physics of the system. Interactions between particle could reduce the localization. In Section 6.4.1, we perform a similar exact-diagonalization calculation that includes interactions and observe that localization persists even in the presence of realistic interactions.

Chapter 6

Relaxation Measurement

Now that I have explained the tools used in this work and characterized the initial state of the gas, I will discuss the dynamics we observed. We measured the hole in the density profile fill in over time. We quantified whether the gas had reached equilibrium with the statistical measure Mardia's B (Section 6.2). We saw the gas thermalized for all lattice depths. The dynamics of the highest lattice depths were very slow, with relaxation times of hundreds of tunneling times (Section 6.3). We considered many possible sources of thermalization (Section 6.4), ruling out lattice scattering and interactions as likely sources of relaxation. We conclude that physics beyond the Bose-Hubbard model may be needed to describe the long-time dynamics of quantum lattice gas systems. The work discussed in this chapter was also described in Reference [4]. Some of the text and figures included in this work were also included in Reference [4]: Copyright 2023 by The American Physical Society.

6.1 Observations of Dynamics

To observe dynamics, we remove the barrier by sigmoidally ramping the barrier beam power to zero over 300 μs. The density distribution was allowed to evolve in the trap and lattice potential for a variable time. After this evolution time, an image was taken *in situ* with variable repumping to control the optical depth and to mitigate imaging artifacts (Section 3.6). We observe that the hole fills in over time (Figure 6.1). We developed a measure, Mardia's B statistic, to quantify the change in the distribution and measure how consistent the distribution is with thermal equilibrium.

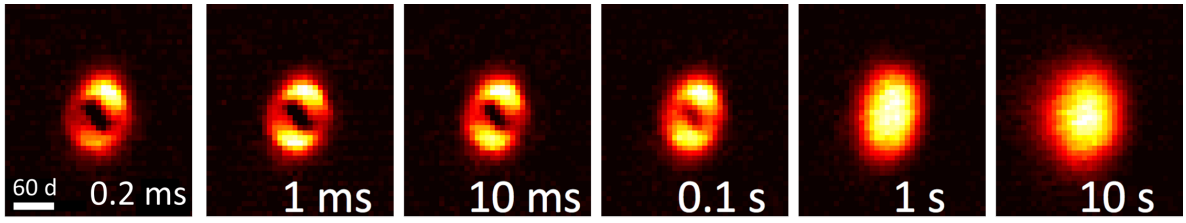

Figure 6.1: Series of images showing relaxation of the gas over time

6.2 Mardia's B Statistic

We use Mardia's B statistic (B) to quantify the dynamical timescale of the change in the density profile and to determine whether the distribution ultimately achieves equilibrium. Mardia's B is a multivariate measure of Gaussianity that is affine-invariant and robust to the overall size, angle, and aspect-ratio of the distribution (which vary over the range of experimental parameters). Since the equilibrium thermal density profile is Gaussian for a trapped gas, Mardia's B is an ideal measure for probing relaxation.

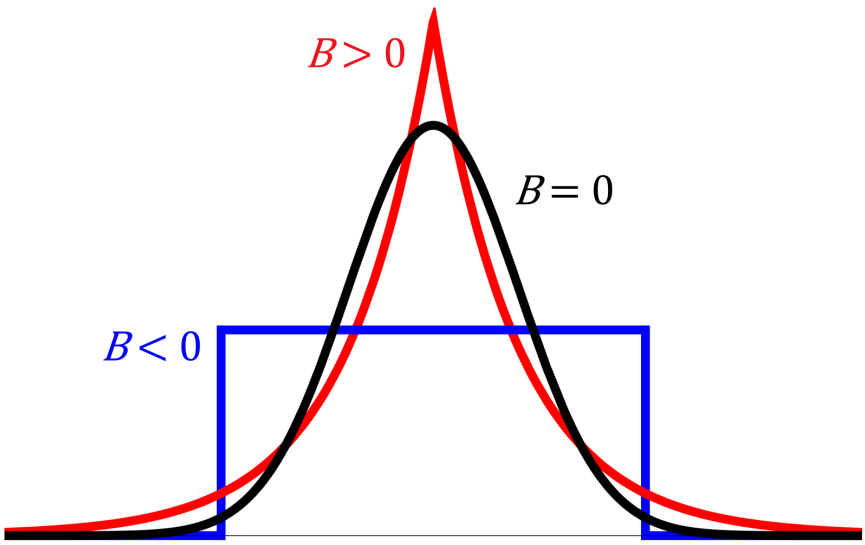

Figure 6.2: Example traces showing the value of B for different kinds of distributions.

Mardia's B is a multi-variate version of kurtosis, which is the fourth moment of a distribution, divided by the second moment squared. Both kurtosis and Mardia's B quantify the shape of the distribution. Put very simply, they measure how much weight is in the tails of the distribution relative to the weight in the center. When the distribution is Gaussian, $B = 0$. For super-Gaussian distributions, where there is more weight in the center of the distribution, $B > 0$ and when the distribution is sub-Gaussian, there is more weight towards the outside of the distribution and $B < 0$ (Figure 6.2). Since B measures the Gaussianity of the distribution independently of its size or aspect ratio, B is an ideal measure to determine whether a harmonically trapped gas is in thermal equilibrium.

Most other measures that have been used to probe density relaxation in strongly correlated lattice gases are insensitive to the shape of the distribution [30–32] as discussed in Section 1.2. Mardia's B measures how closely an image is described by a Gaussian distribution and does not depend on the size of the distribution, so it can be compared across gases with different sizes. This is important as the size of the gas can vary with the lattice depth (as the confinement changes) or the temperature of the gas (which can increase over time). Mardia's B is also insensitive to the aspect ratio and rotation angle of the Gaussian distribution.

Mardia's B for an image is determined according to:

$$B = \frac{1}{8}\sum_{i=1}^{n} w_i \left[\begin{pmatrix} x_i - \bar{x} & y_i - \bar{y} \end{pmatrix} \hat{\Sigma}^{-1} \begin{pmatrix} x_i - \bar{x} \\ y_i - \bar{y} \end{pmatrix} \right]^2 - 1, \tag{6.1}$$

where

$$\hat{\Sigma} = \sum_{j=1}^{n} w_j \begin{pmatrix} x_j - \bar{x} \\ y_j - \bar{y} \end{pmatrix} \begin{pmatrix} x_j - \bar{x} & y_j - \bar{y} \end{pmatrix}, \tag{6.2}$$

and w_i is the normalized weight at pixel i, x_i is the horizontal postion and y_i is the vertical position of pixel i, and $\bar{x}$ is the horizontal centroid and $\bar{y}$ is the vertical centroid [87]. An example point and centroid are illustrated in Figure 6.3.

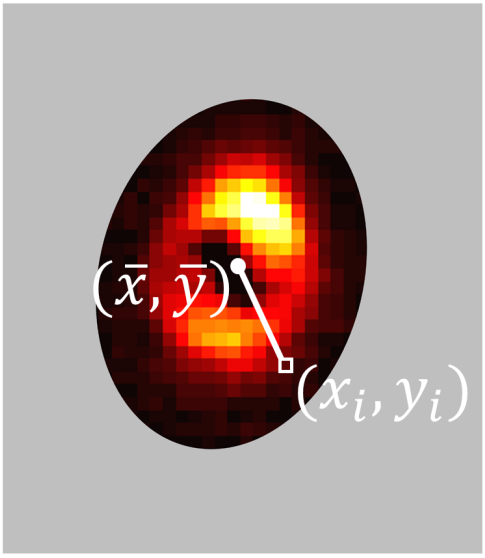

Figure 6.3: Diagram showing points used in Mardia's B calculation.

6.2.1 Masking for Mardia's B

To suppress the impact of imaging noise, we mask the contribution of pixels at large radii, which introduces a small systematic shift in B for an equilibrium Gaussian distribution. Since B is very sensitive to pixels away from the center of the distribution, background imaging noise outside of the atomic distribution can contribute strongly to the value of B. To reduce these effects, a mask was used to calculate B for only the region of the image containing the gas. The mask was determined by fitting the tails of the distribution (where the barrier potential does not influence the density profile) to a rotated Gaussian distribution. Details of this fitting process are explained in Section 6.2.3. The size of the mask must be based on information for each individual image (instead of using one mask for the entire data set) because the size of the gas changes over time (Section 5.3.3). Using a constant mask would introduce systematic error based on the changing size of the distribution.

An elliptical mask is created for each image that removes all pixels outside of the 3σ contour of the fit of a rotated Gaussian distribution to each image. A mask size of 3σ provides a balance that greatly reduces statistical uncertainty in B, while including nearly the entire distribution. The effect of different mask sizes is shown in Figure 6.4. In the figure, Mardia's B was calculated for simulated images of the sum of a large positive Gaussian distribution and a small, negative Gaussian distribution such that the weight $W(\vec{r})$ on the

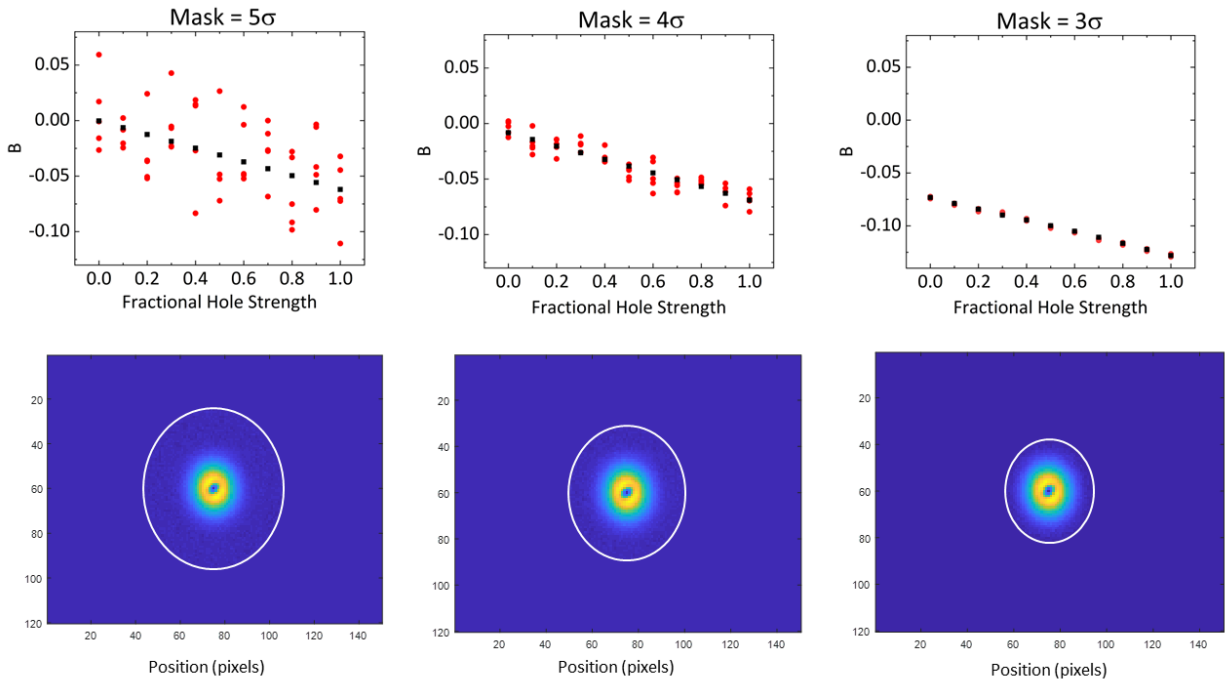

Figure 6.4: The effect of mask size on Mardia's B. Images are simulated distributions of a smaller Gaussian distribution subtracted from a larger Gaussian distribution. Realistic noise was added to the images and Mardia's B was calculated for different mask sizes and fractional hole strengths, $F_{\rm hs}$. Red circles represent B for images with noise, and black squares represent B for images without noise. The example images show the $F_{\rm hs} = 1$ case.

pixel at position $\vec{r}$ is determined by

$$W(\vec{r}) = e^{-\vec{r}^{\,2}/(2\sigma^2_{\rm gas})} - F_{\rm hs} e^{-\vec{r}^{\,2}/(2\sigma^2_{\rm hole})} \tag{6.3}$$

where the size of the positive Gaussian, the "gas", is give by $\sigma_{\rm gas}$ and the size of the negative Gaussian, the "hole", is given by $\sigma_{\rm hole}$. This is a slight simplification of the equation used to generate the images, which allowed for elliptical Gaussian distributions.

The amplitude of the negative Gaussian distribution $F_{\rm hs}$ is the fractional hole strength and was varied between 0 and 1. In images where $F_{\rm hs} = 0$, only a positive Gaussian distribution, the "gas", is present and images where $F_{\rm hs} = 1$ have a completely empty region in the center of the image.

To determine the effect of imaging noise on B, we added realistic noise to the simulated images. Realistic imaging noise was determined by measuring the distribution of optical depth values on pixels in empty regions of real images. The imaging noise was found to be well-approximated by a normal distribution with a standard deviation of 11 mOD. The red points of Figure 6.4 show the value of B for images with realistic imaging noise. For a large mask, the noise in B from the imaging noise is greater than the change in signal of B from the changing strength of the hole.

Since a mask at 3σ cuts off the edge of the tails of the distribution, the mask introduces a small shift in the value of B to -0.07 for a Gaussian distribution. Since the mask is based on the size and angle of the distribution, the masked value of B is uniform for all Gaussian distributions. This was tested for simulated images with a range of aspect ratios and sizes.

