

**Atom-by-atom control and readout for studying
spin-motional dynamics and entanglement
in neutral atom arrays**

by

Brian James Lester

B.S., California Institute of Technology, 2010

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2016

This thesis entitled:
Atom-by-atom control and readout for studying spin-motional dynamics and entanglement
in neutral atom arrays
written by Brian James Lester
has been approved for the Department of Physics

Prof. Cindy A. Regal

Prof. Ana Maria Rey

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Lester, Brian James (Ph.D., Physics)

Atom-by-atom control and readout for studying spin-motional dynamics and entanglement
in neutral atom arrays

Thesis directed by Prof. Cindy A. Regal

In this thesis, I describe early experiments with a new platform that harnesses control over the full quantum state of individual ^{87}Rb atoms to study out-of-equilibrium states of a few atoms placed in tailored optical potentials. We employ an enhanced loading technique that fills each well in 90% of loading attempts, image the configuration of the atoms, and then perform Raman sideband cooling that results in a 90% three-dimensional ground state fraction. Then, after initializing the spin of each atom, we can reconfigure the traps to initialize dynamics in a final optical potential of interest. For example, we can form a double well potential and observe the quantum interference of two atoms tunneling between the wells. Additionally, we have demonstrated the ability to coherently transfer atoms between wells and, by preparing two atoms in opposite spin states, have observed spin-exchange oscillations that periodically entangle the two atoms. I will also discuss plans and ongoing work to combine these capabilities with new techniques to gain more information from systems containing more atoms. In such systems, we wish to study how the spin-motional coupling of independently prepared atoms, in the presence of interactions, will lead to complex dynamics, such as in the Kondo lattice model.

Dedication

This work is dedicated to my parents.

Acknowledgements

Most of all, I want to say a big thank you to Cindy for all of the support and guidance that she has provided over the past six years. I have really enjoyed my time working with and learning from her, and I am very grateful for the freedom that she allowed during the design, testing, and construction of the apparatus.

Of course, the work presented here would not have been possible without the help of everyone I have had the pleasure of working in the lab with! First and foremost, the knowledge and work ethic that Adam brought to the lab was incredible. His ability to bring into focus the most important aspects of the system and get things working at their best was invaluable, but most of all it was a pleasure to work with him for my first five years (and, of course to go grab a Duchess and some wings after a good day in lab). I am also happy to know that the experiment is being transferred to capable hands and I have really enjoyed working with (and, at times, relearning parts of the experiment with) Yiheng, Mark, and Randall. I am grateful for all of the discussions and questions you have brought, and look forward to continuing them! And, although it was shorter, I also enjoyed the time I got to spend working Collin, Junling, Niclas, Kathrin, Ludovic, and Mika over the years.

I also owe a huge thanks to Ana Maria Rey and all the members of her group for their help to fully model and understand both current and future experiments discussed in this thesis. In addition to Ana Maria, I want to specifically thank Michael Wall and Leonid Isaev for the numerous discussions we have had over the past couple of years.

I also want to thank everyone who I shared coffee breaks, beer tastings, and physics discussions with throughout the years! In particular, I want to mention Nir Kampel, who was always up for

discussing potential problems with electronics or general physics, or just going to the Goat to chat and enjoy the company. And, aside from beverages, thanks also to the rest of the “other’ side” of the lab - Bob, Max, Thom, Tom, Ran, and Ben for fostering an great atmosphere in which I was able to live and learn during my time.

And to everyone else I got to share my time with in Boulder, helping me to get my mind out of lab when it was necessary. I don’t think it would be possible to name everyone here, but I will always appreciate the brewery touring, food exploring, and beer brewing adventures - and that is not to mention all of the ski trips - throughout the years with Andrew, Alejandra, Craig, Monica, Daniel, Steven, and Rabin. Thank you for everything!

Contents

Chapter

1	Introduction	1
1.1	Quantum gas assembly	1
1.1.1	Background on control of quantum gases	1
1.1.2	Using single-atom control to build complex many-body states	2
1.2	Work performed as part of this thesis	6
1.3	List of publications	8
2	Experimental Apparatus for Control of Single Atoms	10
2.1	Trapping and Control of Cold Atomic Gases	10
2.1.1	Lasers for cooling and trapping ^{87}Rb	11
2.1.2	Loading directly from a MOT	13
2.1.3	Optical pumping and spin-polarization of ^{87}Rb	14
2.1.4	Single-atom spin addressing	15
2.2	Optical System for Single-Atom Trapping	17
2.2.1	High-numerical aperture lenses	19
2.2.2	Sources of light for optical trapping	20
2.2.3	Optical Rail	22
3	Quantum Gas Assembly: A Platform to Study Tailored Many-Body Systems	25

3.1	Enhanced Loading of Single Atoms	27
3.1.1	Background on light-assisted collisions	27
3.1.2	Experimental implementation of enhanced loading	29
3.1.3	Collisional beam parameters	33
3.1.4	Limitations to the enhanced loading	35
3.2	Manipulation of Optical Tweezers	40
3.2.1	Overall intensity stabilization	40
3.2.2	Acousto-optic control of optical traps	42
3.2.3	Generation of radio-frequency drive	46
3.2.4	Alternative options for arbitrary trapping potentials	46
3.3	Removing entropy from the system	48
3.3.1	Building uniform atom arrays	49
3.3.2	Distilling motional state preparation	50
3.4	Improving readout of the final state of the system	50
3.4.1	Spin-resolved imaging	51
3.4.2	Raman-cooling imaging	52
4	Preparing and Interfering Indistinguishable Atoms	60
4.1	Raman sideband cooling	60
4.1.1	Non-degenerate pulsed Raman sideband cooling	61
4.1.2	Efficient cooling in optical tweezers	64
4.1.3	Raman sideband spectroscopy	66
4.2	Quantum interference of indistinguishable particles	67
4.2.1	Measuring indistinguishability via tunneling in a double well potential	67
4.2.2	Using indistinguishability to improve state preparation	70
5	Entangling Independently-Prepared Neutral Atoms	72
5.1	Dynamical entanglement via spin-exchange	72

5.2	Spin-entangled states in spin-dependent traps	78
5.2.1	Adiabatic preparation of entangled states	80
6	Assembling the Kondo Lattice Model with Optical Tweezers	84
6.1	The Kondo Lattice Model	85
6.1.1	Limiting cases of the Kondo lattice model	87
6.2	The Kondo-Hubbard Model for Atomic Systems	88
6.2.1	Studying the Kondo-Hubbard model with bosons	89
6.3	Implementation with optical tweezers	91
6.3.1	Calculating the expected spin dynamics	92
6.3.2	Weakly interacting regime ($V_{eg} \ll J_e$)	94
6.3.3	Interaction dominating regime ($V_{eg} \gg J_e$)	98
6.3.4	Summary of ideal theoretical investigations	102
	Bibliography	107
	Appendix	
A	Alignment and Characterization of Optical System	116
A.1	External lens testing and alignment procedure	118
A.1.1	Measurement of the point spread function (PSF) via pinhole imaging	118
A.1.2	Wavefront measurement testing	120
A.2	Characterization of the trapping potential	120
B	Derivation of eigenstates in minimal Kondo model	123

Figures

Figure

1.1	Operational procedure for quantum gas assembly	3
2.1	Level diagram for ^{87}Rb and laser addressing schematic	12
2.2	Measurement of the polarization purity of the optical pumping light	15
2.3	Microwave manipulations of the spin state of a single atom	16
2.4	Microwave spectrum taken with light shift beam used to address a single well	17
2.5	Schematic of the optical setup for controlling single atoms	18
2.6	Diagram of optical rail for beam preparation	21
2.7	Schematic of Optical Relay	23
3.1	Operational procedure for quantum gas assembly	26
3.2	Schematic for light-assisted collisions	28
3.3	Enhanced loading summary of results	30
3.4	Loading efficiency versus length of loading sequence	31
3.5	Loading efficiency versus relevant enhanced loading parameters	32
3.6	Depump beam power “resonance”	35
3.7	Effects of trap proximity on loading efficiency	39
3.8	Calibration of the optical tweezer position as a function of drive frequency	44
3.9	Distillation of imperfect motional state preparation	51
3.10	Histogram of photon counts during both PGC imaging and Raman cooling imaging	55

3.11	Imaging during Raman cooling, varying drive strength	58
4.1	Schematic of experimental setup used to perform Raman cooling	62
4.2	Diagram of the Raman cooling procedure	63
4.3	Raman sideband spectra using gaussian pulsed spectroscopy	67
4.4	Schematic of tunneling procedure	68
4.5	Observing two-particle interference while tuning distinguishability	69
4.6	Improving state preparation by optimizing indistinguishability	71
5.1	Generation of spin-entanglement via local spin-exchange	73
5.2	Experimental procedure for the observation of spin exchange and initial data	75
5.3	Verification of spin entanglement after separation of atoms	77
5.4	Spin exchange in traps that have less spin-dependence	80
5.5	Effect of spin-dependent trapping potentials on spin-exchange contrast	81
5.6	Eigenstate energy levels in the presence of differential spin-dependent shifts	82
6.1	Schematic of the Kondo-Hubbard lattice model	86
6.2	Schematic of the Kondo-Hubbard lattice model	90
6.3	Implementation of the microscopic Kondo lattice model with optical tweezers	93
6.4	Single spin basis dynamics in the weakly interacting regime	95
6.5	Joint spin-state basis dynamics in the weakly interacting regime	97
6.6	Alternative measurements of the spin dynamics in weakly interacting regime	99
6.7	Motion of atom in excited band with strong interactions	100
6.8	Spin state dynamics with strong interactions	101
6.9	Dynamics of the Kondo lattice model for realistic trap parameters	104
A.1	Schematic of testing setup to measure the point-spread-function	118
A.2	Schematic of the setup to directly measure wavefront aberrations	121
A.3	Optimizing trap frequencies via alignment of objective lens	122

A.4 Measurement of trap depth via light shift of cycling transition 122

Chapter 1

Introduction

1.1 Quantum gas assembly

1.1.1 Background on control of quantum gases

Modern atomic physics experiments have alternately been gaining impressive experimental control in order to understand properties of complex systems, then using this understanding to create additional experimental tools for a new generation of experiments. The study of quantum gases builds on earlier work developing the techniques to cool and trap clouds of atoms at temperatures well below 1 mK, eventually leading to the observation of Bose-Einstein condensates and the production of a degenerate Fermi gas [1, 2, 3, 4]. Shortly thereafter, these clouds of atoms started to be used to study the complex interactions of ground-state atoms and molecules and, eventually, to the development of techniques to control the strength of interactions, such as the s -wave scattering length via Fano-Feshbach resonances and dipole-dipole interactions via Rydberg excitation [5, 6, 7, 8, 9, 10, 11].

The ability to control interactions and study the effects on bulk properties of clouds of many atoms has been extremely successful and now it has become routine to use this control in the presence of an external potential to study different Hamiltonians with the density and interaction parameters as tunable parameters [12, 13, 14, 15, 16, 17]. Even more recently, there has been a push toward observing these systems with resolution of individual atoms (or sites of a lattice), leading to the development of quantum gas microscopy, where ultracold atoms in an optical lattice

can be observed over a large range of interaction parameters, allowing the direct measurement of correlations in the lattice [18, 19, 20, 21, 22, 23].

In parallel, there has been a significant amount of progress in the trapping, manipulation, and interaction of well-separated atoms, with numerous demonstrations of single and two-qubit control both with trapped atomic ions and neutral atoms via Rydberg blockade [24, 25, 26, 27]. And more recently, these experiments have been building up to larger systems of strongly interacting atoms, while maintaining the single-atom control developed for use in quantum information processing [28, 29, 30, 31]. In particular, these systems have the benefit of being able to initialize each atom independently and, because of the strong interactions, to keep the atoms well-separated throughout the experiment and read out each atom at the end of the experimental sequence. These systems also have the flexibility to generate nearly arbitrary trapping geometries that can be tailored for a variety of experiments and enable the transport of atoms during the course of a single experiment [32, 33, 34, 35].

1.1.2 Using single-atom control to build complex many-body states

The work described in this thesis is a part of the development of a platform that combines the advantages of experiments that maintain individual-atom control in an array of atoms, but which also can be used to build a degenerate quantum gas from the ground-up in a custom optical potential. In particular, we wish to initially set the full quantum state of each atom (including both the spin and motional degrees of freedom) and allow the atoms to delocalize and interact in a projected optical potential. This builds on the capabilities afforded in quantum gas microscope experiments, but turns the process around, such that we could, in principle, build a small degenerate Bose (or Fermi) gas from individually prepared atoms. Schematic representation of this style of experiment, highlighting many of the capabilities it would afford and have been developed in my thesis work, is shown in Figure 1.1.

To start with, atoms are initially loaded into an array of traps using an enhanced loading technique that results in a 90% probability of loading an atom in each trap [36, 37, 38]. After

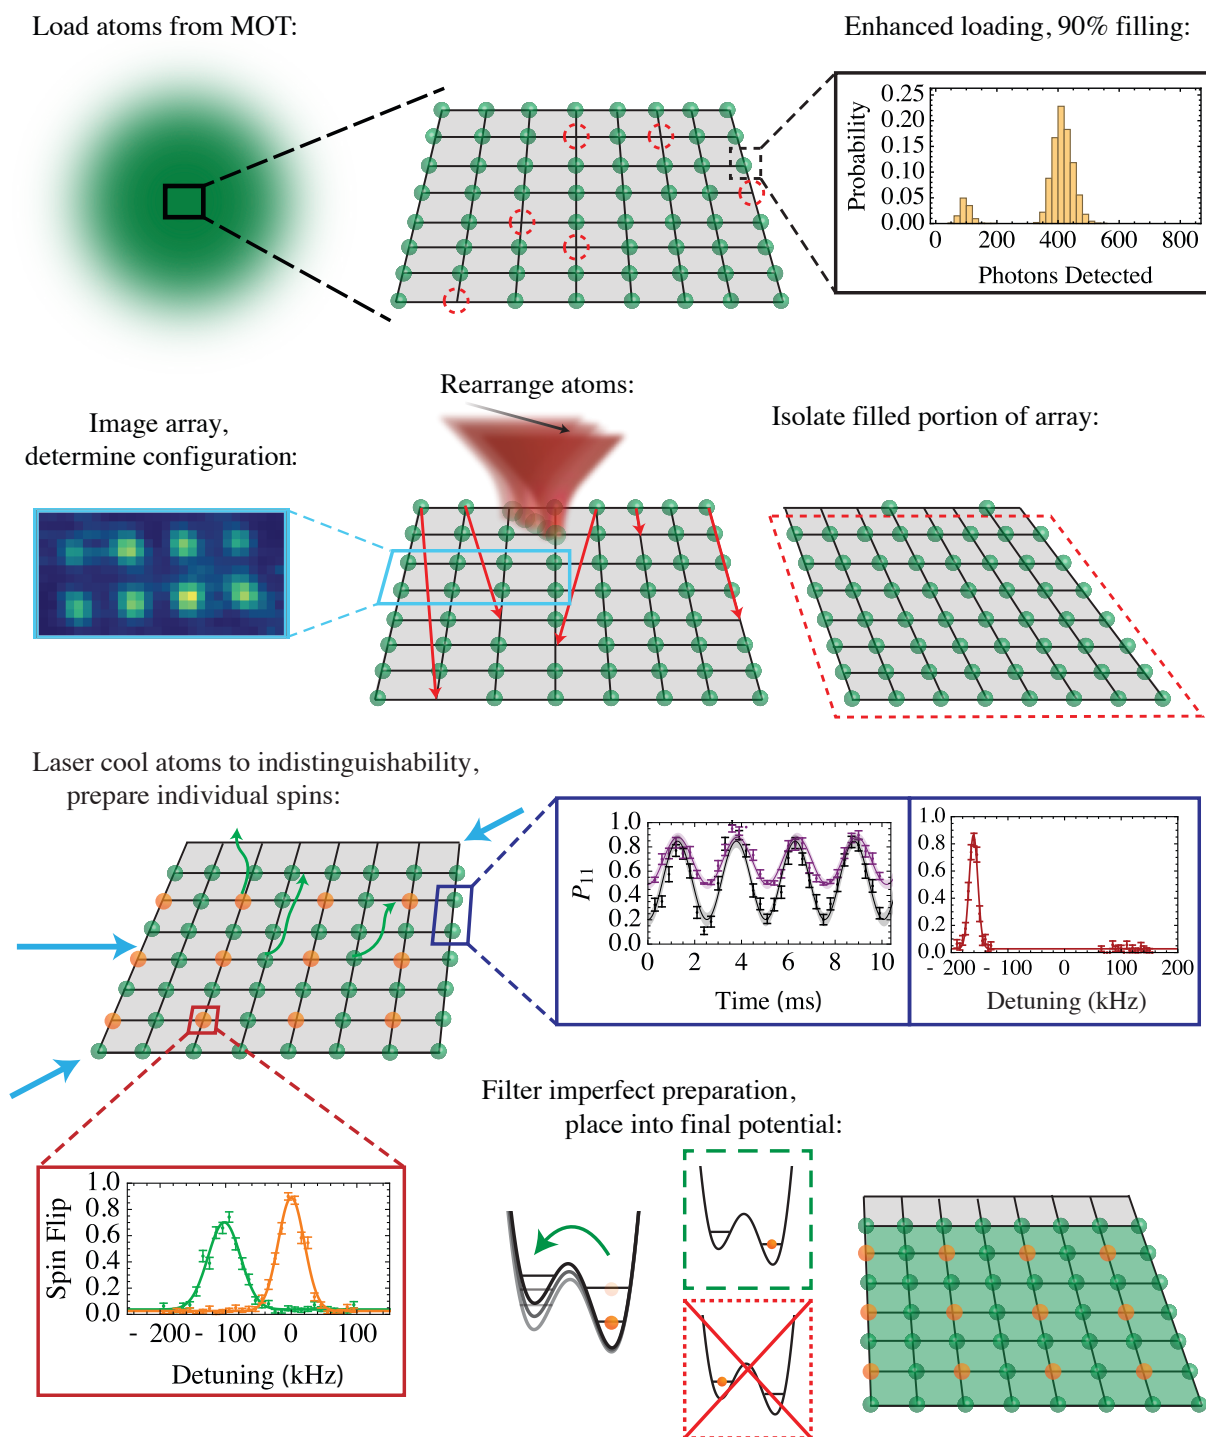


Figure 1.1: A vision of the initialization of experiments via quantum gas assembly, depicting several of the important capabilities that the optical tweezer platform and the work in this thesis will afford.

loading, the array is imaged to measure the occupation of each of the trap, gaining knowledge of how many atoms are available in that run of the experiment, as well as their location in the array of traps. This knowledge can either be used during processing of the data set, when it is useful to performing the same experimental procedure in multiple initial atom configurations, or in real-time to perform different experiments based on the initial state [39, 31]. For the information to be used in real-time, it is likely that rearranging the atoms, either to fill holes in an array, or even loading into a completely separate, customized potential, will allow for faster data taking in configurations that are of most relevance [40, 41, 42]. After constructing a desired pattern of atoms for initializing a particular experiment, the entire array is spin-polarized and Raman sideband cooling is performed, bringing each of the atoms to the three-dimensional ground state of the trapping potential [43, 44]. With all of the atoms in the motional ground state, the spin-state of each atom can then be initialized by taking advantage of the ability to address individual sites with a beam to shift the spin flip resonance [45, 39, 46].

After the entire state-preparation protocol has been performed, the atoms are ready to be placed into the final potential in which the dynamics of interest will be observed. However, this overlooks the possibility to manipulate the system during this procedure in such a way that imperfections in the state preparation can be detected. For example, the process shown in Figure 1.1 will “distill” the motional state preparation by performing an adiabatic passage (discussed more in Section 5.1) in such a way that atoms in an excited motional state will transfer to a second well. This extra well is then separated from the remainder of the system and not included in the final potential; then, by not detecting an atom in the secondary well, it can be inferred that the atom was in motional ground state. The traps containing these fully initialized atoms are then dynamically repositioned to place the atoms into a final potential where the many-body quantum state can evolve. After a period of dynamics, the traps are again rearranged to separate the atoms for readout. This procedure does not need to be identical to the procedure that assembled the state, but the goal is to separate the atoms and to measure both the spin and population dynamics over time.

In our experiment, we have already demonstrated the ability to control all of the internal and external degrees of freedom of a single atom and are actively studying the dynamics of systems of few interacting particles from a pure initial state (in the single-particle basis). In particular, we have shown that we can create indistinguishable atoms via laser cooling, and observe two-particle quantum interference in the dynamics of two atoms tunneling in a double well potential [43, 39]. We also have demonstrated the ability to perform enhanced loading, significantly increasing the probability of uniformly filling arrays of traps. Additionally, the experiments we have performed with two atoms already demonstrate our ability to dynamically reconfigure atoms and transfer between trapping potentials while maintaining the purity of the motional state [38]. We have also used our ability to selectively flip individual spins and coherently transfer atoms between the wells to observe dynamical spin-exchange and verified that the entanglement generated between the two spins is maintained upon separation of the atoms [47]. The ability to image a partially filled array and use optical tweezers to generate a smaller filled array has very recently been demonstrated in one dimension in Ref. [41] and for two dimensions in Ref. [42], and is a technique that we will be implementing as we scale to larger arrays of atoms.

We are currently studying what happens when we combine these capabilities to observe spin-motional coupling in systems of a few interacting atoms. For example, with just three atoms in two bands of a double well potential, we will be able to construct minimum instances of the bosonic Kondo-Hubbard model and observe differences in dynamics as we tune the relative strength of the on-site interactions to the tunneling-rate between the wells. Specifically, in the strongly-interacting regime, we expect to be able to see that, for certain initial states, the tunneling rate will be reduced by a factor of two; this is a direct analog to the observation of mass-enhancement in heavy-fermion materials. On the other hand, for weak interactions, we expect to be able to measure the effect of a nonlocal spin-spin coupling between the separated ground-state atoms mediated by local exchange interactions.

1.2 Work performed as part of this thesis

During my time in Cindy's lab, I have seen this experiment grow from a single table with a recently baked vacuum cell and the laser systems for producing a magneto-optical trap (MOT), to a fully functional optical tweezer system. For my first five years here, I had the pleasure of working closely with Adam Kaufman to help get this system off of the ground and through that experience learned a lot of atomic physics. When I arrived, he was getting the MOT set up and very shortly thereafter began characterizing the details of the setup and implementing the basic systems required for manipulating the internal states of clouds of atoms.

In parallel, my focus was on designing, testing, and setting up significant portions of the high numerical aperture (NA) optical system that both generates our optical tweezer potentials and collects fluorescence scattered from the atom to form an image on our camera. The most important requirement for this system is that the optical traps needed to provide sufficient confinement to enable the ground-state cooling protocol developed in our lab, which means maintaining near-diffraction limited performance, even with the significant challenges presented by mounting the lens outside of the science chamber. We started out this process with a first-generation lens and imaging system that achieved a NA of 0.25, that allowed us to gain experience trapping single atoms and manipulating their state in tight optical traps, not to mention the debugging and characterization of the sources of aberrations in the experimental setup, which provided useful information before moving up to a higher NA system. Once this had been set up and was trapping atoms, I got my hands on our (first) custom made objective lens and was able to start the process of both characterizing and optimizing its performance in an external testing setup. Eventually, it was this lens that was able to provide us with the confinement sufficient for three-dimensional Raman sideband cooling, followed by coherent tunneling between two optical tweezer potentials and quantum interference between two atoms. We have more recently used the system to observe dynamical spin-exchange between the atoms and verify that entanglement is maintained after separation of the atoms.

After helping with the data taking and characterization of the trapping potentials for the

two-particle quantum interference experiment, I was able to take the lead on characterizing our imaging system more thoroughly and demonstrating the ability to image single atoms via light scattered during continuous Raman sideband cooling. While this technique has proven to be less useful for our system (because sub-doppler cooling is so efficient for ^{87}Rb), it is a technique that is being used in several quantum gas microscopes for lighter fermionic atoms, and provided us with a lot of information about the performance of our imaging system [48, 49, 50, 51].

I also worked with a visiting student Niclas Luick to implement an enhanced loading technique that had not previously been used with traps that are as tightly confining as in our system, as well as to demonstrate its utility in loading small arrays of traps. Importantly, we were able to optimize the technique such that we loaded a single atom in $\sim 90\%$ of attempts, while the length of each loading cycle was not significantly increased compared to the standard techniques that achieve closer to 60% loading [38]. This technique is going to play an integral role in improving the data collection efficiency of future experiments where the system size is larger than a few atoms, and is likely to be implemented in similar systems that load an array of traps directly from a thermal cloud of atoms.

Aside from the published work listed in Section 1.3, I have worked to address some of the limitations present in our initial experiments. Specifically, improving the characterization and control of the light used for our Raman cooling procedure, switching to using temporally incoherent light for generating optical tweezers, increasing the dynamic range of our trap intensity control to access trap depths that vary over 4 orders of magnitude, and implementing the use of completely arbitrary waveform generator to allow for arbitrary generation of square arrays of optical tweezers (within the limitations of the optical system). We have additionally tested numerous methods for correcting aberrations of our objective lens *in-situ*, which should lead to better tunneling performance, but none have been to satisfaction. Therefore, we are planning to switch to a new objective lens whose aberrations have been measured and corrected externally.

In the meantime, I have also developed an understanding of the physics and experimental protocols for several upcoming experiments that take advantage of the control that our experiment

has either already demonstrated or is expected to have in the near future. For example, studies of the microscopic implementation of the Kondo experiment, which will be discussed in much more detail in Chapter 6, is an example of an experiment that will take advantage of the improved control of our optical tweezer potentials, as well as the ability to initialize both the spin and motional states of atoms and then coherently transfer the atoms between. I will also discuss our plans to use control over the state-dependence of a double well potential to adiabatically prepare pure singlet or triplet spin states, which could then be separated to use as a source of maximally entangled atoms that is less sensitive to experimental fluctuations than the dynamical entanglement demonstrated in Ref. [47].

Most of the other unpublished ideas are related to gaining more information from our system than in our current experiments. Specifically, we have recently implemented the ability to perform spin-selective imaging of our atoms; this can then be used for post-selecting our data to account for single-atom loss, whereas our current imaging protocol registers loss as a spin-up atom. Additionally, we can use extra tweezers to read out more information about the final state of the system, similar to the method depicted in Figure 1.1, but which can isolate imperfections in other stages of preparation or in the readout of the spin state at the end of an experiment. The ability to use these techniques to gain more information about the state of a system, both before and after a period of evolution is one of the most important features of this platform for quantum gas assembly.

1.3 List of publications

For reference, I am including a list of the published work that I have been involved in during the course of my thesis work.

- (1) A. M. Kaufman, **B. J. Lester**, and C. A. Regal, "Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State", *Phys. Rev. X* **4**, 041014 (2012).
- (2) A. M. Kaufman, **B. J. Lester**, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, "Two-particle quantum interference in

tunnel-coupled optical tweezers”, *Science* **345**, 306 (2014).

- (3) **B. J. Lester**, A. M. Kaufman, and C. A. Regal, ”Raman cooling imaging: Detecting single atoms near their ground-state of motion”, *Phys. Rev. A* **90**, 011804(R) (2014).
- (4) **B. J. Lester**, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, ”Rapid Production of Uniformly Filled Arrays of Neutral Atoms”, *Phys. Rev. Lett.* **115**, 073003 (2015).
- (5) A. M. Kaufman, **B. J. Lester**, M. Foss-Feig, M. L. Wall, A. M. Rey, and C. A. Regal, ”Entangling two transportable neutral atoms via local spin exchange”, *Nature* **527**, 208 (2015).

Chapter 2

Experimental Apparatus for Control of Single Atoms

In this chapter, I will give an overview of our implementation of some of the atomic physics techniques that are made use of in this thesis to initially trap, cool, spin-polarize, and detect atoms in the experiments performed. For all of the experiments in this thesis, we use ^{87}Rb , which has been the workhorse of the cold atomic gas community and along with ^{133}Cs has been the subject of most neutral atom Raman cooling experiments in free space [52, 53, 54, 55, 56]. By choosing this species, we were able to focus our attention on the challenges related to cooling and manipulating ground-state atoms in optical tweezers.

After covering this background information, I will go into more detail on the operational principles of optical tweezers, including the optical setup used to generate our traps, and how they are used specifically for single atom trapping experiments. Additionally, I will detail the methods we have used to carefully align the optics for generating these potentials, as well as methods for characterizing the trapping potentials that are generated using the atoms. Further information on the details of the experimental setup, such as the vacuum system and how particular laser frequencies are generated and referenced to each other, can be found in Adam Kaufman's thesis [57].

2.1 Trapping and Control of Cold Atomic Gases

The experimental setup is built around a single vacuum chamber that is made from 9 fused silica windows attached to a machined quartz frame (assembled by Precision Glassblowing). There is a pair of larger diameter (2.3") windows (that provide a large optical access to the cell), seven

0.9” diameter windows along the octagonal edge of the cell (which act as ports for additional beams to address the atoms), and the final port being occupied by a glass-to-metal transition connecting to both a rubidium source and a 20 liter-per-second ion pump. The experiment is operated by loading atoms from the background vapor into optical traps and then manipulating them all in the same chamber.

Surrounding this glass chamber are four pairs of coils to control the magnetic field environment for atoms located in the center of the vacuum chamber. There are three pairs of coils in the Helmholtz configuration, which provide control over the scalar magnetic field, in addition to a single pair in the anti-Helmholtz configuration, which allows us to apply a tunable magnetic field gradient through the trapping volume. Importantly, we can tune the scalar magnetic field from 0 to approximately 10 Gauss directed along an arbitrary axis. We can also turn up the gradient coils to achieve field gradients suitable for loading a magneto-optical trap of ^{87}Rb .

Aside from the vacuum chamber and magnetic field control, the experimental setup consists of an arrangement of near-infrared (near-IR) laser beams that are used to trap, cool, image, and spin-polarize the atoms. The majority of these lasers are frequency stabilized via an offset-lock from a reference laser (which is locked directly to a rubidium vapor cell), which operates by measuring the beatnote between the laser being locked and the reference, and stabilizes that beatnote to a particular frequency. Details about the lockpoints for the majority of these lasers are given in Adam Kaufman’s thesis, Ref. [57]. Additionally, there is a sawed-off waveguide that delivers microwave radiation to the atoms, allowing us to perform precise rotations of the spin state.

2.1.1 Lasers for cooling and trapping ^{87}Rb

To provide a reservoir of cold atoms, we use standard techniques for generating a magneto-optical trap (MOT) of ^{87}Rb atoms. Cooling and confinement of atoms in a MOT is supplied by the combination of a damping force (provided by doppler cooling) and a restoring force (provided by the magnetic field gradient). A diagram showing the alignment of the MOT beam, as referenced to the vacuum chamber and other beam paths is shown in Figure 2.5.

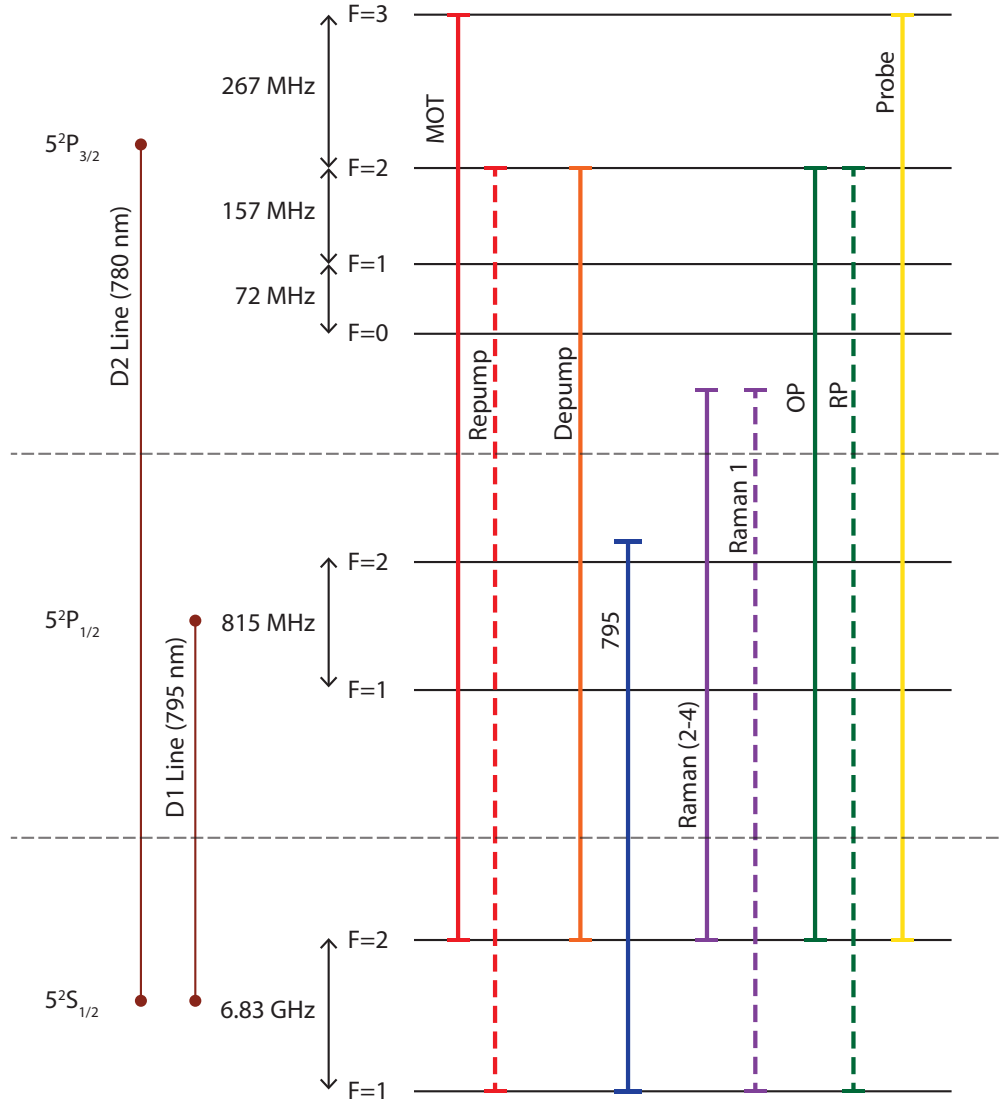
^{87}Rb Level Diagram and Laser Addressing Scheme (*not to scale*):

Figure 2.1: Schematic showing all of the relevant electronic levels in ^{87}Rb and their splittings. Additionally, there are lines indicating the levels being addressed (or detuned from) for each beam used to gain control of the spin and motional degrees of freedom in our experiment: The two red beams indicate the cycling (or MOT, solid) and Repump (dashed) beams used during loading and imaging, the orange and blue lines represent the Depump and 795 beams, respectively, used for enhanced loading (see Section 3.1), the two purple lines are the Raman beams (the dashed beam is RB1 and the solid line represents RB2, 3 and 4), the two green lines are the optical pumping (or OP, solid) and sigma repump (or RP, dashed) beams, and the yellow beam is the probe beam. The dashed purple, green and yellow beams are all on the same beam path, as shown in Figure 2.5.

To generate a three dimensional (3D) trapping potential, we use three beam paths that have a projection on all three dimensions, each path reflected back on itself in the $\sigma^+ - \sigma^-$ polarization configuration. Each beam path is the trap light, which is red-detuned from the cycling transition ($|F, m_F\rangle = |2, 2\rangle \rightarrow |F', m'_F\rangle = |3, 3\rangle$), which is the source of the doppler cooling and generates an optical molasses in the absence of a magnetic field gradient. The magnetic field gradient ensures that scattering rate is minimized at the center of the trap and as the atom moves toward any given beam, that beam is shifted closer into resonance, thus providing a restoring force to the middle of the trap. In addition to the trap light, we need use a repumper beam on two of the beam paths (the paths in the $\hat{x} - \hat{y}$ plane), which is used to keep atoms in the manifold of the cycling transition (the $F = 2$ ground-state manifold). This beam is tuned to be resonant with the $|F, m_F\rangle = |1, 1\rangle \rightarrow |F', m'_F\rangle = |2, 2\rangle$ transition, ensuring any atom off-resonantly scattering out of the $F = 2$ manifold (which occurs roughly once every 10^5 scattering events for our typical detunings) is rapidly pumped back back out of the $F=1$ manifold.