6.2.2 Measuring B for an Equilibrium Gas

To check the validity of Mardia's B as a measure of equilibrium, we measured B for gases known to be in equilibrium (Figure 6.5). Specifically, we prepared gases without the barrier potential present at any stage of evaporation. We added a lattice potential with a variety of lattice depths representative of those used in this work. The value of B was measured after hold times of 5 to 10 seconds. We observed values of B that range from -0.105 to -0.054, which we consider to be the uncertainty band for consistency with equilibrium (gray band in Figures 6.9, 6.10, and 6.11). This is approximately equal to the range that is within two standard deviations of the mean (-0.109 to -0.057). The values of B for equilibrium gases are plotted versus the geometric mean size of the gas ($\sigma = (\sigma_u \sigma_v)^{1/2}$) in Figure 6.5. The value of B does not appear to vary significantly with the size of the gas.

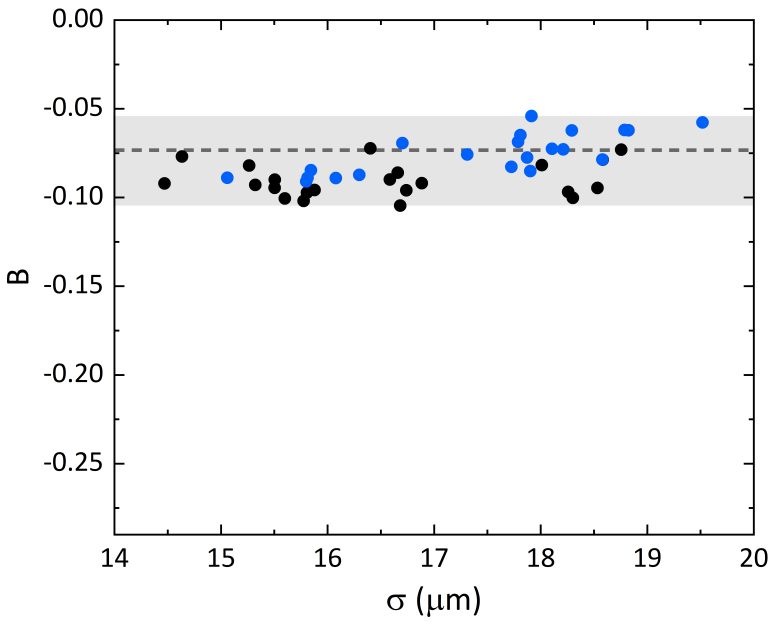

Figure 6.5: Mardia's B for a gas in equilibrium. Gases were prepared without the barrier potential at lattice depths between $s = 0\,E_R$ and $s = 20\,E_R$ and held in the trap for 5 to 10 seconds. Black points were in a trap that was weaker than the trap used to take relaxation data, and blue points were it a trap stronger than the trap used to take relaxation data. The values of B of the equilibrium thermal gases are consistent with the value of B for an ideal Gaussian masked at $3\,\sigma$ (-0.07). This range of values (gray band) was used to define the range of B values that are consistent with a thermalized distribution.

It is possible for sufficiently large gases to have a non-Gaussian distribution in equilibrium, since the confining potential used in this work is made using Gaussian laser beams and therefore deviates from a harmonic potential away from the center of the trap. We used these equilibrium data to verify that the gas was only sampling the harmonic region of the trap. The data in Figure 6.5 were taken in two different traps: one that was stronger ($\omega \approx 60 \times 2\pi$ Hz, blue points) and one that was weaker ($\omega \approx 46 \times 2\pi$ Hz, black points) relative to the trap used for the experiment ($\omega = 54.6 \times 2\pi$ Hz). These equilibrium data include gases with sizes up to 19 μm for both traps. The values of B at this size are consistent with a Gaussian distribution.

In the relaxation data, the size of the gas increased over time after removing the barrier, but was never

larger than the largest sizes measured in these equilibrium data. Specifically, the largest gases seen after removing the barrier potential were $\sigma = 14\,\mu$m (with $\sigma_u = 19\,\mu$m being the largest size in a single direction). Gases in equilibrium in a range of trap depths (which includes the trap depth used for the relaxation measurement) were observed to have a Gaussian distribution, even for large gases in both traps. Since the gases used for the relaxation measurement were not larger than the equilibrium gases measured, we conclude that the gases of the relaxation data were sufficiently small to only samples the harmonic region of the trap.

6.2.3 Image Processing

Finding the Size of the Mask

Calculating Mardia's B for the data involves several steps of image processing in order to reduce the effect of background noise. The first image processing step is to apply the defringing algorithm described in Reference [69].

We also exclude images with a maximum optical depth (defined to be the average of the 5 highest-valued pixels) of more than 1900 mOD (where the optical depth saturation can distort the measured density) or less than 300 mOD (where the imaging noise can pollute the signal). By eliminating images with extreme OD, we are excluding the images with extreme atom number and limiting the study to atomic gases with a slightly smaller range of total atom number than occurs naturally in our evaporation sequence. Images with extreme OD did not appear to have categorically different density distributions than images with typical OD values, and exclusion of these images is unlikely to cause systematic effects.

Once the images are defringed, they are processed with a Matlab program. Using a mask that accurately matches the image is important for calculating an accurate Mardia's B value. In order to find the mask as accurately as possible, we do several stages of fitting the image, where each fit uses results from the previous fit to more accurately find the distribution.[1] An overview of these steps is shown in Figure 6.6.

First, we fit the early-time images where the hold time is zero (specifically less than 0.3 ms) to find the center of the gas and the position and size of the hole. This fitting is done in multiple stages. First, the images (without any masking) are fit to a Gaussian function to find the size and position of the gas. In order to find the typical size and position of the gas, we average the results for all early-time images. This averaging step includes removing outliers with MatLab's `rmoutliers` function, which removes outliers more than three scaled median absolute deviations from the median. This removes all results for a given image if any of the parameters for the image are an outlier. Outlier results are not included in finding the average results of this fitting step, but the outlier images remain in the data-set and are used for further calculations. After removing outliers, we average the remaining fit results to find a starting guess for the next fitter.

Next, we find the center and size of the hole. We fit the early-time images to a sum of two rotated Gaussian distributions. Each Gaussian has its own rotation angle. The fit is constrained based on the results of the first fit of the single Gaussian so that this two-gauss fit returns a positive Gaussian of a similar size to the results of the first fit and a negative Gaussian inside the first one. The exact constraints and guesses used for this fit are available in Appendix A. With this fit, we find the center and size of the hole created by the barrier beam. After removing outliers, the results for every early-time image are averaged and used to create a mask excluding the area within 2 σ_{hole} of the center of the hole. The averaged results of the larger Gaussian are used as guesses for the next fitter.

Next, we loop over all the images and process them using the results of the previous stage of fitting. Each

[1] The code for this software is available in the "image_processing" project shared with the DeMarco Group GitLab.

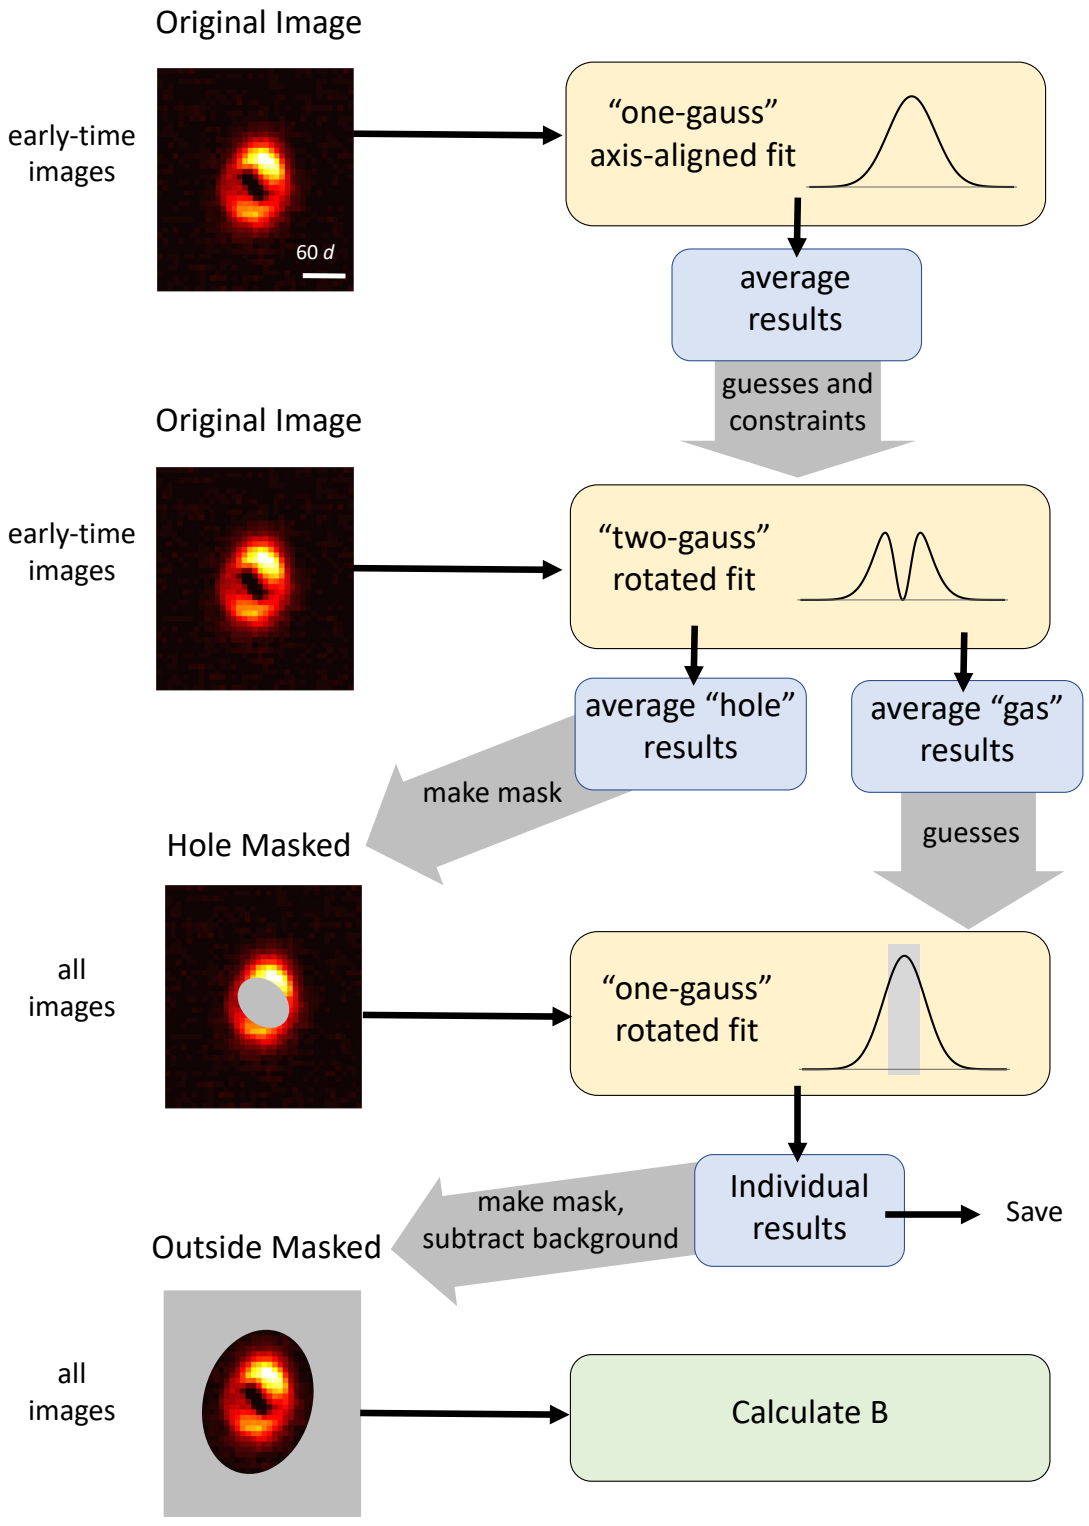

Figure 6.6: Diagram of the series of fitting steps used to create a mask to calculate Mardia's B.

image is masked to exclude the region with the hole, then fit to a rotated Gaussian (where the angle is one of the free parameters) using the results of the early-time fits as starting guesses. The results of this fit to the tails of the gas are used on a per-image-basis to inform the next stage of image processing, namely, calculating Mardia's B.

To calculate Mardia's B, first, the offset given by the fit of the tails of the gas is subtracted from the whole image. A non-zero background would skew the results of the Mardia's B calculation. Next, a mask is created to eliminate everything more than 3σ from the center of the gas according to the fit of the tails. This mask is a rotated ellipse, rotated by the angle given by the fit.

Creating the Mask

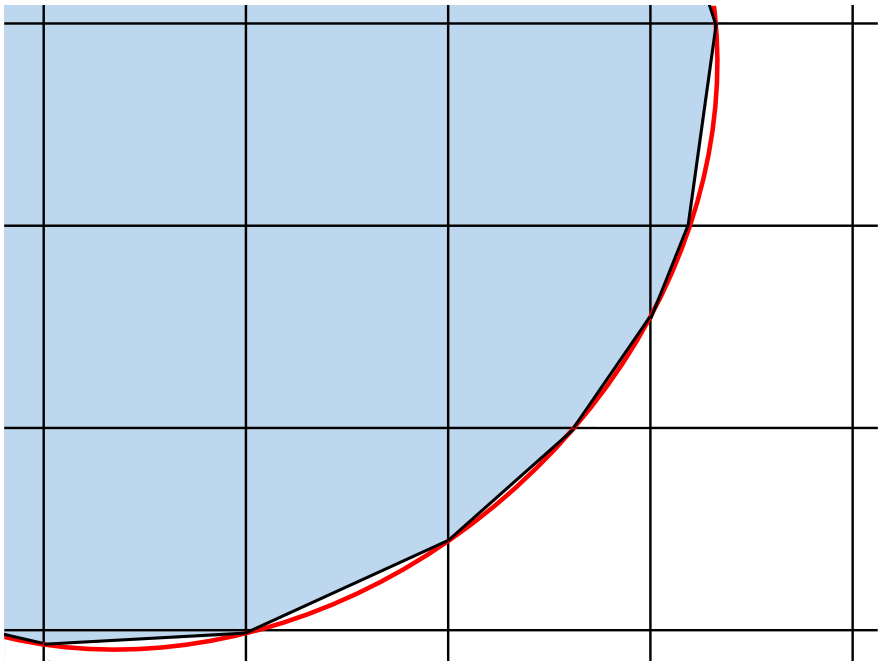

Figure 6.7: Illustration of partial pixel masking. The area of the blue region in each pixel represents the value of the mask at that pixel. A complete square has an area of 1.

The mask partially includes pixels on the boundary of the mask. Since there are a relatively small number of pixels in the image of the gas (the mask has a radius of approximately 10 pixels), changing the mask size by a whole pixel can have a large impact on the value of B. This can lead to discontinuities in the value of B as the mask size increase as the temperature of the gas increases over time. To avoid this effect, pixels that are on the boundary of the mask are partially included. The mask is made up of pixels ranging from zero to one. The pixels that are on the boundary are represented by the fraction of that pixel that is inside the ellipse. Specifically, the area of the triangle, trapezoid, or pentagon that is defined by the points where the ellipse cuts through the edge of the pixel is used to find an approximation of the fraction of the pixel is inside the ellipse. This is shown in Figure 6.7 for an exaggerated example mask. The details of this calculation are given in Appendix B. Since the mask includes fractional weights, when the image is multiplied by the mask pixel-by-pixel, the pixels on the line of the mask are included weighted by the fraction of the pixel that is within the mask. The effect on Mardia's B is approximately the same as upsampling the image and applying a pixel-by-pixel mask to the new image with more pixels. After applying the mask, we normalize the masked

image and calculate Mardia's B as described by Equations 6.1 and 6.2.

Processing the Data

This masking procedure works well as long as the tails of the distribution fit well to a Gaussian distribution, as is true for nearly all the images we collected. However, in less than about 5% of the images, the gas is shifted to one side of the barrier potential. A comparison of these two kind of images can be seen in Figure 6.8. When the gas is shifted to one side of the barrier beam, the distribution that remains after the hole is masked is not Gaussian. The calculated values of B will be very different from those where the hole is in the center of the distribution and the mask, which is based on a fit to a Gaussian distribution, may introduce errors to the calculation. We eliminate these values of B by excluding every image where either the x or y position of the center of the gas (found from the final fitting step where the hole is masked) is more than 2 standard deviations away from the mean position for the data set.

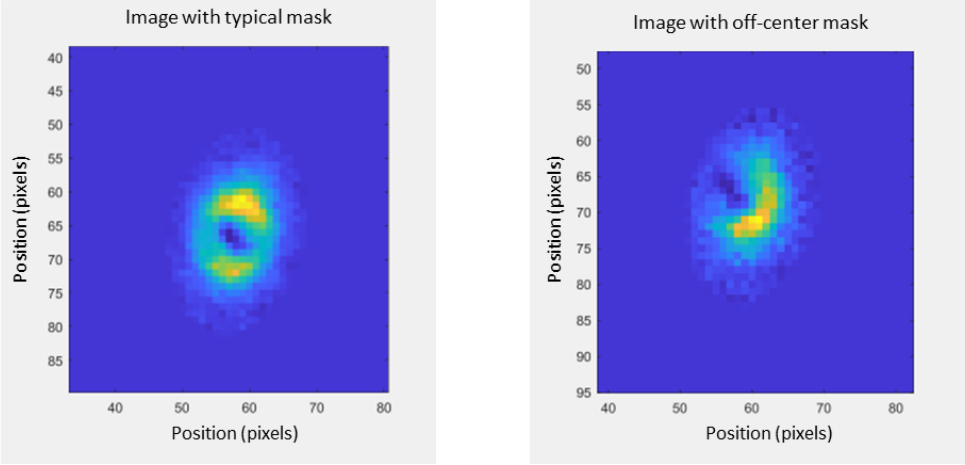

Figure 6.8: Example images showing a normal image (left) and an image where the gas is primarily on one side of the barrier and the mask is off-center from the hole in the density distribution (right).

6.3 Dynamics

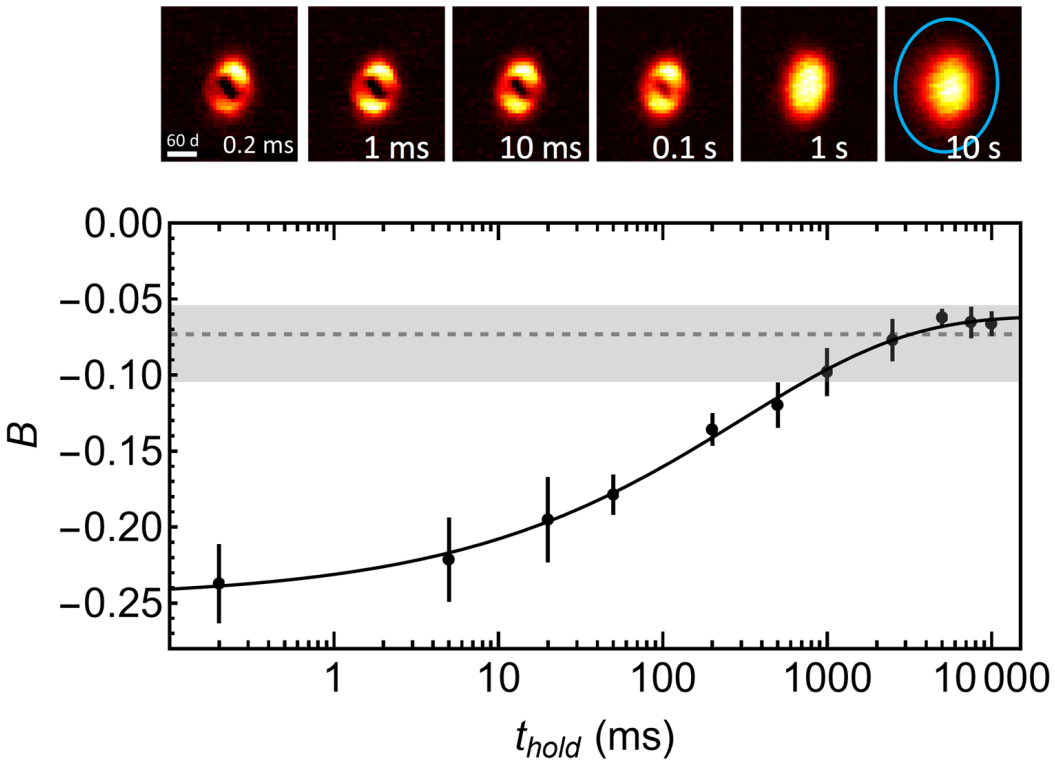

Figure 6.9: Relaxation of the density profile. A series of images taken for $s = 7\,E_R$ is shown, with corresponding measurements of Mardia's B, for different times after the optical barrier potential is removed. Mardia's B increases towards the equilibrium value over long times as the hole disappears from the density profile. The error bars show the standard deviation for the 3–7 measurements averaged at each hold time. The gray dashed line shows B for a Gaussian distribution, and the gray bar displays the range of B values expected for a gas in equilibrium. The elliptical mask used to suppress the effect of imaging noise is superimposed (light blue) for the image at $t_{hold} = 10\,\mathrm{s}$. The solid black line is a fit to a stretched exponential. This figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

We use Mardia's B to quantify how the hole fills in over time. Figure 6.9 shows how B changes over time for sample data taken at $s = 7\,E_R$. Images of the distribution at several hold times are shown above the plot of B. The hole in the density profile disappears over hundreds of milliseconds, as B becomes consistent with equilibrium.

The $s = 7\,E_R$ data is representative of every lattice depth — we observe that the hole in the density profile fills in and B becomes consistent with an equilibrium distribution for every lattice depth. The data for all lattice depths are shown in Figure 6.10. For all lattice depths probed in this work, we find that the time dependence of B fits well to a stretched exponential: $B = B_\infty - A\, e^{-\left(\frac{t}{\tau}\right)^\beta}$, where B_∞ represents the long-time value of B, τ is a time-constant-like parameter, and β is the stretching exponent (Figure 6.9). The timescale of the stretched exponential fit τ is the time when $B = B_\infty - A\,\frac{1}{e}$, the same as a regular exponential fit, for all values of β. A stretched exponential was chosen because it captures the behavior of the data well. The adjusted R^2-values range from 0.87 to 0.97. Neither a regular exponential fit nor a fit to the sum of two exponentials fit the data well at all lattice depths.

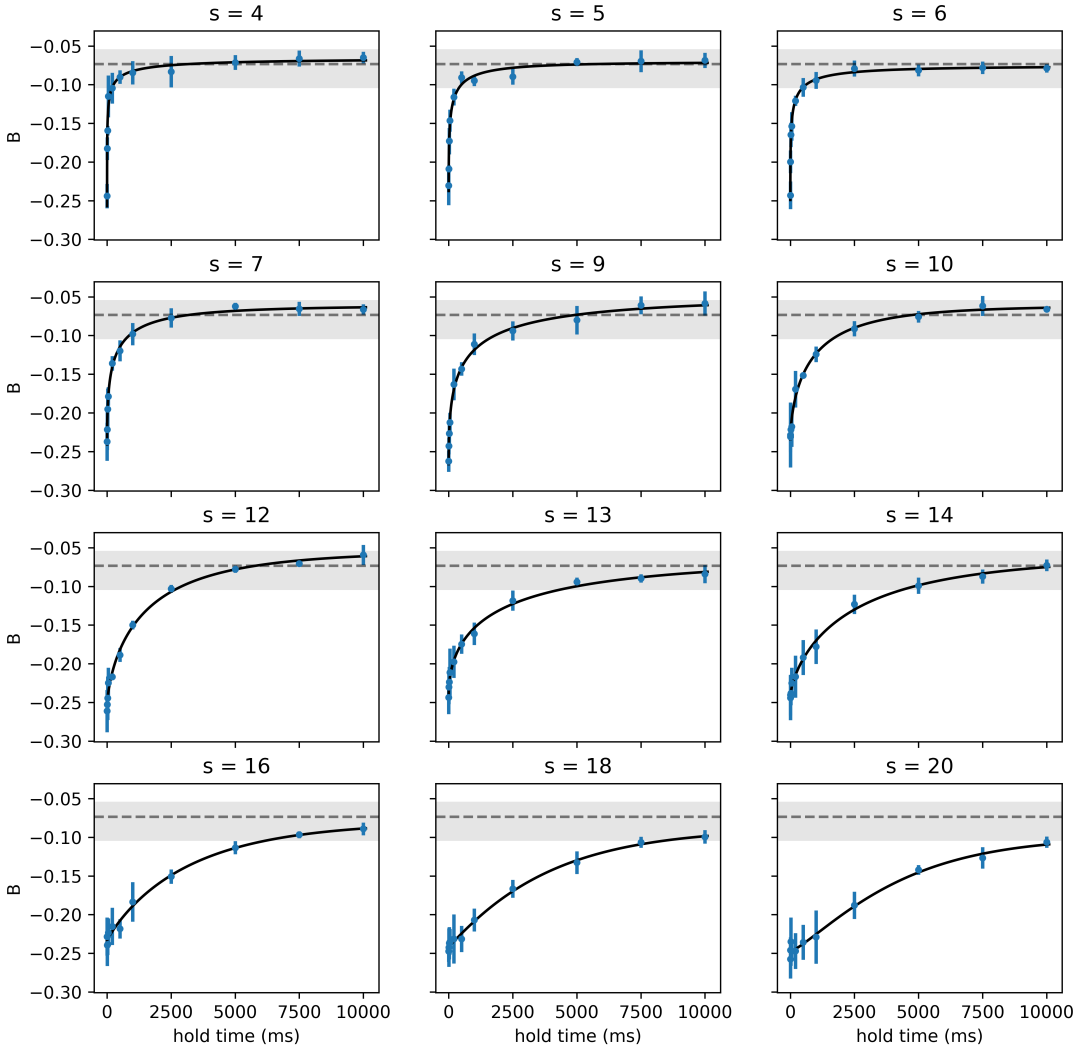

Figure 6.10: The value of B over time for every lattice depth. Error bars are the standard deviation of 3 – 9 measurements. The solid black line is stretched exponential fit to the data. The stretched exponential fit fits the data well for all lattice depths, with adjusted-R^2 values ranging from 0.87 to 0.97. The values of s are in units of E_R.

We observe that equilibrium is achieved at long times for $s = 4 - 20\,E_R$ (Figure 6.11), which includes the regime of complete single-particle localization along at least one lattice direction and nearly complete localization along three directions (Section 5.4). This behavior suggests that Stark localization is disrupted, since the presence of localization of particles along even one direction would prevent thermalization of the density profile.