2.1.2 Loading directly from a MOT

To load single atoms, we simply overlap the optical tweezer potential with the trapping region of our MOT, which provides a reservoir of atoms around the optical tweezer. Then, we perform sub-doppler cooling on the cloud of atoms by further detuning the MOT beams, turning off the magnetic field gradient, and zeroing the total magnetic field. Sub-doppler cooling techniques, such as the polarization-gradient cooling (PGC) that we use, can cool atoms to very near the single-photon recoil limit, which will be very near the ground-state of our deep optical tweezer potentials [58]. In our case, the cooling brings the atoms to roughly $20 \mu\text{K}$, such that any atoms loaded into the trap ends up near the ground-state of the potential.

Additionally, the same light used for cooling will induce light-assisted collisions between pairs of atoms [59], which can result in the loss of both atoms when the kinetic energy gained during the collision is larger than the depth of the trap. Due to the small trapping volume of the optical tweezer and the relatively low-density of our MOT, we are in the collisional blockade regime, where

only a single ^{87}Rb atom can remain in the trap and pairs of atoms are lost [60]. This effect is the basis of parity imaging in many neutral atom experiments [61, 18, 19].

2.1.3 Optical pumping and spin-polarization of ^{87}Rb

An additional beam path is aligned along the typical quantization axis ($-\hat{x}$), as shown in Figure 2.5, for which particular care is taken to achieve a very high σ^+ polarization purity. This beam path contains other frequencies of light, but in order to spin-polarize atoms that are trapped in the optical tweezers, we use two frequencies in particular: The “optical pumping” (OP) beam is tuned to the $F = 2 \rightarrow F' = 2$ transition, while the “sigma repump” (RP) beam is tuned to the $F = 1 \rightarrow F' = 2$ transition. With this pair of beams on, the atom will continue scattering photons until reaching the $|F, m_F\rangle = |2, 2\rangle$ electronic state, which is dark to both of these beams.

This optical pumping procedure is limited by the polarization purity of the beam path, which is very important for the Raman cooling procedure discussed in Section 4.1.2. The fact that the optical pumping procedure is limited by the polarization purity also provides us with a tool for both characterizing and optimizing the polarization purity of this beam. Specifically, we measure the scattering rate of the OP beam with a tilted quantization axis (and thus completely mixed polarization), as shown in Figure 2.2(a), to characterize the scattering rate by measuring how long it takes for the atoms to leave the $F = 2$ manifold. Separately, we can measure how quickly atoms that start in the “dark state” ($|2, 2\rangle$) are scattered out of the $F = 2$ manifold with the quantization in its optimal alignment, which happens because of the residual polarization impurity of the beam (and is thus a measurement of the “darkness” of the dark state). The ratio of these two rates is an estimate of the polarization purity of this beam path [62, 57]. This measurement is shown in Figure 2.2(b), which suggests that the ratio of σ^- and π polarization to σ^+ is better than 1:1250. Note that by fixing the pulse length, this same measurement can be used to optimize the parameters relevant for improving the polarization purity, such as the angle of the quantization axis, by minimizing the survival probability (which would correspond to a lower scattering rate) [57].




Figure 2.2: These two plots are taken using the same experimental procedure, but in (a) the quantization axis is intentionally misaligned, effectively mixing the polarization of the light at the atoms, while in (b) the quantization axis is optimally aligned to the atoms. In both plots, the atom is initially prepared in the $|F, m_F\rangle = |2, 2\rangle$ state before shutting off all other beams, adjusting the quantization axis, and then applying only the OP light for a variable pulse length. Afterward, the probability for the atom to have been depumped to the $F = 1$ manifold is measured and the exponential decay time fitted. The ratio of these decay times ($20 \mu\text{s}$ to 25ms) is an indication of the polarization purity of the beam with the quantization axis well-aligned.

2.1.4 Single-atom spin addressing

After spin-polarizing the atoms in optical tweezers, we need to be able to rotate them to arbitrary states to initialize different experiments. Fundamentally, we need the ability to rotate individual spins in an array of traps, so we take a two-part approach to this: Performing arbitrary global microwave rotations using a microwave drive (coming through the sawed-off waveguide) and localized vector light shifts, which can move the microwave resonance for individual traps by hundreds of kHz [45].

The signal for the global microwave drive is provided by a Rhode and Schwarz microwave generator (SMF100A), which passes through a voltage-controlled variable attenuator, and finally is amplified to around 10 W of microwave power before being coupled into a sawed-off guide that launches the microwave signal toward the atoms in free-space. With this setup, we can drive Rabi oscillations between the $|2, 2\rangle$ and $|1, 1\rangle$ hyperfine states at a rate of $2\pi \times 40 \text{ kHz}$, as shown in Figure 2.3. The variable attenuator also allows us to shape microwave pulses, in particular to achieve Gaussian-profile pulses, which results in a narrower spectral feature at the expense of longer




Figure 2.3: This figure demonstrates our ability to drive coherent rotations between the $|2, 2\rangle$ and $|1, 1\rangle$ spin states of ^{87}Rb by applying a microwave drive at 6.840996 GHz (shifted from the bare transition because of the quantization axis). By varying the length of a square pulse of this resonant microwave drive, we observe coherent oscillations between the two spin states of the atom at a rate of approximately $2\pi \times 40$ kHz.

π pulses. The Gaussian pulse shaping is key to the single-atom addressing because relatively small shifts of the microwave resonance (100 kHz or less) can result in a very small probability to flip the spin of the atom whose resonance is shifted. This is in contrast to the spectral features of a square-pulse, which extend much farther from the resonance and leading to a higher probability of flipping the spin of the atom whose resonance is shifted.

The localized vector light shifts are produced by adding a $\lambda = 785.1$ nm beam (which we call the “light shift”, or LS, beam), which is deflected backwards on the imaging path toward the objective lens, as shown in Figure 2.5. This wavelength is chosen so that with pure σ^- polarization, the $|2, 2\rangle$ state remains unshifted while the $|1, 1\rangle$ level is shifted.¹ This is useful because the energy of the initial spin state is not perturbed, which minimizes the effects of the LS beam (which could provide an additional trapping or anti-trapping potential, if that energy was perturbed). An example spectrum with the light shift beam shifting the resonance of the atom in one of the wells is shown in Figure 2.4, demonstrating that a single spin can be flipped with high efficiency.

In order to align the LS beam to the atom, we can directly measure, and maximize, the vector

¹ Note that this is a change from earlier experiments that used a much closer detuned beam (approximately 50 GHz) and much lower powers to shift the resonance. Importantly, by going to a larger detuning, we can perform longer microwave rotation pulses (including gaussian shaped pulses) to improve the fidelity of our site-selective spin rotations.




Figure 2.4: This is a microwave spectrum taken the light shift beam focused onto one of the wells (green points). The beam induces a large vector light shift of the atom in the right well (red arrow) that allows spin of the unshifted atom (orange) to be flipped with high efficiency and with no significant effect to the spin of the atom shifted out of resonance. The peak of the orange spectrum is less than unity due to a combination of factors including shot-to-shot magnetic field fluctuations and single-atom loss to background, but spin flip probabilities of about 95% are achievable [57]. Additionally, the broadening of the shifted atom resonance is likely due to a small intensity noise on this laser and is not a fundamental limitation.

light shift induced between the $|2, 2\rangle$ and $|1, 1\rangle$ spin states. To measure this, we perform a microwave spin-echo pulse sequence ($\frac{\pi}{2} - \pi - \frac{\pi}{2}$) and turn the LS beam on during the hold between the first two pulses, but not between the second and third pulses. Thus, this is a differential measurement that cancels out any systematic effects due to stray magnetic fields or vector shift induced by the trapping light itself. By varying the alignment of the LS beam during this measurement (using a piezo mirror to adjust the angle), we can maximize the light shift and thus maximize the intensity at the atom we align the beam to. Once aligned, we can tune the power of the LS beam to reliably shift the microwave resonance for the selected atom far enough that the spin is not rotated by the Gaussian pulse, but keeping the power as small as possible to minimize the scattering rate from that beam (and thus minimizing heating or spin-depolarization).

2.2 Optical System for Single-Atom Trapping

Optical tweezers are a type of optical dipole trap that provide very tight (typically sub-wavelength) confinement of atoms. Of particular interest are traps containing only a single atom,

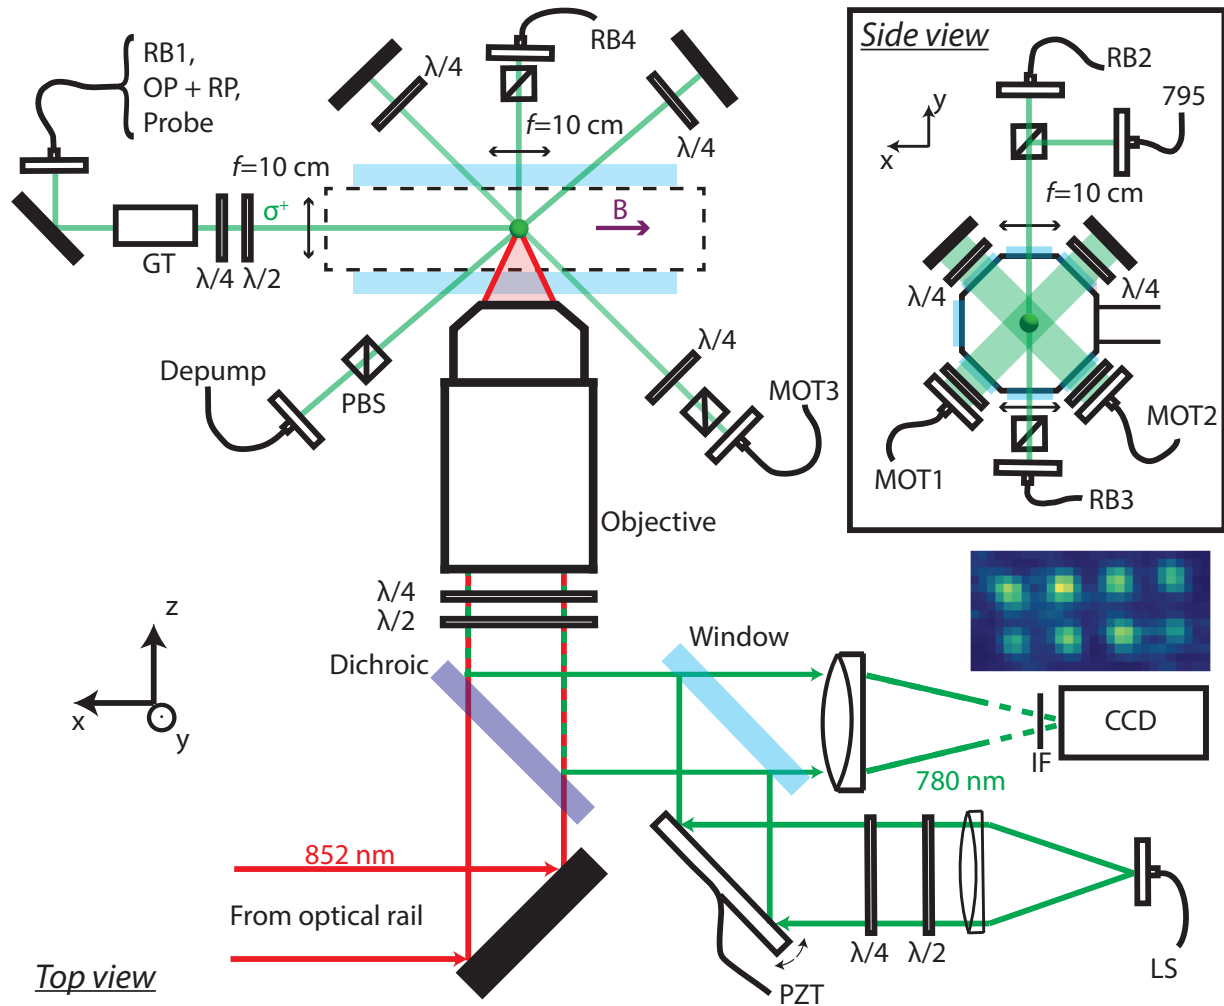


Figure 2.5: Schematic of the optical setup, showing the relative alignment of all the beams for controlling the full quantum state of atoms trapped in optical tweezers formed by the custom objective lens outside of the vacuum chamber. For imaging, the same objective lens is used to collect fluorescence from the atoms, which is then deflected by a dichroic mirror and down an imaging path with an $f = 1$ m lens onto our Andor 897 EMCCD camera. In order to manipulate the spin in a site-selective way, we send another beam back toward the objective lens from the imaging path (via reflection off of a window blank). The remaining beam paths are used for state manipulation, addressing the transitions depicted in the level schematic shown in Figure 2.1.

in principle giving us the control to organize single atoms in arbitrary configurations [32, 63, 35] and even combine separated atoms into the same trapping potential [64, 65, 47]. The simplicity of the optical trap also offers many advantages in that many electronic states can be trapped simultaneously (in contrast to magnetic traps) and the same trap system can generally be used for many different atomic species (possibly with slight modifications to the trapping wavelength).

The optical dipole trap uses light that is far detuned from resonance with any ground-state transitions and relies on the ability to shift the electronic energy levels of the atom without scattering a significant number of photons (which would heat the atoms and destroy any coherence between electronic states). Specifically, we take advantage of the AC Stark shift, which lowers the ground-state energy of the atom when the light is below resonance, or red-detuned, because the light induces a dipole moment that is in phase with the electric field of the light. The confinement of the potential is then formed by the intensity gradient (in all three dimensions) at the focus of the trapping beam, where the intensity of the light is at a maximum and thus the ground-state energy is minimized. Alternatively, light that is far above the same transitions, or blue-detuned, can be used when an intensity minimum is created at the focus (such as with higher order Laguerre-Gauss beams) because, in this case, the light increases the ground-state energy of the atom.

In our system, we use 852 nm light, which is far detuned from both the wavelengths of the D1 and D2 transitions in ^{87}Rb (which are around 795 and 780 nm, respectively). We also collect fluorescence scattered during PGC (at 780 nm), which sets many of the requirements for our optical system. Ideally, the optical system will be diffraction limited for both of these wavelengths and be capable of generating a 5×5 array of diffraction-limited traps, each spaced by $1.5\ \mu\text{m}$.

2.2.1 High-numerical aperture lenses

In the generation of our optical tweezers, the most important element is the high numerical aperture (NA) objective lens that brings the beam of light to a sub-micron focal point, which, when used as an optical dipole trap, provides the tight confinement we want. In order to simplify the system, we use an infinity-corrected objective lens, which will take as input a collimated beam

(and, for imaging, will collimate light scattered from a point source in the focal plane of the lens). Because we want to place the lens outside of the vacuum chamber, there are some constraints in the physical design of the lens, which are set by the dimensions of the glass vacuum chamber and the diameter of the coil mounts. Specifically, the working distance of the objective lens needs to be larger than 20 mm to have some space between the front of the lens and the vacuum window and to preserve the optical access, the half-angle of the cone that the lens can occupy cannot exceed 50° from the center of the vacuum chamber.

We are currently using a custom objective that was designed and built by ASE Optics that provides an effective NA of approximately 0.6. The objective is designed to be achromatic for both 780 and 852 nm and to operate at a NA of 0.64, with a long working distance of ~ 21 mm and correcting for the aberrations introduced by focusing through a 6.35 mm thick fused silica window. The housing of the objective is made out of Ultem because it can be precisely machined and is nonmagnetic, which is desirable because of the proximity of the field coils. The objective was not diffraction limited, but we found that we could get acceptable performance only at a NA of 0.6. Some more details on the characterization of this lens will be given in Section A.1 and there is additional discussion of this particular lens in Ref. [57].

2.2.2 Sources of light for optical trapping

We have used a number of different lasers to produce light for trapping single atoms, but recently we have switched between two high power laser sources: A highly coherent Titanium:Sapphire laser (an M² SolsTiS system) operating at 852 nm and a temporally incoherent laser system consisting of a 1 W tapered amplifier chip (from Eagleyard Photonics) seeded in a double-passed configuration with light from a super-luminescent light emitting diode, or SLED (supplied by Superlum). Both of these systems can reach similar intensity levels at the atoms, which are sufficient for the experiments being performed, but have different advantages. While the SolsTiS system is quite a standard unit, the temporally incoherent source is rather new (we learned about its use in Markus Greiner's lab and decided to try it out based on the success in their system) and so I am




Figure 2.6: Schematic of the optical rail that takes input light from a fiber and then generates a high-quality beam that is used to generate sub micron optical traps. Note that due to the optical relays, the two AODs are in conjugate planes of the optical system, which is then passed to the back focal plane of the 20x telescope, and finally projected onto the back of the objective lens.

adding some of the details of our construction here [66, 67, 68].

Our switch from the SolsTiS system to the temporally incoherent light source was motivated by our concern that reflections off of any of the beam preparation optics, the cell, or even within the objective lens itself, could lead to residual amplitude modulation at the focus of our lens, leading to small (and slow) time-variation of the relative depth between multiple traps. We have some external data suggesting that there was residual drifting that has improved by switching to the incoherent light source, but it is not conclusive and we have yet to verify the improvement with the atoms directly. I am sure more details on this testing will be discussed in an upcoming thesis by our new graduate student, Mark Brown.

The main concern about using a broadband source is that there could be “blurring” of the trap due to the combination of chromatic aberrations, polarization optics working over a narrow wavelength range, and, potentially most importantly, the acousto-optic deflectors (AODs) having a slightly different deflection angle for each wavelength component.

2.2.3 Optical Rail

In order to achieve the expected performance of the optical tweezers, it is important to start with a high-quality and unaberrated beam. The challenge is that we need a beam with a $1/e^2$ radius on the order of 15 mm to fill the input aperture of the long working-distance, high NA objectives that use. Additionally, we want the beam to pass through two acousto-optic deflectors (AODs), which will enable some amount of beam steering (more details on the capabilities afforded by these AODs is given in Section 3.2), which means that the optical system will need several components that are well referenced to each other.

To achieve these goals, we have designed a custom built “optical rail”, which minimizes the number of adjustable degrees of freedom such that it becomes a simpler task to align all of the optical elements using references that are referenced to fixed post locations. Additionally, the rail is designed to be compact (fitting in the area available near the vacuum chamber) with the output beam height fixed to be centered on the optical axis of the objective lens and the vacuum cell. Aside from the mechanical stability and benefits of fixing the location of as many of the optics as possible, the optics on the rail are designed to provide the best performance for the objective lens.

Fundamentally, the optics of the rail consist of a small set of features that give us the flexibility and performance required, which are all included in the diagram shown in Figure 2.6. At the input is a standard fiber launcher setup for a flat-faceted (FC/PC) output fiber connector, which has been anti-reflection (AR) coated to prevent interferences, and where we have taken special care to make sure that the output fiber is centered on the collimation lens. Immediately following this is a polarizing beamsplitter cube (PBS), as well as a $\lambda/2$ waveplate (after the first AOD) that is used to properly align the polarization into each of the AODs (which achieve a higher deflection efficiency when the polarization is transverse to the acoustic wave propagation). After the first AOD is an $f = 4.5$ cm 4- f “optical relay” that brings the deflected beam from the center of the first AOD to the center of the second AOD. An optical relay is simply a 1:1 telescope that maintains the angle and position (up to a sign) of the beam between the input and output planes of the telescope,

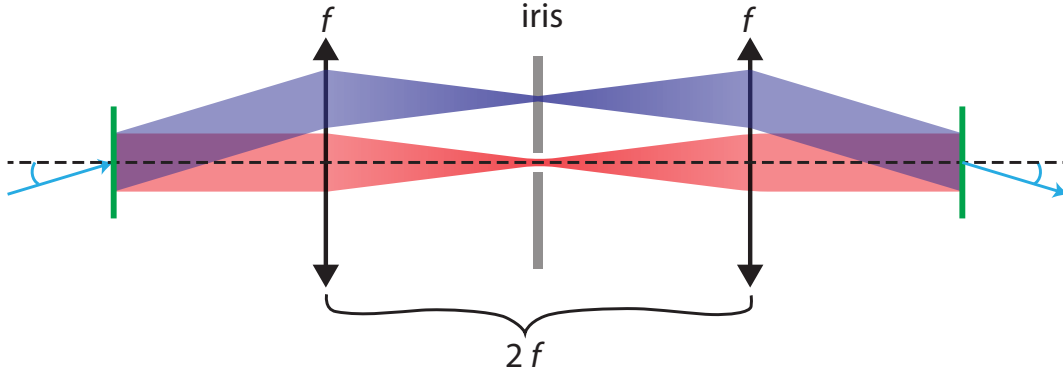


Figure 2.7: Diagram to highlight the utility of the optical relay (also known as a $4f$ -relay or a 1:1 telescope) to keep the beam deflections on-axis through the optical rail (and eventually into the objective lens). Specifically, by positioning two lenses of equal focal length, f , by twice that focal length, any two planes located a total of $d = 4f$ apart (with both lenses in between the two planes) will be conjugate planes. Because the magnification factor is unity, this means that in both planes both the location of the beam and the angle of the beam (with respect to the optical axis) will be the same, up to a sign flip (both the angle and the position will be opposite).

as is demonstrated in Figure 2.6.² Then, there is a longer $f = 15$ cm optical relay that brings the deflections from the second AOD to the back focal plane of the final $M \simeq 20$ telescope that magnifies the beam and brings it to the back of the objective lens.

However, before this final telescope is a set of polarization optics that allow us to pick off a fixed fraction of the total beam power. Specifically, we use the sequence of a PBS (to initially clean up the polarization after the other optics) followed by a $\lambda/2$ waveplate and a second PBS, which together pick off a fixed fraction of the light transmitted by the first PBS. The first PBS may not be necessary, but add stability to long-term drift of the polarization after passing through the two AODs (which are both birefringent and thermally sensitive), two optical relays, and two $\lambda/2$ waveplates. The signal picked off by the second PBS is then used to monitor the power going into the objective lens, which is directly proportional to the trap depth of a single optical tweezer. Further, by feeding back on this measured signal, we can stabilize the trap depth (and use the computer control to ramp this depth to different values during a single experiment).

² These optical relays are particularly important for our beam preparation because of the large magnification telescope, which would take a slightly off-axis beam and cause it to completely miss the objective lens.

In the context of stability, building a completely fixed rail would be ideal, but we do need to make sure there are enough degrees of freedom to achieve the desired performance. Because of the sensitivity of our lenses, this performance requires the precision alignment far beyond what we are likely to achieve with a completely fixed system. For clarity, we have the ability to tune the alignment of the beam axis and to align with references to the physical axis (built-in pinholes) of the optical rail, as well as to tune the spacing between lenses in the optical relays (and the telescope) to maintain a well-collimated beam, which provides the required flexibility to optimize and align the output beam from the rail to the objective lens. A brief outline of how we performed this alignment is included in Section A, where we discuss the full alignment procedure.

Chapter 3

Quantum Gas Assembly: A Platform to Study Tailored Many-Body Systems

As introduced in Section 1.1, the goal of quantum gas assembly is to build up interesting quantum systems with single neutral atoms as the building blocks. There is a tremendous amount of control afforded by the use of optical tweezers, which not only allows us to build these systems from the ground-up, but also gives us the opportunity to access more information about the system we are preparing along the way, as depicted in Figure 3.1. It is the goal of this chapter to discuss some of the methods we have developed to take advantage of this control, as well as to highlight the utility of the information that we have access to, which has been inaccessible, or otherwise not fully utilized, in many earlier quantum gas experiments.

In our experiment, we have already developed the ability to load atoms in tight optical tweezers, rearrange the potentials during the course of a single experiment, laser cool these atoms to the three-dimensional ground state, individually manipulate the spin-state of separated atoms, and the ability to analyze data based on the measured initial and final configurations. The experiments to date have all been with small arrays of traps and this chapter attempts to address some of the challenges that need to be overcome to enable these experiments to expand to more than a few atoms. Specifically, we need to be able to reliably trap a set of atoms that are individually addressable, can be arbitrarily transported and transferred between potentials, and can each be initialized into a pure quantum state of both the spin and motional degrees of freedom.

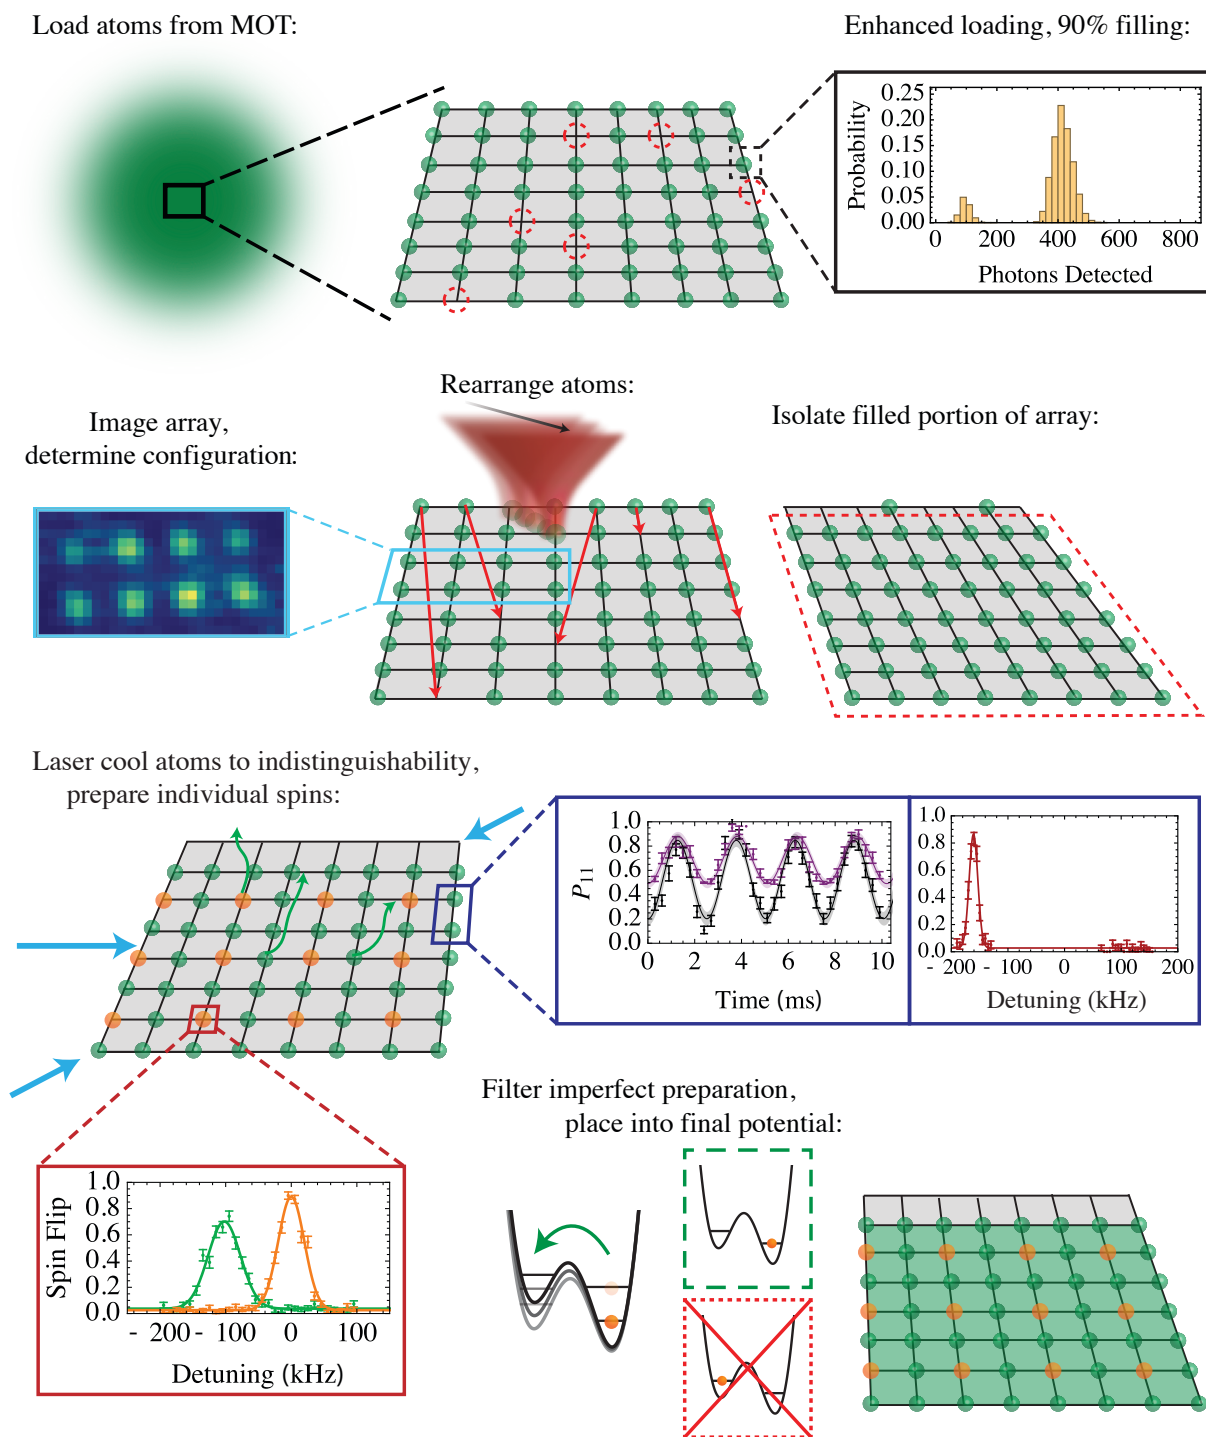


Figure 3.1: A copy of the schematic vision showing the initialization of experiments via quantum gas assembly. Highlighted are a few of the important capabilities that the optical tweezer platform affords, which will be discussed in more detail throughout this chapter.

3.1 Enhanced Loading of Single Atoms

The loading procedure discussed in Section 2.1.2, where atoms are efficiently laser cooled in the vicinity of optical traps, is a reliable method to stochastically load atoms into tightly-confining optical traps. However, the loading efficiency with these small-volume traps is typically limited to 50% [60], with the highest efficiencies achieved in optical tweezer experiments reaching efficiencies of as high as 64% [69, 70]. The understanding of the light-assisted collision process between tightly confined atoms has allowed us to understand these limitations and dramatically improve the efficiency of loading techniques [36, 37, 38].

Note that all previous work studying the use of light assisted collisions to improve the loading efficiency were performed using both larger volume optical traps, where multiple atoms could initially be loaded before initiating light-assisted collisions, and using ^{85}Rb , which has both a different level structure and ground-state scattering properties [71, 36, 37]. Additionally, these studies found that the collisional process to prepare a single atom in the trap required several hundred milliseconds, which would have been prohibitive to increasing the repetition rate of these optical dipole trap experiments. More recently, loading efficiencies around 80% have been demonstrated with ^{85}Rb in a micron-scale optical tweezer and requiring only 30-60 ms collisional pulses [72].

3.1.1 Background on light-assisted collisions

The process of light-assisted collisions is successfully described by onset of resonant optical transition between different molecular (interaction) potentials at specific interatomic distances R_C , as shown in Figure 3.2. The different potentials correspond to the different electronic configurations of the colliding atoms valence electrons. When the electrons of both atoms are in the s -orbital, then the atoms are on an attractive electrostatic C_6/R^6 potential, but when one of the atoms is excited to a p -orbital, the dipolar interaction (which can be positive or negative) varies as C_3/R^3 [59]. The difference in these interactions means that for any reasonable detuning from the bare (single-atom) resonance frequency, there will be a particular distance, defined as the Condon radius, where the




Figure 3.2: a) Interaction potential for two ^{87}Rb atoms colliding in the presence of near-resonant light. For reasonable detunings, there will be an internuclear separation, R_C , where the light is resonant with the excitation between the interaction potentials, resulting in a light-assisted collision. The character of that collision changes depending on the sign of the detuning: For positive (blue) detuning, the excited potential is repulsive, allowing us to control the maximum kinetic energy gained by the pair of atoms undergoing to collision. For negative (red) detuning, the excited potential is attractive. b) Schematic of the relevant energy levels in ^{87}Rb , indicating the detuning (from the light-shifted energy levels) of the beams used to perform enhanced loading (not to scale, for full level diagram see Figure 2.1).

transition becomes resonant with the driving frequency. When atoms are in close proximity the detuning from the bare transition can be quite large (such that the scattering rate from a single atom is low), even while resonantly driving this collisional process.

Of particular relevance for loading via efficient cooling is that the interaction potential to which atoms are excited by the red-detuned light that is used for cooling has a strong curvature, meaning that the atoms will be strongly attracted to each other. The atoms will be accelerated toward each other on the long-range attractive potential, but the “quasi-molecule” can emit a photon (now of a different frequency) bringing both atoms back to the s -orbital and giving the pair of atoms a large amount of kinetic energy. This is the dominant process that will result in the loss of both atoms from the trap during cooling in the collisional blockade regime. This pair-loss effect is what leads to the typically 50% loading efficiency (which can be pushed to 60%, in some cases)

and to parity imaging [70, 38].

Fortunately, there also exists a repulsive interaction potential that is resonant when the light is blue-detuned (higher in energy) from the bare atomic transition. Importantly, atoms on the repulsive interaction potential cannot gain more kinetic energy during the collision than the detuning of the beam (converted into energy units) because the potential curves asymptotically approach the bare levels. This additional control, combined with our knowledge (and control over) the depth of the optical trap is used to preferentially expel a single atom at a time (rather than pairs of atoms). Specifically, by setting the detuning to be equal to the depth of the trap, the maximum kinetic energy gained by the pair of atoms is only enough for a single atom to be lost! This enhancement is demonstrated in Figure 3.3(a), where the collisional process allows us to improve the probability of fully filling of an array with four sites by a factor of four compared to our best PGC loading procedure. While the process is still stochastic, this is a drastic improvement for beginning to scale to larger arrays of atoms.

3.1.2 Experimental implementation of enhanced loading

However, there remain some complications to make the repulsive light-assisted collisions useful for enhanced loading in the experiment: First, we need to be sure that the rate at which photons are scattered from a single atom in the trap is small enough that a single atom in the trap is not lost. Additionally, we must be sure that the scattering/collision rate induced by the blue-detuned collisional photons is much larger than that of any other light that is red-detuned of the transition. In the end, these effectively require us to use very deep trapping potentials so that a collisional beam with detuning equal to the trap depth results in a single-atom scattering rate that is quite small, even with the power necessary to rapidly drive these collisions. In our traps, using such a deep trap has the side-effect of shifting the D2 transitions used for the MOT well out of resonance once the atoms are loaded into the trap; this helps in balancing the scattering rates between the residual MOT light and the blue-detuned collisional light.¹

¹ I do not think that this is actually necessary for the enhanced loading to work, but it does help to order the scattering rates as desired. In principle, one should be able to use a trap that uses a magic wavelength for a particular

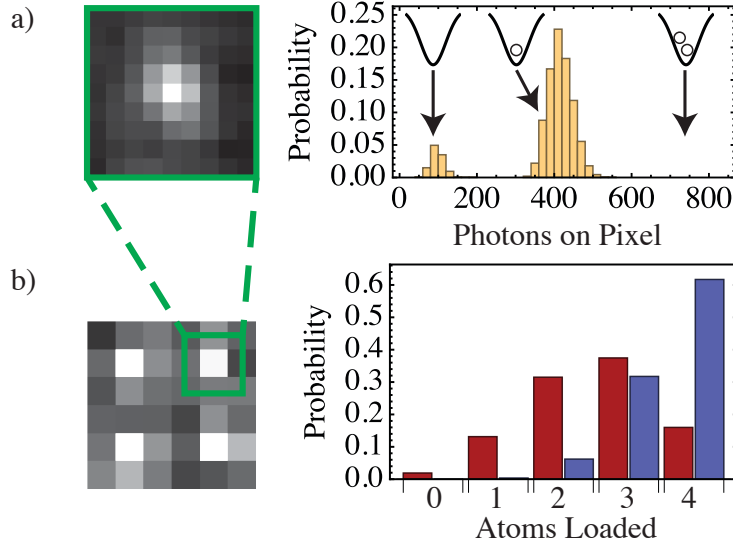


Figure 3.3: This figure shows the results of the enhanced loading. a) For a single site, we can look at the number of photons collected and plot a histogram of the number of photons collected during an image after each loading attempt; we see two peaks in the histogram corresponding to zero and one atoms loaded into the trap, with a much larger (approximately 90%) probability for an atom to be loaded. b) For a small array of optical tweezers, we find that we can achieve similar loading efficiency at each site. On the right, we compare the probability to load n atoms into the 4 site array when using enhanced loading (blue) and PGC loading (red). Specifically, we already see over a factor of 4 enhancement in the probability to load 4 atoms compared to using standard PGC loading.

As mentioned, the most important concern for implementing this enhanced loading scheme is to consider the rates of the various processes that are happening. Because we have such small-volume optical potentials (and are thus in the collisional blockade regime), we cannot rely on loading many atoms and then sequentially applying a collisional pulse to reduce the remaining number to 1. Thus, we need near-resonant (in free-space) beams on to generate a reservoir of atoms that are being cooled into the trap.² In our case, these are simply the MOT beams, which address the D2 $F = 2 \rightarrow 3'$ cycling transition and the D2 $F = 1 \rightarrow 2'$ repump transition. These beams are sufficient for loading a single atom in 50-60% of trials, but we will add two additional beams that are transition and design a scheme around that trap which can have all the rates work out. However we do not currently have a trap in which to try this.

² In principle, this could be a blue-detuned set of beams, which could help alleviate some of the requirements for minimizing two-atom loss events (and possibly improving the final loading efficiency further), but this would require adding another set of cooling beams on the MOT beam paths (or adjusting the locking setup to allow us to tune farther blue-detuned than we currently can go), which prohibited exploring this possibility.

tuned to address atoms that are already loaded into the optical tweezers: The main beam is used to induce the collisions as described above and is blue-detuned from the D1 $F = 1 \rightarrow 2'$ transition. The final beam is nearly resonant (in the trap) with the D2 $F = 2 \rightarrow 2'$ depumping transition. For a discussion of the reasoning behind the choices of these transitions, refer to Section 3.1.4.1.

For simplicity, we choose to add these collisional beams during the MOT loading stage, which means we don't need to add a multi-stage loading procedure that is long and complicated. This procedure results in one atom remaining in the trap more often than our standard MOT loading procedure because the primary loss mechanism is single-atom loss induced by the blue-detuned light-assisted collisions. As seen in Figure 3.4(b), the loading efficiency saturates to around 90% in about 170 ms, this time being primarily limited by the loading time of our MOT (which could be improved, but without larger scale changes to the setup would require sacrificing the trap lifetime) rather than any effects related to the blue-detuned collisions.

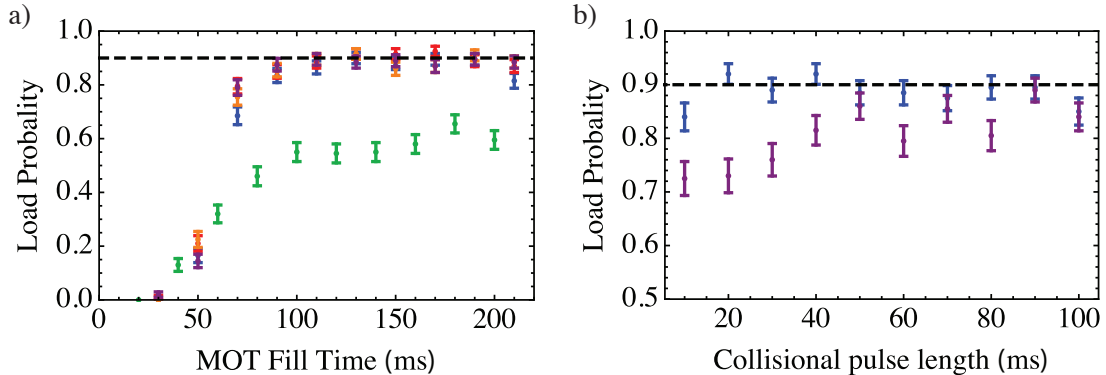


Figure 3.4: The loading process is split into two sections: The MOT fill time, which provides the reservoir of atoms to load into the optical tweezers, and the collisional pulse, which ensures that there is never more than a single atom at the end of the loading sequence. a) Here, we measure the loading probability while varying the MOT fill time for a fixed collisional pulse length of 35 ms. The probability for each well in a 2×2 array (blue, red, orange, and purple) is plotted along with the data taken when loading a single well via PGC (green). b) Here, we look at the effect of the collisional pulse length on the final loading probability. The blue points were taken using the optimal collisional beam power, while the purple points were taken with 30% of that beam power. The black dashed lines indicate a 90% loading efficiency.

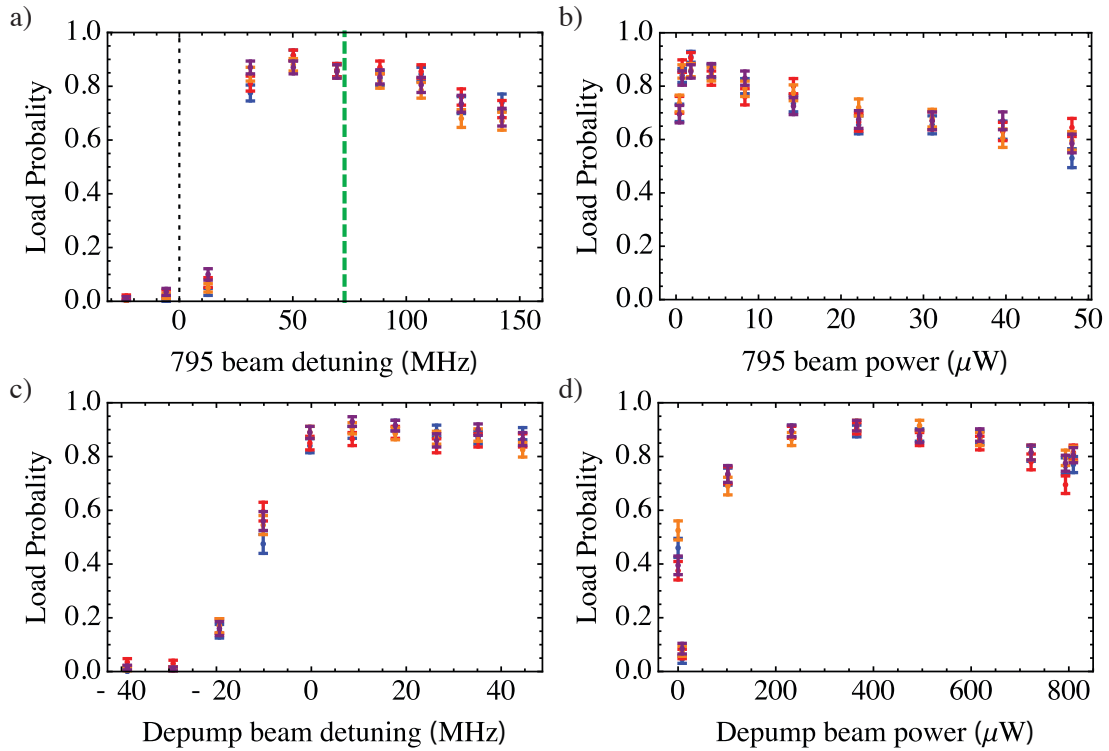


Figure 3.5: a) Loading efficiency as a function of the “795 beam” detuning with the trap depth highlighted as a green dashed line. The optimum detuning is slightly less than the trap depth, and larger detunings results reduced efficiency due to pair loss. b) Loading efficiency as a function of the 795 beam power. c) Loading probability as a function of the depump beam detuning. d) Loading probability as a function of the depump beam power, indicating the relatively high power required for enhanced loading. Note the drop in efficiency at low, but nonzero, powers (see Figure 3.6). In all of these figures, the colors (blue, red, orange, and purple) represent the measured loading efficiency for each of the four sites of the 2×2 array shown in Figure 3.3(b).

3.1.3 Collisional beam parameters

For our particular implementation of the enhanced loading procedure, we use our standard MOT beam configuration, as outlined in Section 2.1.1 and add the two additional beams discussed here. The primary collisional beam, called the “795 beam”, is detuned from the D1 transition (which is around 795 nm) and is aligned with the atoms along the same path as the top Raman beam (discussed in Section 4.1.1). The polarization of the 795 beam is linear and aligned along the propagation axis of the optical tweezer (orthogonal to the polarization of the top Raman beam). The second additional beam, called the “depumping beam”, is used to quickly depump atoms that have been loaded into the optical tweezer to the $F = 1$ manifold so that it becomes dark to the MOT light and can be addressed by the 795 beam. This beam is effectively unpolarized due to the beam being incident to the optical tweezer axis at an approximately 50° angle, and being retroreflected in the $\text{lin}\perp\text{lin}$ configuration.

The 795 beam is the primary driver of these light assisted collisions, so it is not surprising that both the power and detuning are important parameters for achieving the best possible loading efficiencies. The effect of these parameters is highlighted in Figure 3.5(a,b), which show a clear optimum in both parameters. It is important to note that these parameters are coupled: The heating of single atoms due to off-resonant scattering of the light limits both the loading at higher-power and smaller-detuning. However, the global optimum will occur when the detuning is very nearly equal to the trap depth, and then the power is tuned such that the loss due to heating of single atoms is negligible (at least in comparison to the cooling rate from the far-off resonant MOT beams). Fundamentally, this is because we are maximizing the rate at which the blue-detuned light assisted collisions will occur, which allows this rate to compete favorably to the pair-loss induced by the MOT beams that are required to initially cool atoms into the optical tweezers. This behavior is consistent with earlier work in larger volume dipole traps where the collision rate was smaller (and which were not in the collisional blockade regime, so more than one atom could be initially loaded into their traps), and thus much longer collisional pulses at lower powers were required to minimize

single-atom loss [37], but because we are in the collisional blockade regime, both of the light-assisted collision rates are much faster and collisional pulses only need to be tens of milliseconds before we are guaranteed to be left with either one or zero atoms, which is demonstrated in Figure 3.4(b).

While the role of the depumping beam is conceptually simple, it plays an important part in promoting the processes we care about during our loading procedure. Importantly, we find that there is a rapid drop in loading performance when moving outside of the range of values that work well for the detuning and power of this beam, as demonstrated in Figure 3.5(c,d). For the detuning, problems only occur when detuning red of the light-shifted transition (meaning, red-detuned of the transition for atoms loaded into the trap) up to at least +50 MHz detuning. However, for the power, we find that there is a broad “sweet-spot” where optimal loading occurs. Surprisingly, if the power is too low, we see that the atoms are always kicked out of the trap, in addition to problems when the power is too high (likely due to the scattering rate becoming high enough to heat single-atoms out of the trap).

This low power problem is highlighted in Figure 3.6(a), which plots the loading probability as we vary the depump beam power for a fixed detuning. Notice that at very low powers (or when the beam is off), we get approximately 50% loading efficiency, which is equivalent performance to not having a blue-detuned collisional beam on because, without depumping, the atoms are unlikely to undergo blue detuned collisions. However, as the power is increased, the loading efficiency rapidly drops to very nearly to 0, before increasing again and saturating at much higher powers. The resulting conclusion is that **we need very large depumping beam powers** to achieve optimal loading efficiencies, however the exact reason is not fully understood. It is likely that it is related to the competition of rates between various scattering processes, but what is particularly interesting is that this is a single-particle effect!

That this is a single particle effect is highlighted in the data shown in Figure 3.6(b), where we see two curves that both exhibit the same dip in the survival: The black points are data taken where two atoms are placed in the same well before applying the collisional light, while the blue points demonstrate the same effect when the same experiment starts with only a single atom. Because

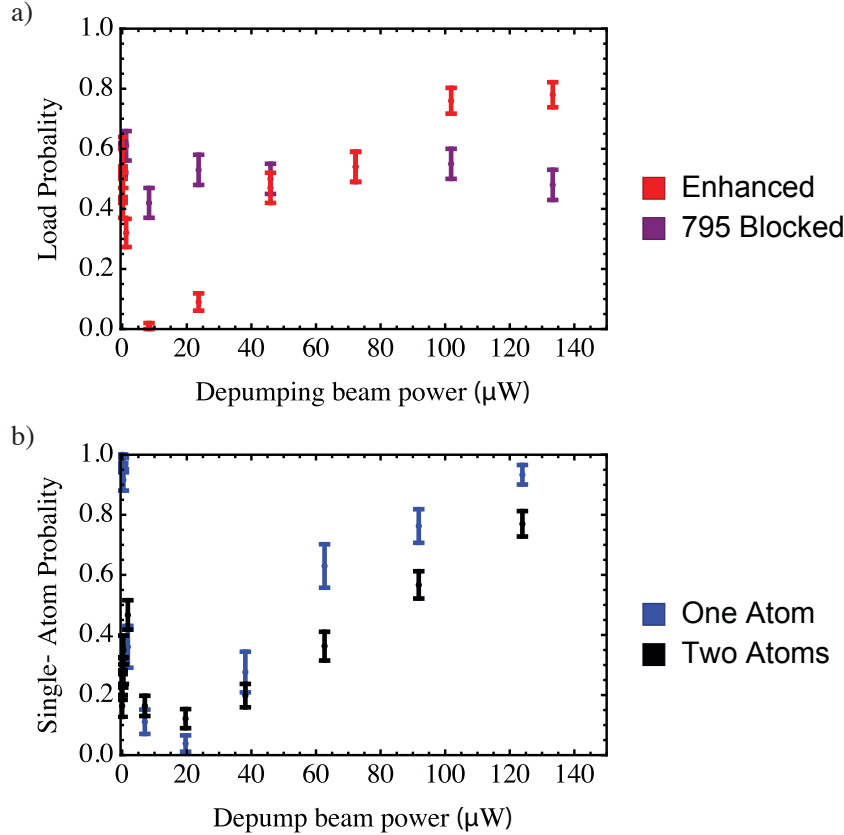


Figure 3.6: a) Measurement of the loading efficiency as a function of depump beam power with finer resolution than in Figure 3.5(d), comparing the behavior seen when the enhanced loading procedure is performed as usual (red) and when the 795 beam is blocked (purple). b) Measurement of the probability for one atom to remain in a single trap, given that it was initially prepared with one (blue) or two (black) atoms before a collisional pulse (consisting of only the 795 and depumping beams) was applied to the trap.

this is a single particle effect, it is worth mentioning that this could be due to the exact orientation of the beams we are using in our setup; we do know that effect requires *both* the 795 and depumping beams to be on in order for atoms to be kicked out.

3.1.4 Limitations to the enhanced loading

In the best case, we are currently limited to an approximately 90% loading efficiency using blue-detuned light assisted collisions, but its worth asking two questions: What are the limitations

to improving this optimum rate? And are there any limits when scaling up to larger arrays of traps?

3.1.4.1 Limitations to the maximum loading efficiency

As has been the theme throughout this section, there will be a fundamental limitation to the loading efficiency given by the competition of the two-atom loss rate to the single-atom loss rate when multiple atoms remain in the trap. The fundamental limitation here comes from the level structure of the atom being loaded into the traps and the relevant interaction potentials for undergoing light assisted collisions [37]. In ^{87}Rb we benefit from well-isolated hyperfine structure and the ability to, ideally, isolate a single interaction potential, but this is never perfect for any real case. Without changing atomic species, this means one has to carefully consider the choice of beam(s) used for driving these light assisted collisions to minimize the total probability of off-resonant red-detuned, far blue-detuned, or spin-changing light-assisted collisions from occurring, which could each result in both atoms being lost from the trap. Because these can never be completely turned off, the important consideration then becomes maximizing the rate of the desired single-atom loss channel in comparison to these two-atom loss channels.

Specifically, our scheme in ^{87}Rb is chosen to utilize the D1 $F = 1 \rightarrow 2$ transition because it has no nearby higher-energy transitions that could lead to driving collisions on an attractive interaction potential. Additionally, the excited hyperfine structure is well-resolved, which means that we are an extra 816 MHz off-resonant from any blue-detuned collisions that would lead to large gain in energy (which means that the atoms need to be much closer together to drive this transition, so the probability is much smaller).³ An additional benefit of using this transition is that there trap induced light shifts of both the ground and excited state energy levels are mainly scalar shifts, which means that the detuning is the same for all Zeeman sublevels and, because we

³ To be explicit, even for quite large blue-detunings there will be an internuclear separation where the transition will be resonant. For such large detunings, this distance is very small and thus the probability of driving this transition in comparison to the desired transition is likely small. However, when this additional detuning becomes on the order of the detuning from the desired transition (i.e., the trap depth), which could certainly be the case with different transitions, other atomic species, or by going to much larger trap depths, then this concern becomes much more relevant.

are not spin polarized during loading, this is particularly important for efficient loading [73, 74].

Another major limitation to controlling single versus two particle loss comes from the uncertainty related to the initial conditions of colliding atoms. In particular, if the temperature of the atoms is not well known, then the spread in initial energies will lead to a higher probability of both atoms being lost, even when colliding in the desired channel. If the energy is too high to begin with, then it is possible for both atoms to be kicked out of the trap, even when the energy *added* is only enough for one atom to be removed. This also cannot be mitigated by using a smaller detuning: If you err on the side of not enough energy, then there will be a larger probability of neither atom getting kicked out, which means you start with a much higher initial energy during the next collision.

An additional limitation to the maximum loading efficiency will be set by the competition of the single-atom loss rate (when only a single atom remains in the trap) to the length of the collisional pulse required to ensure that there are never multiple atoms remaining in the trap. This source of single-atom loss will fundamentally be limited by the background (vacuum) collision rate, but for now is negligible and is mainly limited by the heating rate associated with scattering the collisional beam photons from the atom in the absence of cooling beams. Standard methods can be used to mitigate this (larger detunings or lower powers) by lowering the scattering rate, but this has to be balanced with the primary limitations dictated by the balance of the light-assisted collision rates. In the end, the combination of both of these concerns points to going to deeper traps such that larger detunings (as long as they are not so large that the rate of collisions leading to two atom loss becomes significant) can be used to optimize the single-atom loss rate, and then higher powers can also be used (thus improving the ratio of the light-assisted collision rates) without increasing the scattering rate off of single atoms.

It is certainly worthwhile to pursue an even better overall loading efficiency, which in our case would mean going to larger trap depths (and thus larger blue-detunings), but this pushed us into two other limitations that complicate efforts to significantly improve the loading efficiency. The first are a set of technical limitations because going to much higher trap depths requires a

proportional amount more optical power per trap, as well as a proportionally larger tuning range for any lasers that need to be resonant with atoms loaded into the trap. Additionally, when going to larger trap depths, the initial temperature of any atoms initially loaded into the trap will likely increase because of the larger light shifts at the center of the trap, which, as discussed, also limits the maximum loading efficiency. In the end, we used several different trap depths ranging from 2 to 4 mK and found that above 3 mK the efficiency saturated around 90%.

3.1.4.2 Limitations when scaling to larger arrays

While our movable tweezers mitigate some of the challenges associated with loading overlapping optical tweezers, it is important to consider proximity effects that become relevant when scaling to larger arrays of traps with a finite accessible field of view (as discussed in Section 3.2.2.1). In particular, even once we implement the ability to rearrange atoms after loading, the ability to load atoms with high efficiency will be critical for scaling up to larger arrays. However, this loading procedure relies on our ability to selectively induce light-assisted collisions between atoms in the same trap without affecting atoms in nearby traps. In the end, we do find that the loading efficiency becomes compromised as the traps are brought into close proximity, as demonstrated in Figure 3.7 [38].

The experimental loading probability drops off quickly as the spacing of two wells is reduced below $\sim 1.7 \mu\text{m}$ [Fig. 3.7(a)], at which point the barrier between the two wells is reduced. By comparing the outcome to a separate experiment where we load two atoms into one side of the double well potential before applying a collisional pulse, we observe that this leads to an increase in the probability of hopping of at least one of the atoms to the second well. Importantly, we note that these collisions resulting in the atom hopping to the other well reduces the probability of completely removing a single atom during each collision event. Using the numbers measured in the experiment, we can perform a simple Monte-Carlo simulation to estimate the loading efficiency we would expect if we simply increase the hopping fraction. We find that the reduction in final loading efficiency is reasonable and compare the results of this Monte-Carlo to two sets of loading

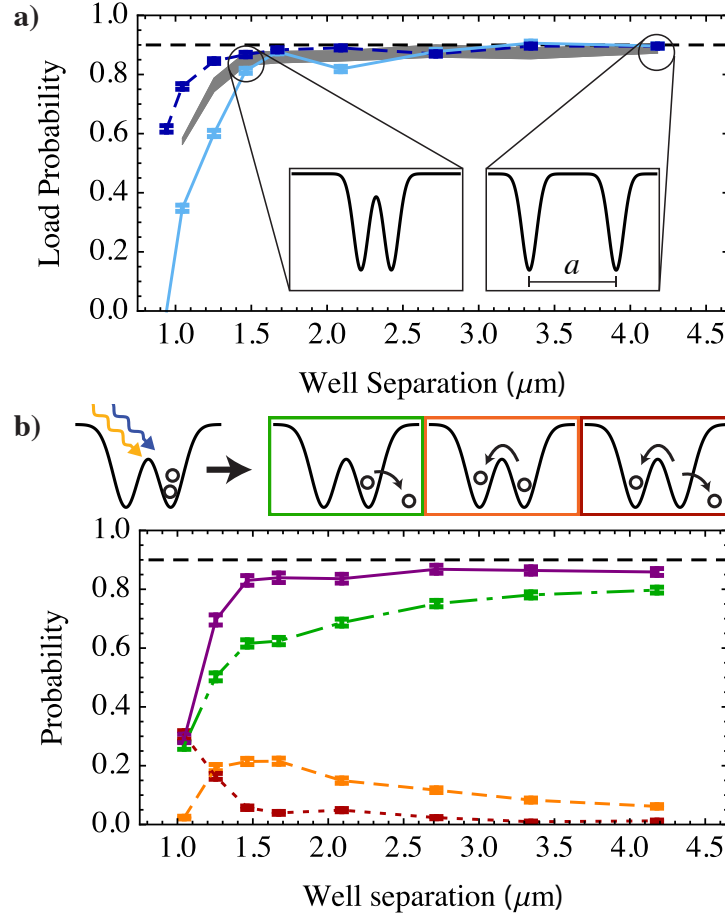


Figure 3.7: Effect of tweezer spacing on loading efficiency. a) Maximum loading probability achieved (per tweezer) after optimizing the loading parameters at each well spacing (dark-blue dashed line) and the loading probability when using the isolated-tweezer beam parameters at each spacing (light-blue solid line) versus the spacing between the center of two tweezers a . The loading probability calculated from a Monte Carlo simulation of the loading process (gray band) using the measured probabilities of the four possible outcomes [three of which are shown in part (b)] during each collision event and neglecting any single-particle effects. (b) Controlled experiment studying the effect of collisional light on two trapped atoms in proximity to a second trap. Measured probabilities of a single atom remaining in the right well (dash-dotted green line), one atom remaining in each well (dashed orange line), a single atom remaining in the left well (dotted red line), and the total probability for an atom to remain in the right well (solid purple line) as a function of the spacing between the optical tweezers. Note that this image is reproduced from Ref. [38].

data measured (with different collisional and cooling beam parameters) in Figure 3.7(a).

3.2 Manipulation of Optical Tweezers

As discussed in the Section 3.1, we use relatively deep traps (typically around 3.5 mK depth) that are far apart (meaning they are not overlapping) in order to promote processes that enhance the probability of losing a single atom at a time for all of the traps in an array. After this loading process is complete, we need to dynamically modify the trapping potentials so that the atoms locations can be imaged, the spin and motional degrees of freedom prepared, and then the atoms transferred into a configuration relevant for the experiment of interest. Within our current experiments, this can be achieved with the ability to change three quantities: 1) The spacing between neighboring rows or columns in a rectangular array, 2) The relative intensity between neighboring rows or columns (the “bias”), 3) The total optical power used to generate all of the potentials. The remainder of this section discusses how we gain control over the remaining two quantities within our experimental setup, as well as how it can be improved or extended for future experiments.

3.2.1 Overall intensity stabilization

To control and stabilize the depth of the optical tweezer array, we control the total optical power entering the rail for beam preparation. Specifically, this sets the overall depth scale for the traps, but its important to note that it does not control the actual depth of any given optical tweezer (which is modified by the total number of traps and their relative intensities, as discussed further in Section 3.2.3). In the case of our system, there are two concerns for the intensity stabilization: First, we need to be sure that the shallow science potentials are consistent shot-to-shot (and thus there is no long-term drift in the absolute power), such that the tunneling and interaction parameters remain constant. Second, we need to be sure that there is no resonant intensity noise on the trap light that will lead to heating of the atoms (and thus lead to distinguishability between our atoms, as discussed in Chapter 4).

The optical power entering the rail is controlled by modifying the RF power sent to an

acousto-optic modulator (AOM) that deflects light into the optical fiber at the beginning of the optical setup described in Section 2.2.3. This intensity is stabilized with standard feedback electronics via the power measured at the pickoff just before the beam expanding telescope in the optical rail. There are two main considerations for designing this feedback system: We need to maintain a stable intensity at DC, up to about 10 kHz in bandwidth (for ramping the intensity during experiments on millisecond timescales) without adding noise at frequencies resonant with parametric excitation processes (in the range of 10-300 kHz). Additionally, we need to be able to stabilize the intensity over a very large dynamic range of approximately four orders of magnitude (e.g., ranging from just a few microwatts to 40 milliwatts of power).

Standard techniques can easily achieve the desired bandwidth requirements, so I will not go into details about the feedback controller (typically a loop filter containing some amount of proportional and integral gain). As a part of my thesis work, it was important to go beyond the limitations of trap depth (both to go to deep traps for enhanced loading and shallower traps for more reliable tunneling) that were limited by the dynamic range of our previous detection system. Namely, we were previously completely saturating a linear photodetector for cooling trap depths, and could not go to shallower traps because we were running into the noise floor of the photodetector at tunneling depths.

However, to achieve a large dynamic range is a bit of a challenge and is typically limited by how the response of the system changes at various setpoints. In particular, we have begun using a logarithmic transimpedance amplifier, which has a very large dynamic range (in principle, the AD8304 can operate over 8 orders of magnitude in photocurrent). Additionally, to preserve our ability to use a standard feedback controller, we use a “linear-in-dB” variable gain amplifier that linearizes the response of the system [75]. Unfortunately, the response of this system changes quite dramatically as a function of the DC offset, preventing us from simply setting the loop filter as if this was a truly linear system. Currently, we have found a way to tune the system to give acceptable performance at most signal levels, but are working on ways to improve this using adaptive feedback systems, which will be discussed in more detail in future work.

Another method that we have explored (and may come back to in the future) is to use more than a single linear photodetector, with separate photodetectors optimized for different, overlapping, dynamic ranges. This has the benefit of a lower noise floor than a single logarithmic amplifier, and returns to a “simple” feedback system for each segment of the dynamic range. However, the complication is that we need a smarter controller than can either take more than a single photodiode input, which should be possible by switching to the use of a digital feedback controller. Alternatively, one could use an external control system to switch between separate control loops, each containing a single photodetector and feedback system.

3.2.2 Acousto-optic control of optical traps

The actual optical setup that is used to generate the optical tweezer potentials was described in detail in Section 2.2.3, so we will not go through the entire optical system here. However, it is important to note that the acousto-optic deflectors (AODs) located in the optical rail are in the Fourier plane of the trapping potentials (and are in conjugate planes with each other and the entrance aperture of the high numerical aperture objective lens).⁴ Importantly, this means that changing the beam deflection angle with either of these AODs directly translates to motion of the tweezer in the focal plane of the lens. Further, because the AODs are oriented orthogonal to each other, we have the ability to steer the optical tweezer potential in both X and Y at the focus by changing the deflection angles in each of the AODs. However, simply changing the angle of a single beam is not enough of a reason to use AODs; in principle, we could do the same thing with piezo mirrors at the same locations with minimal change to the optical rail.

The use of AODs is necessary because of the ability to change the deflection angles very rapidly, as well as to deflect multiple beams simultaneously (and to change the number of deflections in real time), all without physically changing the optical setup to scale to larger arrays. Some

⁴ The main difference between an acousto-optic deflector (AOD) and an acousto-optic modulator (AOM) is in its design: An AOD is designed to optimize for deflection efficiency and to have a large bandwidth over which the deflection remains efficient, which allows for efficient beam steering. An AOM is optimized for speed, such that a beam can quickly be turned on and off, or have its power modulated, but the frequency is expected to remain relatively constant in order to maintain the optimum deflection efficiency.

alternative options we have considered are discussed in more detail in Section 3.2.4. To demonstrate the utility of this system, we will start by reviewing the operational principles and important characteristics of these AODs.⁵

Acousto-optic devices operate by driving a crystal, whose index of refraction is particularly sensitive to stress, with a piezoelectric transducer and thus coupling an acoustic wave into the material. This acoustic wave leads to a spatially varying index of refraction, which is to say that it imprints a periodic index grating into the crystal. When the input angle of an optical beam is matched to the (acoustic) wavevector of this grating, then efficient diffraction of the input beam to a different mode (with angle related to the frequency of the acoustic wave and the material properties of the crystal itself) can be achieved. Because this deflection angle is directly dependent on the frequency used to drive the AOD, then by changing this driving frequency we can tune the deflection angle, and thus the location of the tweezer at the focus of our objective lens. Additionally, this is simply a Fourier optics problem, so by overlapping multiple frequencies into the AOD, we create a more complex diffraction grating with multiple spatial frequency components, and thus we end up with multiple deflections out of the AOD. Using the imaging magnification calibrated in Section A.1.1, and by measuring the spacing of the images of two atoms as a function of the frequency difference between the two traps, we can calibrate the change in position as a function of the frequency into the AOD, an example of which is shown in Figure 3.8.

3.2.2.1 Technical design considerations

The ultimate limit to our ability to position optical traps at the focus of the objective is set by the bandwidth of the AOD, the specific material properties of the crystal, and the details of the optical system used to move the beam from the AODs to the focal plane (specifically, the total

⁵ The AODs we use are Gooch and Housego model #: 46080-3-.85-LTD. They have 3 mm aperture longitudinal-mode TeO₂ crystals (4.2 mm/ μ s acoustic velocity), which allows for the use of a collimated beam in generating deflections (which is important to put the AOD in the Fourier plane of the focus of the lens). Importantly, they can achieve about a 90% deflection efficiency into the first mode at the peak response of 80 MHz (the specification is for > 70%, but measured 88 and 92% for the two in the rail when set up in an external test). They are designed to have a 30 MHz RF bandwidth, which achieves a ± 3 mRad deflection angle about the 16 mRad deflection angle of the 1st order at 80 MHz (and all at 852 nm).




Figure 3.8: Plot of the measured spacing between two atoms held in two traps as a function of the frequency difference between the traps. This *in-situ* measurement of the spacing calibration gives $(0.204 \pm 0.003) \mu\text{m}/\text{MHz}$ and is consistent with both external measurements and the expectation based on the properties of the AO crystal and the magnification of the rail telescope.

magnification of the telescopes and the effective focal length of the objective lens). Our optical system is designed so that we can, in principle, scale up to a 5×5 array of tweezers, each spaced by $1.5 \mu\text{m}$, which is sufficiently spaced for resolved imaging (although may not be sufficient for enhanced loading with the current aberrations). This is a somewhat conservative choice made to ensure that we are not limited by the frequency difference between adjacent deflections in partially overlapping traps.⁶

This is an extremely important consideration: The frequency difference corresponding to the desired minimum trap spacings should be much larger than the expected trap frequencies that might be used at these spacings. In our system, we chose to use the AODs discussed in the previous section (which result in a relatively small available field of view) because we are able to get to $0.5 \mu\text{m}$ spacing (on the order of the minimum spot radius we might be able to achieve) at $>2 \text{ MHz}$ in frequency difference, which is more than an order of magnitude larger than the trap frequencies we would encounter at these spacings.⁷ The reason this is an important consideration is that these beams

⁶ Recall that the frequency of the beam deflected by the AOD will be shifted by the acoustic frequency, so two adjacent traps will have optical frequencies that differ by the difference these frequencies.

⁷ If the frequency of modulation of the trap (due to interference of beams at slightly different frequencies) is large compared to the trapping frequency of the potentials, then there is unlikely to be significant heating induced by this modulation of the trap.

will be copropagating and spatially coherent, so the frequency difference between these beams will lead to a beating of the intensity of the trapping potential wherever the potentials overlap! In particular, we need to bring our traps close enough together to couple atoms between individual wells, which requires the traps to overlap without significantly affecting the atomic temperature. This requirement is in contrast to other experiments using optical tweezers, such as experiments with Rydberg atoms, where the atoms remain separated and the overlap between the traps can be kept to a minimum.

If desired, the maximum array size could be improved in a number of ways, but all of them will require some modification of the optical system. For a relatively small change, adjusting the beam size within the AOD would give a slightly smaller trap size, which could allow us to use more spots within the same field of view (at the expense of more power being clipped into the objective lens). Alternatively, a larger initial beam size would allow us to use a slightly smaller magnification in the final telescope, resulting in a larger deflection range (at the expense of a larger trap spacing for a given frequency difference), which would enable larger arrays of traps. However, for even larger arrays we would probably need to use a different AOD whose crystal properties or bandwidth enabled an even larger increase the range of deflection angles available (and thus a larger number of resolvable spots at the focus of the objective).

As you can see, it is important to balance these design considerations when building the system initially, but when properly designed, we transfer the complication of controlling the optical potentials into control of the drive signal for these AODs, which is simply an electronics problem (and excellent control can be achieved, as discussed in Section 3.2.3). Specifically, by changing the frequencies in each of the AODs, we can move rows and columns of an array of tweezers. Further, by changing the relative driving power in each of these frequencies, we adjust the deflected intensity of the traps in each of these rows or columns (thus, tuning the bias between rows or columns). This does not give completely arbitrary control over individual traps in larger arrays, but it will be sufficient in our typical use cases (and, for one-dimensional arrays it does give completely arbitrary control).

3.2.3 Generation of radio-frequency drive

We have historically used a few different methods to generate the appropriate frequencies to drive the AODs, starting with voltage controlled oscillators and then function generators with more built-in functionality. These were convenient and relatively simple, but each channel could only provide a single trap and, unless we sacrificed the stability of the signals, we were unable to trap positions in real-time.⁸ Further, in order to get more than a single trap, we needed to use multiple output channels and combine them onto the same line using external RF components (e.g., a power combiner), which means there were more components before the amplifier, which could lead to interferences, and the scheme is not arbitrarily scalable!

With the goal of improving both scalability and stability of the RF intensities, the source of the RF signals is now a single (dual-channel, one for each AOD) arbitrary waveform generator (AWG). This AWG has a large onboard memory (2 GB), which can store a set of completely arbitrary waveforms (at 400 MS/second) along with a script that determines the order of the waveforms.⁹ Critically, the script will repeat a base waveform until the generator receives a trigger to initiate some action, then hold at another waveform before being triggered to go to another position.