The thermalization of Stark localization in this system is unexpected, particularly for the largest lattice depths where localization in nearly complete. Above the threshold of localization, dynamics are generally expected to remain frozen [44]. To our knowledge, relaxation of fully Stark-localized states has not been observed previously. Although there has been relatively limited work measuring the long-time dynamics of Stark-localized systems [31, 32] and weak relaxation effects may not have been discernible on shorter times scales. In Section 6.4, we discuss possible causes of the observed relaxation. Although we observe the system to reach equilibrium, the dynamics are much slower than would be expected for non-localized particles.

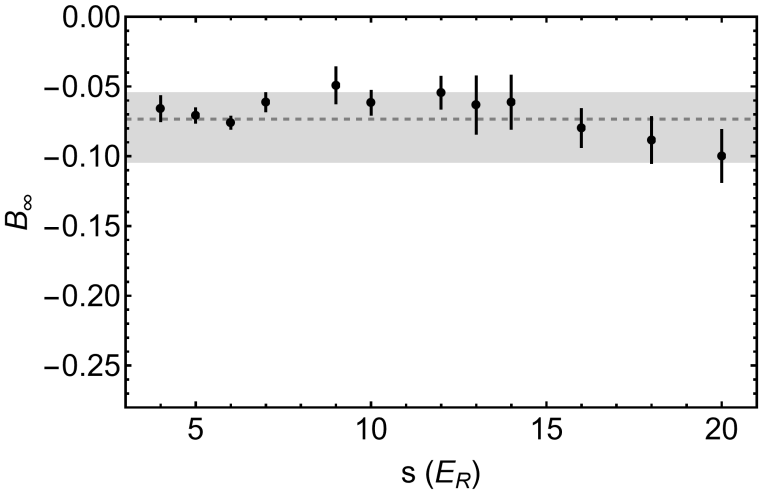

Figure 6.11: Long-time value of Mardia's B determined from the stretched-exponential fits shown in Figure 6.10. Mardia's B reaches a value consistent with equilibrium for all lattice potential depths s. The error bars show the fit uncertainty. The dashed line shows the value of B for a Gaussian distribution, and the gray bar is the range of B for an equilibrated gas determined by measurements of equilibrated thermal gases. This figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

The timescale τ for establishing equilibrium grows with increasing lattice depth and exceeds several seconds at the highest lattice depth, as shown in Figure 6.12(a). This trend is consistent with previous quasimomentum relaxation measurements of our lab group [75], which revealed more rapid relaxation at higher lattice depths. Faster relaxation of momentum implies slower equilibration of the density distribution, since the self-diffusion constant is proportional to the time-integrated velocity autocorrelation function [88]. The timescale for thermalization, τ, rapidly increases for $s > 6\,E_R$, which is the regime for which we observed violation of the Mott-Ioffe-Regel criterion [75]. In this regime, doublon binding and un-binding also become energetically suppressed [66, 89, 90], leading to slow dynamics for atoms located on the same site to break apart and for two initially separated atoms to tunnel onto the same site.

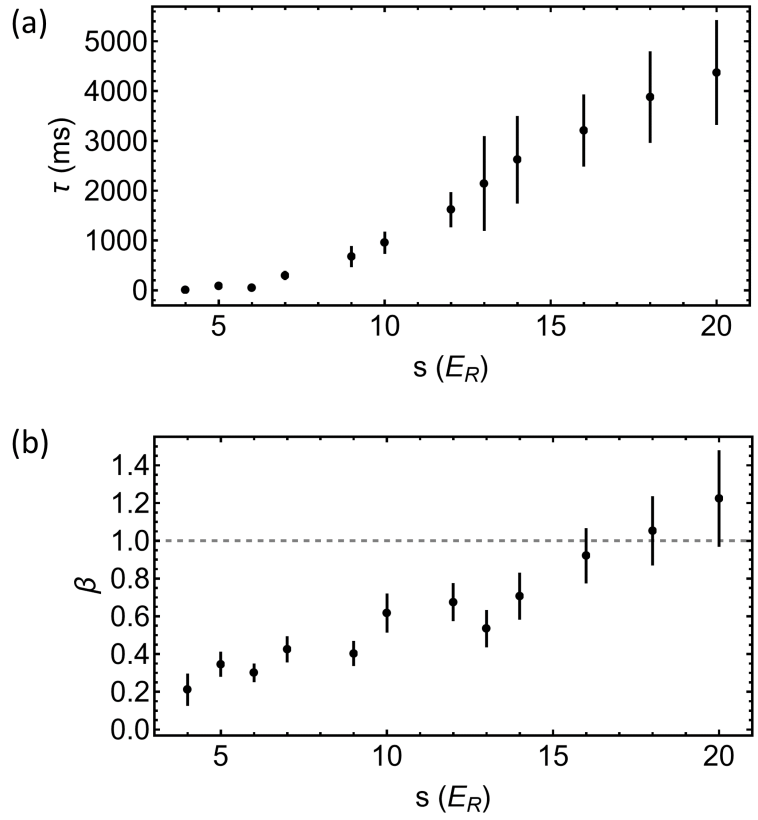

Figure 6.12: The relaxation time τ and stretching exponent β determined from fits of Mardia's B for varied lattice potential depth s. (a) The relaxation time rapidly increases above $s = 6\,E_R$, exceeding several seconds for the highest lattice potential depths. (b) The stretching exponent exhibits non-exponential relaxation for $s < 15\,E_R$. The dashed line shows agreement with exponential behavior. This figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

We also observe that equilibration is non-exponential, except at the highest lattice depths sampled here. Figure 6.12(b) shows how the stretching exponent, β, changes with s. The expectation for gases that are weakly interacting or diffusive is exponential relaxation and $\beta = 1$ [88]. However, for $s < 15\,E_R$, we observe sub-diffusive behavior (i.e., $\beta < 1$), which was also measured in a thermal two-dimensional tilted Fermi-Hubbard system [32] and for a Bose-Einstein condensate confined in a quasi-periodic lattice in the presence of repulsive interactions [19].

In an unconstrained lattice, a particle can be thought of as undergoing a random walk as it randomly tunnels to neighboring lattice sites. This leads to diffusive behavior, where the average distance of a particle from its starting position goes like $\sqrt{t}$. For an out-of-equilibrium system, equilibration is achieved as particles collide and exchange energy. This relaxation can be shown to occur exponentially. Additional constraints on the particles can cause relaxation to be sub-diffusive and exhibit relaxation that is non-exponential. Sub-diffusive behavior is often seen in systems near a localization transition. Previous theoretical and experimental work has shown that a dramatic slowdown of relaxation and non-exponential behavior can be induced by localization [44, 54, 91–94].

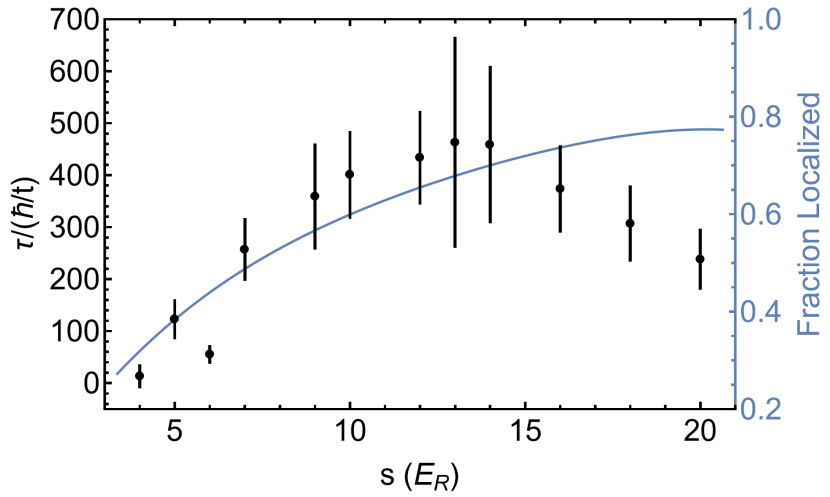

Figure 6.13: Relaxation time normalized to the tunneling time $\hbar/t$ for varied lattice potential depths s. The error bars are determined by fit uncertainty. The blue solid line shows the fraction of particles localized in every direction for the initial state, where localization is defined as exclusion of the wave function from the central lattice site. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

Even without localization, dynamics are expected to be slower at higher lattice depths. As the lattice depth increases, the tunneling between neighboring sites t becomes weaker and the single-particle characteristic tunneling time $\hbar/t$ becomes longer. To separate localization from the suppression of tunneling as the lattice depth is increased, we show the measured relaxation time τ normalized to the tunneling time $\hbar/t$ in Figure 6.13.

When scaled by the tunneling time, the relaxation time increases with increasing lattice depth until approximately $s = 12\,E_R$ where it becomes relatively constant and then begins to fall for the highest lattice

depths. The exact significance of this trend is hard to interpret. Until approximately $s = 14\,E_R$, the normalized relaxation time approximately follows the fraction of localized particles. This suggests that localized states are responsible for the long relaxation times, and equilibration is the result of some effect that delocalizes these states. The decrease in the scaled relaxation times at high lattice depths may be explained by a relaxation effect with a long relaxation time that does not scale with tunneling time. Since the tunneling time increases with lattice depth, a constant relaxation time would appear to decrease when scaled by the tunneling time.

We observe thermalization at all lattice depth, with a relaxation time that increases with increasing lattice depth. The cause of this relaxation remains unclear. In the next section, we consider many possible sources of relaxation.

6.4 Possible Sources of Relaxation

There are several candidates that may disrupt localization. On-site interactions between atoms could allow atoms to exchange energy and relax to thermal equilibrium. Scattering of photons from the light field of the lattice could create a path to delocalization. The semi-classical non-interacting dynamics of the system could create a density distribution that is indistinguishable from equilibrium (discussed in Reference [4]). Atoms could populate excited bands where localization is weaker. Lastly, beyond-Bose-Hubbard terms could introduce a path to equilibration. In this section, we discuss each of these potential effects and their likelihood of explaining the thermalization we observe in the system. We conclude that beyond-Bose-Hubbard physics may be needed to explain the relaxation over the very long-time scales that we observe.

6.4.1 Interaction

The most obvious source of relaxation to consider is on-site interactions. The most significant interaction effect is the Hubbard energy, U, which captures the effect of collisions between particles on the same lattice site. We used an exact diagonalization calculation of an interacting Hamiltonian to estimate the effects of U on the localization in our system. We find that the fraction of particle in localized states is slightly modified by the interaction at low lattice depths and not influenced by interaction at high lattice depths. This suggests that the Hubbard U term does not disrupt localization at the low densities and high lattice depths that we probe.

This calculation is similar to the one done in Section 5.4.3. However, in this case, a set of basis states with multiple particles is necessary to capture the effect of interactions.[2] We created a set of basis states in the Fock space consisting of all possible configurations of two particles populating a one-dimensional, 60-site lattice centered in a harmonic potential. A few example basis states are:

$$|1\rangle = \begin{pmatrix} 2 \\ 0 \\ ... \\ 0 \\ 0 \end{pmatrix} \quad |2\rangle = \begin{pmatrix} 1 \\ 1 \\ ... \\ 0 \\ 0 \end{pmatrix} \quad |3\rangle = \begin{pmatrix} 1 \\ 0 \\ 1 \\ ... \\ 0 \end{pmatrix} \quad |60\rangle = \begin{pmatrix} 1 \\ 0 \\ ... \\ 0 \\ 1 \end{pmatrix} \quad |1829\rangle = \begin{pmatrix} 0 \\ 0 \\ ... \\ 1 \\ 1 \end{pmatrix} \quad |1830\rangle = \begin{pmatrix} 0 \\ 0 \\ ... \\ 0 \\ 2 \end{pmatrix}. \tag{6.4}$$

In this calculation, the number of the basis state no longer corresponds to the lattice site that is occupied,

[2] This calculation was done in "Two Particle Stark MBL.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

since there are two particles in each state. We limit the calculation to two particles on a 60 site lattice to keep the basis set reasonably small. This calculation has 1830 states in the basis set. Two particles on a 300 site lattice (as used for the single particle case) would have 45 150 states in the basis set, which would require more computational resources than we had available. The limitation in the size of the basis set limits how accurately the calculation describes our system, which occurred in a larger lattice. The density of the system may also be different. The average density (density-weighted density described in Section 5.2.2) in our distribution is 0.06 atoms per lattice site. The average density in this calculation is approximately 0.03 atoms per lattice site.

The Hamiltonian used for this calculation is the standard Bose-Hubbard Hamiltonian with a harmonic potential (Equation 2.31),

$$H = -t \sum_{\langle i,j \rangle} (\hat{b}_i^\dagger \hat{b}_j + h.c) + \sum_i \frac{U}{2} \hat{n}_i(\hat{n}_i - 1) + \sum_i \Omega r_i^2 \hat{n}_i, \tag{6.5}$$

where i and j are lattice indices, $\langle i,j \rangle$ represents a sum over nearest neighbors, $\hat{b}_i^\dagger$ and $\hat{b}_i$ represent the creation and annihilation operators, and $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$ gives the number of atoms at site i. The tunneling t, the interaction U, and the harmonic confinement Ω all vary with the lattice depth s as described in Section 2.2.6.

We found the eigenstates of the Hamiltonian and categorized the eigenstates as localized or non-localized. Similar to the single-particle exact-diagonalization calculation, we defined localized states as those with less than 0.001% weight on the central lattice site. The choice of threshold is arbitrary. The results vary less than 10% over three orders of magnitude in the chosen threshold value, and the trend is consistent over this range.