3.2.4 Alternative options for arbitrary trapping potentials

Many methods for generating projected optical potentials have been explored, including the use of time-averaged potentials [33], as well as spatial light modulators (SLMs) [76, 35] and digital micromirror devices (DMDs) [77], but none of these can alter the potential with sufficient bandwidth to fully replace the AODs in our system. Therefore, the AODs described above are an important and necessary component of the real-time control required for loading individual atoms,

⁸ In order to change the trap position in real-time, we needed to use a DAC voltage to control the frequency of a VCO (or, equivalently, as an input to certain function generators), which resulted in a less stable and noisier output RF frequency.

⁹ While we designed the system around an AWG that is designed (and specified) to work at 400 MS/s, the generator we are using (a National Instruments PXIe-5451) cannot actually meet these full specifications while also using triggers in scripting mode, so we have been operating at 350 MS/s as we work to find an option that will fulfill all of our specifications.

performing ground-state cooling, and observing dynamics of atoms tunneling between optical potentials. However, there are limitations (arbitrary two-dimensional control, aberration correction) that have caused us to consider alternatives for generating the final “science potential” used to observe dynamics, where we are most affected by these limitations, an dI will outline some of the basic advantages and disadvantages of these systems here.

There are two main capabilities that we do not currently have in our system that lead us to search for alternatives to AODs. First is the ability to project nearly arbitrary two-dimensional potentials onto the atoms, as have been demonstrated in other groups [76, 35, 77], which is not possible with static potentials constructed with AODs arranged in a square array (and time-averaged potentials are hard when the trap frequencies are as large as we have in our tweezers). Second, these holographic beam shaping methods immediately provide the ability to correct for nearly arbitrary wavefront aberrations in the optical system, which will distort the focus, weaken the tightness of the traps, and lead to asymmetric overlap of the optical potentials.

A liquid-crystal spatial light modulator (LC SLM) is a device that has many pixels (for the case of the one we were looking at, 512×512), each of which applies a voltage across that segment of the LC display. There are multiple ways to operate these devices, but in our case we want to operate in phase-only modulation mode, which means aligning the polarization of the light in the direction that is modulated by the birefringent liquid crystal material.¹⁰ This is important very important feature of their operation, because a small misalignment of the polarization to this axis directly leads to a small polarization rotation of the light, which couples to amplitude modulation after going through polarizing optics farther down the beam path. Importantly, one technical detail of this style of SLM is that it has to refresh the voltage on each pixel at a rate of about 1 kHz, which we found coupled to a small, but measurable, amplitude modulation at 1 kHz. Unfortunately, after carefully characterizing a test SLM with the best achieved in-situ alignment, we found that the

¹⁰ An LC SLM operated by applying a different voltage at each pixel located in a grid across a thin sheet of liquid crystal. The different voltage leads to different alignment of the LC molecules, which changes the effective index of refraction of that portion of the SLM. The modulators that we were looking into worked in reflection, where the light passed through this sheet to liquid crystal twice. The spatial variation of the index of refraction is what is then used to either holographically generate patterns or gratings, or to shift the wavefront to correct for wavefront aberrations.

residual amplitude modulation was not small enough to perform tunneling experiments; the refresh rate becomes near resonant with the trap frequency for very shallow traps and an increased atom loss rate was observed. This has led us to consider other more static options that will allow us to observe dynamics of ground-state atoms.

A digital micromirror device (DMD) is a device with many individual mirrors that can be turned on or off and which can be used to define a diffraction pattern or image, depending on the optical setup (most DMD chips come from projectors and so have standard screen resolutions, e.g. 1024×768) [77]. This overcomes the limitation of heating due to the refresh rate because the mirrors remain fixed for any given pattern. However, the efficiency of the deflection into the desired mode is very small (on the order of 5%), which means that we will need a laser that is a factor of 20 larger for any given array size. Additionally, while the DMD is relatively fast (in comparison to SLMs), the update rate is still below (or on the order of) the trap frequency of our traps, so we would only be able to use this as a static diffraction grating (that could be changed experiment-to-experiment).

3.3 Removing entropy from the system

After the loading procedure is complete, we immediately take a fluorescence image by scattering the light used for PGC cooling, which allows us to determine how many atoms have been loaded into the array [78, 79, 80, 18]. This image provides information about the number of atoms that have been loaded into an array and their locations, which can be used to post-select experiments based on the atomic configuration or even to rearrange the atoms in real-time. As discussed in Section 1.1, this is a key technique for quantum gas assembly in that we can gain information about the system that does not lead to atom loss and then can use this information during final state preparation to construct a state of interest using the information gained [41, 42].

Specifically, this process of gaining information about the configuration of atoms before performing state preparation and then an experiment with knowledge of the initial configuration, allows us to remove entropy from the system that would otherwise be inaccessible. For instance, experi-

ments that start with a Bose condensed gas and generate an array by going to the Mott insulating regime of the Bose-Hubbard model cannot measure whether there are imperfections in the array before running an experiment, which is a significant source of entropy in many of these systems. In our experiment, this means that we can measure the configuration of the atoms, then perform laser cooling to the ground state of the optical potentials, produce an array of (independently prepared) ground-state atoms that can then be studied in much the same way as Mott insulators have been studied in quantum gas microscopes in recent years [18, 19, 81, 22].

3.3.1 Building uniform atom arrays

In past experiments, the information we have gained has strictly been used to separate data taken in a single run into separate datasets; data starting with a single atom in the left well of a double-well potential versus data starting with one atom in each well, for instance [43, 47]. However, this becomes increasingly inefficient as the number of wells used in a given experiment is increased and a significant subset of the possible configurations does not provide useful data. Using enhanced loading, as discussed in Section 3.1, improves the data collection rate, but given the capabilities of our setup that have been discussed in Section 3.2, we can do better!

Specifically, we can use the information about the number of atoms available to perform different experiments depending on what we have. For instance, if we are loading 3 wells for a particular experiment (e.g., one described in Chapter 6), we can perform certain calibration experiments when only one or two of the atoms were loaded, and then only perform the “real” experiment when all three atoms are present. We can even go a step further by using the configuration information provided by the same initial image to rearrange the atoms into a specific desired configuration for different experiments, as has recently been demonstrated in one dimension in Ref. [41] and two dimensions in Ref. [42]. Importantly, this style of rearranging will allow us to start each experiment with precisely the desired configuration of atoms, as long as the required number of atoms was initially loaded.

3.3.2 Distilling motional state preparation

Aside from the entropy associated with a lack of knowledge of the presence of atoms when a system is assembled, there is also entropy associated with the imperfect ground state preparation. When planning techniques for future studies with larger arrays of atoms and, in particular, methods to gain as much information as possible about the state of our system, it became clear that we could use our ability to transfer atoms in different motional states between traps to selectively isolate atoms that are not in the ground state of the optical trap, thereby performing a “controlled distillation” of the quantum state.

While improving the ground-state preparation will always be the goal, some amount of imperfection is inevitable and my vision for future experiments is to use our ability to dynamically rearrange the optical tweezers to read out imperfect preparation. Specifically, by introducing an empty trap near a trap containing an atom that has been cooled, we can pull off atoms that are located in excited states by dynamically shifting the bias between the two wells (a process that will be described in more detail in Chapter 5). Importantly, this extra well can then be pulled off to the side and read out at the end of the experiment; an atom in one of these wells indicates that the cooling was imperfect and the result of that run of the experiment should be ignored.¹¹

In this process, depicted in Figure 3.9, we are purifying the motional state of the single atom by gaining as much information as we can. Thus, with this new technique, we will be able to improve our ability to generate indistinguishable atoms after postselection, which is important not only for the experiments described in this thesis, but also for all future experiments where quantum statistics play a significant role in the dynamics of the system.

3.4 Improving readout of the final state of the system

So far, the theme of this chapter has been to initially prepare the system as well as we can, then gain as much information as possible about the system we do prepare before performing the

¹¹ An alternative is to simply drop these extra wells after pulling off atoms in excited states. In this case, we are transferring imperfection in the cooling to effective loss of that atom, so for experiments where we want to post-select on atom survival, we lose no significant information in this process.




Figure 3.9: This is a simple schematic showing how we would be able to further purify the motional state of a single atom in a tweezer before using that atom in an experiment. By using a second tweezer to pull excited motional levels out of the original trap, we can distill the state and postselect the data from experiments where the atom was in the ground state.

experiment. Now, we turn our attention to trying to extract the most information about the state at the end of an experiment. Specifically, at the end of all experiments, the atoms are separated into resolved traps and some form of spin readout process is performed. Typically, this consists of pushing out one of the spin states and performing standard fluorescence imaging on the remaining population, but two alternative imaging methods that have been implemented will be discussed here.

3.4.1 Spin-resolved imaging

For future experiments, it will be much more important for us to be able to resolve the individual spin states at the end of an experiment compared to keeping the atoms, but only being sensitive to the population. In particular, we would like to expand the imaging sequence to take a picture of the location of all “spin-up” atoms, immediately followed by the location of all “spin-down” atoms because this will give us more information about the system compared to kicking out

all atoms that are in one spin state and imaging the location of the remaining atoms. The ability to know whether an atom was actually in the spin-up state or lost due to imperfect preparation (or a collision with a thermal atom) makes a big difference when analyzing data. Specifically, this spin-sensitive imaging provides additional information during each run of the experiment to postselect the data on runs of the experiment where known errors did not occur (such as loss or imperfect spin preparation), giving us a better signal-to-noise in the final dataset.

Aside from the benefits related to postselection, one version of this imaging method is, in principle, lossless, meaning that we can read out the spin of the atom without losing the atom from the trap [82, 83]. In this case, we also would benefit by being able to reuse atoms in several experiments before reloading, saving a significant amount of time because our loading procedure currently dominates our cycle time. There are two keys to properly implementing this method: High polarization purity to preserve the cycling transition (scattering photons only off of one of the spin states). And a large detection efficiency, since we can only scatter a few hundred photons off of the atom before it will be heated out of the trap. In our case, we require the high polarization purity during the optical pumping cycle for Raman cooling, so we can easily achieve this by adding another beam along the same optical path. Additionally, the large collection efficiency is “for free” with the high numerical aperture lens we are using to collect the fluorescence. The limitation is then only the signal-to-noise, which is limited by the readout noise of the EMCCD camera we are using to detect the collected photons.

3.4.2 Raman-cooling imaging

The work discussed in this section was detailed in Ref. [48] and much of the information below has been reproduced or adapted from this work.

Imaging ^{87}Rb via the fluorescence from standard PGC cooling light, as discussed in Section 2.1.2, is extremely efficient. The signal to noise is very high, even for relatively short images, and the temperature of the atom scattering these photons typically remains below $50\ \mu\text{K}$. However, there may be situations where it would be beneficial to keep the atom very near the ground state

of the traps (or if using a different atomic species where standard fluorescence imaging is not as efficient) and for these cases one can consider collecting and imaging the scattered photons during a period of continuous Raman sideband cooling. This work in this section was published in Ref. [48] and several groups have more recently used this in quantum gas microscope experiments [84, 49, 50].

In such situations, efficient continuous Raman sideband cooling can be used to maintain a significant ground-state fraction of trapped atoms and during the cooling process there are photons scattered that can be collected to form an image. However, the operational principle of sideband cooling relies on the existence of a dark state, composed of a spin degree of freedom and the motional ground state, that is decoupled from the cooling light. Thus, to achieve a sufficient signal to noise during a period of imaging requires some method by which the dark state is tunably compromised to allow continuous photon scattering. There are numerous methods which could be chosen, but for this work we applied a parametric modulation to the trapping potential to couple the atom out of the dark state during the imaging process. By varying the parameters of this drive, we increased the detection fidelity and speed of the imaging procedure when a smaller ground-state fraction could be tolerated.

The full optical tweezer apparatus is shown in Figure 2.5, with the most important elements highlighted in Figure 4.1(a) for implementing imaging via Raman cooling [Fig. 4.2] [43, 85, 52, 53, 86]. The efficiency of the cooling cycle is determined by our ability to optically pump back to the initial spin state while preserving the reduced motional state, which is related to the Lamb-Dicke parameter $\eta^{OP} = x_0 k$, where $x_0 = \sqrt{\hbar/2m\omega}$ is the oscillator length for a particle of mass m with trap frequency ω , and k is the wavenumber of the optical pumping (OP) light. For our optical tweezer traps in this work, $U/k_B = 1.1$ mK and the radial (axial) trap frequency is $\omega_r = 2\pi \times 140$ kHz ($\omega_a = 2\pi \times 30$ kHz), which gives $\eta_r^{OP} = 0.16$ ($\eta_a^{OP} = 0.35$) [43]. This work uses continuous Raman cooling, where the OP beams remain on during the entire procedure [39].

Each experimental cycle begins with a single atom that has been stochastically loaded by overlapping a magneto-optical trap (MOT) with an optical tweezer trap [60, 70]. After 10 ms of polarization-gradient cooling (PGC), the atom is cooled to near the bottom of the trap, with a final

temperature of $\sim 15 \mu\text{K}$ [58, 87, 43]. The trap occupancy is then measured with a standard PGC fluorescence image [78, 43], which works very well for ^{87}Rb . We use this as a pre-selection image to indicate the presence of an atom for the studies we present. After the pre-selection image, two more images are taken in each experimental cycle. The second image is the exposure during continuous Raman cooling. The final image is another PGC fluorescence image (the post-selection image) that is used in combination with the pre-selection image to distinguish experiments where the atom was lost. For all of our data analysis, the recorded signal is the number of photons detected in a single pixel on which the atomic signal is centered. This represents $\sim 1/3$ of the total detected signal. While the best SNR is achieved when all of the signal is binned onto a single pixel, we operate the camera with a smaller pixel size to fully separate the signal from atoms in different traps. However, for absolute comparison of the measured scattering rates (discussed below), we integrate the signal on the surrounding 3×3 pixels to ensure that the majority of detected photons are counted.

Figure 3.10(a) shows an example of the signal from 25 ms of PGC fluorescence imaging, which provides a useful benchmark for imaging single neutral atoms [78, 79]. For this comparison, we use a single pair of PGC beams in the $\sigma^+ - \sigma^-$ configuration [shown in Fig. 4.1(a)] with 0.61 mW of power in a beam with a gaussian waist of 1.66 mm¹². This configuration minimizes interferometric intensity gradients near the atom, which reduces fluctuations in the atomic scattering rate [58]. We red detune the PGC light 38 MHz from the trap-shifted cycling transition, which corresponds to 13 MHz from the free-space transition [73, 88, 89, 90]. Taking into account the 6% collection efficiency of our optical system, the measured scattering rate during PGC imaging is $(2.6 \pm 0.4) \times 10^5$ photons s^{-1} . This is reasonable in comparison to the estimate of 4.5×10^5 photons s^{-1} [71], which assumes perfect alignment of the peak intensity of the PGC beams to the atom, balanced circularly-polarized beams, and the ideal steady-state spin distribution for an atom at rest in one-dimensional (1D) PGC light [58].

During Raman cooling imaging, 3D cooling must be performed even if our controlled paramet-

¹² Note that the pre- and post-selection images used a six-beam PGC configuration, which we find achieves a comparable SNR.

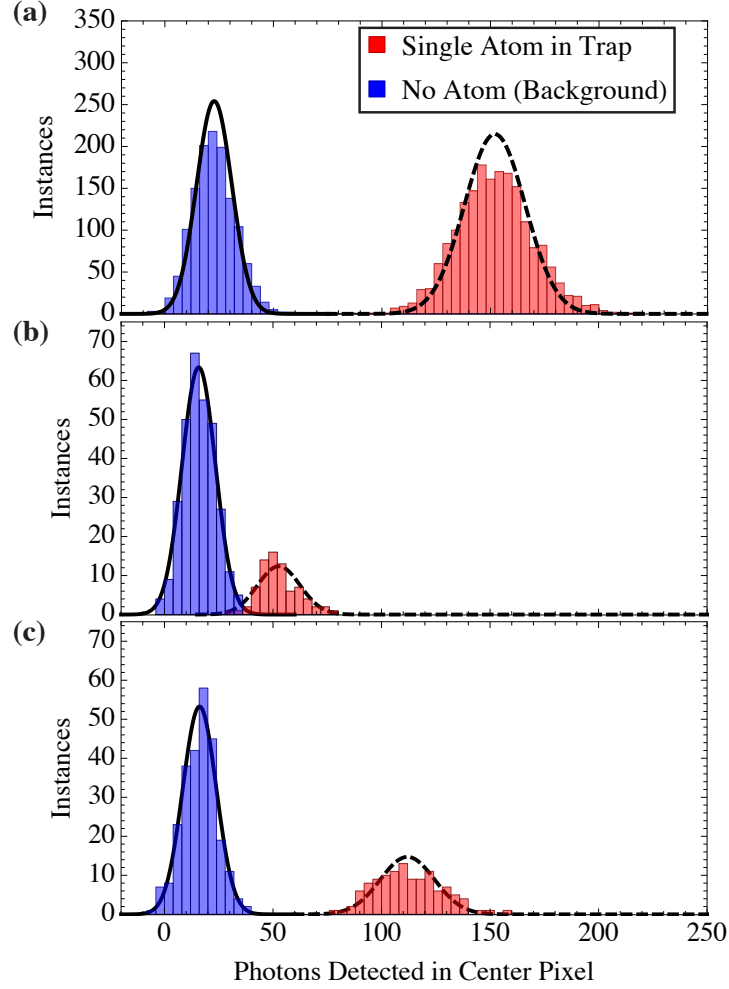


Figure 3.10: Photon detection histograms. The red (blue) bars on the right (left) indicate images taken with (without) an atom in the trap, as determined by the pre-selection image. The dashed (solid) black gaussian curve approximates the expected distribution for the atomic signal (background). (a) Fluorescence signal from 25 ms of PGC imaging. (b) Signal collected during 1 second of continuous Raman cooling. The presence of an atom is triggered using the pre- and post-selection PGC images. (c) same as in (b), but now applying a sinusoidal parametric drive with $\Delta U/U \sim 0.055$, resonant with the radial trap frequency, to the optical tweezer potential.

ric drive affects only one motional axis. This requirement arises from the optical pumping process, which can heat any motional axis via spontaneous emission. We use three beams to simultaneously perform continuous Raman cooling along both the radial (RB1+RB2) and axial (RB1+RB4) dimensions of the trap. For the radial cooling, we take advantage of the near degeneracy of the two radial trap frequencies of our trap to cool with only one beam pair and still achieve the necessary

(albeit slower) cooling along both dimensions.

Figure 3.10(b) shows a histogram of the number of photons detected during 1 second of continuous Raman sideband cooling. The red (blue) data represent experiments where an atom was present (absent) in both the pre- and post-selection images, from which we infer that there was an atom present (absent) during the entire period of continuous Raman cooling. Selecting experiments where the atom is not lost allows us to directly characterize the imaging procedure, independent of imperfections (such as the vacuum lifetime) that are specific to our setup. We note that the signal peak is smaller than the background peak because the trap loading rate was smaller than experiments in Fig. 3.10(a) (roughly 25% instead of 60%) due to a lower MOT density during loading. This does not affect the results other than to reduce the available statistics for the signal data points compared to background.

In principle, we should expect no signal when no parametric drive is applied to the trap because the atom should remain in the dark state and thus scatter no photons. The main contributions to the signal in Fig. 3.10(b) are from off-resonant carrier [$\Delta n=0$ in Fig. 4.1(b)] Raman transitions, which initiate an optical pumping cycle, and, to a lesser extent, residual coupling of the dark state to the optical pumping beams. Off-resonant carrier transitions are relevant in these experiments because the ratio of the carrier Rabi rate (the transition linewidth) to the trap frequency (the detuning from the transition) is non-negligible, leading to a finite population transfer. The measured scattering rate for the data in Fig. 3.10(b) is $(1.9 \pm 0.3) \times 10^3$ photons sec^{-1} . We estimate the maximum scattering rate from off-resonant carrier transitions to be 2×10^3 photons s^{-1} by summing the scattering rates from two 1D calculations of the axial and radial rates.

The fidelity of the imaging is determined by a combination of the signal amplitude and the ability to separate the signal and background distributions. The expected standard deviation in the number of detected photons is a combination of shot noise (Poisson statistics) and the readout noise of our detector ($\sigma_{RO} = 6.6 \frac{\text{photons}}{\sqrt{\text{pixel}}}$). The expected distributions for the atomic signal (background) are approximated by the dashed (solid) gaussian curves on the histograms in Fig. 3.10. Based on the overlap integral of two normalized gaussian distributions with means and standard deviations

from the data in Fig. 3.10(b), we estimate a minimum detection error rate of 1.4×10^{-2} for 1 second of integration time [82]¹³. Note that this is a measure of the achievable SNR, but the effects of atom loss are removed based on the pre- and post-selection images.

To improve the signal detected during continuous Raman cooling, we apply a parametric drive by modulating the total optical power generating the optical tweezer trap at twice the radial trap frequency ($2\omega_r = 2\pi \times 280$ kHz). This drives the atom out of the motional ground state, allowing the cooling cycle to restart, and thus more OP photons are scattered during the exposure. Figure 3.10(c) clearly shows an increase in the atomic signal during Raman cooling imaging with a parametric drive applied to the trapping potential. The increased signal corresponds to a reduction of the minimum detection error rate to $< 10^{-4}$ for the same 1 second integration time. For comparison, we estimate that 20 ms of our PGC imaging, which has a much higher scattering rate, can achieve the same minimum detection error rate.

The Raman cooling imaging procedure can be simply understood as the competition between a variable parametric excitation rate and the (constant) continuous Raman cooling. The competition of these rates rapidly yields an equilibrium ground-state fraction, and hence non-dark fraction, which largely determines the measured scattering rate. This simple description is effective because the parametric drive does not directly interfere with the cooling and the Raman cooling cycle occurs on sub-millisecond timescales compared to the 1 second image.

The average number of photons detected during Raman cooling imaging is observed to increase as a function of the applied parametric drive amplitude [red circles in Fig. 3.11(a)]. Hence, increasing parametric drive amplitude reduces the equilibrium ground-state fraction, allowing us to control the atomic fluorescence rate. The atom signal can be compared to the average number of background photons detected (blue squares) measured at each drive amplitude. The amplitude of the parametric drive is defined as the ratio of the peak-to-peak variation of the trap depth ΔU to the trap depth in the absence of a drive U .

¹³ The minimum detection error rate δ is half the overlap integral of the two normalized gaussian distributions, which requires an optimal choice of threshold to achieve. The fidelity of the measurement is then defined to be $\mathcal{F} = 1 - \delta$ (i.e., $\delta = \epsilon_B = \epsilon_D$ in Ref. [82]).

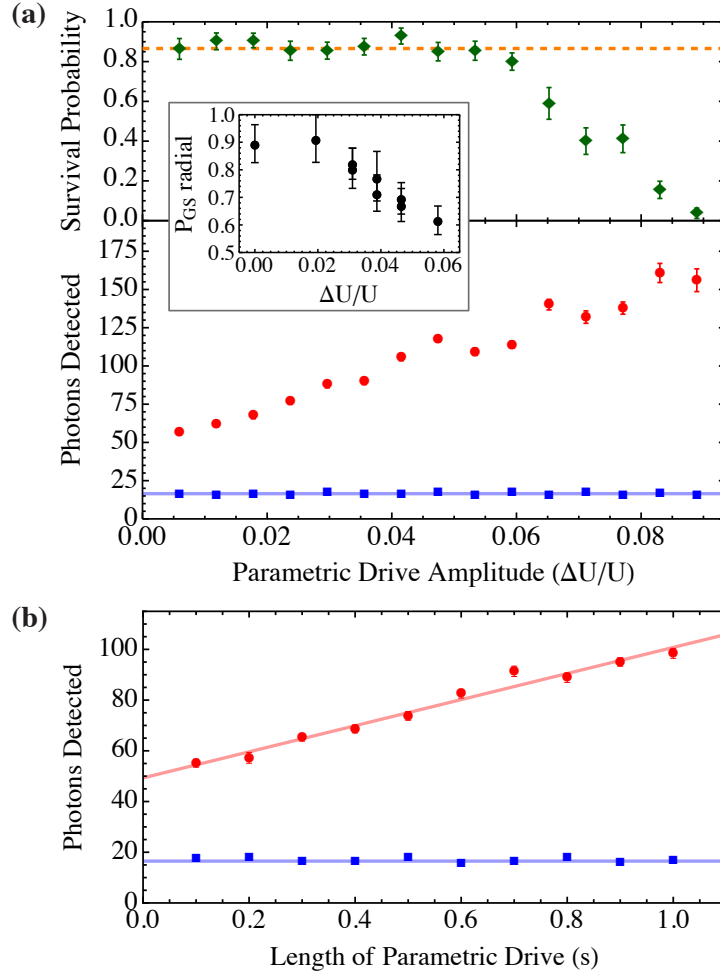


Figure 3.11: (a) Varying the amplitude of the parametric drive applied to the optical tweezer potential during 1 second of Raman cooling imaging. The green diamonds are the atom survival probability, as determined using the pre- and post-selection images, and the orange dashed line is the measured vacuum lifetime limit. The red circles (blue squares) are the signal collected with (without) an atom in the trap during the entire imaging period. The blue line indicates the average background value of 16.5 detected photons. The inset displays the equilibrium radial ground-state fraction (P_{GS} radial) measured via Raman sideband spectroscopy immediately following a period of Raman cooling imaging [43]. (b) Signal as a function of the parametric drive length, for fixed total imaging length (1 s) and drive amplitude ($\Delta U/U \sim 0.04$). The red line is a linear fit to the data. All error bars indicate the standard error of the measurement.

Up to intermediate drive amplitudes, we characterize the equilibrium radial ground-state fraction during imaging by performing radial sideband thermometry immediately after a period of Raman cooling imaging. We verify that the equilibrium ground-state fraction varies with the applied

parametric drive amplitude [inset of Fig. 3.11(a)]. For the lowest drive amplitudes ($\Delta U/U \lesssim 0.02$), thermometry is limited by the resolution of our sideband spectroscopy [43], suggesting a radial ground-state fraction $P_{GS} > 80\%$. We further find that a significant ground-state fraction $P_{GS} \geq 55\%$ is maintained for drive amplitudes up to $\Delta U/U \sim 0.06$ (near the maximum drive without loss). The scattering rate observed at this drive amplitude is $(4.9 \pm 0.7) \times 10^3$ photons s^{-1} . We estimate an expected scattering rate of $\leq 7 \times 10^3$ photons s^{-1} by weighting the scattering rates from each of the motional excited states undergoing Raman cooling, assuming a thermal distribution corresponding to the measured motional occupation at $\Delta U/U \sim 0.06$ [43].

Further, the survival probability of the atom is measured as a function of the parametric drive amplitude [green diamonds in Fig. 3.11(a)] by recording the probability for an atom to appear in the post-selection image, given that it was present in the pre-selection image. The orange dashed line represents the expected survival probability due to the vacuum lifetime ($\tau = 7.1$ s) over the course of the 1.023 seconds between the pre- and post-selection images. The atom loss is initially consistent with background collisions, but at larger drive amplitudes the parametric excitation rate exceeds the Raman-cooling rate and the atom is quickly heated out of the trap.

In Fig. 3.11(b), we perform 1 second of continuous Raman cooling and apply a parametric drive for a fraction of the image. The constant drive amplitude of $\Delta U/U \sim 0.04$ realizes an equilibrium radial ground-state fraction of $\geq 70\%$ [inset of Fig. 3.11(a)]. The measured data (red circles) are well fit by a linear increase (red line), which is further evidence that the radial ground-state fraction rapidly reaches equilibrium. For the same drive amplitude, we estimate that a minimum detection error of $< 10^{-2}$ can be reached after about 0.5 seconds of imaging. This is comparable to the site-resolved imaging used in recent optical lattice experiments [18, 19].

Chapter 4

Preparing and Interfering Indistinguishable Atoms

In Chapter 3, we reviewed the methods by which we can generate an array of atoms for use in an experiment, as well as ideas for how to read out these atoms efficiently, gaining as much information as possible about the initial state prepared and the final state after an experiment for use in the data analysis. However, we left out a crucial step for performing real experiments: Bringing the atoms to the ground motional state of the traps before initiating dynamics. In this Chapter, I will review the basics of Raman sideband cooling and how we have implemented it to cool single atoms trapped in optical tweezers [43]. The Raman cooling is an important technical aspect of the experiment, for which I provide details here that are not necessarily in any of our other published work, though reviewed some in Adam Kaufman's thesis [57]. This work led to the observation of two-particle quantum interference of independently prepared atoms, verifying the ability to make atoms indistinguishable via laser cooling alone. This work is described in Ref. [39, 57] and, hence, I do not cover the details here. However, I do briefly discuss the application of our ability to measure the indistinguishability of two atoms to improving our initial state preparation.

4.1 Raman sideband cooling

Reliably bringing atoms to the motional ground state of a trapping potential is a hard problem. In many ultracold atomic gas experiments, a large cloud of atoms is trapped, cooled to the limit of traditional laser cooling techniques (in the few μK regime), and then initiating some form of evaporative cooling reach sufficiently small temperatures for Bose condensation (or the genera-

tion of a degenerate Fermi gas) with the remaining atoms. While this is effective and has led to significant progress in the field of ultracold gases, it is inefficient both in terms of the number of ground-state atoms versus number loaded and in terms of the length of the cooling process, and is not amenable to quantum gas assembly as envisioned in Chapter 1 [1, 2, 3, 4, 12, 91, 92, 18, 19].

For cooling single atoms to the ground state of tightly confining traps, Raman sideband cooling has been extremely successful, and quite generally applicable, being used for trapped ions [93], atoms in optical lattices [52, 53, 54, 56, 86, 94], optical cavities [95, 96, 97], and optical dipole traps [43, 44]. The operational principle is to drive transitions that manipulate both the spin and the motional degrees of freedom, decreasing the motional quantum number in each step, and eventually reaching a state that is dark to the cooling light, which occurs in the motional ground-state of the trapping potential. Below, I will briefly review the form of pulsed Raman sideband cooling that we use in our system, but there are many other forms of Raman sideband cooling, such as continuous sideband cooling and projection sideband cooling, in addition to a large amount of detailed derivations that I am not going to repeat in this thesis because they are covered very thoroughly elsewhere, in particular in Refs. [98, 99, 57].

4.1.1 Non-degenerate pulsed Raman sideband cooling

First, note that much of the content of the next two sections is detailed in Refs. [57] and [43], but certain details have changed since the original publication and are discussed here. Here, we focus on the current implementation of a nondegenerate, pulsed Raman cooling scheme. Briefly, this is a two-step process that alternates between driving a coherent, energy-reducing spin-flip and then driving spontaneous transitions to bring the atom back to its original spin state, while maintaining the reduced motional level. By repeating such a procedure many times, the population accumulates in the ground state of the trap with the original spin state.

The key component of this procedure is the ability to drive a coherent, energy-reducing spin-flip; this is achieved using a two-photon Raman process, where two far-off-resonant beams have a frequency difference that (nearly) matches the microwave transition frequency for the spin-flip

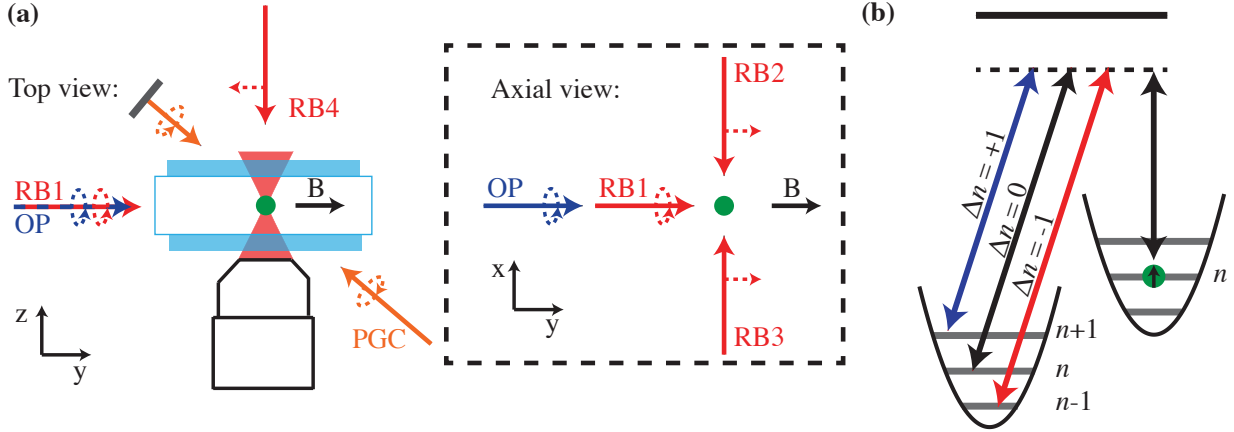


Figure 4.1: a) The content of this diagram is contained in Figure 2.5, but here we highlight the beams used during Raman cooling (RB1-RB4, and OP), as well as the beam used for both sub-Doppler cooling and imaging (PGC). The PGC beam is used to initially cool the atoms such that the temperature of the atoms is amenable to Raman cooling. The polarizations of each of these beams is also indicated in the figure with a dashed line (or circle). b) Diagram depicting two ground-state spin levels the beams used to perform a Raman transition. The black vertical arrow represents one of the linearly polarized beams (RB2-RB4), and is far-detuned from the excited state energy level. The three arrows on the right indicate three possible transitions selected by the σ^+ Raman beam (RB1), all of which flip the spin of the atom, but address transitions between different motional states (red for $\Delta n = -1$, black for $\Delta n = 0$, and blue for $\Delta n = +1$).

that we want to drive. These two beams then cause the atom to undergo a virtual process, where a photon is absorbed from one beam and emitted into the other, changing the spin-state of the atom in the process. As depicted in Figure 4.1(b), by adjusting the frequency difference between the beams, we can simultaneously change the motional state of the atom in the trap, meaning that for particular frequency differences, we drive coherent, energy-reducing, spin-flips that are the first step in the Raman cooling procedure. Then, the spontaneous process that brings the atom back to its original spin state is simply the optical pumping that was described in Section 2.1.3.

Figure 4.1(a) highlights the relevant beams for performing the Raman transitions in our setup. In particular, we highlight that the optical pumping beams (OP) are aligned along the quantization axis and are copropagating with the Raman beam, RB1, which is used when addressing all three axes (this ensures that all three beams have the same high polarization-purity). The remaining three Raman beams (RB2, RB3, and RB4) are all aligned orthogonal to RB1 and the three pairs are




Figure 4.2: This is a schematic representation of the pulsed Raman cooling procedure. a) Along a single axis, this shows the pulse sequence and how the population of atoms is brought to a lower temperature. by repeating a single Raman cooling cycle repeatedly (each two-step sequence is highlighted by the dashed box around it). b) Pulse sequence demonstrating how we sequentially cool all three axes (each completing a full Raman cycle before the others have a chance to accumulate population in the three dimensional ground state of the trap).

used to address three orthogonal axes of the trap. Because these beams are pointing along different axes, there is a net momentum kick imparted to the atom during a two-photon Raman transition, which then allows us to couple different motional states of the atom. Assuming a nonzero total momentum kick, then we can selectively address transitions that change the motional state by a known number of quanta (typically either $\Delta n = \pm 1$ for the blue and red sidebands, or 0 for the carrier transitions) while flipping the spin of the atom, as depicted in Figure 4.1(b).