Once we had determined which eigenstates were localized, we needed to determine what fraction of the particles were in a localized eigenstate. Unlike the single-particle case, there is not a straight-forward way to map the two-particle states to single-particle lattice-site occupation, and we were unable to calculate the localization fraction for the initial density profile. Instead, we found an average localization fraction over the final (equilibrium) distribution by taking the thermal average of the eigenstates at the final, measured temperature. The probability of finding the system at temperature T in eigenstate $|E_n\rangle$ is

$$P_n = \frac{e^{-E_n/(k_B T)}}{\sum_n e^{-E_n/(k_B T)}}, \tag{6.6}$$

where E_n is the energy of eigenstate $|E_n\rangle$. The localization fraction l_{1D} is found by weighting the localization of each eigenstate l_n by the probability of that eigenstate:

$$l_{1D} = \sum_n P_n l_n. \tag{6.7}$$

This one-dimensional average approximates the fraction of states that are localized in the equilibrium distribution for each lattice depth based on the measured temperatures of the long-time images.

We did this calculation with both the correct interaction strength, U, for each lattice depth and for $U = 0$. We find that interactions modify the fraction localized at low lattice depths, but do not affect the localization at high lattice depths (Figure 6.14). Interactions reduce the thermally averaged fraction of localized eigenstates for the two-particle Hamiltonian by 20% at $s = 4\,E_R$. This effect is suppressed at higher lattice depths. For $s = 10\,E_R$, interactions have a 1% effect, and the effect of U at $s = 20\,E_R$ is insignificant. The weak effect of interaction may, in part, be due to the relatively small presence of doubly occupied sites.

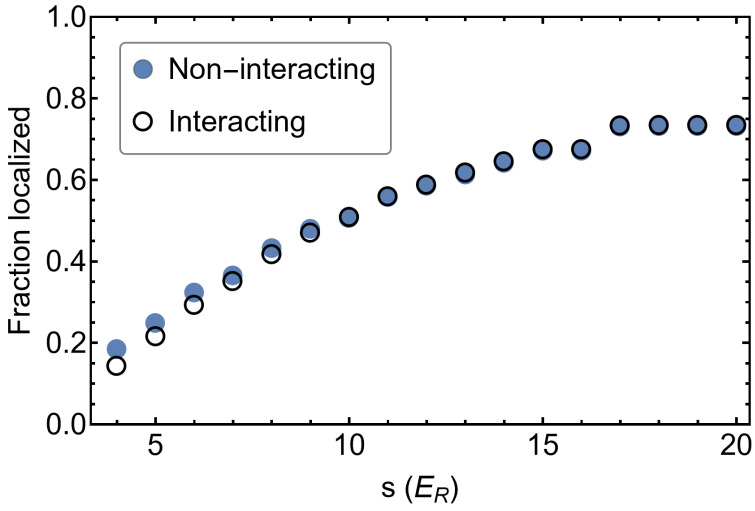

Figure 6.14: The fraction of states localized for interacting (open black) and non-interacting (closed blue) particles in equilibrium as calculated for the thermally-averaged eigenstates of a one-dimensional, 60-site lattice in a Bose-Hubbard Hamiltonian. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

In the final state at $4\,E_R$, 5% of atoms are in a site with greater than single occupation. At $20\,E_R$, only 2% of atoms are in a site that is more than singly occupied. This calculation shows that localization persists even in the presence of interactions, suggesting that interactions are not responsible for the observed thermalization.

6.4.2 Lattice Scattering as a Source of Relaxation

As discussed in Section 5.3.3, the lattice light can off-resonantly scatter and heat the atoms. This interaction with the light field could act as a bath to allow the atoms to thermalize. To determine the influence of lattice-light heating, we measured the relaxation rate τ at fixed s but different lattice wavelengths.

If the lattice scattering was the dominant source of relaxation, we would expect the relaxation times to be proportional to the inverse of the heating rate. The heating rate for lattice scattering is given by Equation 5.16 and shown in Figure 6.15 for a range of wavelengths near the lattice wavelength used in this work. The heating rate depends on the detuning from resonance. The nearest resonance transition occurs at 795 nm and the heating rate changes quickly near this wavelength.

We measured the relaxation time in a $10\,E_R$ lattice for λ =803 nm, λ =812 nm, and λ =825.16 nm lattices (Figure 6.16). We found that the relaxation time changes very little with lattice depth, changing by 15% over the measured range. However, the inverse of the heating rate, which is proportional to the scattering-induced relaxation time, changes by a factor of 4 over this range.[3] The gray dashed line in Figure 6.16 shows the inverse heating rate scaled to be equal to τ at $\lambda = 812$ nm. The lack of agreement between the scattering-induced relaxation time and the measured relaxation time suggests that the lattice light is likely not responsible for thermalization.

We used a bootstrap analysis to estimate the probability of our measured data arising as a random

[3]This was calculated in "Scattering Rate.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

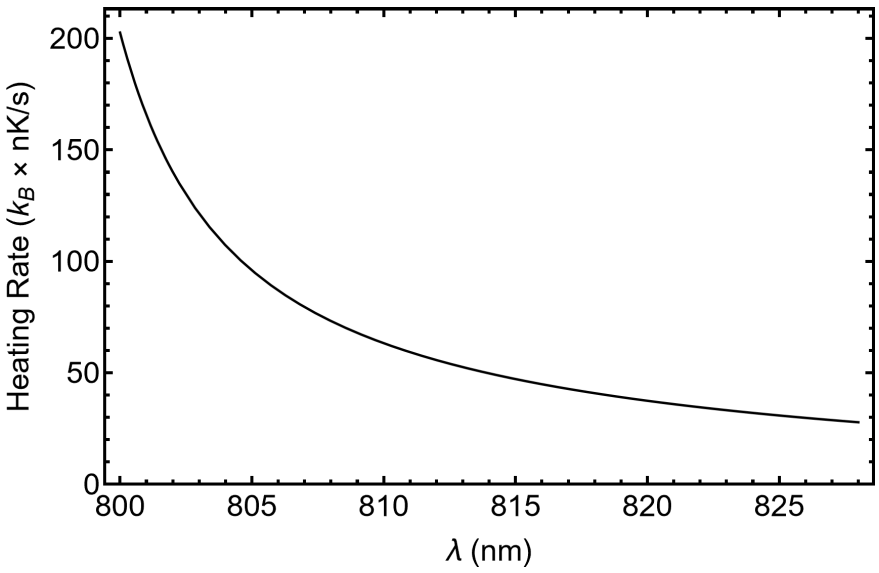

Figure 6.15: The predicted heating rate from the lattice light for a $s = 10\,E_R$ lattice over a range of lattice-light wavelengths.

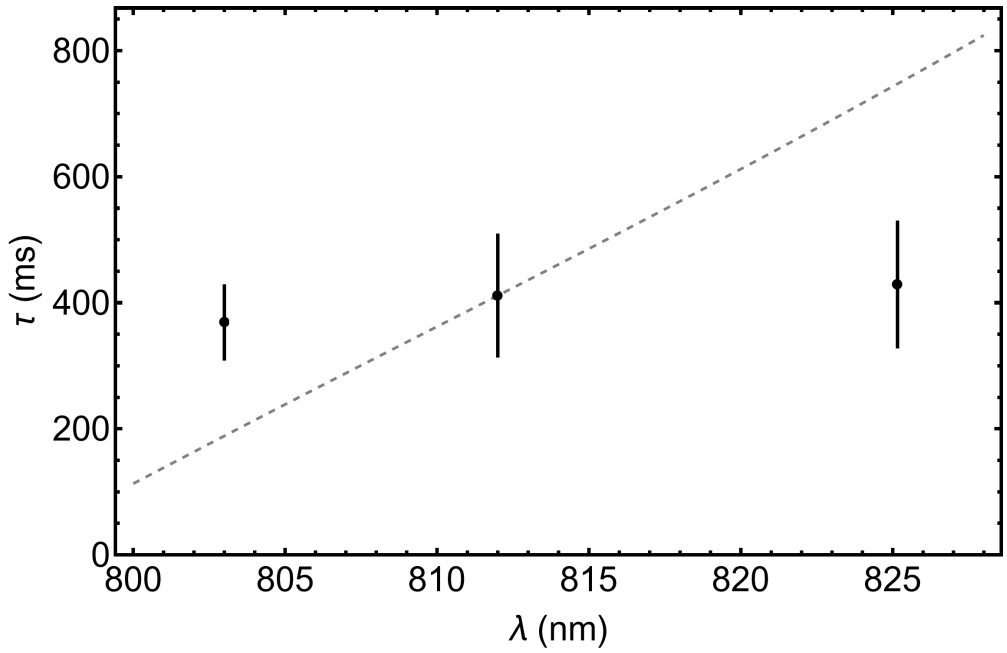

Figure 6.16: Relaxation time at $s = 10\,E_R$ lattice for varied lattice wavelength. The lattice potential depth was kept fixed by tuning the lattice laser optical power. The gray, dashed line shows the dependence of the inverse lattice heating rate on wavelength. The error bars are determined by fit uncertainty. Based on a bootstrap analysis, the measured slope (2.9 ± 0.8) ms/nm is inconsistent with the predicted scaling at greater than the 99.99% confidence level. A modified version of this figure appears in Reference [4]: Copyright 2023 by The American Physical Society.

instance of the distribution of relaxation times, τ, predicted by the scattering rate. We created 2×10^7 sets of three points, one at each measured wavelength, sampled from a normal distribution centered at the τ value predicted by the scattering rate with a standard deviation given by the average uncertainty of the data. We then fit these points to find the slope. We found that 0.00014% of samples had a slope smaller than the largest slope consistent with a fit to the data (3.7 ms/nm, which is the slope returned by the fit plus the fit uncertainty in the slope). Approximately twice as many points, 0.00027%, were less than the maximum slope drawn from the top of the top point's error bar and the bottom of the bottom point's error bar (9.8 ms/nm). It is extremely unlikely that our data are consistent with a relaxation time that is proportional to the inverse of the lattice heating rate. We conclude that heating induced by the lattice light is not the dominant source of relaxation. The other light source present at the atoms is the 1064 nm light of the optical dipole trap. The predicted scattering from this wavelength and intensity of light is one-tenth of the scattering from the lattice light at $s = 20\,E_R$ and is constant across all lattice depths. Therefore, we conclude that scattering from the 1064 nm light is also not a significant source of relaxation.

6.4.3 Excited Band Atoms

Another possible source of relaxation is mixing of excited-band states. For all calculations done this far, we have taken the approximation that all particles are in the ground band of the lattice. However, it is possible that some of the atoms are in an excited state, since a fraction of the atoms in the thermal distribution have as much potential energy as the band gap between the ground and first excited bands of the lattice.

The exact band structure for a three-dimensional lattice is complicated, but we can find an approximation of the excitation levels of a three-dimensional lattice using the one-dimensional band structure along each of the lattice directions. To estimate the population of particles that could be in an excited state, we first found the band gap between the ground and excited bands through exact diagonalization in one dimension, as described in Section 2.1.2. We make the approximation that the band is flat, and measure the band gap as the distance between the centers of the bands, which are relatively flat for high lattice depths. For $s = 20\,E_R$, the gap is $E_{\text{gap}} = 1290\,k_B\times$ nK. Since the lattice potential is separable, the band index of the excitation level n_i of a particle is independent in each lattice direction i.

n_x	n_y	n_z	$E_{\text{excitation}}(n_x, n_y, n_z)$
0	0	0	0
1	0	0	E_{gap}
0	1	0	E_{gap}
0	0	1	E_{gap}
1	1	0	$2\ E_{\text{gap}}$
1	0	1	$2\ E_{\text{gap}}$
0	1	1	$2\ E_{\text{gap}}$
1	1	1	$3\ E_{\text{gap}}$

Table 6.1: The total energy E for a particle in excitation level n_i for each of the i lattice directions.

In three dimensions, there are eight total states that involve only the ground or first excited state of the three one-dimensional lattices, as shown in Table 6.1. This approximation of the band energies neglects the bandwidth and considers any state with the same number of excitations to have the same energy, instead of the range of energies that could be accommodated in the band. We did an atomic limit calculation considering each of these possible states. We created a grid of lattice sites in the positive octant of the trap, with sites up

to 70 lattice sites from the center. We used the cylindrically symmetric harmonic trapping potential $V(x, y, z)$ defined by Equation 5.21. Adding the potential energy from the trap to the energy from the band excitation gives the total energy per a particle at each lattice site in each excited state,

$$E(x, y, z, n_x, n_y, n_z) = V(x, y, z) + E_{\text{excitation}}(n_x, n_y, n_z). \tag{6.8}$$

We then use Maxwell Boltzmann statistics to find the probability of each state at temperature T,

$$P(x, y, z, n_x, n_y, n_z) = \frac{e^{-E(x,y,z,n_x,n_y,n_z)/(k_B T)}}{\sum e^{-E(x,y,z,n_x,n_y,n_z)/(k_B T)}}, \tag{6.9}$$

which is probability $P(x, y, z, n_x, n_y, n_z)$ of finding a particle on lattice site (x, y, z) in excitation state (n_x, n_y, n_z). These probabilities determine the population in the excited bands. We calculated the fraction in the higher energy bands for a gas for the strongest lattice depth measured ($s = 20\, E_R$) at the temperature (700 nK) observed at the longest hold time (10 s)[4]. We determine that 30% of the atoms are in the first excited band, and 5% are in the second excited band. The bandwidth of the first excited band, $0.24\, E_R$, is larger than the ground band, but still small relative to the gradient away from the center of the trap. For $s = 20\, E_R$, the gradient from the harmonic trap is greater than the first excited bandwidth by 11 sites away from the center in the vertical direction and 14 sites away from the center in the horizontal direction. For the equilibrium density distribution at 700 nK, approximately 70% of the atoms in the first excited band are localized to a single lattice site. We conclude that, although excited-band population is possible, excited-band atoms would also experience localization and the presence of excited-band atoms is unlikely to explain the observed thermalization.