A schematic representation of the full cooling procedure is shown in Figure 4.2(a). After loading of the optical tweezer with a single ^{87}Rb atom, sub-doppler cooling is performed to lower the initial temperature of the atom such that the atom starts near the bottom of the trap; this is important for ensuring that the cooling procedure is effective, as discussed in Section 4.1.2. Then, before the Raman cooling procedure is initiated, optical pumping is performed to set the spin state of the atom to the $|F, m_F\rangle = |2, 2\rangle = |\uparrow\rangle$ Zeeman sublevel. Then, a pair of Raman beams (RB1 plus RB2, RB3, or RB4) is used to drive a coherent rotation of the spin to the $|1, 1\rangle = |\downarrow\rangle$ Zeeman sublevel, while changing the motional level by $\Delta n = -1$ (to be explicit, this is a Raman transition between the states $|\uparrow, n\rangle \rightarrow |\downarrow, n-1\rangle$), as shown in Figure 4.2. The spin state is then optically pumped back to the original spin state $|\uparrow\rangle$, while preserving the reduction in the motional state,

and the final two steps can be repeated many times to ensure that the atoms end up in the ground state at the end of the process. Note that this procedure is shown for a single motional axis, which is determined by the choice of the pair of Raman beams and the relative frequency difference, but we actually switch between the different axes for each cycle to bring the atom to the three-dimensional ground state.

Near the end of a successful Raman cooling procedure, the majority of atoms are in the “dark state”, meaning that the atom is not affected by the any of the light applied. Specifically, when applying light for a Raman transition that is designed to reduce the motional state by one quanta during the spin flip, then no resonant transition will occur when you are already in the ground state of the trap. Additionally, because of the design of the optical pumping process, the $|\uparrow$ spin state is already dark to the optical pumping light. Therefore, the state $|\uparrow, 0\rangle$ remains the dark state for the entire Raman cooling procedure.

4.1.2 Efficient cooling in optical tweezers

One of the most important considerations for the successful implementation of ground-state Raman sideband cooling is to have sufficiently strong harmonic confinement of the atom (or at least approximately harmonic for the relevant energy scales). Specifically, this confinement is needed in order to resolve transitions between individual motional states, as well as to ensure the scattering of a small number of photons (during optical pumping) will not result in the atom’s motional state being increased more often than it is reduced. The parameter describing used to describe the confinement of a trap is called the Lamb-Dicke parameter [100]. The Lamb-Dicke parameter is defined as $\eta = k_r x_0$, with k_r being the recoil momentum imparted to the atom by the light addressing a given transition and x_0 the harmonic oscillator length of the ground state wavefunction. Equivalently, this can be written as $\eta^2 = \frac{\omega_r}{\omega_0}$, with ω_r the recoil frequency (where the energy gained by the atom emitting a photon is given by $E_r = \hbar\omega_r$) and ω_0 is the trap frequency of the harmonic potential.

Note that the probability to preserve the reduced motional state is dependent on the average

number of photons scattered during the optical pumping stage, the Lamb-Dicke parameter η^{OP} for scattered optical pumping photons, and the motional state that the atom is in when scattering photons. There is a fundamental minimum to the number of photons scattered during optical pumping that is set by the decay pathways from the excited state. Conservatively, we assume 3 scattering events per optical pumping cycle, which would say that if we increase the energy level of the atoms less than once every three scattering events, then we will have a net cooling effect during each cycle. Importantly, if there are no mechanisms for heating or scattering out of the dark state, then, if a net cooling is achieved, arbitrarily high ground-state fractions can be generated with a long enough period of cooling.

However, there is always some limit to the fidelity of the dark state, either from a mechanism that can directly heat the atom out of the ground state, or due to the polarization of the optical pumping light not being perfectly σ^+ and causing an undesired spin flip. Heating of the atom can come from the elastic scattering of photons, possibly from the trap or Raman light, or from intensity noise modulating the trap at a frequency resonant with motional state changes. In either case, these will never be completely gone, but can be minimized with relatively standard techniques. On the other hand, preventing unwanted spin flips from the optical pumping light is more challenging because this light is designed to bring the atom back to the $|\uparrow\rangle$ spin state quickly, and therefore has a relatively high scattering rate of $\Gamma_{\text{OP}} \sim 100$ kHz. The “darkness” of the dark state then depends on our ability to prevent unwanted scattering events, which relies on high polarization purity. To be specific, from the $|\uparrow\rangle = |2, 2\rangle$ spin state, the OP light (from Figure 2.1) would be resonant with transitions to both the $|F', m'_F\rangle = |2, 2\rangle$ and $|2, 1\rangle$ states, if the polarization were not purely σ^+ . As discussed in Section 2.1.3, we are able to achieve a polarization purity to better than a part in 1000, which means the scattering rate out of the dark state is less than 100 Hz.

As mentioned above, the probability of the atom heating during a photon scattering event is also dependent on the motional state of the atom, which we can quantify by defining an effective Lamb-Dicke parameter as $(\eta_{\text{eff}}^{\text{OP}}) = (2\bar{n} + 1) (\eta^{\text{OP}})^2$. This modification simply means that to achieve efficient Raman sideband cooling, we need to start with a low thermal occupation before initiating

the Raman cooling process which the in-trap PGC cooling provides.

The coupling between the motional states is also parameterized by a Lamb-Dicke parameter, but in this case the momentum is replaced by Δk , the amount of momentum imparted to the atom during a Raman transition (the difference in the momenta of a photon absorbed from one beam and emitted into the other). Therefore, the Raman Lamb-Dicke parameter is defined as $\eta^R = \Delta k_R x_0$. This parameter determines the ability to drive sideband transitions with a given set of beams and, along with the motional occupation, defines the relative rate at which sideband transitions will be driven in comparison to the carrier ($\Delta n = 0$) transition.¹ The strength of this coupling will affect the spectra resolution of the sidebands from the carrier transitions because it changes the ratio of the Rabi rates, and thus the widths of the spectral features. This also affects the final cooling efficiency because a reduction in the spectral resolution leads to a higher probability of driving an off-resonance carrier transition (instead of a sideband). This directly leads to the scattering of optical pumping photons without the benefit of reducing the motional state first, directly leading to heating of the atom.

This highlights another important concern for maximizing the efficiency of the cooling procedure, in that improving the spectral resolution between the sidebands and the carrier is beneficial. Specifically, there are two things we have changed that improved this in our setup: Changing the alignment of the axial Raman beam (thus *increasing* the Lamb-Dicke parameter η^R) and using Gaussian shaped pulses to narrow the spectral extent of each pulse.

4.1.3 Raman sideband spectroscopy

After performing a sequence of Raman sideband cooling, we can estimate the motional state occupancy (or, more specifically, the ground-state fraction) by performing Raman sideband spectroscopy. In this procedure, we perform a π -pulse for the blue sideband (assuming the sideband Rabi rate associated with atoms from the ground-state) while varying the relative frequency be-

¹ This also has some significant consequences, where atoms with too high a motional occupation will undergo a 2π rotation rather, getting some fraction of the population stuck in higher harmonic oscillator levels. In the end, this enforces a similar requirement that we need to start with a low enough initial temperature.




Figure 4.3: These plots show (a) the radial and (b) the axial Raman sideband spectrum taken after a three-dimensional cooling process that has been optimized by improving the indistinguishability of the atoms, as discussed in Section 4.2.2. The dashed black lines indicate the expected location of the red (energy reducing) sideband based on the measured trap frequencies. The lack of a red sideband (on the right) indicates a large ground state fraction. In (b), the central peak is the overdriven carrier transition, which adds noise that partially overlaps with the location of both the sidebands.

tween the two Raman beams. The resulting spectra are shown in Figure 4.3 for the radial and axial dimensions after the cooling procedure. In both of the spectra, the energy-reducing red sidebands are not noticeable, suggesting that the atoms are in the ground-state of the trapping potential. The large asymmetry between the sidebands (and especially the absence of the red sideband) is suggestive of the large ground state fraction after this cooling process.

4.2 Quantum interference of indistinguishable particles

The details of our experiment that led to the observation of two-particle quantum interference are detailed extensively in Refs. [39] and [57], hence I will only briefly introduce the experimental sequence and relevant observable before discussing how we have more recently used the results of this experiment as a tool for optimizing our initial preparation procedure.

4.2.1 Measuring indistinguishability via tunneling in a double well potential

A schematic of the tunneling procedure is shown in Figure 4.4, along with a representative set of recent single-particle tunneling data that has been post-selected on the single atom not being

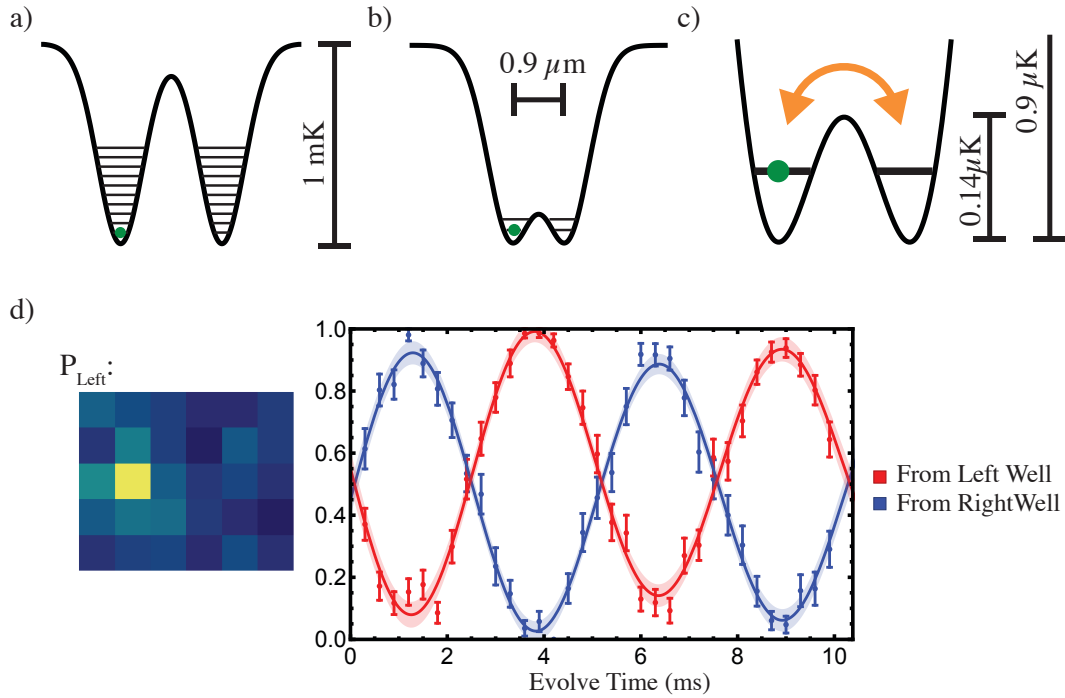


Figure 4.4: Schematic of the tunneling procedure used to observe interference between indistinguishable atoms. a) Atoms are loaded into $\sim 1 \text{ mK}$ deep traps that are separated by $1.57 \mu\text{m}$ to perform Raman cooling to the motional ground-state. b) The traps are dynamically reconfigured to both partially overlap and be shallow enough that there is a tunnel coupling between the sides of the double well potential, in the example shown this means a $0.9 \mu\text{m}$ separation and $0.9 \mu\text{K}$ single-well depth. c) The atom is allowed to evolve in this final trap configuration for a period of time, during which the atom can coherently tunnel between the two wells. d) The traps are separated and imaged to measure the population in each well. An example of data taken with the trap parameters listed above is shown, where the probability for the atom to end in the left well is shown as a function of time for a single atom initially in the left (red) and right (blue) wells.

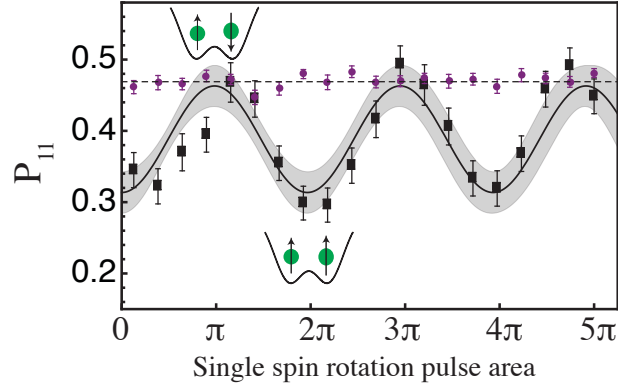


Figure 4.5: In this figure, we observe two-particle quantum interference as a function of the distinguishability of the two atoms. The black points on this plot show the measured probability for two atoms to end up in separate wells after tunneling for a period that splits each atom into both wells. The two-particle quantum interference is the reduction of this probability below the purple points. The purple points show the expected probability for distinguishable atoms to end up in separate wells, calculated from the single particle data (as shown in Figure 4.4), which would ideally go to 0.5. In this figure, we are specifically varying the spin state of one of the atoms before this tunneling procedure occurs. When the atoms are in opposite spin states, they are then distinguishable and there is no interference “dip” below the distinguishable limit, but when they are the same spin state, we do see this interference.

lost before the end of the experiment (as discussed in Section 3.4.1). The coherent tunneling of single atoms between the wells forms the basis of our observation of quantum interference between two atoms. In this sequence, two optical tweezers are stochastically loaded when separated and the trapped atoms are prepared in the motional ground state, as described in Section 4.1.1, before repositioning the wells to initiate a period of tunnel-coupling. After a variable period of evolution, the wells are separated and the population measured.

When two atoms are initially loaded, one in each of the two tweezers, then they will simultaneously experience a tunnel-coupling between the two wells, and periodically both atoms will be split evenly between the two wells. However, because they started in opposite wells, the phases between the two atomic wavepackets will be different, which leads to quantum interference that results in the atoms always being observed in the same well when the atoms are indistinguishable in all other degrees of freedom [101, 39, 57, 102]. Figure 4.5 demonstrates the change in the nature of the interference between two atoms split evenly between the two wells, as the distinguishabil-

ity between the two atoms (in the form of the initial spin state) is tuned via site-resolved spin addressing.

4.2.2 Using indistinguishability to improve state preparation

As demonstrated in Figure 4.5, the depth of the two-particle interference between particles evenly split between the two wells of a double-well potential is a good measure of the distinguishability of the two atoms. Therefore, because the goal of our initial state preparation is to put all atoms in the motional ground state of their respective traps, we can use this measure of the distinguishability to optimize our state preparation. This is equivalent to the experiment demonstrated in Figure 4.5, except that instead of purposefully rotating one of the spins (and not the other), we will change the global parameters used during state preparation, such as the cooling frequencies and Raman beam powers, and choose to use the values which give a minimum in the interference.

An example of the signal from this optimization procedure, as applied to the parameters used during the Raman cooling procedure, is shown in Figure 4.6. Specifically, for this run, the axial Raman beam frequency used during cooling was varied to find where the axial cooling rate was maximal. After tuning this and several other cooling parameters, we achieved the sideband spectra shown in Figure 4.3.




Figure 4.6: In this figure, we are observing the quantum interference between two atoms as we vary the axial Raman beam frequency used during the initial state preparation. We can clearly see that the interference is maximized for a particle frequency for Axial sideband cooling. This allows us to set the exact frequency used to the best of my abilities, and is often the best signal around for optimizing any given Raman cooling parameter.

Chapter 5

Entangling Independently-Prepared Neutral Atoms

We have already seen that the quantum statistics for identical particles can have drastic effects on the dynamics of indistinguishable atoms in a double well potential. However, in this chapter we will study effects in a system where we purposefully prepare distinguishable atoms and allow them to interact in a single tweezer. In particular, we observe how the quantum statistics of the atoms imposes additional restrictions on the total quantum state, which leads to the correlated degrees of freedom will become coupled in such a way that they can no longer be described independently [103]. Specifically, by preparing two atoms in different motional states of a single trap and with opposite spin states, the symmetrization of the total wavefunction will lead to spin-exchange dynamics, which periodically generated spin-entanglement between the two atoms [39].

We note that this type of coupling forms the basis of many exciting areas of research, including quantum information protocols with quantum dots (where dynamics similar to those observed here are used to entangle and manipulate qubits), to effects observed in unconventional superconductors, which will be discussed further in Chapter 6. The information presented here thus lays the foundation for the discussions in Chapter 6, but certain details will be left out and can be found in Ref. [47] and [57].

5.1 Dynamical entanglement via spin-exchange

The work discussed in this section was published in Ref. [47] and for any additional information see discussions there, as well as in Ref. [57].

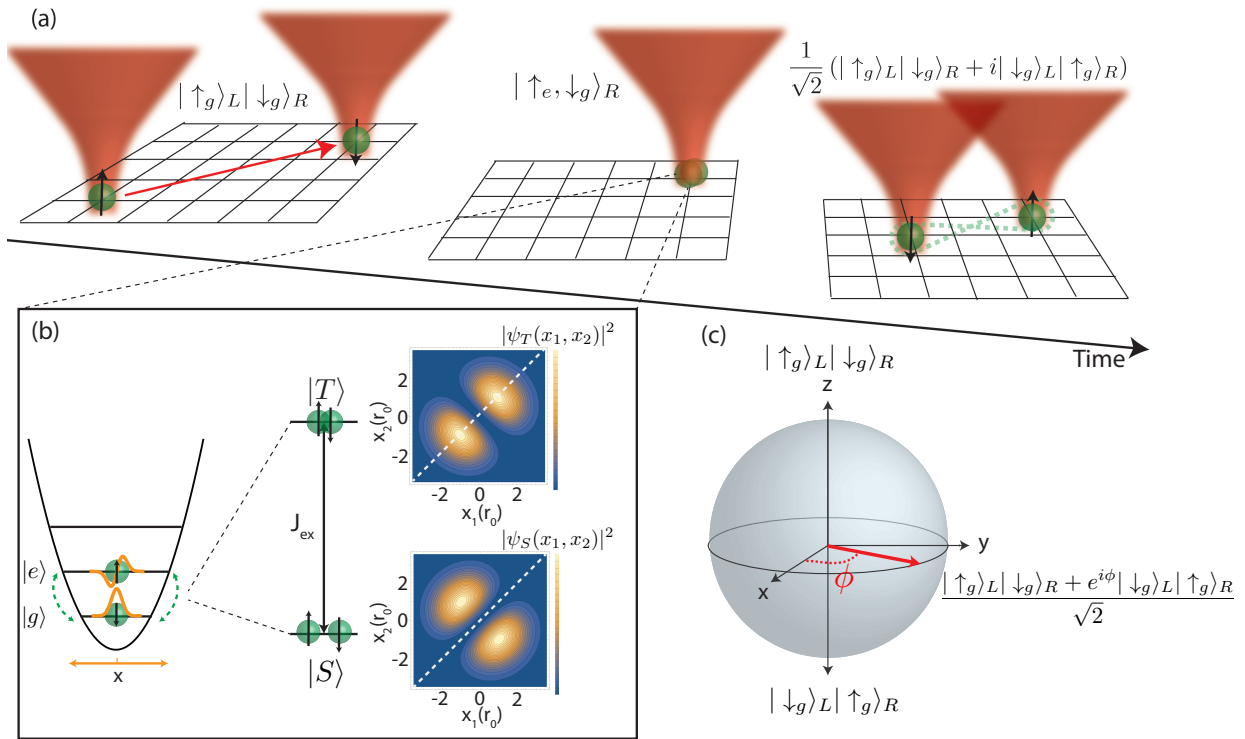


Figure 5.1: a) Starting on the left, two atoms are loaded into separate optical tweezers and prepared in their motional ground state and in opposite spin states. Then, the tweezers are reconfigured to b) The on-site contact interactions shift the energy of triplet spin-state (which will be in a symmetric spatial wavefunction), but due to the symmetrization of the total wavefunction, there is no contact interaction in the singlet state (which must be in an antisymmetric spatial configuration). The differential phase accumulation between these spin states thus leads to the spin exchange dynamics that we wish to observe. (Note that this figure is adapted from Ref. [47].)

In this section, I will discuss our demonstration of the use of local interactions to entangle the spins of two atoms that are initially trapped in separate wells. When the two atoms are placed in the same well, the symmetry of the two-particle spatial wavefunction determines the interaction energy of the particles. However, assuming the atoms are identical in all other degrees of freedom, the spatial symmetry of the two-particle wavefunction directly determines the spin-symmetry, which means that the singlet and triplet spin wavefunctions will experience different interaction strengths, as shown in Figure 5.1. It is this difference in interaction strength that leads to a differential phase accumulation between the singlet and triplet spin states, resulting in the exchange of spin between the two atoms.

In the experiment, we can observe this effect by preparing an initial state with two atoms in two optical tweezers, bringing both to the ground state, and then flipping the spin of one of the atoms. The remainder of the procedure is depicted in Figure 5.2, where we sweep the wells together in an asymmetric configuration such that we are not resonant with tunneling, and slowly reconfigure this tilt to couple the ground motional state of the left well to the first excited motional state of the right well [Fig. 5.2], adiabatically transferring the atom into the right well in a process I will refer to as an “adiabatic passage”. After a period of evolution, this process is reversed to separate the two atoms, and the spin state of the atoms in the two separate wells is measured.

The observed dynamics are shown in Figure 5.2, exhibiting anti-correlated oscillations between the $|\uparrow\rangle_e |\downarrow\rangle_g$ (green) and $|\downarrow\rangle_e |\uparrow\rangle_g$ (purple) spin configurations. We also verify that, as we vary the depth of the well during the exchange period (and only during this hold to ensure that the remaining process is identical), the oscillation rate changes in a way that is consistent with the expected Hubbard parameters [Fig. 5.2(c)], as calculated using a full three-dimensional model of our system, as discussed in Refs. [39] and [104].

These measurements show correlations in single-particle spin states. However, to use the entanglement generated in this process for future quantum information protocols, we must verify that the phase of the entangled state is preserved after separation of the particles. The entanglement verification protocol for separated atoms is summarized in Figure 5.3(a). For explanatory

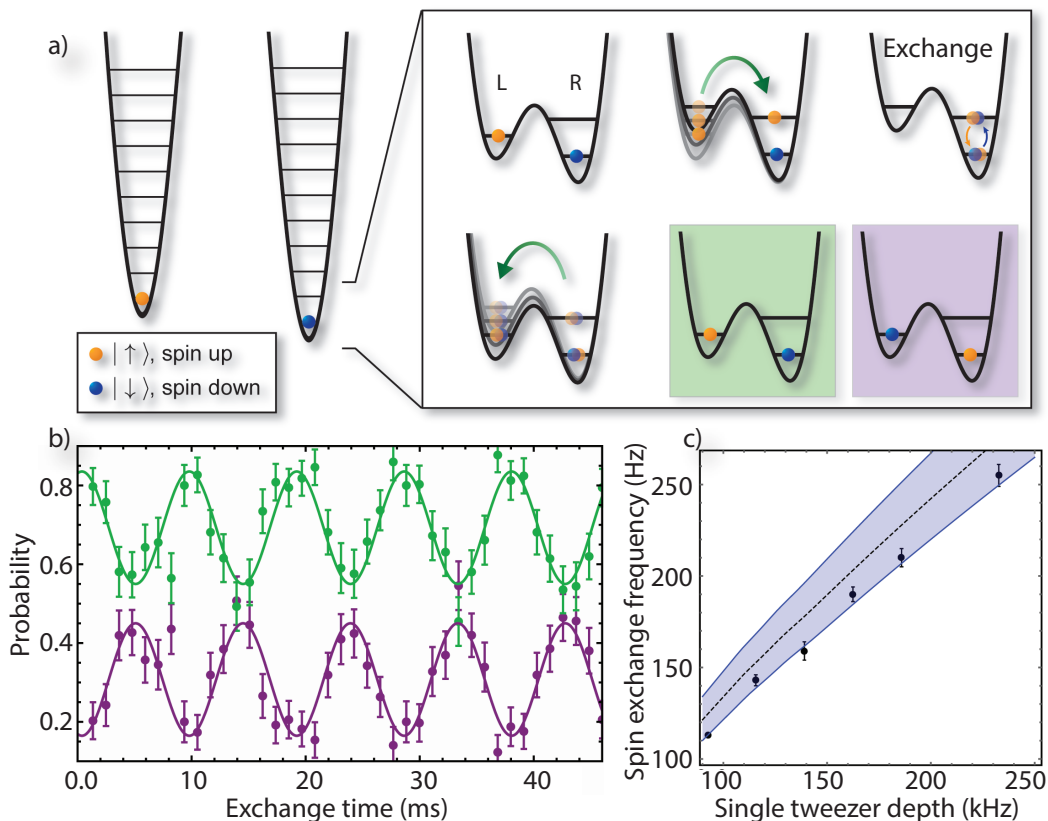


Figure 5.2: a) Schematic representation of the experimental procedure for observing spin exchange oscillations. Starting with atoms in separate wells and in opposite spin states, we then proceed by rearranging the traps to couple the atom from the ground-state of the left well to the first excited state of the right well. Then, we let the system evolve for some period of time before separating the wells and observing the final spin configuration of the system. b) Time-evolution of the spin configurations over time. We see the anti-correlated oscillations of the spin states that is indicative of spin-exchange dynamics. c) Varying the depth of the trap during the period of evolution, we see that the oscillation rate varies in agreement with the theory calculated for the trap parameters being used. (Note that this figure is adapted from Ref. [47].)

purposes, we first focus on the case when the particles are separated after an exchange time of $t_{\text{ent}} = n\pi/2J_{\text{ex}}$ where n is an odd integer. The entangled state after the second adiabatic passage is $|\psi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|\downarrow\rangle_L|\uparrow\rangle_R \pm i|\uparrow\rangle_L|\downarrow\rangle_R)$, omitting from now on the ground-state (g) motional subscripts to simplify notation. The $|\psi_{\pm}\rangle$ states correspond to the gray and orange Bloch vectors, respectively, in Figure 5.3(b). We then apply a magnetic-field gradient that imposes a difference, $\delta\hbar$, in the $|\uparrow\rangle \leftrightarrow |\downarrow\rangle$ single-atom-transition energy between the left and right optical tweezer. By applying the gradient for a time t_g , a transformation $|\psi_{\pm}\rangle \rightarrow \frac{1}{\sqrt{2}}(|\downarrow\rangle_L|\uparrow\rangle_R \pm ie^{i\delta t_g}|\uparrow\rangle_L|\downarrow\rangle_R)$ is achieved. As a function of t_g , the state rotates between the singlet [pink in Figure 5.3(b)] and triplet (blue) with frequency δ . We then apply a global $\pi/2$ pulse in the $\{|\uparrow\rangle, |\downarrow\rangle\}$ sub-space. This pulse maps the singlet back to itself, while it maps the triplet to a Bell state $\frac{i}{\sqrt{2}}(|\uparrow\rangle_L|\uparrow\rangle_R + |\downarrow\rangle_L|\downarrow\rangle_R)$. Therefore, by measuring the probability that the spins are aligned or anti-aligned as a function of t_g , we can observe singlet-triplet oscillations whose amplitude characterizes the two-particle coherence. We quantify this probability with the parity $\Pi(t_g) = \sum_j P_j(-1)^j$, where P_j is the likelihood to measure j atoms in the spin-down state [105, 106, 27, 26]. The parity is equivalently the projection of the Bloch vector in Figure 5.3(b) onto the x -axis of the effective two-particle Bloch sphere (in the $S_z = 0$ subspace) prior to the $\pi/2$ -pulse, and hence the gradient is essential because, though entangled, the states $|\psi_{\pm}\rangle$ (gray, orange) exhibit zero parity after application of a $\pi/2$ -pulse.

We demonstrate the outcome of the verification protocol on the state $|\psi_+\rangle$ in Figure 5.3(c). We plot $\Pi(t_g)$ after the microwave $\pi/2$ pulse, and observe oscillations in the parity signal as the gradient time t_g is scanned. The contrast of these oscillations is consistent with what is expected given the exchange oscillation contrast in Figure 5.2, and non-vanishing parity oscillation would certify entanglement in the ideal case of perfect spin preparation. However, we have imperfect spin preparation, and the erroneous spin populations outside the $\{|\downarrow\rangle_L|\uparrow\rangle_R, |\uparrow\rangle_L|\downarrow\rangle_R\}$ manifold could lead to parity oscillations even in the absence of entanglement. We have derived a condition on the parity oscillation contrast that is necessary and sufficient to certify entanglement and is experimentally the most accessible way to see there is entanglement in our system (for a full derivation of this condition, see Ref. [47]). We relate the measured parity contrast, C , to the

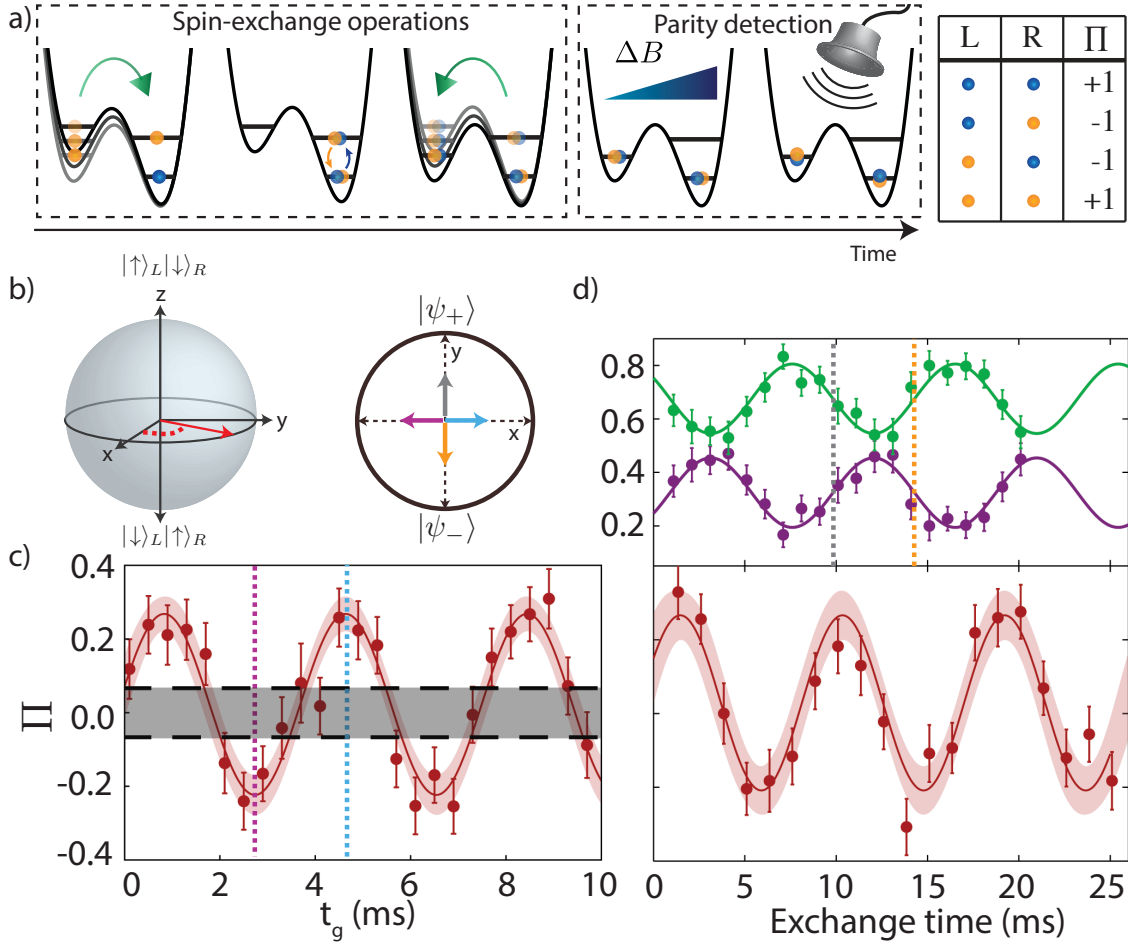


Figure 5.3: a) Procedure for performing spin exchange, followed by the entanglement verification procedure, which consists of apply a magnetic field gradient to the separated atoms and then performing a spin parity readout sequence. b) Two-particle Bloch sphere representing the $S_z = 0$ subspace of the spins. The eigenstates of the spin-exchange interaction are along the x -axis (singlet and triplet), while the eigenstates in separate wells and in the presence of a gradient are along the z -axis. Then, spin exchange becomes a σ_x rotation, while the gradient evolution is a σ_z rotation. Then, upon application of a microwave $\pi/2$ pulse, we map the projection of the Bloch vector onto the x -axis to parity. c) Evolution of the parity after performing spin-exchange for a time that should result in an entangled state and then varying the length of the gradient evolution. The gray bar represents the calculated oscillation bound of oscillation amplitude that would verify the presence of spin entanglement. d) Varying the length of spin-exchange evolution *before* performing the readout sequence for a fixed gradient evolution time. (Note that this figure is adapted from Ref. [47].)

measured probabilities ($P^{\uparrow\uparrow}, P^{\downarrow\downarrow}$) that the spins are erroneously prepared in the same spin-state: If $C > C_{\text{bnd}} = 4(P^{\uparrow\uparrow}P^{\downarrow\downarrow})^{1/2}$, then the state is entangled. By directly measuring the spin populations and their associated uncertainty, we ascertain $C_{\text{bnd}} = 0.133(25)$ as indicated by the dashed lines in Figure 5.3(c). The observed parity oscillation contrast $C = 0.49(4)$ exceeds C_{bnd} by more than 7σ , certifying the presence of entanglement in the final state of the separated spins. We verify entanglement without correcting the measured parity for experimental imperfections, such as single atom loss due to background collisions.

While in Figure 5.3(c) we varied the parity detection parameters via t_g , in Figure 5.3(d) we measure the dependence of the parity on the exchange time at fixed t_g , thereby observing oscillations as the exchange interactions periodically entangle and unentangle the two atoms. We fix t_g in the parity detection such that the entangled state $|\psi_+\rangle$ [gray lines in Figure 5.3(b,d)] is rotated to a peak in Π , corresponding to the creation of the triplet [blue lines in Figure 5.3(b,c)]. Because this t_g amounts to a $\pi/2$ rotation about the z -axis of the Bloch sphere, it will also rotate $|\psi_-\rangle$ to the singlet, which corresponds to maximal negative parity. In the lower panel of Figure 5.3(d), we show how the parity measured under these conditions oscillates at the exchange frequency $J_{\text{ex}}/(2\pi\hbar)$. For comparison, in the upper panel, we show the measured exchange oscillations (purple, green) without the parity detection. At the linear points of the exchange oscillations, one expects maximal entanglement corresponding to states $|\psi_+\rangle$ (gray) and $|\psi_-\rangle$ (orange) and thus the extremal points of the parity. At the minima and maxima of the exchange oscillations, the atoms are unentangled and the parity vanishes.

5.2 Spin-entangled states in spin-dependent traps

In the verification protocol discussed in the previous section, we purposefully break the degeneracy of the $|\uparrow_L, \downarrow_R\rangle$ and $|\downarrow_L, \uparrow_R\rangle$ states by using a magnetic field gradient between the two wells, which initiates a σ_z rotation on the effective Bloch sphere representing the $S_z = 0$ subspace of the two atoms (as shown in Fig. 5.3). The breaking of this degeneracy is useful for such operations, but when not controlled or zeroed at other times, the fidelity of the initial state preparation can be

degraded due to the splitting of the eigenstates, which should be degenerate during the adiabatic passage procedure.

This becomes especially important when considering that our trapping potential during this state preparation is intentionally unequal between the two wells in order to couple atoms between the ground band in the left well to the first excited band in the right well (as shown in Fig. 5.2). In this situation, any imperfection in the state independence of a single trap (meaning, the differential energy level shifts between the $|\uparrow\rangle$ and $|\downarrow\rangle$ single particle states on a single site), will directly turn into a position-dependent spin-selective shift between the two wells, which is equivalent to adding a magnetic field gradient. In particular, an abundance of σ^+ (or σ^-) polarization in the trap light will act as an effective magnetic field for atoms trapped in that potential, but because we are purposefully making the intensity of light unequal between the traps, this turns into an effective magnetic field gradient.