6.4.4 Beyond Hubbard Terms

Another source of delocalization and thermalization may be terms beyond the Bose-Hubbard expansion. These represent the full physics of atoms undergoing s-wave collisions in a continuous sinusoidal potential. The largest of these terms is next-nearest-neighbor tunneling. The timescale associated with the next-nearest-neighbor tunneling energy is shown in Figure 6.17. Across the range of lattice depths we probe, this timescale is within an order of magnitude of τ, implying that this effect may play a role in equilibration. The timescale associated with the next-largest beyond Bose-Hubbard term, nearest-neighbor interactions, corresponds to 90 ms at $s = 4\, E_R$ and 97 000 ms at $s = 20\, E_R$, and is therefore likely too small to contribute. These terms are discussed in more detail in Section 2.2. Although the Bose-Hubbard model has been successful in describing the short-time dynamics of bosons in a lattice potential, a more complete picture may be needed to explain the long-time dynamics of these systems.

Experimental studies exploring beyond Bose-Hubbard physics have been limited, but the extended Bose-Hubbard model has been considered numerically. The phase diagram of the extended Bose-Hubbard model (including nearest-neighbor interactions) was explored numerically in Reference [95] and [96]. Reference [97] considers the effects of the addition of both nearest-neighbor interaction and next-nearest-neighbor tunneling on the energy-level spacing for the Bose-Hubbard model. Reference [98] explores a scheme to use the excited states of atoms in an optical potential to realize long-range interactions. Consideration of whether exited-band–induced long-range interactions are present in this work may be warranted.

To our knowledge, previous experimental realization of the extended Bose-Hubbard model has been

[4]This calculation was done in "Excited State Density Distribution.nb" in `Z:\Pubs\Laura Wadleigh\thesis code`.

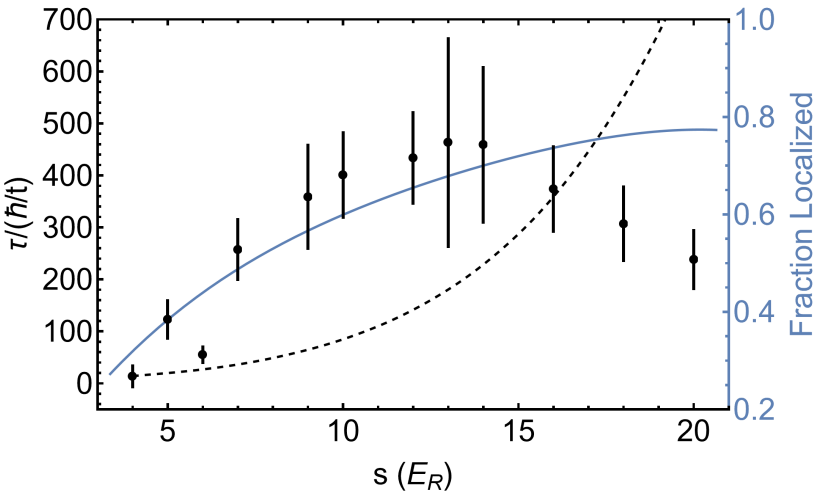

Figure 6.17: The next-nearest-neighbor tunneling time t_{nnn} normalized to $\hbar/t$ is shown along with the relaxation time normalized to the tunneling time $\hbar/t$ (black points) and the fraction of particles localized in every direction for the initial state, where localization is defined as excluded from the central lattice site (blue line).

limited to particles with dipolar interactions, such as magnetic erbium atoms [99] and semiconductor dipolar excitons [100]. This work observes thermalization of a localized lattice gas without dipolar interactions. This thermalization is not due to scattering of lattice light, suggesting that extended Bose-Hubbard effects may also be relevant to realizations of the Bose-Hubbard model without dipole interactions.

The energy scales of extended Bose-Hubbard terms are small in this system. Therefore, the effects of these terms may only be discernible over long-time scales. If localization is to be considered a truly stable state, such long-time effects must be understood. More modeling of the continuous lattice potential and the effect of extended Bose-Hubbard terms is necessary to better understand localization in these systems.

Chapter 7

Conclusion

7.1 Conclusion

In this work, we created Stark-localized states in a thermal lattice bose gas with barrier beam. We measured equilibration of nominally localized states over very long timescales (equivalent to hundreds of tunneling times) using a new probe of equilibration for trapped gases that is sensitive to the shape of the entire density distribution. Through direct measurement, we have shown that the equilibration timescales are largely insensitive to coupling to the environment through lattice scattering. We conclude that the Bose-Hubbard model may be insufficient to describe the long-time dynamics of this system. A model that includes the continuum lattice potential may be necessary to describe the low-energy effects that can influence long-time dynamics. This observation questions the stability of localization in other lattice-gas systems. More work is needed to understand the nature of localization in interacting, lattice-gas systems.

7.2 Future Work

The new tools described in this work can be used to study more than the thermal relaxation dynamics discussed in this thesis. We developed the ability to manipulate and measure the atomic density distribution on a much smaller scale than was previously possible in this apparatus. This is broadly useful to explore the dynamics of a variety of quantum states.

The next experiment planned for this apparatus is to measure the diffusion dynamics in superfluid and Mott insulator states. Instead of using the barrier beam to create a hole in the density profile, the beam has been modified to apply a spatially selective spin flip to the atoms. An additional beam has been coupled into the beam path of the barrier beam so that it can apply a Raman transition to the atoms. In this configuration, both beams are detuned from the 795 nm D1 transition, and they are shifted 59 MHz apart to excite a Raman transition from the $F = 1$, $m_F = -1$ to the $F = 1$, $m_F = 0$ hyperfine state. The atoms in either spin state can be selectively imaged. This effectively "spin-tags" the atoms in the center of the gas.

We will measure the dynamics of a spin impurity using this technique. Since the spin-flip transition does not disrupt the density profile, the gas remains in equilibrium, and we can measure diffusion constants in a thermal gas by observing the initial spin perturbation expand over time. We can also measure the diffusion constant in other quantum states, such as a superfluid, Mott insulator, or Bose glass.

References

[1] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, "Colloquium: Many-body localization, thermalization, and entanglement", Reviews of Modern Physics **91**, 021001 (2019).

[2] F. Alet and N. Laflorencie, "Many-body localization: An introduction and selected topics", en, Comptes Rendus Physique, Quantum simulation / Simulation quantique **19**, 498–525 (2018).

[3] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, "Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms", en, Nature **415**, 39–44 (2002).

[4] L. Wadleigh, N. Kowalski, and B. DeMarco, "Interacting Stark localization dynamics in a three-dimensional lattice Bose gas", Physical Review A **107**, 043325 (2023).

[5] R. Nandkishore and D. A. Huse, "Many-Body Localization and Thermalization in Quantum Statistical Mechanics", Annual Review of Condensed Matter Physics **6**, 15–38 (2015).

[6] J. M. Deutsch, "Eigenstate thermalization hypothesis", en, Reports on Progress in Physics **81**, 082001 (2018).

[7] E. Altman, "Many-body localization and quantum thermalization", en, Nature Physics **14**, 979–983 (2018).

[8] S. Bhattacharjee and A. Dutta, "Quantum thermal machines and batteries", en, The European Physical Journal B **94**, 239 (2021).

[9] J. Q. Quach, K. E. McGhee, L. Ganzer, D. M. Rouse, B. W. Lovett, E. M. Gauger, J. Keeling, G. Cerullo, D. G. Lidzey, and T. Virgili, "Superabsorption in an organic microcavity: Toward a quantum battery", Science Advances **8**, eabk3160 (2022).

[10] D. Rossini, G. M. Andolina, and M. Polini, "Many-body localized quantum batteries", Physical Review B **100**, 115142 (2019).

[11] C. M. Chandrashekar and T. Busch, "Localized quantum walks as secured quantum memory", en, Europhysics Letters **110**, 10005 (2015).

[12] R. Peña, T. H. Kyaw, and G. Romero, "Stable Many-Body Resonances in Open Quantum Systems", en, Symmetry **14**, 2562 (2022).

[13] D. A. Abanin and Z. Papić, "Recent progress in many-body localization", Annalen der Physik **529**, 1700169 (2017).

[14] A. Lagendijk, B. v. Tiggelen, and D. S. Wiersma, "Fifty years of Anderson localization", en, Physics Today **62**, 24–29 (2009).

[15] P. W. Anderson, "Absence of Diffusion in Certain Random Lattices", Phys. Rev. **109** (1958).

[16] S. S. Kondov, W. R. McGehee, W. Xu, and B. DeMarco, "Disorder-Induced Localization in a Strongly Correlated Atomic Hubbard Gas", Physical Review Letters **114**, 083002 (2015).

[17] M. Rispoli, A. Lukin, R. Schittko, S. Kim, M. E. Tai, J. Léonard, and M. Greiner, "Quantum critical behaviour at the many-body localization transition", en, Nature **573**, 385–389 (2019).

[18] B. Chiaro, C. Neill, A. Bohrdt, M. Filippone, F. Arute, K. Arya, R. Babbush, D. Bacon, J. Bardin, R. Barends, S. Boixo, D. Buell, B. Burkett, Y. Chen, Z. Chen, R. Collins, A. Dunsworth, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, M. Harrigan, T. Huang, S. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, E. Lucero, J. McClean, X. Mi, A. Megrant, M. Mohseni, J. Mutus, M. McEwen, O. Naaman, M. Neeley, M. Niu, A. Petukhov, C. Quintana, N. Rubin, D. Sank, K. Satzinger, T. White, Z. Yao, P. Yeh, A. Zalcman, V. Smelyanskiy, H. Neven, S. Gopalakrishnan, D. Abanin, M. Knap, J. Martinis, and P. Roushan, "Direct measurement of nonlocal interactions in the many-body localized phase", Physical Review Research **4**, 013148 (2022).

[19] E. Lucioni, B. Deissler, L. Tanzi, G. Roati, M. Zaccanti, M. Modugno, M. Larcher, F. Dalfovo, M. Inguscio, and G. Modugno, "Observation of Subdiffusion in a Disordered Interacting System", Physical Review Letters **106**, 230403 (2011).

[20] B. Min, T. Li, M. Rosenkranz, and W. Bao, "Subdiffusive spreading of a Bose-Einstein condensate in random potentials", Physical Review A **86**, 053612 (2012).

[21] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, "Observation of many-body localization of interacting fermions in a quasirandom optical lattice", Science **349**, 842–845 (2015).

[22] S. Iyer, V. Oganesyan, G. Refael, and D. A. Huse, "Many-body localization in a quasiperiodic system", Physical Review B **87**, 134202 (2013).

[23] P. Sierant and J. Zakrzewski, "Challenges to observation of many-body localization", Physical Review B **105**, 224203 (2022).

[24] D. A. Abanin, J. H. Bardarson, G. De Tomasi, S. Gopalakrishnan, V. Khemani, S. A. Parameswaran, F. Pollmann, A. C. Potter, M. Serbyn, and R. Vasseur, "Distinguishing localization from chaos: challenges in finite-size systems", Annals of Physics **427**, 168415 (2021).

[25] Y.-T. Tu, D. Vu, and S. D. Sarma, "Avalanche stability transition in interacting quasiperiodic systems", Physical Review B **107**, 014203 (2023).

[26] R. Sahay, F. Machado, B. Ye, C. R. Laumann, and N. Y. Yao, "Emergent Ergodicity at the Transition between Many-Body Localized Phases", Physical Review Letters **126**, 100604 (2021).

[27] S. Gopalakrishnan and S. A. Parameswaran, "Dynamics and transport at the threshold of many-body localization", en, Physics Reports, Dynamics and transport at the threshold of many-body localization **862**, 1–62 (2020).

[28] D. Vu, K. Huang, X. Li, and S. Das Sarma, "Fermionic Many-Body Localization for Random and Quasiperiodic Systems in the Presence of Short- and Long-Range Interactions", Physical Review Letters **128**, 146601 (2022).

[29] Q. Guo, C. Cheng, H. Li, S. Xu, P. Zhang, Z. Wang, C. Song, W. Liu, W. Ren, H. Dong, R. Mondaini, and H. Wang, "Stark Many-Body Localization on a Superconducting Quantum Processor", Physical Review Letters **127**, 240502 (2021).

[30] W. Morong, F. Liu, P. Becker, K. S. Collins, L. Feng, A. Kyprianidis, G. Pagano, T. You, A. V. Gorshkov, and C. Monroe, "Observation of Stark many-body localization without disorder", en, Nature **599**, 393–398 (2021).

[31] S. Scherg, T. Kohlert, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger, "Observing non-ergodicity due to kinetic constraints in tilted Fermi-Hubbard chains", en, Nature Communications **12**, 4490 (2021).

[32] E. Guardado-Sanchez, A. Morningstar, B. M. Spar, P. T. Brown, D. A. Huse, and W. S. Bakr, "Subdiffusion and Heat Transport in a Tilted Two-Dimensional Fermi-Hubbard System", Physical Review X **10**, 011042 (2020).

[33] M. Schulz, C. A. Hooley, R. Moessner, and F. Pollmann, "Stark Many-Body Localization", Physical Review Letters **122**, 040606 (2019).