Thus, having an extremely pure linear trap polarization (we prefer it to be π polarized for other reasons, but for this purpose other linear polarizations can work) is very important to maintain the degeneracy of the $|\uparrow_L, \downarrow_R\rangle$ and $|\downarrow_L, \uparrow_R\rangle$ states when initiating spin-exchange dynamics. To achieve the largest spin-exchange contrast, we add a quarter-wave plate to the trapping light and use it to minimize the spin dependence of the trapping light. Specifically, we measure the microwave resonance for the $|2, 2\rangle \rightarrow |1, 1\rangle$ transition as a function of trap depth to get a measure of the spin-dependence of the traps and then adjust the quarter-wave plate to minimize the shift in the resonance as a function of trap depth. For reference, the data shown in Fig. 5.2 was taken when the shift in the microwave from a zero-depth trap to that in a $k_B \times 1$ mK depth trap of approximately $h \times 120$ kHz. After adjusting the quarter-wave plate, we were able to null the state dependence of the trap to better than a $h \times 1$ kHz shift between the free-space transition and a $k_B \times 1$ mK depth, as demonstrated by the microwave spectra in Fig. 5.4(b). The resulting spin exchange dynamics are shown in Fig. 5.4(a) and have a contrast more than 50% larger than that seen without polarization correction.

To confirm that these changes are correlated, we used the same quarter-wave plate to make

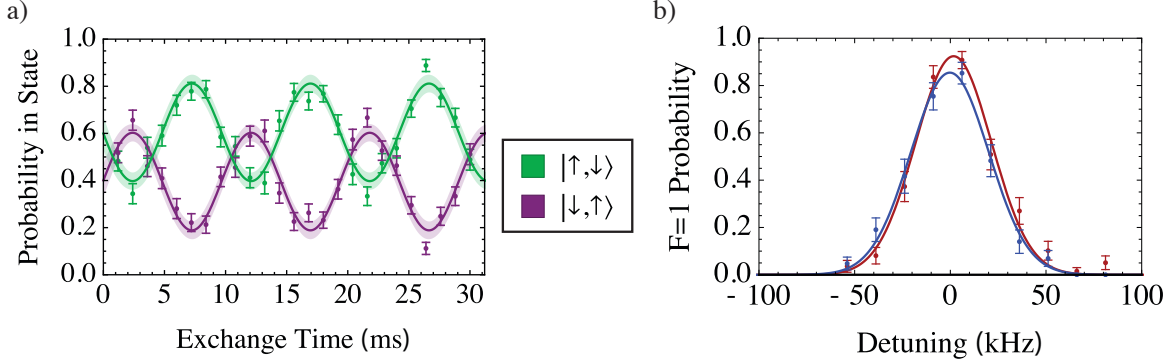


Figure 5.4: a) Observation of spin-exchange dynamics in a double well potential after compensating for imperfect trap polarization and, therefore, improving the state dependence between the two wells during the spin exchange procedure. Notice that after correcting for the spin dependence of the trapping potential, we have improved the spin exchange contrast by over 50%. b) Microwave spectra taken for both deep (red) and shallow (blue) traps, showing that there is no longer a significant vector light shift induced by the trapping light. For comparison, this same measurement for the original data shown in Figure 5.2, these peaks were separated by over $h \times 120$ kHz.

the polarization worse. Specifically, we tune the shift from the zero-depth transition to be $h \times 140$ kHz, which translates into a roughly $h \times 70$ Hz difference in the differential shift of the spin states between the two traps during the adiabatic passage process (at tunneling depths). The resulting spin-exchange dynamics are shown in Fig. 5.5, demonstrating a reduction of contrast by 50% from the optimized dynamics. The reduction observed is consistent with a picture where breaking the degeneracy of these states results in an improper transfer of state population from the state $|\uparrow_L, \downarrow_R\rangle$ into an equal superposition of the localized singlet and triplet spin-states. Thus, using this state dependence, we preferentially end up in one of the singlet or triplet spin states.

5.2.1 Adiabatic preparation of entangled states

After observing how large this effect can be, we consider a method that could take advantage of such a state dependence to reliably generate pure entangled states, without relying on the dynamics of spin-exchange. In Figure 5.6, we plot the eigenenergies of the system of two atoms in a tilted double-well potential, both with an external magnetic field gradient (which introduces a spin-dependent shift between the two wells) and without the magnetic field (such that all spin-states are

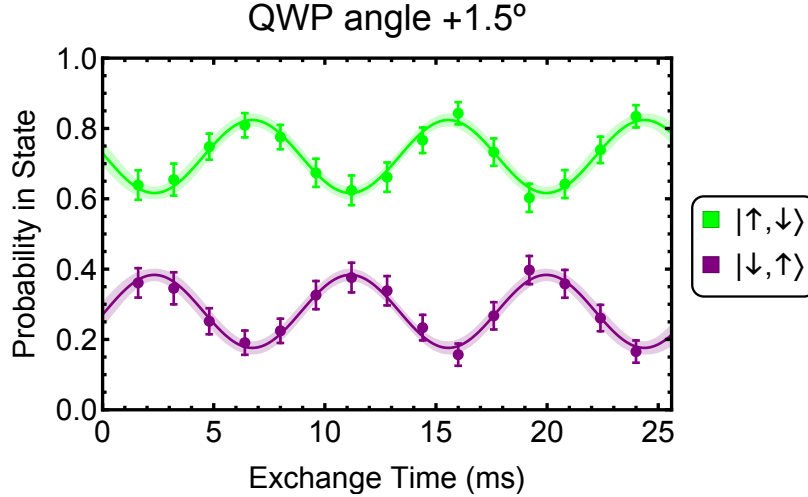


Figure 5.5: A plot of the spin exchange dynamics in the presence of a small circular component of the trap polarization, which is achieved by turning a $\lambda/4$ waveplate by a small angle from its optimum position, such that the tilt between the two wells leads to breaking of the degeneracy of $|\uparrow\rangle_e |\downarrow\rangle_g$ and $|\downarrow\rangle_e |\uparrow\rangle_g$ by $\sim h \times 70$ Hz, which is on the order of the interaction strength in the same well. We can see that this has the effect of reducing the contrast of the spin exchange oscillations by roughly a factor of two compared to the optimum case, where this shift is 0.

degenerate when the atoms are in separate wells). We notice that with the magnetic field gradient on, the $|\uparrow\rangle_L |\downarrow\rangle_R$ ($|\downarrow\rangle_L |\uparrow\rangle_R$) eigenstate is adiabatically connected to the singlet (triplet) spin state localized in the right well. Thus, we can transfer from an easy-to-prepare state ($|\uparrow\rangle_L |\downarrow\rangle_R$) into a maximally entangled state by performing an adiabatic passage.

The goal is then to separate the two atoms without projecting back into the original state. Thus, if we can shut off the state-dependent shift between the two wells, a singlet spin state should not preferentially map into one of the spin eigenstates (due to their being degenerate). This procedure is depicted in Fig. 5.6, where in part (a) we use an external magnetic field gradient to preferentially transfer from $|\uparrow\rangle_L |\downarrow\rangle_R$ into a singlet state and then, in part (b), we turn off the gradient and reverse the adiabatic passage, resulting in a singlet spin state, in separate wells.

The requirements for this procedure are simply that we need to be able to break the degeneracy of the $|\uparrow\rangle_L |\downarrow\rangle_R$ and $|\downarrow\rangle_L |\uparrow\rangle_R$ spin configurations on-demand, as well as the ability to perform the transfer of atoms at a rate that is much slower than the spin-exchange coupling. While we

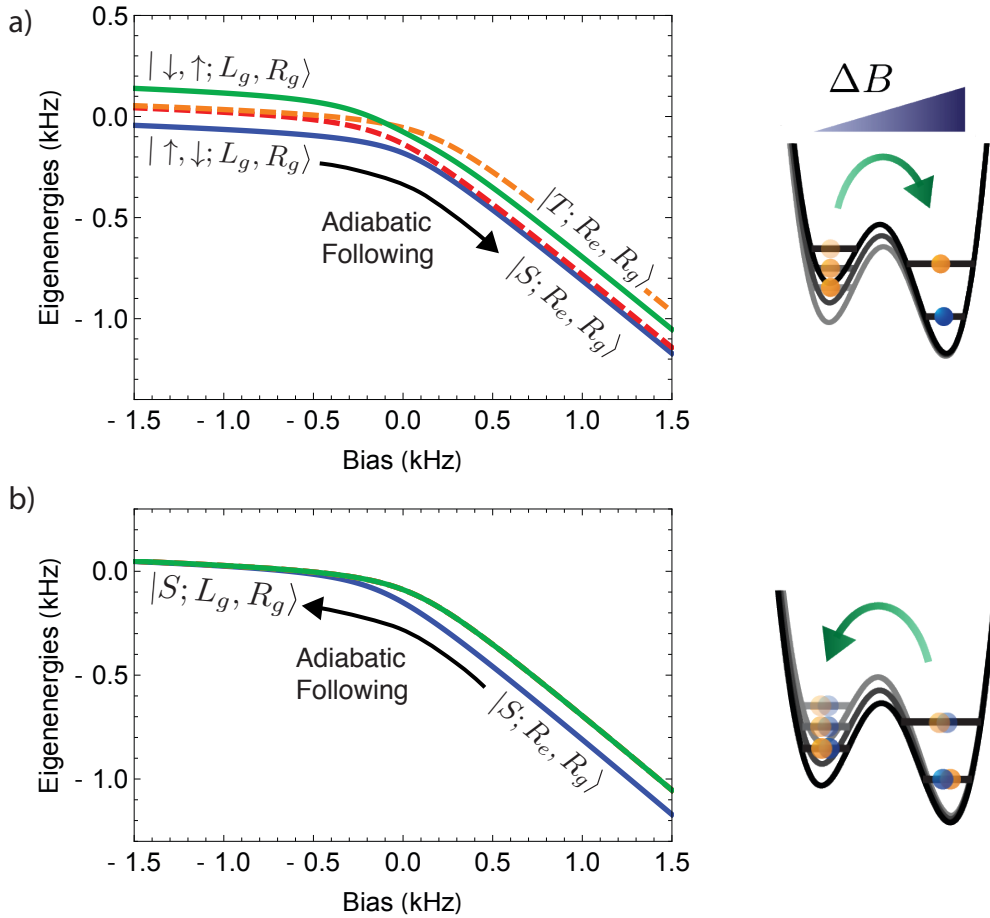


Figure 5.6: a) Diagram showing the eigenstate energy levels as a function of double-well bias, highlighting the adiabatic passage procedure (as depicted in Figure 5.2). In particular, this shows that by sufficiently breaking the degeneracy of the $|\uparrow\rangle_e |\downarrow\rangle_g$ and $|\downarrow\rangle_e |\uparrow\rangle_g$ states in separated wells, for instance by applying a large magnetic field gradient, we can adiabatically follow a single eigenstate to a pure singlet or triplet state with both atoms in the same well. b) Importantly, if we can then shut off the state dependence and separate the traps, then we would be able to separate the pure singlet or triplet states after the initial preparation.

demonstrated this effect by changing the polarization of the trap light, it is unlikely that doing this in real-time is going to be stable or reproducible, but we can easily use the magnetic field gradient that is used in the verification protocol of Section 5.1 and achieve sufficient splitting; this is the magnitude of the spin-dependence that was assumed in the calculation used to make Figure 5.6. Alternatively, we may also prefer to use the vector light shift beam localized to one of the wells (at much lower intensity than is used for microwave spin flips) because it will provide us with much

larger shifts and faster actuation than the magnetic field gradient, which may become necessary to increase the speed while simultaneously limiting the coupling to other states during the transfer.

In the end, this procedure will likely be slower than using dynamical spin exchange to entangle the atoms, but may be much more reliable, as it is not sensitive to many day-to-day fluctuations of the trapping potential. Additionally, if we truly want to make a pure singlet or triplet for use in other simulations, as will be discussed in Section 6.3.3, then this may be the most reliable method of state preparation.

Chapter 6

Assembling the Kondo Lattice Model with Optical Tweezers

In this chapter, I will discuss ongoing work that will combine many of the capabilities discussed in this thesis toward the goal of implementing a microscopic realization of the bosonic Kondo-Hubbard model. A future demonstration of this control would be a significant realization of the utility of this platform for studying the interplay of the spin and motional degrees of freedom in systems of a few atoms with local interactions, when the individual degrees of freedom are no longer separable and thus the full quantum state must be considered. It will be particularly interesting to observe how the symmetrization of the many-body wavefunction in the presence of local interactions manifests in the out-of-equilibrium dynamics of states assembled from a initially pure single-atom states.

First, we will introduce the Kondo lattice model along with some background and a discussion of the utility of this model for explaining interesting phases of matter in heavy-fermion systems [107]. Then, we will take a closer look at the bosonic analog of the model and, in particular, its minimal version that we propose to implement in optical tweezers. For this system, we will take a look at some preliminary simulations we have been performed to understand the observables related to interactions of interest and discuss some of the challenges related to implementation, as well as highlighting regimes in which we will, realistically, be able to work.

6.1 The Kondo Lattice Model

The development of the Kondo lattice model has a long history in condensed matter physics, starting with the understanding of the “Kondo effect”, which describes the observations that metals doped with small fractions of magnetic impurity atoms had a minimum in their resistance as a function of temperature.¹ In this system, the exchange interaction between localized f -shell electrons (bound to the impurity atoms) and the conduction band electrons result in the creation of localized quasiparticles when the thermal energy is not enough to overcome the binding of the energetically favorable singlet states to the localized electrons.² However, the presence of these localized states reduce the mobility of the conduction electrons and therefore increase the resistance at low temperature [108].

The Kondo lattice model is an extension of this model of the Kondo effect, with many of these impurities arranged in an ordered fashion (typically a square lattice) to study the interplay of the competing local interactions with the quantum statistics and motion of conduction band electrons [109, 107]. This model has the minimum ingredients to explain the complex behaviors present in heavy-fermion materials, where the spin and motional degrees of freedom of electrons are coupled via the Coulomb interaction [110, 111]. In particular, the propagation on a lattice of the type quasiparticles that are described in the Kondo effect can explain the very large effective masses that these heavy-fermion materials exhibit. The specific form of this model imagines an array of spins that are fixed, which can interact with mobile electrons that can move between (and interact with) spins localized in different locations, as shown in the schematic representation given in Figure 6.1(a). The interaction Hamiltonian is given by

$$H_{\text{KLM}} = -t \sum_{\langle i,j \rangle, \sigma} \left(c_{i\sigma}^\dagger c_{j\sigma} + c_{j\sigma}^\dagger c_{i\sigma} \right) + J \sum_i \mathbf{S}_i^f \cdot \mathbf{S}_i^c \quad (6.1)$$

Where here we use the notation defined in Ref. [109], where $c_{i,\sigma}$ ($c_{i,\sigma}^\dagger$ is the annihilation (creation)

¹ These observations were at low-temperature, where the resistance of a normal (not superconducting) metal would approach a constant minimum value with power-law dependence.

² This description assumes that the exchange interaction is negative, which will shift the singlet state to lower energies for fermions. However, for impurities where the interaction coefficient is positive, the resistance decreases faster at low temperature, as mentioned in Ref. [108].

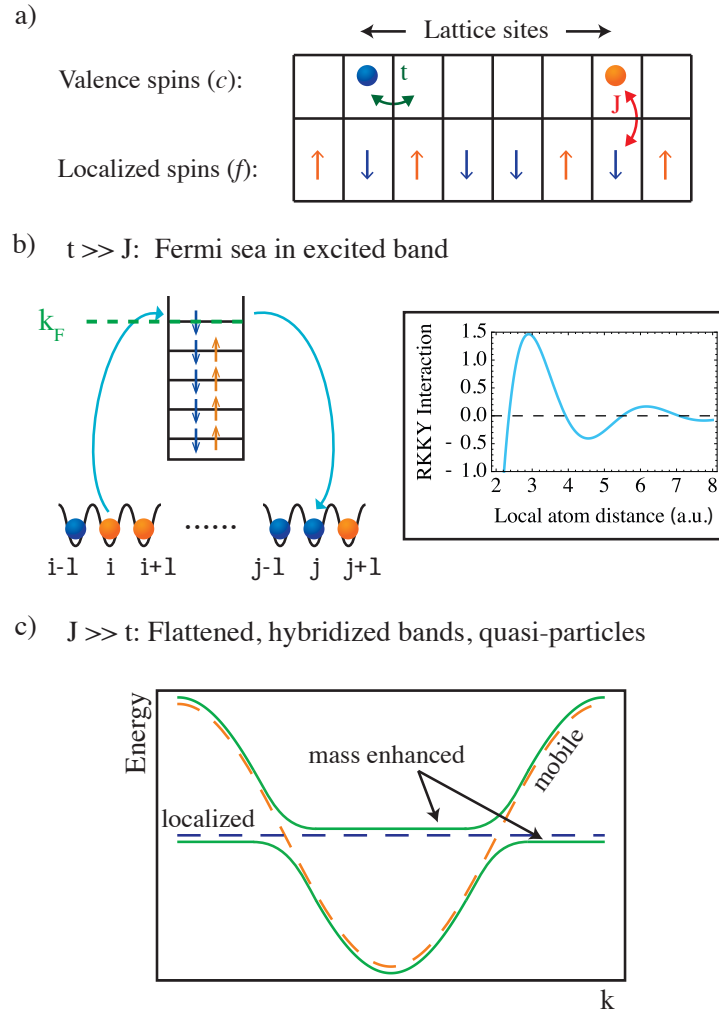


Figure 6.1: a) Diagram to depict the Kondo lattice model written out in Equation 6.1. Notice that the model contains only two terms: The on-site spin-interaction term J (indicated in red), and the tunneling term t (indicated in green). Note that this is a general model that does not specify the underlying realization, but is meant to describe, for example, a solid material where each localized atom contains an inner-shell electron that interacts with the mobile valence electrons. b) Schematic showing how, in the weakly-interacting regime, the perturbative local interactions can lead to a second-order coupling between distant fixed spins via virtual excitation of electrons above the Fermi sea. This virtual excitation leads to a distance scale that is set by the Fermi momentum, k_F (indicated by the green dashed line). c) Depiction of the band structure in the strongly interacting regime, where the flat band representing the fixed spins, and the cosinusoidal band represents the mobile electrons. The effective mass of the quasiparticles formed in this system (which is proportional to the slope of the band) will become very large when the lattice filling factor is near $1/2$, such that there can be a single atom per site in the excited band (and thus all ground spins are paired into singlets).

operator for the conduction band spins on site i , \mathbf{S}_i^c is the total spin operator for the conduction band spins on site i , and \mathbf{S}_i^f is the total spin operator for the localized spins (defined by the creation operator f_i^\dagger) on site i . In this model, J represents the interaction energy between the localized and conduction band spins, while t represents the energy scale for tunneling of the conduction band spins between sites. It is important to note that this notation will differ from notation used in Section 6.2 and beyond (see text after Equation 6.2 for details), but I want to use the original notation in this introduction.

6.1.1 Limiting cases of the Kondo lattice model

In this model, it is interesting to consider the two limiting regimes of strong and weak interactions compared to the tunneling term ($J \gg t$ and $J \ll t$, respectively), where interesting phases of the system develop. In particular, I have already mentioned the strongly-interacting regime, where the formation of quasiparticles on the lattice leads to measurements of the effective mass of the conduction band electrons becoming extremely large (up to a factor of 1000 larger than the bare mass of an electron) [112, 113]. In this regime, the strong interaction between the localized spins on the lattice and the conduction band electrons results in a hybridization of the band structure that results in a flattening of regions of the dispersion relation, as depicted in Figure 6.1(b). As described in band theory, the flattening of this band is a signature of the increasing effective mass of the electron quasiparticles because the effective mass is defined by the local curvature of the dispersion relation (specifically, the second derivative of the energy with respect to the momentum) for a given electron momentum.

In the weakly interacting regime, the interactions are a perturbative effect, so it is appropriate to expand the Hamiltonian in orders of the on-site interaction energy. At second-order, the ground-band atoms interact with each other in the same manner as that described by the Ruderman-Kittel-Kasuya-Yosida (RKKY) interaction that was originally derived when studying the magnetic ordering of nuclear spins in a metal [114, 115, 116]. Specifically, this effective interaction between the localized spins will lead to the long-range ordering of these spins in a way that depends heavily

on the electrons in the conduction band because the process that drives these interactions requires an effective flipping of a conduction band electron. However, to flip a conduction band electron, it has to virtually be excited out of the fermi-sea, propagate to the other local spin, and then flip back; this process, as depicted in Figure 6.1(c), sets the energy (and length) scale of this interaction, which is the Fermi energy E_F (or the inverse of the Fermi momentum k_F). Thus, at low filling fractions, this direct coupling of the localized spins leads to a ferromagnetic interaction, while at larger filling fractions there is an oscillatory coupling between local spins that depends on the separation of these spins. This second-order interaction between the localized spin states is thought to stabilize long-range ordering of the localized spins in materials that have weaker interactions between the local and conduction band electrons [111].

6.2 The Kondo-Hubbard Model for Atomic Systems

Despite decades of research invested into studying the Kondo lattice model, its general phase diagram remains unknown even for a simple case of a 2D square lattice [111], which points to the complexity of the spin-motional coupling, especially as the array size, and number of particles, is increased. Hence, any insight into its physical properties obtained within a controlled setting is valuable. Such an environment is offered by cold atomic gases that have proven to be excellent candidates for studying condensed matter models [18, 117, 19, 81]. In our work, we will leverage the additional control afforded with the optical tweezer platform, this seems to be a particularly interesting model to study with our system.

The Kondo-Hubbard model is a representation of the Kondo lattice model that can be replicated in a relatively simple system, but describes the complex behaviors present in certain strongly correlated materials where the spin and motional degrees of freedom of electrons are coupled [110, 111]. The interaction Hamiltonian for the Kondo-Hubbard model, as implemented in a two-band lattice, is

$$H_{KHM} = -J_e \sum_{\langle i,j \rangle} b_i^\dagger b_j + \frac{U_e}{2} \sum_i (n_i^b)^2 + 2V_{eg} \sum_i \mathbf{S}_i^a \cdot \mathbf{S}_i^b \quad (6.2)$$

Here a_i^\dagger (a_i) are the creation (annihilation) operators for atoms in the ground-band at site i , b_i^\dagger (b_i) are the same for atoms in the excited-band, n_i^b is the operator indicating the number of atoms in the excited band at site i , $\langle i, j \rangle$ indicates summation over all pairs of neighboring wells i and j , \mathbf{S}_i^σ is the total spin operator for atoms in the σ band at site i , and the coefficients J_g , U_g , J_e , U_e , and V_{eg} set the relative tunneling rates and on-site interaction energies for the system.³ Note also that to be general, we have added a term for on-site interactions within the excited band, because multiple bosonic atoms in the same state can spatially overlap (unlike the case of fermions in Equation 6.1, where only one particle per spin state can exist in a single band at each site).

The Hamiltonian in Equation 6.2 also assumes that each site in the ground-band can have only a single atom per site, that we can neglect the tunneling of atoms within the ground-band of the lattice, and that we can neglect the presence of any band-changing collisions (which would require taking into account higher bands of the lattice). Note that the Hamiltonian is written with simplified notation that may suggest a one-dimensional chain of atoms, but many systems of interest will actually be two- or three-dimensional, in which case the indices i and j simply represent coordinates in a higher-dimensional array.⁴

6.2.1 Studying the Kondo-Hubbard model with bosons

Of course, in our system, we are using spin-1/2 bosons, rather than the fermions considered for condensed matter systems. However, many of the same effects can be seen in the bosonic Kondo-Hubbard model, and the understanding of these interactions gained with bosons will only give more information to aid in understanding the fermionic case (which has been studied extensively in condensed matter systems, albeit without the level of control afforded to cold atom

³ The notation I am using here will be familiar to those used to working with cold, bosonic quantum gases. However, the notations differ quite dramatically in the condensed matter literature, where the exchange energy (here V_{eg}) is written as J and the tunneling rate (here J_e) is given by t . Because I am writing this with the simplified model we work with in our system (and in comparison to past work with effective Hubbard parameters), I have chosen to stick with this notation, but want to include a warning because this can lead to significant confusion.

⁴ While the model works with no other modifications, care must be taken in real implementations of the model to take into account degeneracies between different axes (the first excited band of the lattice have different motional states along each dimension) or different couplings in each dimension (if you break the degeneracy, then the tunnel coupling will likely be altered as well).

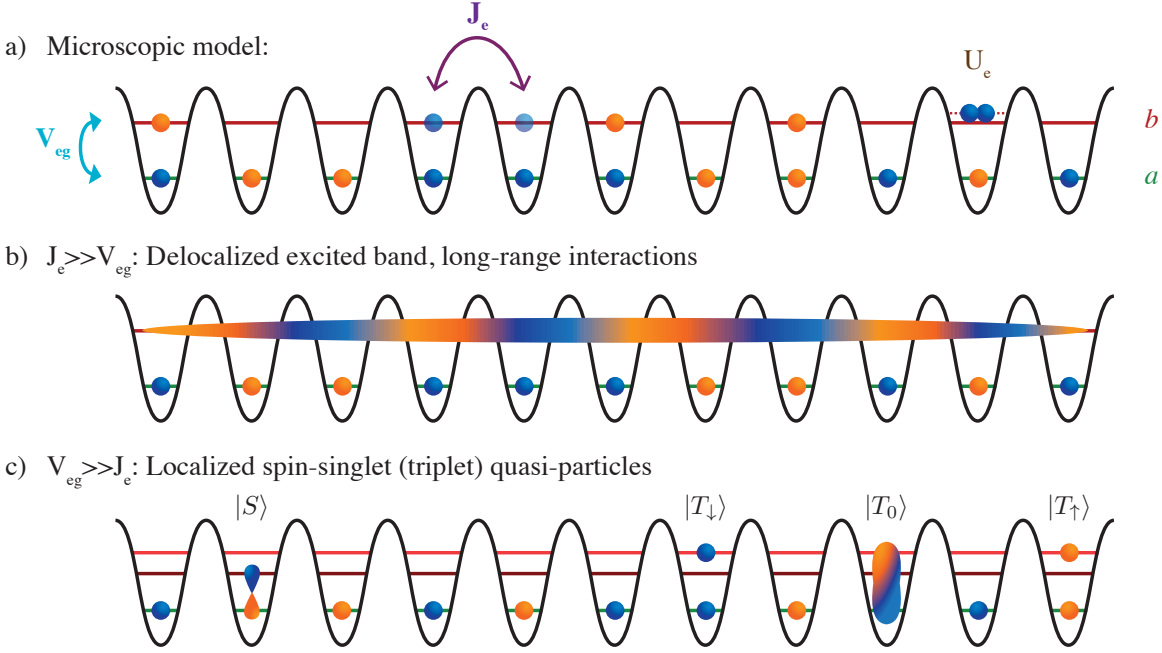


Figure 6.2: a) Schematic showing the components of the Kondo lattice model, as implemented in a two-band optical lattice. The diagrams here as shown in one dimension, but can be extended to two or three dimensions, where the full phase diagrams remain unsolved. In this model, the excited band (red) contains atoms that can tunnel around the lattice potential with tunneling coefficient J_e . The ground band of the lattice (green) contain atoms that are effectively fixed (the tunneling rate J_g , is much smaller than all other timescales in the experiment), but which can interact with the excited band particles through the on-site exchange interaction, characterized by the interaction coefficient V_{eg} . b) When the tunneling timescale, J_e dominates, then the excited band particles effectively become delocalized, which will lead to mediated exchange interactions between the localized ground band atoms. For fermions, the different quasimomenta of the atoms in the excited band can lead to periodic spin-coupling terms, analogous to the RKKY interaction term. c) When the interaction strength, V_{eg} , dominates over the tunneling, then the energies related to different on-site spin configurations (singlet vs. triplet) will be split. In this situation, the excited band tunneling rate will be modified based on these splitting. This modification to the tunneling rate is an analog to the heavy fermion physics observed certain condensed matter systems.

experiments) [118, 119, 120, 121]. The fundamental difference in these two models comes from how the quantum statistics affect the configurations of atoms (and the distribution of spin states) in the excited band: Due to Pauli blocking, there can be no more than two fermionic atoms on each lattice site (and they would have opposite spin), while bosons will happily bunch many atoms onto the same site. Thus, the number of particles per site will depend heavily on the ratio of the

interaction energy U_e to the tunneling rate J_e (in addition to the density of carriers) for the bosonic case. For example, when $J_e \gg V_{eg}$, this results in the magnetic ordering becoming ferromagnetic no matter the filling fraction of the lattice, as opposed to the spin ordering in the fermionic case that depends on the filling fraction (via the Fermi momentum k_F), where at low filling fractions the ordering is ferromagnetic, but at higher filling fractions becomes oscillatory (as described by the RKKY interaction, introduced above) [120, 121]. Importantly, this means that we can apply studies of low filling-fractions with bosons to the general problem, but certain behaviors that are observed in condensed matter systems at higher filling fractions will not be apparent with bosons.

6.3 Implementation with optical tweezers

Typical theoretical studies of this model assume are carried out for infinite systems using. However, in our case, we have the ability to directly study small arrays of atoms placed into a lattice that is projected into the focal plane of our objective lens. In such systems, the Hilbert space is already very large, such that detailed theoretical study becomes very challenging, if not impossible. For initial experiments with just a few atoms, this will be accomplished by reorganizing and overlapping the optical tweezers themselves, but could also involve transferring the atoms to an arbitrary potential generated via a DMD, as discussed in Section 3.2.4. Using the techniques for quantum gas assembly, we have the ability to create and study out-of-equilibrium states of the system (by preparing a pure state in well-separated potentials and quenching into the final potential to initiate dynamics) and then by pulling the atoms back to the separated potentials, we can follow the dynamics of the system from initial state to final state.

For our first experiments, we will concentrate on the minimum requirements to implement the microscopic Kondo lattice Hamiltonian, which is to use three atoms in a double well potential. Two of these atoms will occupy the ground-state of the two wells and the third atom will be mobile and in an excited band of the potential; for the purposes of this discussion, we will use the first radial excited band along the axis connecting the two wells. As in previous work, we will encode the spins in the ground state hyperfine levels of ^{87}Rb , and for the study of this system will prepare

the total spin of the system to have a projection of $S_z = \pm 1/2$ (meaning the spin of one atom will be opposite the other two atoms). The specifics of how we prepare these atoms, as well as the shape of the final potential we use, will have drastic effects on the dynamics of the system.

The process for initializing these states is depicted in Figure 6.3. Part (a) shows the standard steps to trap a set of atoms and initialize the spin and motional states before assembling the state of interest from single atoms, but the crucial differences are highlighted in Figure 6.3(b-d), where both the shape of the final trapping potential and the method for initiating dynamics are altered (and in the case of (d), a different spin configuration is used, as discussed in Section 6.3.3). While these minimum realizations of the Kondo lattice model can be solved analytically, the observation of these dynamics in our system will serve as a verification of the techniques used to study spin-motional coupling with optical tweezers, and provides a stepping stone to studying larger systems where an analytic solution may not be available.

6.3.1 Calculating the expected spin dynamics

Before discussing the different parameters regimes that we want to study, as well as how we can implement them with our system, I want to recognize the productive collaboration we have had with Ana Maria Rey's group for studying this system. In particular, Leonid Isaev derived the analytic result for a simplified three-spin model, which has been extremely useful for quickly exploring different parameter regimes, and Michael Wall has performed full three-dimensional calculations of the potential to extract effective Hubbard parameters that we can use in these calculations [104]. The plots in the remainder of this chapter that show the expected spin dynamics are generated using the analytic solutions to the restricted three-spin model that we are implementing by taking advantage of the simplifying assumption that the ground-state atoms are fixed in their respective wells, as well as knowledge of the symmetries associated with the trapping potential and the spin interactions.⁵

Many of these systems were also numerically solved by propagating the initial states in the full

⁵ This derivation is contained in a set of notes written by Leonid Isaev and is briefly summarized in Appendix B.




Figure 6.3: Diagram showing how to implement the microscopic Kondo lattice model with optical tweezers. a) Starting with three isolated atoms, we bring all atoms to the ground state of their optical tweezers and prepare the spins of each atom according to the requirements for any given experiment (one specific preparation is shown). The traps are then repositioned to form two double well potentials and, as was done for a single pair in Section 5.1, we perform an adiabatic passage that transfers atoms in the outer wells into the excited band of the inner wells. Note that this same process results in motional distillation on the two ground-band atoms (as described in Section 3.3.2), which will improve the final state preparation. b,c) After the initial state preparation from (a), we can rearrange the two central wells (of the four total wells) and initiate dynamics a central double well potential. For different trap parameters, we will initiate the dynamics in this potential differently: b) In the weakly-interacting regime, we will want to delocalize the excited atom, which is achieved by performing half of an adiabatic passage after the wells are close. c) In the strongly-interacting regime, we start with a localized excited particle, and thus will sweep the traps together at zero bias to initiate dynamics.

Hilbert space via the Schrödinger equation (with the full Hamiltonian given in Equation 6.2). The numerical propagation produced nearly identical results, aside from the presence of the superexchange interaction between the ground band atoms, which is neglected in the analytic solution. However, the timescale for superexchange is much longer than all other dynamical timescales of interest, which justifies neglecting that term initially. While the analytic solution is much faster for this particular system, the parallel development of these two models is important because the analytic solution is specific to this system and will not hold for larger systems of atoms; that is not to say an analytic solution cannot be found, but it becomes less likely (and certainly much more complex to connect the eigenstates of the system back to the experimentally prepared states). On the other hand, the numerical propagation can be very simply extended to larger arrays with more atoms, where the complexity of the system is dramatically increased.

6.3.2 Weakly interacting regime ($V_{eg} \ll J_e$)

To understand the parameter space of this problem, we will first look at extreme cases, where one energy scale dominates, and observe the resulting dynamics. In these cases, with the appropriate choice of the initial state we can effectively restrict ourselves to a smaller subspace of the system, where the observed dynamics have clear interpretations. In the case where the interaction energy V_{eg} is small compared to the tunneling rate J_e , the separated subspaces are the delocalized states of the excited atom, one that has positive (and the other negative) inversion symmetry; these are the eigenstates of a single atom under a tunnel coupling, as discussed in Ref. [39]. Each of these subspaces is spanned by three spin states of the two ground-state atoms (assuming no external forces that can change \mathbf{S}_z^{tot}). For now, we will work in the subspace with $\mathbf{S}_z^{tot} = -\frac{1}{2}$ and use the single-particle measurement basis for the ground-state spins: $|\uparrow_L, \downarrow_R\rangle$, $|\downarrow_L, \uparrow_R\rangle$, and $|\downarrow_L, \downarrow_R\rangle$. Figure 6.4(a,b) shows the population of the $|\uparrow_L, \downarrow_R\rangle$ and $|\downarrow_L, \uparrow_R\rangle$ spin configurations over time. We can immediately see that the spin dynamics in this system are relatively complex, so let's look closer at the Hamiltonian to see how to extract the information we want.