[34] L. Zhang, Y. Ke, W. Liu, and C. Lee, "Mobility edge of Stark many-body localization", Physical Review A **103**, 023323 (2021).

[35] P. Zhang, "Subdiffusion in strongly tilted lattice systems", Physical Review Research **2**, 033129 (2020).

[36] G. Zisling, D. M. Kennes, and Y. Bar Lev, "Transport in Stark many-body localized systems", Physical Review B **105**, L140201 (2022).

[37] E. van Nieuwenburg, Y. Baum, and G. Refael, "From Bloch oscillations to many-body localization in clean interacting systems", Proceedings of the National Academy of Sciences **116**, 9269–9274 (2019).

[38] E. V. H. Doggen, I. V. Gornyi, and D. G. Polyakov, "Stark many-body localization: Evidence for Hilbert-space shattering", Physical Review B **103**, L100202 (2021).

[39] E. V. H. Doggen, I. V. Gornyi, and D. G. Polyakov, "Many-body localization in a tilted potential in two dimensions", Physical Review B **105**, 134204 (2022).

[40] X. Wei, X. Gao, and W. Zhu, "Static and dynamical Stark many-body localization transition in a linear potential", Physical Review B **106**, 134207 (2022).

[41] D. S. Bhakuni, R. Nehra, and A. Sharma, "Drive-induced many-body localization and coherent destruction of Stark many-body localization", Physical Review B **102**, 024201 (2020).

[42] V. Khemani, M. Hermele, and R. Nandkishore, "Localization from Hilbert space shattering: From theory to physical realizations", Physical Review B **101**, 174204 (2020).

[43] P. P. Mazza, G. Perfetto, A. Lerose, M. Collura, and A. Gambassi, "Suppression of transport in non-disordered quantum spin chains due to confined excitations", Physical Review B **99**, 180302 (2019).

[44] R. Yao and J. Zakrzewski, "Many-body localization of bosons in an optical lattice: Dynamics in disorder-free potentials", Physical Review B **102**, 104203 (2020).

[45] D. Emin and C. F. Hart, "Existence of Wannier-Stark localization", Physical Review B **36**, 7353–7359 (1987).

[46] E. E. Mendez and G. Bastard, "Wannier-Stark Ladders and Bloch Oscillations in Superlattices", Physics Today **46**, 34–42 (1993).

[47] H. Ott, E. de Mirandes, F. Ferlaino, G. Roati, V. Türck, G. Modugno, and M. Inguscio, "Radio Frequency Selective Addressing of Localized Particles in a Periodic Potential", Physical Review Letters **93**, 120407 (2004).

[48] A. M. Rey, G. Pupillo, C. W. Clark, and C. J. Williams, "Ultracold atoms confined in an optical lattice plus parabolic potential: A closed-form approach", Physical Review A **72**, 033616 (2005).

[49] D. McKay, M. White, and B. DeMarco, "Lattice thermodynamics for ultracold atoms", Physical Review A **79**, 063605 (2009).

[50] T. Chanda, R. Yao, and J. Zakrzewski, "Coexistence of localized and extended phases: Many-body localization in a harmonic trap", Physical Review Research **2**, 032039 (2020).

[51] M. Serbyn, D. A. Abanin, and Z. Papić, "Quantum many-body scars and weak breaking of ergodicity", en, Nature Physics **17**, 675–685 (2021).

[52] J. Feldmeier, P. Sala, G. de Tomasi, F. Pollmann, and M. Knap, "Anomalous Diffusion in Dipole- and Higher-Moment Conserving Systems", Physical Review Letters **125**, 245303 (2020).

[53] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, "Exploring the many-body localization transition in two dimensions", en, Science **352**, 1547–1552 (2016).

[54] H. P. Lüschen, P. Bordia, S. S. Hodgman, M. Schreiber, S. Sarkar, A. J. Daley, M. H. Fischer, E. Altman, I. Bloch, and U. Schneider, "Signatures of Many-Body Localization in a Controlled Open Quantum System", Physical Review X **7**, 011034 (2017).

[55] C.-L. Hung, X. Zhang, N. Gemelke, and C. Chin, "Slow Mass Transport and Statistical Evolution of an Atomic Gas across the Superfluid–Mott-Insulator Transition", Physical Review Letters **104**, 160403 (2010).

[56] N. Ashcroft and D. Mermin, *Solid State Physics* (Thomson Learning, Inc., 1976).

[57] P. A. Russ, "Compressibility of a bosonic Mott-insulator in a disordered optical lattice", Thesis (University of Illinois Urbana-Champaign, 2020).

[58] M. White, "Ultracold Atoms in a Disordered Optical Lattice", PhD thesis (2009).

[59] C. Pethick and H. Smith, *Bose-Einstein condensation in dilute gases*, en (Cambridge University Press, Cambridge ; New York, 2002).

[60] D. A. Steck, *Rubidium 87 D Line Data*, tech. rep. (Los Alamos National Laboratory, Sept. 2001).

[61] D. Mckay, "Quantum simulation in strongly correlated optical lattices", PhD thesis (University of Illinois Urbana-Champaign, 2012).

[62] J. Li, Y. Yu, A. M. Dudarev, and Q. Niu, "Interaction broadening of Wannier functions and Mott transitions in atomic BEC", en, New Journal of Physics **8**, 154 (2006).

[63] O. Dutta, M. Gajda, P. Hauke, M. Lewenstein, D.-S. Lühmann, B. A. Malomed, T. Sowiński, and J. Zakrzewski, "Non-standard Hubbard models in optical lattices: a review", en, Reports on Progress in Physics **78**, 066001 (2015).

[64] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, "Boson localization and the superfluid-insulator transition", Physical Review B **40**, 546–570 (1989).

[65] V. Gurarie, L. Pollet, N. V. Prokof'ev, B. V. Svistunov, and M. Troyer, "Phase diagram of the disordered Bose-Hubbard model", Physical Review B **80**, 214519 (2009).

[66] R. Sensarma, D. Pekker, E. Altman, E. Demler, N. Strohmaier, D. Greif, R. Jördens, L. Tarruell, H. Moritz, and T. Esslinger, "Lifetime of double occupancies in the Fermi-Hubbard model", en, Physical Review B **82**, 224302 (2010).

[67] Y. Khorramzadeh, F. Lin, and V. W. Scarola, "Boson core compressibility", Physical Review A **85**, 043610 (2012).

[68] D. Chen, "Non-equilibrium dynamics of ultracold atoms in optical lattices", PhD thesis (University of Illinois Urbana-Champaign, 2015).

[69] C. Meldgin, "Dynamics of ultracold bosons in a disordered optical lattice", PhD thesis (University of Illinois Urbana-Champaign, 2016).

[70] P. della Porta, C. Emili, and S. J. Hellier, *Alkali Metal Generation an dGas Evolution from Alkali Metal Dispensers*, Technical Report TR 18 ().

[71] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, "Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential", Physical Review A **79**, 063631 (2009).

[72] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, *Optical dipole traps for neutral atoms*, Feb. 1999.

[73] A. S. Poulsen, "Calibration of a 3D Optical Lattice", en,

[74] C. Meldgin, U. Ray, P. Russ, D. Chen, D. M. Ceperley, and B. DeMarco, "Probing the Bose glass–superfluid transition using quantum quenches of disorder", en, Nature Physics **12**, 646–649 (2016).

[75] D. Chen, C. Meldgin, P. Russ, B. DeMarco, and E. Mueller, "Disappearance of quasiparticles in a Bose lattice gas", Physical Review A **94**, 021601 (2016).

[76] D. C. McKay and B. DeMarco, "Cooling in strongly correlated optical lattices: prospects and challenges", en, Reports on Progress in Physics **74**, 054401 (2011).

[77] D. C. McKay, C. Meldgin, D. Chen, and B. DeMarco, "Slow Thermalization between a Lattice and Free Bose Gas", Physical Review Letters **111**, 063002 (2013).

[78] D. McKay and B. DeMarco, "Thermometry with spin-dependent lattices", en, New Journal of Physics **12**, 055013 (2010).

[79] B. Goodwin, *What is the difference between the FFT and Huygens PSF?*, en-US, Mar. 2021.

[80] C. Robens, S. Brakhane, W. Alt, F. Kleißler, D. Meschede, G. Moon, G. Ramola, and A. Alberti, "High numerical aperture (NA = 0.92) objective lens for imaging and addressing of cold atoms", EN, Optics Letters **42**, 1043–1046 (2017).

[81] P. S. Considine, "Effects of Coherence on Imaging Systems", EN, JOSA **56**, 1001–1009 (1966).

[82] E. C. Kintner and R. M. Sillitto, "Edge-ringing in Partially Coherent Imaging", Optica Acta: International Journal of Optics **24**, 591–605 (1977).

[83] V. K. Pustovalov and V. A. Babenko, "Optical properties of gold nanoparticles at laser radiation wavelengths for laser applications in nanotechnology and medicine", en, Laser Physics Letters **1**, 516–520 (2004).

[84] C. Robens, "Testing the Quantumness of Atom Trajectories", en,

[85] J. Lowney, *MATLAB-For-PICAM*, Mar. 2014.

[86] H. T. Grahn, *Semiconductor Superlattices: Growth and Electronic Properties*, en (World Scientific, 1995).

[87] K. V. Mardia, "Measures of Multivariate Skewness and Kurtosis with Applications", Biometrika **57**, 519–530 (1970).

[88] F. Reif, *Fundamentals of Statistical and Thermal Physics*, en (Waveland Press, Jan. 2009).

[89] N. Strohmaier, D. Greif, R. Jördens, L. Tarruell, H. Moritz, T. Esslinger, R. Sensarma, D. Pekker, E. Altman, and E. Demler, "Observation of Elastic Doublon Decay in the Fermi-Hubbard Model", en, Physical Review Letters **104**, 080401 (2010).

[90] K. Winkler, G. Thalhammer, F. Lang, R. Grimm, J. Hecker Denschlag, A. J. Daley, A. Kantian, H. P. Büchler, and P. Zoller, "Repulsively bound atom pairs in an optical lattice", en, Nature **441**, 853–856 (2006).

[91] A. Signoles, T. Franz, R. Ferracini Alves, M. Gärttner, S. Whitlock, G. Zürn, and M. Weidemüller, "Glassy Dynamics in a Disordered Heisenberg Quantum Spin System", Physical Review X **11**, 011011 (2021).

[92] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, K. Boothby, P. Bunyk, C. Deng, C. Enderud, S. Huang, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, T. Lanting, R. Li, T. Medina, R. Molavi, R. Neufeld, T. Oh, I. Pavlov, I. Perminov, G. Poulin-Lamarre, C. Rich, A. Smirnov, L. Swenson, N. Tsai, M. Volkmann, J. Whittaker, and J. Yao, "Phase transitions in a programmable quantum spin glass simulator", Science **361**, 162–165 (2018).

[93] J. Choi, S. Choi, G. Kucsko, P. C. Maurer, B. J. Shields, H. Sumiya, S. Onoda, J. Isoya, E. Demler, F. Jelezko, N. Y. Yao, and M. D. Lukin, "Depolarization Dynamics in a Strongly Interacting Solid-State Spin Ensemble", Physical Review Letters **118**, 093601 (2017).

[94] G. Kucsko, S. Choi, J. Choi, P. C. Maurer, H. Zhou, R. Landig, H. Sumiya, S. Onoda, J. Isoya, F. Jelezko, E. Demler, N. Y. Yao, and M. D. Lukin, "Critical Thermalization of a Disordered Dipolar Spin System in Diamond", Physical Review Letters **121**, 023601 (2018).

[95] D. Rossini and R. Fazio, "Phase diagram of the extended Bose–Hubbard model", en, New Journal of Physics **14**, 065012 (2012).

[96] T. Ohgoe, T. Suzuki, and N. Kawashima, "Ground-state phase diagram of the two-dimensional extended Bose-Hubbard model", Physical Review B **86**, 054520 (2012).

[97] C. Kollath, G. Roux, G. Biroli, and A. M. Läuchli, "Statistical properties of the spectrum of the extended Bose–Hubbard model", en, Journal of Statistical Mechanics: Theory and Experiment **2010**, P08011 (2010).

[98] V. W. Scarola and S. Das Sarma, "Quantum Phases of the Extended Bose-Hubbard Hamiltonian: Possibility of a Supersolid State of Cold Atoms in Optical Lattices", Physical Review Letters **95**, 033003 (2005).

[99] S. Baier, M. J. Mark, D. Petter, K. Aikawa, L. Chomaz, Z. Cai, M. Baranov, P. Zoller, and F. Ferlaino, "Extended Bose-Hubbard models with ultracold magnetic atoms", Science **352**, 201–205 (2016).

[100] C. Lagoin, U. Bhattacharya, T. Grass, R. W. Chhajlany, T. Salamon, K. Baldwin, L. Pfeiffer, M. Lewenstein, M. Holzmann, and F. Dubin, "Extended Bose–Hubbard model with dipolar excitons", en, Nature **609**, 485–489 (2022).