In this regime, we are looking for evidence of a second-order, direct-coupling term between

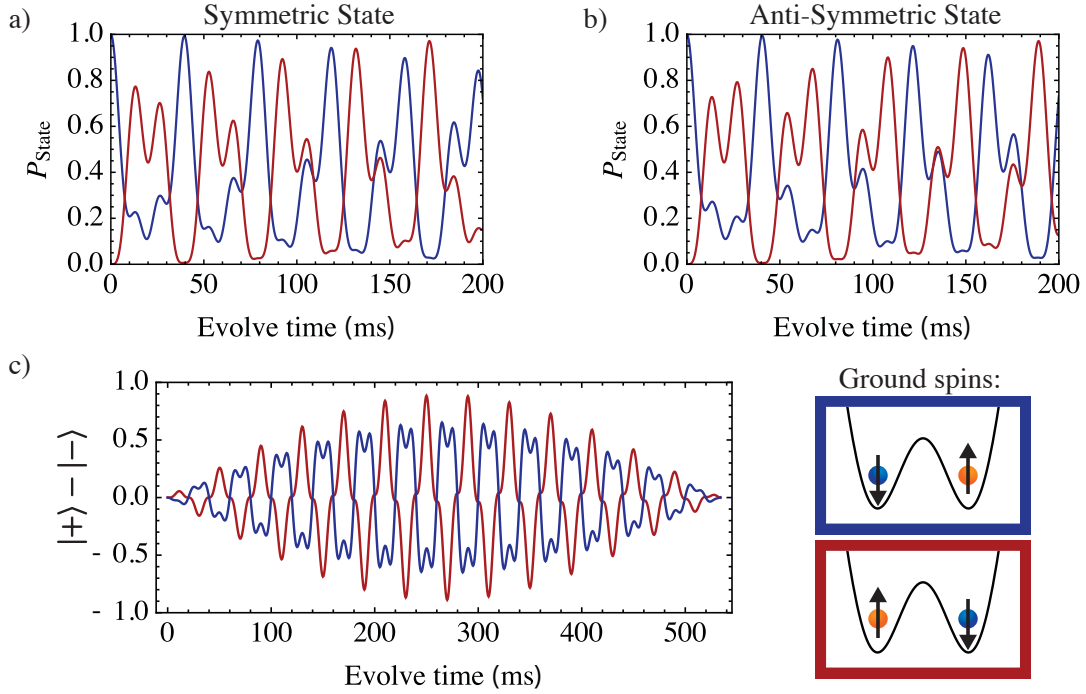


Figure 6.4: a) Plot of the population in the $|\uparrow_L, \downarrow_R\rangle$ (red) and $|\downarrow_L, \uparrow_R\rangle$ (blue) ground state spin configurations as a function of time in the weakly interacting regime. The initial state was in the $|\downarrow_L, \uparrow_R\rangle$ ground-state spin configuration (the ground-state spin configuration indicated in the blue box), with the excited band atom in the symmetric parity state. b) Plot of the same spin population dynamics, but not with the excited band atom in the anti-symmetric parity state. c) Looking at the difference in the time evolution of the spin populations between the symmetric ($|+\rangle$) and antisymmetric ($|-\rangle$) parity states of the excited band atom. The envelope of these oscillations is a manifestation of the direct spin-interaction term in Equation 6.3.

the ground-state spins (e.g., $\mathbf{S}_R \cdot \mathbf{S}_L$), which is due to the third atom being in the excited band. The effective spin Hamiltonian, up to second-order in the perturbative spin-interaction, is written in Equation 6.3 where \mathbf{S}_e represents the spin- $\frac{1}{2}$ operator for the excited atom and the sign of the second term is positive (negative) for the symmetric (antisymmetric) delocalized state of the excited atom. In order to derive this Hamiltonian, we look at how the presence of an atom in the excited band perturbs the spin eigenstates of the ground-band atoms.⁶ Notice that the second

⁶ The symmetric and antisymmetric delocalized states of the excited atom are the parity eigenstates related to the tunnel-coupling between the two wells. A localized atom in one well of the excited band becomes a superposition of these two parity states, and since they have different energies, the phase evolution results in tunneling. The addition of interactions further splits these manifolds, and it is a difference in the sign of this splitting, as indicated in Equation 6.3, that we want to measure in this experiment.

term of the effective Hamiltonian is simply the Heisenberg model on a triangle, which includes the direct coupling term that we wish to detect. The problem is that the dynamics will be dominated by the direct coupling between the two ground-state atoms and the excited state atom at a rate $\propto V_{eg} \gg J_e$. Thus, to isolate the relevant signal we will use the change of sign and compare the dynamics between identical initial spin states with opposite parity of a delocalized excited atom, as shown in Figure 6.4(c); this will reveal the shift of the energy levels due to the second-order coupling term.

$$H_{eff} = 2V_{eg} \mathbf{S}_e \cdot (\mathbf{S}_R + \mathbf{S}_L) \pm \frac{V_{eg}^2}{4J_e} [\mathbf{S}_R \cdot \mathbf{S}_L + \mathbf{S}_e \cdot (\mathbf{S}_R + \mathbf{S}_L)] \quad (6.3)$$

Specifically, we start by observing the dynamics of the ground state spins for both parity states, as shown in Figure 6.4(a,b). The simulated evolution is performed using “idealized” parameters of $V_{eg} = h \times 50$ Hz and $J_e = h \times 1000$ Hz, to ensure that the perturbative model is justified. The initial spin state is prepared using the combination of the pure single-atoms spin states $|\downarrow_R, \uparrow_L; \downarrow_e\rangle$, and we measure joint probability for the ground state atoms to have swapped to the state $|\uparrow_R, \downarrow_L\rangle$ during the evolution. Figure 6.4(c) shows the difference in the dynamics between the two initial states and it is the slowly varying envelope signal that signifies the presence of the second term in Equation 6.3. The faster oscillation is the result of a direct coupling between multiple spin states, which we will explore in more detail below, and because it is not a pure sinusoidal term, indicates that there are at least three levels involved in these dynamics.

While this initial state (and measurement basis) will be the *easiest* to use experimentally, we will look at the more natural spin basis of the singlet $[\frac{1}{\sqrt{2}}(|\uparrow_L, \downarrow_R\rangle - |\downarrow_L, \uparrow_R\rangle)]$, triplet $[\frac{1}{\sqrt{2}}(|\uparrow_L, \downarrow_R\rangle + |\downarrow_L, \uparrow_R\rangle)]$, and down $(|\downarrow_L, \downarrow_R\rangle)$ states to understand the spectrum of the energy levels. This information also gives us the information to explore ways to expose a cleaner signal, even if it is harder to implement in the experiment. Specifically, we take a closer look at the case where we initialize a triplet state between the two ground-band atoms in Figure 6.5, measuring the probability to be in the down state as a function of time. Importantly, this gives a single frequency of oscillation,

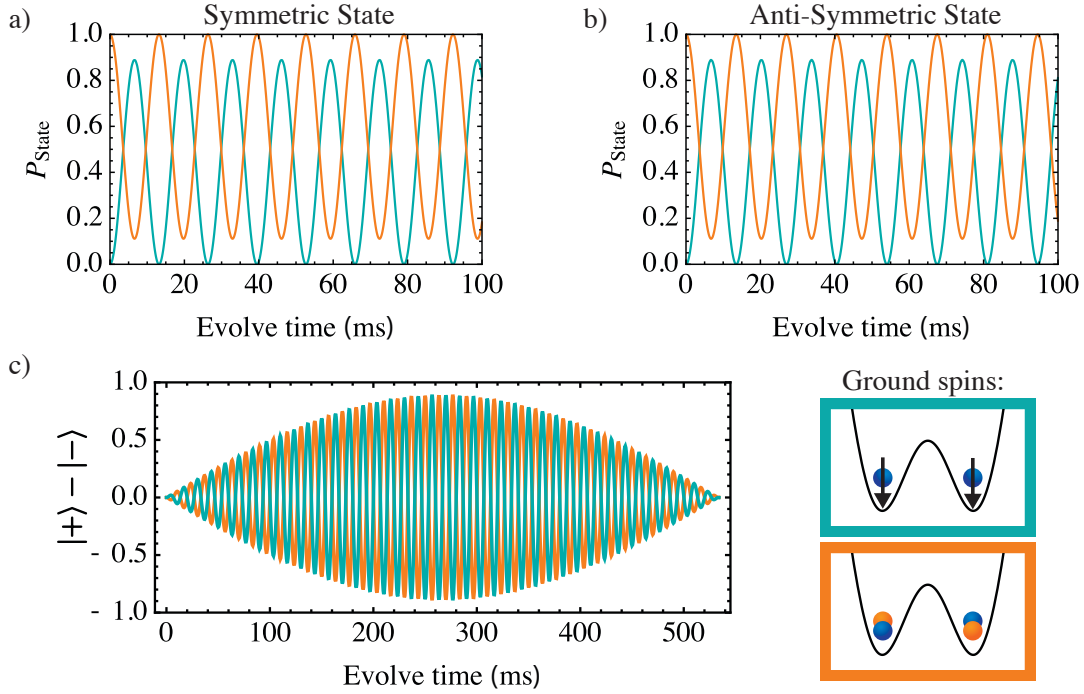


Figure 6.5: These are an analogous set of plots to those in Figure 6.4, but now both measuring and initializing in the typical basis for a two-spin system. Note that the singlet state $[\frac{1}{\sqrt{2}}(|\uparrow_L, \downarrow_R\rangle - |\downarrow_L, \uparrow_R\rangle)]$ is an eigenstate and will remain constant on these plots. Hence we have chosen to initialize the two ground band atoms in the triplet state $[\frac{1}{\sqrt{2}}(|\uparrow_L, \downarrow_R\rangle + |\downarrow_L, \uparrow_R\rangle)]$, indicated in the orange box) and are plotting the triplet and down ($|\downarrow_L, \downarrow_R\rangle$, indicated in the cyan box) state fractions over time. This plot demonstrates that when preparing the right initial state (and measuring in the same basis), we can simplify the dynamics such that we are effectively measuring the difference in exchange frequency for initial states with the excited atom in the symmetric (a) or antisymmetric (b) parity state. The difference signal (c) shows the same envelope as in Figure 6.4, but the frequency of this demonstrates that the frequency of this envelope is simply the difference in frequencies of the oscillations in (a) and (b).

which is a measure of the splitting between the triplet and down states, which will be changed as the sign of the second term in Equation 6.3 is changed. Importantly, this means that rather than subtracting amplitudes, which are affected by noise and dephasing experimentally, we would be able to measure a change in the frequency of oscillations to detect the presence of this second-order term and this would likely be a more robust signal to the noise typically encountered in our experiments.

While simply measuring the difference of frequencies is sufficient to detect the shifts (and thus the presence of the second-order interaction), we can also directly look at the analogous

difference in the dynamics, as shown in Figure 6.5(d), which exhibits the same envelope as the dynamics in Figure 6.4. Further, we observe that when initially preparing the singlet state, we are in an eigenstate of the spin-interaction Hamiltonian, so we will not observe any dynamics. This explains the more complex dynamics in Figure 6.4 as being due to the interference of all three of the singlet, triplet, and down states. It also suggests that, while the oscillation amplitudes will be reduced by a factor of two, we can measure the down state when starting with the initial state $|\downarrow_R, \uparrow_L; \downarrow_e\rangle$, as shown in Figure 6.6. This gives us the option of measuring the difference in oscillation frequency even when starting with the easier-to-prepare initial state, which may prove to be useful experimentally.

6.3.3 Interaction dominating regime ($V_{eg} \gg J_e$)

At the other extreme, we can observe dynamics in our system when the interaction energy between the atoms dominates the tunneling rate. In this regime, there is a separation of the energy spectrum for the spin states within a single well, so we choose the natural basis in the $\mathbf{S}_z^{tot} = -\frac{1}{2}$ spanned by the singlet $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle - |\downarrow_g, \uparrow_e\rangle)]$, triplet $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle + |\downarrow_g, \uparrow_e\rangle)]$, and down ($|\downarrow_g, \downarrow_e\rangle$) states. In this section, we will also simulate the dynamics using “idealized” parameters, but now with the scales reversed: $J_e = h \times 50$ Hz and $V_{eg} = h \times 1000$ Hz. In this regime, we are looking for evidence of “mass enhancement” of the excited band particle, which will manifest itself as the slowing of the tunneling rate of the excited band atom. Thus, we initialize states with the excited atom localized to one of the wells, as depicted in Figure 6.3(d) and observe the tunneling of this atom for various spin configurations. In particular, we first measure the baseline tunneling rate by using a spin-polarized state ($\mathbf{S}_z^{tot} = -\frac{3}{2}$) to get a baseline for the excited band tunneling in the absence of interactions; these dynamics are shown in blue on in Figure 6.7. Equivalently, we can measure the tunneling of the excited atom in the absence of ground-band atoms and would observe the same tunneling rate.

In comparison, we initiate dynamics after preparing the initial spin states of a localized singlet (green) and triplet (purple) states in the right well of the potential. These each have significantly

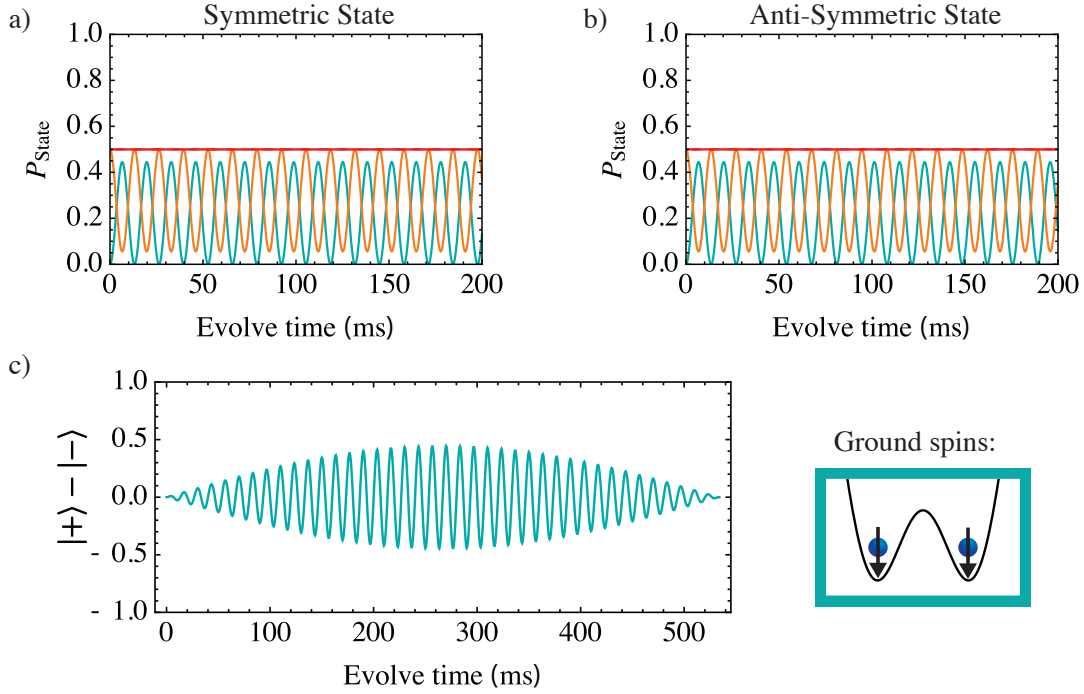


Figure 6.6: This is just a proof-of-principle that by mixing the preparation and measurement bases, we can actually simplify the state preparation, while also potentially improving our ability to read out the shift by measuring the change in oscillation frequencies. In this set of plots, we initialize the $|\uparrow_L, \downarrow_R\rangle$ state, as in Figure 6.4, while measuring the down state $|\downarrow_L, \downarrow_R\rangle$ (cyan). While we sacrifice a factor of two in the amplitude of the signal because half our state is stuck in the singlet, shown in red in (a) and (b) (which has the side effect of preventing us from easily measuring the triplet state as well), we may win because we are less sensitive to further amplitude reduction due to the dephasing of the state (as long as the dephasing time is long enough that we can get a precise enough measurement of the frequency).

different dynamics from the baseline tunneling rate, but importantly the singlet state exhibits pure oscillation at half the frequency of the bare tunneling. This is precisely what we expect in the Kondo lattice model with a single mobile spin, and is a manifestation of the preference to generate localized singlet states [109]. The dynamics in the case of the localized triplet state is more complex due to the coupling between the triplet and down states (as seen in Figure 6.8), but qualitatively we see that there appears to be dynamics occurring at both rates, but the atom never fully transfers to the other well.

While not experimentally feasible, in the simulation we can also look at the spin states in




Figure 6.7: Probability for the excited band atom to be in the left well (highlighted in red) as a function of time, plotted for three different initial spin configurations, and artificially choosing the parameters to be deep in the interaction dominated regime ($J_e = h \times 50$ Hz, $V_{eg} = h \times 1000$ Hz). In all three initial states, the excited atom is initially localized to the right well, but the initial spin states change: In blue, we are in the fully spin-polarized case, where the excited atom tunnels between the wells at rate $2J_e$, as if there are no ground band atoms. In green, we show the case where the two atoms in the right well form a spin singlet state $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle - |\downarrow_g, \uparrow_e\rangle)]$, which tunnels at full contrast but at half the rate of a bare atom. This reduction in the tunneling rate can be translated to an effective mass enhancement factor of 2. In purple, we show the excited atom dynamics when the atoms in the right well form a triplet state $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle + |\downarrow_g, \uparrow_e\rangle)]$, and we can see that there are multiple frequencies due to the mixing of the triplet manifold states, as discussed in Figure 6.8.

the transferred well to get a better understanding of how the states evolve over time. Specifically, in Figure 6.8 we plot the probability of the three relevant spin states localized to the right well (depicted at the bottom of the figure) as a function of time for different initial spin configurations. In all four cases, the excited atom is initially localized to the right well and the total spin is set to $S_z^{tot} = -\frac{1}{2}$; we then vary only the combined spin state of the two atoms the right well (and, if relevant, the spin of the atom in the left well to preserve the total spin). In (a,b), we show the dynamics for atoms starting in the triplet and down state, respectively, demonstrating that the two states are coupled as the atom tunnels between the wells. In particular, notice that when the atoms starts in the triplet state, it tends to transfer first to the down state in the right well, and vice-versa; it also is the interference of these couplings that leads to the imperfect transfer between the two wells. In part (c), we show the spin dynamics when we start with a singlet spin state, which

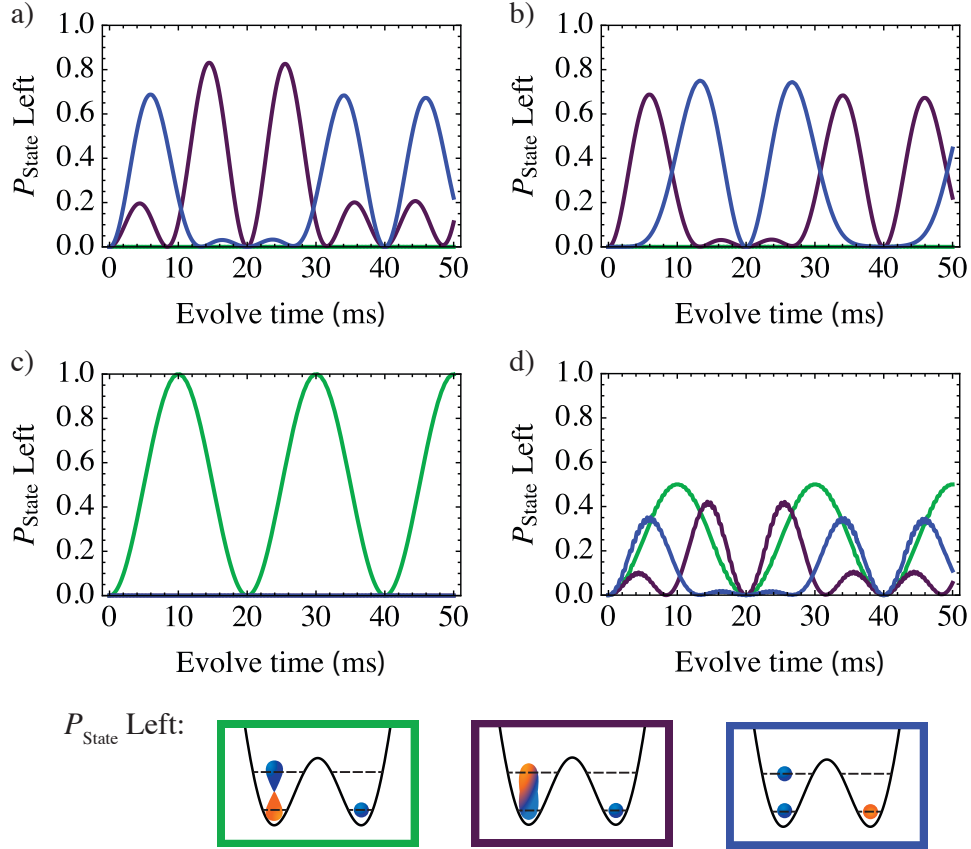


Figure 6.8: These plots show the evolution of the two-atom spin states in the left well for the same parameters as in Figure 6.7. This information gives some insight into the different behaviors for varied spin configurations. a) However, if we start with a triplet $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle + |\downarrow_g, \uparrow_e\rangle)]$, we can see that the probability to be in the triplet and $|\downarrow_g, \downarrow_e\rangle$ state is mixed over time. (Note that these states are both in the triplet manifold, but for clarity in notation I only refer to the $S_z = 0$ state as the triplet state.) b) A similar evolution occurs when we start in the $|\downarrow_g, \downarrow_e\rangle$ state, but with slightly different mixing. c) When starting with a singlet state $[\frac{1}{\sqrt{2}}(|\uparrow_g, \downarrow_e\rangle - |\downarrow_g, \uparrow_e\rangle)]$, the singlet is transferred between the wells at unit peak probability and at a single frequency. d) Now, if we start with the state $|\uparrow_g, \downarrow_e\rangle$, the dynamics will be a combination of the singlet and triplet states, meaning we project into all three spin states. This complex dynamics also includes some fast oscillations that are due to the large splitting between the singlet and triplet manifolds.

demonstrates that the singlet state transfers completely; this suggests that the localized singlet can be thought of as a quasiparticle that tunnels between the wells, which has an effective mass twice that of the bare atom in the excited band. Additionally, we show that starting with a mixed initial spin state, $|\uparrow_g, \downarrow_e\rangle$, we see a superposition of the same dynamics, but it is important to note that, in

the measurement basis, we would see very fast spin-exchange dynamics whenever there are nonzero components of both the singlet and triplet states.

6.3.4 Summary of ideal theoretical investigations

The above simulations have been useful to understand how the different interactions affect the dynamics observed in the minimum instance of the bosonic Kondo-Hubbard model (two sites with three atoms). Importantly, we see that it is possible to observe interesting dynamics related to both extremes of the model, given the ability to arbitrarily set the interaction and tunneling energy scales. Additionally, we note that the observables in such systems are experimentally accessible. Unless noted above, the simulations shown were either projecting out the population distribution in the excited band between the two wells (for the mass-enhanced regime), or the observed ground-band spin states in the (observable) basis of $|\uparrow\rangle$ and $|\downarrow\rangle$ on individual sites.

With knowledge that these interesting dynamics can be observed in our system, it is also worth considering the state initialization to understand a specific experimental procedure that would result in this observation. In particular, we must be conscious of other interactions and dynamics that will prevent us from quenching into the pure initial state that we want. In general, this means we must always have spin eigenstates during the preparation (before initiating the dynamics); for example, in Figure 6.3(c), we would not be able to initialize to opposite-spin atoms in the left pair of wells because they would undergo spin-exchange during the sweep that initializes the tunneling dynamics of the excited band atom. Additionally, when delocalizing an atom in the excited band, the adiabatic ramp would need to be slow compared to the tunneling rate, but fast compared to the exchange rate, setting a bound on the separation of energy scales that we can experiment with and dictating an optimal rate for the best possible state preparation.

With this in mind, I would aim for two specific implementations of the experiment, one exploring each of these extremes. To observe the ground-state spin coupling in the weakly interacting regime, I think simply measuring the probability for both atoms to be in the same spin state, as shown in Figure 6.6, will give the cleanest signal. While the signal amplitude is smaller than it

could be, there are only two frequencies of oscillation, which will simplify the analysis; the main question is what the absolute timescale will be in comparison to the tunneling dephasing.

The most promising experiment, in my opinion, is to first implement the adiabatic preparation of a singlet state, as discussed in Section 5.2 (which is a spin eigenstate), but leave the atoms in the same well after shutting off the magnetic field gradient. Then, by sweeping the spacing of the middle two wells [as shown in Figure 6.3(c)], we can initiate mass-enhanced tunneling dynamics, which can be compared to the exact same procedure, but with spin-polarized atoms (if all spins are the same, we would observe the bare tunneling rate), which should be a factor of two slower. In this experiment the state preparation is more complicated, but the ability to separate the challenge of preparation from the initialization of dynamics, combined with the very clear signal of the tunneling rate, makes this an attractive first experiment to demonstrate the coupling of the spin and motional degrees of freedom in our system.

6.3.4.1 Experimental limitations and available parameters

Now that we have an idea of the signals we are looking for, it's important to discuss the prospects for actually observing them in our experiment with the accessible trap parameters, rather than the idealized parameters used in the simulations. In principle, we can tune the shape of the trapping potential to cover the range of parameters extending all the way from the weakly to strongly interacting regimes (achieving ratios of $\frac{V_{eg}}{J_e}$ from 0.1 up to 10) simply by tuning the overall trap depth and spacing between the traps. Generally speaking, there are two limitations that we need to balance: The fidelity for preparing and initializing dynamics with the desired state, and the stability of the trapping potential setting the interaction parameters in the Hamiltonian. For example, we can make $V_{eg} \gg J_e$; however, the absolute value of each of these will be small because we can only reasonably achieve $V_{eg} \sim h \times 100$ Hz (without going to extremely large trap depths) and then would need to adjust the trap spacing to make the tunneling rate even smaller. However, with these parameters, the dynamics of interest will occur over hundred's of milliseconds, during which time noise in the system (e.g., fluctuations of the magnetic field or the bias between the




Figure 6.9: a) Spin dynamics shown for realistic trap parameters that should allow us to observe the ground-state spin coupling in the weakly interacting regime. Specifically, this simulation was run for a trap separation of 900 nm and single-well depth of $h \times 75$ kHz, trap parameters that are expected to result in the interaction parameters $V_{eg} = h \times 48$ Hz and $J_e = h \times 310$ Hz. b) Excited atom dynamics for a different set of parameters that should be sufficient to observe dynamics in the strongly interacting regime. Specifically, the trap separation is 905 nm and the single-well trap depth is $h \times 150$ kHz, which should result in the interaction parameters $V_{eg} = h \times 93$ Hz and $J_e = h \times 38$ Hz.

wells) will likely cause dephasing of the system.

Fortunately, both of these regimes have very clear observables, which should allow us to see the effects described above, even without long coherence times or moving far into the limits of strong or weak interactions. Specifically, in the weakly interacting limit we simply need to see a deviation of the difference between the dynamics with a symmetric and antisymmetric excited atom state (or, equivalently, a difference in the frequency of oscillation of the probability for the ground atoms to be in the $|\downarrow_R, \downarrow_L\rangle$ state, as shown in Figure 6.6). And in the limit of strong interactions, we simply need to see a single oscillation to see the difference compared to the bare tunneling rate. Examples of the dynamics we hope to observe (in the absence of dephasing), for parameters that approach both of these regimes, are shown in Figure 6.9.

For observing the dynamics in the weakly interacting regime, we plan to set the spacing between the tweezers that define the final double-well potential to 900 nm and a single-well trap depth of $h \times 75$ kHz (using the same definitions given in Ref. [39]). Based on calculations taking into account the full three dimensional shape of our tweezer potentials, this configuration is expected

to achieve the effective Hubbard parameters parameters $V_{eg} = h \times 48$ Hz and $J_e = h \times 310$ Hz, putting us very close to the weakly interacting regime. Figure 6.9(a) shows the signal observed on shorter timescales when initiating the ground atoms in the $|\downarrow_R, \uparrow_L\rangle$ state and measuring the probability to end in the $|\downarrow_R, \downarrow_L\rangle$ over time (as demonstrated in Figure 6.6). In this parameter regime, there are two experimental concerns that might reduce, or even wash out, the signal we are searching for: First, we must reliably be able to prepare the delocalized excited atom states. The plan for achieving this is to simply perform an adiabatic passage halfway and stop when the wells are in resonance. We have received positive results with some initial testing of this procedure in the experiment, but need to confirm that it is a coherent process (and not simply thermal) before we can be confident it will work reliably. Additionally, the timescale for the second-order coupling is very long, which likely means that the dephasing will occur before we see anything close to a full oscillation. However, we note that any significant deviation between the oscillation frequencies (or, equivalently, the deviation of the difference in spin probabilities from zero) is sufficient to verify the second order coupling. So, as shown in Figure 6.9, with dephasing times in the tens of milliseconds (as we can typically achieve for tunneling experiments, and even longer is possible for exchange experiments) we should be able to observe this deviation.

On the other hand, to observe dynamics in the strongly interacting regime, we plan to set the spacing between the tweezers to 905 nm with the single-well trap depth of $h \times 150$ kHz, which should provide the effective Hubbard parameters parameters $V_{eg} = h \times 93$ Hz and $J_e = h \times 38$ Hz. As can be seen in Figure 6.9(b), atoms in this trap clearly exhibit tunneling the characteristic slowing of the excited particles motion, however with some additional fast oscillations on imposed on the tunneling. In this case, the fast oscillations are not going to be problematic for differentiating between the two cases. However, there are also two concerns in this regime: First, we need to be able to generate a pure singlet on a single site before initiating the dynamics. This is another experiment that we think is possible, as discussed in Section 5.2, but the fidelity of this operation needs to be high, which has not been verified. And second, the dephasing time need to be longer than a single oscillation, meaning ~ 30 ms or longer. While we have certainly achieved these dephasing times

in specific trap configurations, it is not yet clear that this particular trap configuration can reach those dephasing times because the larger depth will increase the absolute amplitude of the noise compared to the tunneling rate (which may cause rapid dephasing).

Bibliography

- [1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of Bose-Einstein condensation in a dilute atomic vapor,” *Science* **269**, 198 (1995).
- [2] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, “Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions,” *Phys. Rev. Lett.* **75**, 1687 (1995).
- [3] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. Van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose-Einstein Condensation in a Gas of Sodium Atoms,” *Phys. Rev. Lett.* **75**, 3969 (1995).
- [4] B. DeMarco and D. S. Jin, “Onset of Fermi Degeneracy in a Trapped Atomic Gas,” *Science* **285**, 1703 (1999).
- [5] J. L. Roberts, N. R. Claussen, S. L. Cornish, and C. E. Wieman, “Magnetic Field Dependence of Ultracold Inelastic Collisions near a Feshbach Resonance,” *Phys. Rev. Lett.* **85**, 728 (2000).
- [6] T. Johnson, E. Urban, T. Henage, L. Isenhower, D. Yavuz, T. Walker, and M. Saffman, “Rabi Oscillations between Ground and Rydberg States with Dipole-Dipole Atomic Interactions,” *Phys. Rev. Lett.* **100**, 113003 (2008).
- [7] K.-K. Ni, S. Ospelkaus, M. H. G. de Miranda, A. Pe’er, B. Neyenhuis, J. J. Zirbel, S. Kotochigova, P. S. Julienne, D. S. Jin, and J. Ye, “A High Phase-space-density Gas of Polar Molecules,” *Science* **322**, 231 (2008).
- [8] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman, “Observation of Rydberg blockade between two atoms,” *Nat. Phys.* **5**, 110 (2009).
- [9] A. Gaëtan, Y. Miroshnychenko, T. Wilk, A. Chotia, M. Viteau, D. Comparat, P. Pillet, A. Browaeys, and P. Grangier, “Observation of collective excitation of two individual atoms in the Rydberg blockade regime,” *Nat. Phys.* **5**, 115 (2009).
- [10] T. Lompe, T. B. Ottenstein, F. Serwane, A. N. Wenz, G. Zürn, and S. Jochim, “Radio-Frequency Association of Efimov Trimers,” *Science* **330**, 940 (2010).
- [11] R. Schmidt, H. R. Sadeghpour, and E. Demler, “Mesoscopic Rydberg Impurity in an Atomic Quantum Gas,” *Phys. Rev. Lett.* **116**, 105302 (2016).

- [12] 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,” *Nature (London)* **415**, 39 (2002).
- [13] R. Jördens, N. Strohmaier, K. Günter, H. Moritz, and T. Esslinger, “A Mott insulator of fermionic atoms in an optical lattice,” *Nature (London)* **455**, 204 (2008).
- [14] L. Tarruell, D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger, “Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice,” *Nature (London)* **483**, 302 (2012).
- [15] S. Krinner, D. Stadler, D. Husmann, J.-P. Brantut, and T. Esslinger, “Observation of quantized conductance in neutral matter,” *Nature (London)* **517**, 64 (2014).
- [16] R. A. Hart, P. M. Duarte, T.-L. Yand, X. Liu, T. Paiva, E. Khatami, R. T. Scalettar, N. Trivedi, D. A. Huse, and R. G. Hulet, “Observation of antiferromagnetic correlations in the Hubbard model with ultracold atoms,” *Nature (London)* **519**, 1 (2015).
- [17] 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 (2016).
- [18] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, “A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-regime Optical Lattice,” *Nature (London)* **462**, 74 (2009).
- [19] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, “Single-atom-resolved fluorescence imaging of an atomic Mott insulator,” *Nature (London)* **467**, 68 (2010).
- [20] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, and S. Kuhr, “Observation of Correlated Particle-Hole Pairs and String Order in Low-Dimensional Mott Insulators,” *Science* **334**, 200 (2011).
- [21] P. Schauß, M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch, “Observation of spatially ordered structures in a two-dimensional Rydberg gas,” *Nature (London)* **491**, 87 (2012).
- [22] R. Islam, R. Ma, P. M. Preiss, M. E. Tai, A. Lukin, M. Rispoli, and M. Greiner, “Measuring entanglement entropy in a quantum many-body system,” *Nature (London)* **528**, 77 (2015).
- [23] M. F. Parsons, A. Mazurenko, C. S. Chiu, G. Ji, D. Greif, and M. Greiner, “Site-resolved observations of antiferromagnetic correlations in the Hubbard model,” arXiv:1605.02704 (2016).
- [24] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland, “Demonstration of a Fundamental Quantum Logic Gate,” *Phys. Rev. Lett.* **75**, 4714 (1995).
- [25] F. Schmidt-Kaller, H. Häffner, M. Rieve, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, “Realization of the Cirac-Zoller controlled-NOT gate,” *Nature (London)* **422**, 408 (2003).
- [26] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, “Demonstration of a Neutral Atom Controlled-NOT Quantum Gate,” *Phys. Rev. Lett.* **104**, 010503 (2010).

- [27] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, “Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade,” *Phys. Rev. Lett.* **104**, 010502 (2010).
- [28] J. W. Britton, B. C. Sawyer, A. C. Keith, C. C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, “Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins,” *Nature (London)* **484**, 489 (2012).
- [29] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C. C. J. Wang, J. K. Freericks, and C. Monroe, “Emergence and Frustration of Magnetism with Variable-Range Interactions in a Quantum Simulator,” *Science* **340**, 583 (2013).
- [30] K. M. Maller, M. T. Lichtman, T. Xia, Y. Sun, M. J. Piotrowicz, A. W. Carr, L. Isenhower, and M. Saffman, “Rydberg-blockade controlled-not gate and entanglement in a two-dimensional array of neutral-atom qubits,” *Phys. Rev. A* **92**, 022336 (2015).
- [31] H. Labuhn, D. Barredo, S. Ravets, S. de Leseleuc, T. Macri, T. Lahaye, and A. Browaeys, “Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models,” *Nature (London)* **534**, 667 (2016).
- [32] J. Beugnon, C. Tuchendler, H. Marion, A. Gaëtan, Y. Miroshnychenko, Y. R. P. Sortais, A. M. Lance, M. P. A. Jones, G. Messin, A. Browaeys, and P. Grangier, “Two-dimensional Transport and Transfer of a Single Atomic Qubit in Optical Tweezers,” *Nat. Phys.* **3**, 696 (2007).
- [33] B. Zimmermann, T. Müller, J. Meineke, T. Esslinger, and H. Moritz, “High-resolution Imaging of Ultracold Fermions in Microscopically Tailored Optical Potentials,” *New J. Phys.* **13**, 043007 (2011).
- [34] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland, “Coherent Diabatic Ion Transport and Separation in a Multizone Trap Array,” *Phys. Rev. Lett.* **109**, 080502 (2012).
- [35] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, “Single-Atom Trapping in Holographic 2D Arrays of Microtraps with Arbitrary Geometries,” *Phys. Rev. X* **4**, 021034 (2014).
- [36] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen, “Near-deterministic Preparation of a Single Atom in an Optical Microtrap,” *Nat. Phys.* **6**, 951 (2010).
- [37] A. V. Carpentier, Y. H. Fung, P. Sompet, A. J. Hilliard, T. G. Walker, and M. F. Andersen, “Preparation of a single atom in an optical microtrap,” *Laser Phys. Lett.* **10**, 125501 (2013).
- [38] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, “Rapid Production of Uniformly Filled Arrays of Neutral Atoms,” *Phys. Rev. Lett.* **115**, 073003 (2015).
- [39] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, “Two-particle quantum interference in tunnel-coupled optical tweezers,” *Science* **345**, 306 (2014).