Appendix A

"Two-Gauss" Fitting Function

This is the fitting function used to find the hole location in the gas. The function uses the results of an earlier fit, `meanFirstFits`, to define guesses and constraints for this fitter.

```
function fitterResults = runTwoGaussFreeAngFit(inputImage, meanFirstFits)

fit1Offset = meanFirstFits(1);
fit1pkOD = meanFirstFits(2);
fit1xcenter = meanFirstFits(3);
fit1xsigma = meanFirstFits(4);
fit1ycenter = meanFirstFits(5);
fit1ysigma = meanFirstFits(6);
fit1bkgdtiltx = meanFirstFits(7);
fit1bkgdtilty = meanFirstFits(8);
fit1avesigma = (fit1xsigma+fit1ysigma)/2;

pkOD1 = fit1pkOD;
pkOD2 = -fit1pkOD;
theta1 = 0;
theta2 = 50;

pkODBoundFrac = 3;
sigmaBoundFrac = 0.5;
pk2ODBoundFrac = 0.9;
center2BoundOffsetX = fit1xsigma;
center2BoundOffsetY = fit1ysigma;
sigma2ScaleFact = 0.2;
sigma2 = fit1avesigma * sigma2ScaleFact;
difHoleSigmaFactor = 0.2;
sigma2LBoundFrac = 0.1;
sigma2UBoundFrac = 3;
```

```
fitname = 'twogauss_freeang';

% guess                          lowerbound                       upperbound
fitInput=...
[fit1Offset,                     fit1Offset,                      fit1Offset;...
 pkOD1,                          pkOD1*(1-pkODBoundFrac),         pkOD1*(1+pkODBoundFrac);...
 fit1xcenter,                    fit1xcenter,                     fit1xcenter;...
 fit1xsigma,                     fit1xsigma*(1-sigmaBoundFrac),   fit1xsigma*(1+sigmaBoundFrac);...
 fit1ycenter,                    fit1ycenter,                     fit1ycenter;...
 fit1ysigma,                     fit1ysigma*(1-sigmaBoundFrac),   fit1ysigma*(1+sigmaBoundFrac);...
 theta1*pi/180,                  -pi/4,                           pi/4;...
 0,                              0,                               0;...
 0,                              0,                               0;...
 pkOD2,                          pkOD2*(1+pk2ODBoundFrac),        pkOD2*(1-pk2ODBoundFrac);...
 fit1xcenter,                    fit1xcenter-center2BoundOffsetX,fit1xcenter+center2BoundOffsetX;...
 sigma2*(1+difHoleSigmaFactor), sigma2*sigma2LBoundFrac ,        sigma2*sigma2UBoundFrac;...
 fit1ycenter,                    fit1ycenter-center2BoundOffsetY,fit1ycenter+center2BoundOffsetY;...
 sigma2*(1-difHoleSigmaFactor), sigma2*sigma2LBoundFrac ,        sigma2*sigma2UBoundFrac;...
 theta2*pi/180,                  0-pi/10,                         pi/2+pi/10];

guess  = fitInput(:,1);
lbound = fitInput(:,2);
ubound = fitInput(:,3);

[fits, sresults] = newfitter(inputImage,fitname, ...
                             'guess', guess, ...
                             'lbound', lbound, ...
                             'ubound', ubound);

fitterResultsFullAngle = [fits.rawresult];
if(fitterResultsFullAngle(15) > pi/2 - pi/10)
    newAngle = fitterResultsFullAngle(15) - pi/2;
    originalSigmaA = fitterResultsFullAngle(12);
    originalSigmaB = fitterResultsFullAngle(14);
    fitterResultsFullAngle(15) = newAngle;
    fitterResultsFullAngle(12) = originalSigmaB;
    fitterResultsFullAngle(14) = originalSigmaA;
end

fitterResults = fitterResultsFullAngle;
```

Appendix B

Rotated Partial-Pixel Mask

This is `createRotatedMaskPartial.m` from the image processing code. It describes the exact function we used for creating an elliptical mask with partial-pixel masking.

```
function rotEllipseMask = createRotatedMaskPartial(x, y, centerx, centery, a, b, theta)

% This fuction makes a matrix of the fractional area of a pixel
% contained within the specified ellipse.
% x and y are the matrices of their respective coordinates

% This sign change is for consistancy in the sign convention with other parts of the code
theta = -theta;

xminus = x - 0.5;
xplus = x + 0.5;
yminus = y - 0.5;
yplus = y + 0.5;

% a and b are the radii
function ellipseN = ellipseNums(x,y)
    ellipseN =((x-centerx)*cos(theta)+(y-centery)*sin(theta)).^2/a^2 + ...
            (-(x-centerx)*sin(theta)+(y-centery)*cos(theta)).^2/b^2;
end

function solutions = xellipseAt1(centerx, centery, theta, a, b, y)
    xellipsePlus  = (b^2*centerx*cos(theta)^2 + a^2*centerx*sin(theta)^2 + ...
            a^2*b^2*(-(centery^2*cos(theta)^4 - b^2*cos(theta)^2 - a^2*sin(theta)^2 + ...
            centery^2*sin(theta)^4 + y^2*cos(theta)^4 + y^2*sin(theta)^4 - ...
            2*centery*y*cos(theta)^4 - 2*centery*y*sin(theta)^4 + ...
            2*centery^2*cos(theta)^2*sin(theta)^2 + 2*y^2*cos(theta)^2*sin(theta)^2 - ...
            4*centery*y*cos(theta)^2*sin(theta)^2)/(a^2*b^2))^(1/2) - ...
            a^2*centery*cos(theta)*sin(theta) + b^2*centery*cos(theta)*sin(theta) + ...
```

```
            a^2*y*cos(theta)*sin(theta) - ...
            b^2*y*cos(theta)*sin(theta))/(sin(theta)^2*a^2 + ...
            cos(theta)^2*b^2);
    xellipseMinus = (b^2*centerx*cos(theta)^2 + a^2*centerx*sin(theta)^2 - ...
            a^2*b^2*(-(centery^2*cos(theta)^4 - b^2*cos(theta)^2 - a^2*sin(theta)^2 + ...
            centery^2*sin(theta)^4 + y^2*cos(theta)^4 + y^2*sin(theta)^4 - ...
            2*centery*y*cos(theta)^4 - 2*centery*y*sin(theta)^4 + ...
            2*centery^2*cos(theta)^2*sin(theta)^2 + 2*y^2*cos(theta)^2*sin(theta)^2 - ...
            4*centery*y*cos(theta)^2*sin(theta)^2)/(a^2*b^2))^(1/2) - ...
            a^2*centery*cos(theta)*sin(theta) + b^2*centery*cos(theta)*sin(theta) + ...
            a^2*y*cos(theta)*sin(theta) - ...
            b^2*y*cos(theta)*sin(theta))/(sin(theta)^2*a^2 + ...
            cos(theta)^2*b^2);
    solutions = [xellipsePlus, xellipseMinus];
end

function solutions = yellipseAt1(centerx, centery, theta, a, b, x)
    yellipsePlus = (a^2*centery*cos(theta)^2 + b^2*centery*sin(theta)^2 + ...
            a^2*b^2*(-(centerx^2*cos(theta)^4 - a^2*cos(theta)^2 - b^2*sin(theta)^2 + ...
            centerx^2*sin(theta)^4 + x^2*cos(theta)^4 + x^2*sin(theta)^4 - ...
            2*centerx*x*cos(theta)^4 - 2*centerx*x*sin(theta)^4 + ...
            2*centerx^2*cos(theta)^2*sin(theta)^2 + 2*x^2*cos(theta)^2*sin(theta)^2 - ...
            4*centerx*x*cos(theta)^2*sin(theta)^2)/(a^2*b^2))^(1/2) - ...
            a^2*centerx*cos(theta)*sin(theta) + b^2*centerx*cos(theta)*sin(theta) + ...
            a^2*x*cos(theta)*sin(theta) - ...
            b^2*x*cos(theta)*sin(theta))/(cos(theta)^2*a^2 + ...
            sin(theta)^2*b^2);
    yellipseMinus = (a^2*centery*cos(theta)^2 + b^2*centery*sin(theta)^2 - ...
            a^2*b^2*(-(centerx^2*cos(theta)^4 - a^2*cos(theta)^2 - b^2*sin(theta)^2 + ...
            centerx^2*sin(theta)^4 + x^2*cos(theta)^4 + x^2*sin(theta)^4 - ...
            2*centerx*x*cos(theta)^4 - 2*centerx*x*sin(theta)^4 + ...
            2*centerx^2*cos(theta)^2*sin(theta)^2 + 2*x^2*cos(theta)^2*sin(theta)^2 - ...
            4*centerx*x*cos(theta)^2*sin(theta)^2)/(a^2*b^2))^(1/2) - ...
            a^2*centerx*cos(theta)*sin(theta) + b^2*centerx*cos(theta)*sin(theta) + ...
            a^2*x*cos(theta)*sin(theta) - ...
            b^2*x*cos(theta)*sin(theta))/(cos(theta)^2*a^2 + sin(theta)^2*b^2);
    solutions = [yellipsePlus, yellipseMinus];
end

ellipseNumsxpyp = ellipseNums(xplus, yplus);
ellipseNumsxpym = ellipseNums(xplus, yminus);
ellipseNumsxmym = ellipseNums(xminus, yminus);
ellipseNumsxmyp = ellipseNums(xminus, yplus);
```

```
rotEllipseMaskxpyp = ellipseNumsxpyp < 1;
rotEllipseMaskxpym = ellipseNumsxpym < 1;
rotEllipseMaskxmyp = ellipseNumsxmyp < 1;
rotEllipseMaskxmym = ellipseNumsxmym < 1;

rotEllipseMask = ...
    (rotEllipseMaskxpyp + rotEllipseMaskxpym + rotEllipseMaskxmyp + rotEllipseMaskxmym)/4;

[fracPixRows,fracPixColumns] = find(rotEllipseMask>0 & rotEllipseMask < 1);

for jj = 1: length(fracPixRows)
    rowInd = fracPixRows(jj);
    colInd = fracPixColumns(jj);
    xminus0 = xminus(rowInd, colInd);
    xplus0 = xplus(rowInd, colInd);
    yminus0 = yminus(rowInd, colInd);
    yplus0 = yplus(rowInd, colInd);
    pixCornersList = [...
        rotEllipseMaskxpyp(rowInd, colInd) rotEllipseMaskxpym(rowInd, colInd) ...
        rotEllipseMaskxmym(rowInd, colInd) rotEllipseMaskxmyp(rowInd, colInd) ...
        rotEllipseMaskxpyp(rowInd, colInd) ...
    ];
    cornerCoords = [
        xplus0 yplus0; ...
        xplus0 yminus0; ...
        xminus0 yminus0; ...
        xminus0 yplus0 ...
    ];
    difAroundPix = diff(pixCornersList);
    xpIntersections = NaN;
    ymIntersections = NaN;
    xmIntersections = NaN;
    ypIntersections = NaN;
    pixFrac = NaN;
    for inx=1:length(difAroundPix)
        if difAroundPix(inx) ~= 0
           switch inx
               case 1
                   % xplus yplus yminus
                   solutions = yellipseAt1(centerx, centery, theta, a, b, xplus0);
                   xpIntersections = solutions(yminus0 < solutions & solutions < yplus0);
                   if length(xpIntersections) > 1
```

```
                pixFrac = 1;
            end

        case 2
            % yminus xplus xminus
            solutions = xellipseAt1(centerx, centery, theta, a, b, yminus0);
            ymIntersections =solutions(xminus0 < solutions & solutions < xplus0);
            if length(ymIntersections) > 1
                pixFrac = 1;
            end

        case 3
            % xmiuns yminus yplus
            solutions = yellipseAt1(centerx, centery, theta, a, b, xminus0);
            xmIntersections =solutions(yminus0 < solutions & solutions < yplus0);
            if length(xmIntersections) > 1
                pixFrac = 1;
            end

        case 4
            % yplus xminus xplus
            solutions = xellipseAt1(centerx, centery, theta, a, b, yplus0);
            ypIntersections =solutions(xminus0 < solutions & solutions < xplus0);
            if length(ypIntersections) > 1
                pixFrac = 1;
            end
    end
  end
end

intersections = [xpIntersections, ymIntersections, xmIntersections, ypIntersections];
if isnan(pixFrac)
    switch sum(pixCornersList(1:4))
        case 1
            inCornerCoords = cornerCoords(pixCornersList(1:4),:);
            if isnan(xpIntersections)
                ydist = abs(xmIntersections-inCornerCoords(2));
            else
                ydist = abs(xpIntersections-inCornerCoords(2));
            end
            if isnan(ypIntersections)
                xdist = abs(ymIntersections-inCornerCoords(1));
            else
```

```
                xdist = abs(ypIntersections-inCornerCoords(1));
            end
            pixFrac = (ydist*xdist)/2;
        case 2
            if isnan(xpIntersections)
                aveIntersections = (ymIntersections + ypIntersections)/2;
                if pixCornersList(1)
                    %xp is in the ellipse
                    pixFrac = abs(xplus0 - aveIntersections);
                else
                    %xm is in the ellipse
                    pixFrac = abs(xminus0 - aveIntersections);
                end
            else
                aveIntersections = (xmIntersections + xpIntersections)/2;
                if pixCornersList(1)
                    %yp is in the ellipse
                    pixFrac = abs(yplus0 - aveIntersections);
                else
                    %ym is in the ellipse
                    pixFrac = abs(yminus0 - aveIntersections);
                end
            end
        case 3
            outCornerCoords = cornerCoords(~pixCornersList(1:4),:);
            if isnan(xpIntersections)
                ydist = abs(xmIntersections-outCornerCoords(2));
            else
                ydist = abs(xpIntersections-outCornerCoords(2));
            end
            if isnan(ypIntersections)
                xdist = abs(ymIntersections-outCornerCoords(1));
            else
                xdist = abs(ypIntersections-outCornerCoords(1));
            end
            pixFrac = 1 - (ydist*xdist)/2;
    end
end
rotEllipseMask(rowInd, colInd) = pixFrac;
end

end
```