- [40] D. S. Weiss, J. Vala, A. V. Thapliyal, S. Myrgren, U. Vazirani, and K. B. Whaley, “Another Way to Approach Zero Entropy for a Finite System of Atoms,” *Phys. Rev. A* **70**, 040302(R) (2004).
- [41] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletić, M. Greiner, and M. D. Lukin, “Cold Matter Assembled Atom-by-Atom,” arXiv:1607.03044 (2016).
- [42] D. Barredo, S. de Leseleuc, V. Lienhard, T. Lahaye, and A. Browaeys, “An atom-by-atom assembler of defect-free 2D atomic arrays,” arXiv:1607.03042 (2016).
- [43] A. M. Kaufman, B. J. Lester, and C. A. Regal, “Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State,” *Phys. Rev. X* **2**, 041014 (2012).
- [44] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletic, and M. D. Lukin, “Coherence and Raman sideband cooling of a single atom in an optical tweezer,” *Phys. Rev. Lett.* **110**, 133001 (2013).
- [45] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, “Single-spin addressing in an atomic Mott insulator,” *Nature (London)* **471**, 319 (2011).
- [46] H. Labuhn, S. Ravets, D. Barredo, L. Béguin, F. Nogrette, T. Lahaye, and A. Browaeys, “Single-atom addressing in microtraps for quantum-state engineering using Rydberg atoms,” *Phys. Rev. A* **90**, 023415 (2014).
- [47] A. M. Kaufman, B. J. Lester, M. Foss-Feig, M. L. Wall, A. M. Rey, and C. A. Regal, “Entangling two transportable neutral atoms via local spin-exchange,” *Nature (London)* **527**, 208 (2015).
- [48] B. J. Lester, A. M. Kaufman, and C. A. Regal, “Raman cooling imaging: Detecting single atoms near their ground state of motion,” *Phys. Rev. A* **90**, 011804(R) (2014).
- [49] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner, “Site-Resolved Imaging of Fermionic ${}^6\text{Li}$ in an Optical Lattice,” *Phys. Rev. Lett.* **114**, 213002 (2015).
- [50] L. W. Cheuk, M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein, “Quantum-Gas Microscope for Fermionic Atoms,” *Phys. Rev. Lett.* **114**, 193001 (2015).
- [51] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross, “Microscopic Observation of Pauli Blocking in Degenerate Fermionic Lattice Gases,” *Phys. Rev. Lett.* **115**, 263001 (2015).
- [52] S. E. Hamann, D. L. Haycock, G. Klose, P. H. Pax, I. H. Deutsch, and P. S. Jessen, “Resolved-Sideband Raman Cooling to the Ground State of an Optical Lattice,” *Phys. Rev. Lett.* **80**, 4149 (1998).
- [53] H. Perrin, A. Kuhn, I. Bouchoule, and C. Salomon, “Sideband Cooling of Neutral atoms in a Far-detuned Optical Lattice,” *Europhys. Lett.* **42**, 395 (1998).

- [54] V. Vuletic, C. Chin, A. J. Kerman, and S. Chu, “Degenerate Raman Sideband Cooling of Trapped Cesium Atoms at Very High Atomic Densities,” *Phys. Rev. Lett.* **81**, 5768 (1998).
- [55] A. J. Kerman, V. V. Vuletic, C. Chin, and S. Chu, “Beyond Optical Molasses: 3D Raman Sideband Cooling of Atomic Cesium to High Phase-Space Density,” *Phys. Rev. Lett.* **84**, 439 (2000).
- [56] D.-J. Han, S. Wolf, S. Oliver, C. McCormick, M. T. DePue, and D. S. Weiss, “3D Raman Sideband Cooling of Cesium Atoms at High Density,” *Phys. Rev. Lett.* **85**, 724 (2000).
- [57] A. M. Kaufman, Laser-cooling atoms to indistinguishability: Atomic Hong-Ou-Mandel interference and entanglement through spin exchange, Ph.D. thesis, University of Colorado at Boulder (2015).
- [58] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,” *J. Opt. Soc. Am. B* **6**, 2023 (1989).
- [59] J. Weiner, V. S. Bagnato, S. Zilio, and P. S. Julienne, “Experiments and theory in cold and ultracold collisions,” *Rev. Mod. Phys.* **71**, 1 (1999).
- [60] N. Schlosser, G. Reymond, and P. Grangier, “Collisional Blockade in Microscopic Optical Dipole Traps,” *Phys. Rev. Lett.* **89**, 023005 (2002).
- [61] M. T. DePue, C. McCormick, S. L. Winoto, S. Oliver, and D. S. Weiss, “Unity occupation of sites in a 3D optical lattice,” *Phys. Rev. Lett.* **82**, 2262 (1999).
- [62] A. Steffen, W. Alt, M. Genske, D. Meschede, C. Robens, and A. Alberti, “Note: In situ measurement of vacuum window birefringence by atomic spectroscopy,” *Rev. Sci. Instrum.* **84**, 126103 (2013).
- [63] D. Barredo, S. Ravets, H. Labuhn, L. Béguin, A. Vernier, F. Nogrette, T. Lahaye, and A. Browaeys, “Demonstration of a Strong Rydberg Blockade in Three-Atom Systems with Anisotropic Interactions,” *Phys. Rev. Lett.* **112**, 183002 (2014).
- [64] K. Eckert, J. Mompart, X. X. Yi, J. Schliemann, D. Bruss, G. Birkl, and M. Lewenstein, “Quantum computing in optical microtraps based on the motional states of neutral atoms,” *Phys. Rev. A* **66**, 042317 (2002).
- [65] D. Hayes, P. S. Julienne, and I. H. Deutsch, “Quantum Logic via the Exchange Blockade in Ultracold Collisions,” *Phys. Rev. Lett.* **98**, 070501 (2007).
- [66] V. Bolpasi and W. von Klitzing, “Double-pass tapered amplifier diode laser with an output power of 1 W for an injection power of only 200W,” *Rev. Sci. Instrum.* **81**, 113108 (2010).
- [67] R. Ma, Engineered potentials and dynamics of ultracold quantum gases under the microscope, Ph.D. thesis, Harvard University (2014).
- [68] P. M. Preiss, Atomic Bose-Hubbard Systems with Single-Particle Control, Ph.D. thesis, Harvard University (2015).
- [69] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais, and A. Browaeys, “Light-assisted collisions between a few cold atoms in a microscopic dipole trap,” *Phys. Rev. A* **85**, 062708 (2012).

- [70] P. Sompet, A. V. Carpentier, Y. H. Fung, M. McGovern, and M. F. Andersen, “Dynamics of two atoms undergoing light-assisted collisions in an optical microtrap,” *Phys. Rev. A* **88**, 051401(R) (2013).
- [71] D. A. Steck, “Rubidium 87 D Line Data,” available online at <http://steck.us/alkalidata> (revision 2.1.4, 23 December 2010).
- [72] Y. H. Fung and M. F. Andersen, “Efficient collisional blockade loading of a single atom into a tight microtrap,” *New J. Phys.* **17**, 073011 (2015).
- [73] F. Le Kien, P. Schneeweiss, and A. Rauschenbeutel, “Dynamical polarizability of atoms in arbitrary light fields: general theory and application to cesium,” *Eur. Phys. J. D* **67**, 92 (2013).
- [74] Y. H. Fung, A. V. Carpentier, P. Sompet, and M. Andersen, “Two-Atom Collisions and the Loading of Atoms in Microtraps,” *Entropy* **16**, 582 (2014).
- [75] A. Peng, Quantum Gas Microscope With Optical Lattice, Ph.D. thesis, Harvard University (2010).
- [76] S. Bergamini, B. Darquié, M. Jones, L. Jacubowicz, A. Browaeys, and P. Grangier, “Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator,” *J. Opt. Soc. Am. B* **21**, 1889 (2004).
- [77] P. Zupancic, P. M. Preiss, R. Ma, A. Lukin, M. E. Tai, M. Rispoli, R. Islam, and M. Greiner, “Ultra-precise holographic beam shaping for microscopic quantum control,” *Opt. Express* **24**, 13881 (2016).
- [78] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, “Sub-poissonian loading of single atoms in a microscopic dipole trap,” *Nature (London)* **411**, 1024 (2001).
- [79] K. D. Nelson, X. Li, and D. S. Weiss, “Imaging Single Atoms in a Three-dimensional Array,” *Nat. Phys.* **3**, 556 (2007).
- [80] M. Karski, L. Förster, J. M. Choi, W. Alt, A. Widera, and D. Meschede, “Nearest-Neighbor Detection of Atoms in a 1D Optical Lattice by Fluorescence Imaging,” *Phys. Rev. Lett.* **102**, 053001 (2009).
- [81] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, and S. Kuhr, “Single-site- and single-atom-resolved measurement of correlation functions,” *Appl. Phys. B* **113**, 27 (2013).
- [82] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais, and A. Browaeys, “Free-Space Lossless State Detection of a Single Trapped Atom,” *Phys. Rev. Lett.* **106**, 133003 (2011).
- [83] M. J. Gibbons, C. D. Hamley, C.-Y. Shih, and M. S. Chapman, “Nondestructive Fluorescent State Detection of Single Neutral Atom Qubits,” *Phys. Rev. Lett.* **106**, 133002 (2011).
- [84] Y. S. Patil, S. Chakram, L. M. Ayccock, and M. Vengalattore, “Nondestructive imaging of an ultracold lattice gas,” *Phys. Rev. A* **90**, 033422 (2014).
- [85] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, “Laser Cooling to the Zero-Point Energy of Motion,” *Phys. Rev. Lett.* **62**, 403 (1989).

- [86] L. Forster, M. Karski, J.-M. Choi, A. Steffen, W. Alt, D. Meschede, A. Widera, E. Montano, J. H. Lee, W. Rakreungdet, and P. S. Jessen, “Microwave Control of Atomic Motion in Optical Lattices,” *Phys. Rev. Lett.* **103**, 233001 (2009).
- [87] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, “Energy Distribution and Cooling of a Single Atom in an Optical Tweezer,” *Phys. Rev. A* **78**, 033425 (2008).
- [88] M. S. Safronova and U. I. Safronova, “Critically evaluated theoretical energies, lifetimes, hyperfine constants, and multipole polarizabilities in ^{87}Rb ,” *Phys. Rev. A* **83**, 052508 (2011).
- [89] B. Arora and B. K. Sahoo, “State-insensitive trapping of Rb atoms: Linearly versus circularly polarized light,” *Phys. Rev. A* **86**, 033416 (2012).
- [90] C.-Y. Shih and M. S. Chapman, “Nondestructive light-shift measurements of single atoms in optical dipole traps,” *Phys. Rev. A* **87**, 063408 (2013).
- [91] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto, “Controlled Exchange Interaction Between Pairs of Neutral Atoms in an Optical Lattice,” *Nature (London)* **448**, 452 (2007).
- [92] S. Trotzky, P. Cheinet, S. Fölling, M. Feld, U. Schnorrberger, A. M. Rey, A. Polkovnikov, E. A. Demler, M. D. Lukin, and I. Bloch, “Time-resolved Observation and Control of Superexchange Interactions with Ultracold Atoms in Optical Lattices,” *Science* **319**, 295 (2008).
- [93] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, “Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy,” *Phys. Rev. Lett.* **75**, 4011 (1995).
- [94] X. Li, T. A. Corcovilos, Y. Wang, and D. S. Weiss, “3D Projection Sideband Cooling,” *Phys. Rev. Lett.* **108**, 103001 (2012).
- [95] P. Maunz, T. Puppe, I. Schuster, N. Syassen, P. W. H. Pinkse, and G. Rempe, “Cavity Cooling of a Single Atom,” *Nature (London)* **428**, 50 (2004).
- [96] A. D. Boozer, A. Boca, R. Miller, T. E. Northup, and H. J. Kimble, “Cooling to the Ground State of Axial Motion for One Atom Strongly Coupled to an Optical Cavity,” *Phys. Rev. Lett.* **97**, 083602 (2006).
- [97] A. Reiserer, C. Nölleke, S. Ritter, and G. Rempe, “Ground-State Cooling of a Single Atom at the Center of an Optical Cavity,” *Phys. Rev. Lett.* **110**, 223003 (2013).
- [98] A. Kerman, Raman Sideband Cooling and Cold Atomic Collisions in Optical Lattices, Ph.D. thesis, Stanford University (2002).
- [99] L. Förster, Microwave control of atomic motion in a spin dependent optical lattice, Ph.D. thesis, Universität Bonn (2010).
- [100] D. J. Wineland, C. Monroe, D. M. Meekhof, B. E. King, D. Leibfried, W. Itano, J. C. Bergquist, D. Berkeland, J. J. Bollinger, and J. Miller, “Quantum state manipulation of trapped atomic ions,” *Proc. R. Soc. A* **454**, 411 (1998).

- [101] C. K. Hong, Z. Y. Ou, and L. Mandel, "Measurement of subpicosecond time intervals between two photons by interference," *Phys. Rev. Lett.* **59**, 2044 (1987).
- [102] R. Lopes, A. Imanaliev, A. Aspect, M. Cheneau, D. Boiron, and C. I. Westbrook, "Atomic Hong-Ou-Mandel experiment," *Nature (London)* **520**, 66 (2015).
- [103] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard, "Coherent Manipulation of Coupled Electron Spins in Semiconductor Quantum Dots," *Science* **309**, 2180 (2005).
- [104] M. L. Wall, K. R. A. Hazzard, and A. M. Rey, "Effective many-body parameters for atoms in nonseparable Gaussian optical potentials," *Phys. Rev. A* **92**, 013610 (2015).
- [105] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe, "Experimental entanglement of four particles," *Nature (London)* **404**, 256 (2000).
- [106] D. Liebfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jenlenkovic, C. Langer, T. Rosenband, and D. J. Wineland, "Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate," *Nature (London)* **422**, 412 (2003).
- [107] A. C. Hewson, The Kondo Problem to Heavy Fermions, Cambridge studies in magnetism: 2 (Cambridge University Press, 1997).
- [108] J. Kondo, "Resistance Minimum in Dilute Magnetic Alloys," *Prog. Theor. Phys.* **32**, 37 (1964).
- [109] H. Tsunetsugu, M. Sigrist, and K. Ueda, "The ground-state phase diagram of the one-dimensional Kondo lattice model," *Rev. Mod. Phys.* **69**, 809 (1997).
- [110] P. Coleman, "Heavy Fermions: electrons at the edge of magnetism," *Handbook of Magnetism and Advanced Magnetic Materials* (2007).
- [111] P. Coleman and A. H. Nevidomskyy, "Frustration and the Kondo Effect in Heavy Fermion Materials," *Journal of Low Temperature Physics* **161**, 182 (2010).
- [112] K. Andres, J. E. Graebner, and H. R. Ott, "4f-Virtual-Bound-State Formation in Ce Al₃ at Low Temperatures," *Phys. Rev. Lett.* **35**, 1779 (1975).
- [113] F. Steglich, J. Aarts, C. D. Bredl, W. Lieke, D. Meschede, W. Franz, and H. Schäfer, "Superconductivity in the Presence of Strong Pauli Paramagnetism: Ce Cu₂Si₂," *Phys. Rev. Lett.* **43**, 1892 (1979).
- [114] M. Ruderman and C. Kittel, "Indirect Exchange Coupling of Nuclear Magnetic Moments by Conduction Electrons," *Phys. Rev.* **96**, 99 (1954).
- [115] T. Kasuya, "A theory of metallic ferro- and antiferromagnetism on Zener's model," *Prog. Theor. Phys.* **16**, 45 (1956).
- [116] K. Yosida, "Magnetic Properties of Cu-Mn Alloys," *Phys. Rev.* **106**, 893 (1957).

- [117] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner, “Probing the Superfluid-to-Mott Insulator Transition at the Single-Atom Level,” *Science* **329**, 547 (2010).
- [118] L. M. Duan, “Controlling ultracold atoms in multi-band optical lattices for simulation of Kondo physics,” *Europhys. Lett.* **67**, 721 (2004).
- [119] M. Foss-Feig, M. Hermele, and A. M. Rey, “Probing the Kondo lattice model with alkaline-earth-metal atoms,” *Phys. Rev. A* **81**, 051603 (2010).
- [120] M. Foss-Feig and A. M. Rey, “Phase diagram of the bosonic Kondo-Hubbard model,” *Phys. Rev. A* **84**, 053619 (2011).
- [121] T. Flottat, F. Hebert, V. G. Rousseau, R. T. Scalettar, and G. G. Bartrouni, “Bosonic Kondo-Hubbard model,” *Phys. Rev. B* **92**, 035101 (2015).

Appendix A

Alignment and Characterization of Optical System

Alignment of an optical system should be a straightforward task, where the goal is simply to align the light to the optical axis of each element and set the spacings properly to maintain collimation, where appropriate, as well as generating images that are in focus. However, this procedure can be complicated and, in particular for high numerical aperture (NA) systems, the sensitivity of the system performance to this alignment may be extremely high. In our system, we have worked to separate these problems into the alignment of our optical rail system, described above, and then the alignment of the output beam to both the lens and the cell (which must be aligned with respect to each other).

For the optical relay rail, depicted in Figure 2.6, we have designed a system where the optical axis is mechanically defined via a precision machined rail, on which we mount fixed-height optics holders. This mechanical definition of the optical axis works well for the low-NA optics that are used on the rail (aside from the initial collimation lens, the shortest focal length lens is $f = 15$ mm for an input beam of $w_0 = 0.68$ mm, which would be an effective $\text{NA} < 0.05$) and simplifies the setup significantly. In particular, because the optical axis is mechanically defined, we can add reference pinholes along the two telescope arms of the rail (the first with $M = 1$, the second with $M = 20$), as indicated in Figure 2.6, and use the mirrors to align the collimated beam along this axis before adding the lenses to these paths. Once the beam is aligned through these references, the lenses can be mounted to generate an output beam that is collimated and that will be directed toward the cell and objective lens using a single mirror.

In fact, the most critical alignment in the system is that of the objective lens to the window through which the beam is focused (and scattered light imaged). This is a fundamental challenge imposed when focusing a high NA lens through a thick (6.35 mm) window because the objective is designed to compensate the optical path lengths for each part of the beam passing symmetrically through this window, but when the cone of rays is tilted through the window, these path lengths become asymmetric and there will be a mismatch of the phase-fronts at the focus. This mismatch primarily results in the third-order aberration coma, but also includes higher-order terms and, for real misaligned systems, will likely introduce other aberrations such as astigmatism. Because of this sensitivity to angle, it can be challenging to achieve a sufficiently high-precision alignment of the optical axis to the beam using transmission of the light through the objective, so we have instead turned to using a beam reflected off of the flat front-facet of the lens as one of our alignment references.

Unfortunately, simply adding this alignment reference to the others would result in over-constraining the system and, in the case of our lenses, these alignment references are often not fully consistent. In particular, this is likely due to the known presence of aberrations in the objective lens (discussed more in Section A.1, below), which makes the optimal alignment of the lens not occur when the optical axis is aligned to the cell. In this case, we typically choose two references that are reproducible and can be referred to an alignment that give optimal performance in an external testing setup, for example: 1) Setting a specific angle between the front facet of the lens and the cell window through which it is focusing (ideally this would be 0° , but as discussed below does not have to be) and 2) That a small beam transmitted through the lens remains on-axis at a given point behind the lens. Unfortunately, for aberrated systems where the angle of the lens is nonzero, this second metric will not be unique (in that it will likely depend on exactly where the reference point is located behind the lens), but the important thing is that it gets the lens very close to centered on the beam, since there is a larger tolerance to transverse displacement of the objective than to tilting of the lens.¹

¹ In principle, a better alignment reference would be to ensure that a small beam passing through the objective




Figure A.1: In this testing setup, we are directly imaging the point spread function of the objective lens (and thus the imaging system used in the experiment) by illuminating a sub-wavelength (350 nm diameter) aperture with 852 nm light. This results in a uniform illumination of the objective lens aperture as a spherical wave diverges from the aperture. The beam that is collimated and then imaged onto a CCD camera (using the standard imaging system from the experiment) will exhibit the best-focus that can be generated by the imaging system for the alignment of the imaging system. This provides the opportunity to adjust the alignment of the objective lens *in-situ* to optimize the performance, and then measure the optimal alignment using external references (which can then be used to align the objective to the cell).

A.1 External lens testing and alignment procedure

We have relied on two particular types of external testing to optimize the performance of our optical system, as well as to estimate the remaining aberrations present in the system. Importantly, we were able to use these testing setups to determine the proper method for aligning the optical system and to achieve the best optical tweezer performance, which has since been verified using by measuring trap properties with the atoms while varying parameters of the alignment.

A.1.1 Measurement of the point spread function (PSF) via pinhole imaging

The first measurements we performed (and the basis for our external characterization of the ASE lens) was to measure the point spread function (PSF) of the objective lens by imaging a sub-wavelength pinhole (which is small enough that we can effectively consider it a point-source) through the objective and onto a CCD camera. A simple schematic of this setup is shown in [Figure A.1](#). The beam propagates parallel to the optical axis in the absence of the objective (meaning it is going through the optical axis of the system, even if that is tilted and imparts a displacement on the beam). However, experimentally observing this is harder because it requires using multiple reference points along the path after transmission through the lens, and making sure that the offset of the beam from the center of each reference is identical at different points.

Figure A.1. The main benefit of this setup is that we can directly observe the optimal performance achievable with the objective lens, assuming a completely aberration-free source, and can change the alignment *in-situ* while directly observing the change in the focal shape.

As in the experiment itself, the alignment of this system is critical. However, if we set up the remainder of the system properly, we can leave the degrees of freedom relevant to the objective lens free to be manipulated while knowing that we continue to observe the effective spot size. Briefly, this alignment starts by defining the optical axis from the source to the CCD camera and aligning that optical axis to be normal to the window that acts as our cell. Next, we will align the imaging lens both to the optical axis, and ensure that the focal plane is in the plane of the CCD array. Then, we will add the objective lens and roughly align it to the optical axis (in this case, we use an iris between the objective and the imaging lens, as well as the location of the beam on the CCD to ensure that the transmitted beam follows the optical axis). Finally, we add a relatively small focal length lens before the window, which will focus light onto the back of the pinhole, and use it to recollimate (as much as possible) the beam after the objective lens. Then, we can simply add the pinhole at the focus between the two lenses, and position it by aligning it to the optical axis (when its image on the CCD is where the initial beam was located) and, at this point, we are measuring the point-spread-function (PSF) with the lens optical axis of the lens aligned to the window. However, to optimize the performance, we can adjust the combination of the lens alignment and pinhole location (using both to keep the focal spot at the same location on the camera) to find an alignment that gives a focus with minimum aberrations.

At the end of this procedure, once we have found the optimum performance of the lens, we can remove the pinhole and focusing lens and measure the angle of the lens with respect to the window for use in the alignment procedure in the experiment. Note that there are a couple of ways to measure this angle, but the optimum will use a long focal length lens, placed roughly one focal length away from the lens, and then the separation between the two foci will indicate the relative angle with very high precision. This is also the method that we use in the experimental setup to make sure the alignment is identical.

This setup also has the benefit of quantifying the magnification of our imaging system, which allows us to calibrate the spacing of our traps in situ. Specifically, we have used different test targets that consist of multiple pinholes arranged in close proximity, with known distances between these apertures (measured with a scanning electron microscope). In particular, our final calibration for the ASE objective was been performed with a test target that has two apertures, each 350 nm in diameter, separated by 6.861 μm , giving a magnification of 48.8(5) for the imaging system (the infinity-corrected objective and an $f = 1$ m achromat).

A.1.2 Wavefront measurement testing

More recently, we have put together a few different setups that are able to directly measure the wavefront aberrations introduced by the objective lens. For fully understanding and attempting to correct the wavefront aberrations present, this is the best method to isolate effects due to the objective lens itself. Previously, we tried to perform measurements with a Thorlabs Wavefront Sensor, but the results were unreliable. However, it is possible to set up a Michelson interferometer with the objective lens in the arm under test, which produces an interference pattern containing the wavefront aberrations introduced by the objective lens.

A.2 Characterization of the trapping potential

After setting up the system in-situ, we are able to use the atoms to characterize certain properties of the trapping potential. This is not a direct measurement of any aberrations, but a local probe of the harmonic potential that is generated for the atoms. The goal of this alignment procedure is, again, to achieve the highest harmonic confinement for a given power into the back of the lens, which will likely correspond to a minimum in the overall aberrations.

The key part of this procedure is to be able to measure the trap depth that the atom experiences, which is a measure of the intensity of light at the focus of the trap, as well as the trap frequency that the atoms experience. Combined, we can estimate the size of the trap, given a few assumptions about the intensity profile based that is based on the clipping of the input beam

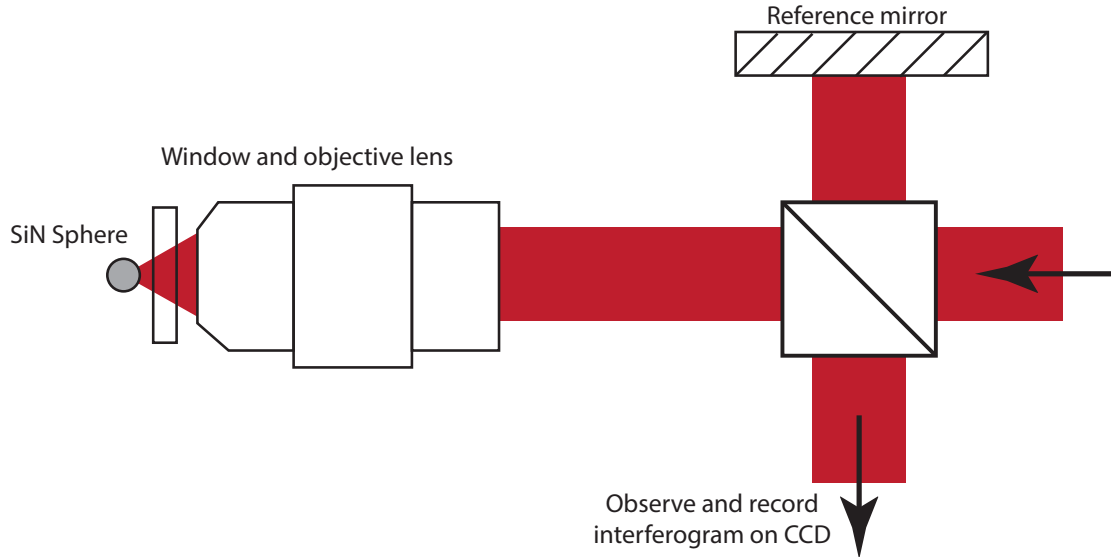


Figure A.2: In this testing setup, an interference pattern that contains information about the wavefront aberrations of the objective lens is generated using a Michelson interferometer. The reference path is defined with a reference mirror whose flatness has been verified on a commercial Zygo interferometer. Thus, the interference pattern arises solely from aberrations introduced on the lens arm of the interferometer and appear with twice the magnitude because the beam passes through the lens twice. By recording the interference pattern and calculating the Fourier transform, we can extract the wavefront curvature, and thus have full knowledge of the wavefront aberrations that need to be corrected for optimal performance.

and the expected PSF from our Zemax model. Note that the main role of the Zemax model here is used to estimate an effective gaussian waist that is representative of the actual size of the trap, given that for a clipped gaussian beam, the trap frequency is noticeably lower than that of an ideal gaussian beam of similar extent (meaning, if you fit the actual profile with a gaussian profile, a trap with that waist would have trap frequencies larger by more than 10%).

In the end, we estimate that we have an effective Gaussian waist of $w_0 = 0.7(1) \mu\text{m}$ based on the measured trap depth, trap frequency, and ratio of the harmonic component of the PSF calculated by Zemax (which is an approximation of the intensity profile) to that of a gaussian fit to the same calculated PSF. Notably, this estimate is consistent with the waist used in the full three-dimensional theory calculations given in Ref. [39] to understand the interactions and tunneling rates in a double-well potential.




Figure A.3: This figure shows how we can optimize the trapping frequencies that the atom experience by adjusting the alignment of the lens to the cell, as well as the collimation of the beam into the back of the objective lens. In particular, we note that the angle optimum angle of the lens with respect to the cell window is consistent with the optimum angle measured in the external testing setup.

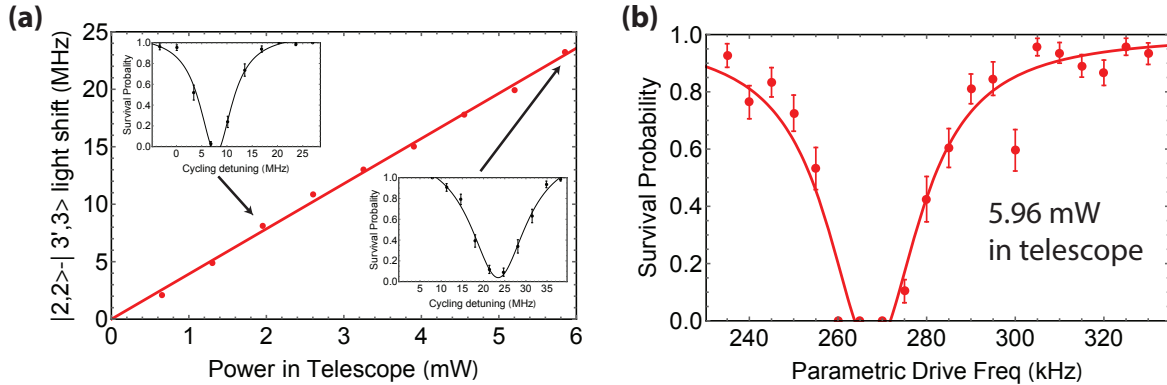


Figure A.4: This figure shows the method by which we characterize the intensity of the light at the atoms, where we measure the cycling transition resonance at different trap depths, then match the slope of this line with the expected light shift as a function of intensity of light at the atom. The key capability here is to use a pure σ^+ polarized beam, which simplifies the analysis of the light shift. In this case, we assume a pure linear polarized trapping potential, which is a reasonable assumption based on the additional vector light shifts being at most 100 kHz at full depth (as shown in Figure 5.4).

Appendix B

Derivation of eigenstates in minimal Kondo model

On the following pages, I am reproducing notes written by Leonid Isaev (from Ana Maria Rey's group), in which the exact eigenstates are derived for the model system composed of a double-well potential with one fixed spin in the ground-band of each well, as well as a single atom in the excited band of the double-well potential. The eigenstates are derived using arguments of symmetry, which lead to conserved quantities in the form of the total spin of the system, as well as the parity the wavefunction under exchange of the two well indices. Note that the notation for fermionic systems (e.g., the Kondo lattice model) was used when writing these notes, but for reference the energy scales are given by $t = J_e$ and $V = 2V_{eg}$.

The exact solutions can then be used to perform fast simulations with arbitrary interaction and tunneling parameters from any initial state (by projecting into the eigenbasis of the system, propagating in time, and then projecting back into observable basis states); this was the process used in simulations shown in Chapter 6. Importantly, this allows us to rapidly change the parameters of the simulation and calculate the change in the expected dynamics (as compared to numerically solving the Schrodinger equation, which is computationally intensive).

After deriving the exact eigenstates, the effective Hamiltonian given in Equation 6.3 is found using second-order perturbation theory in the weakly-interacting limit ($V \ll t$, in this notation, or $V_{eg} \ll J$ in the notation of Section 6.2).

Kondo box and geometric frustration with cold atoms

MODEL

We consider a double-well system shown in Fig. 1(a). The lowest levels are strongly localized, while the higher band splits due to particle tunneling between the wells [see Fig. 1(b)].

The double-well is populated with three spin- $\frac{1}{2}$ atoms: one in the excited band, and two localized in the lowest levels of individual wells. This system can be modeled by a Hamiltonian:

$$H = t \sum_{\sigma} (c_{1\sigma}^{\dagger} c_{2\sigma} + c_{2\sigma}^{\dagger} c_{1\sigma}) + V (\mathbf{S}_1 \cdot \mathbf{s}_{c1} + \mathbf{S}_2 \cdot \mathbf{s}_{c2}). \quad (1)$$

The first term describes the hopping (hybridization) between wells 1 and 2 with an amplitude t which we assume to be real and positive. The operator $c_{i\sigma}^{\dagger}$ creates a particle (the statistics does not matter) in the i -th well with spin projection σ : $|1_{i\sigma}\rangle = c_{i\sigma}^{\dagger} |\text{vac}\rangle$ ($|\text{vac}\rangle$ is an empty state). The second term is the exchange coupling between local atoms in each well and the mobile atom. We assume that $V \geq 0$ but this is not a crucial assumption because we will not be interested in the ground state properties. \mathbf{S}_i is the local spin- $\frac{1}{2}$ operator in i -th well, and $\mathbf{s}_{ci} = \frac{1}{2} \sum_{\alpha\beta} \sigma_{\alpha\beta} c_{i\alpha}^{\dagger} c_{i\beta}$ is the spin of the mobile particle (σ is the vector of Pauli matrices). Below we will use implicit summation over repeated spin indices.

EIGENSTATES OF THE MODEL HAMILTONIAN

The Hamiltonian (1) has two important symmetries. First, there is the spin $SU(2)$ continuous symmetry that guarantees conservation of the total spin $\mathbf{S}_{\text{tot}} = \mathbf{S}_1 + \mathbf{S}_2 + \mathbf{s}_{c1} + \mathbf{s}_{c2}$ of the system. S_{tot} can be either $\frac{1}{2}$ or $\frac{3}{2}$. The second one is a discrete ‘‘lattice’’ symmetry that corresponds to a mirror plane interchanging the wells (i.e. $1 \leftrightarrow 2$) and leads to the conservation of parity \hat{P} : $[\hat{H}, \hat{P}] = 0$. The parity transformation can be defined as $\hat{P} \mathbf{S}_1 \hat{P} = \mathbf{S}_2$ and $\hat{P} c_{1\sigma}^{\dagger} \hat{P} = c_{2\sigma}^{\dagger}$. Since $\hat{P}^2 = 1$, the parity operator has two eigenvalues $P = \pm 1$ (no hat). The corresponding three-particle states are schematically shown in Fig. 1(c) and (d). With the aid of these symmetries one can easily determine all eigenstates of the model (1).

In the following it will be convenient to diagonalize the hopping term in H by introducing symmetric (f) and antisymmetric (a) combinations of the mobile particle:

$$\begin{pmatrix} f_{\sigma} \\ a_{\sigma} \end{pmatrix} = \frac{1}{\sqrt{2}} (c_{1\sigma} \pm c_{2\sigma}), \quad (2)$$

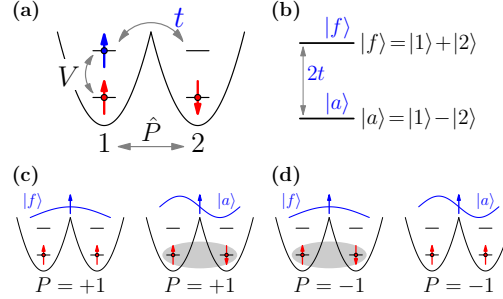


FIG. 1. Panel (a) The double-well setup with two local atoms (red arrows) and one mobile atom (blue arrow). The parity operator \hat{P} interchanges the two wells. (b) Hybridization splitting $2t$ of the mobile band. $|f\rangle$ and $|a\rangle$ denote symmetric and antisymmetric states (2). (c) Schematic illustration of different states with $P = +1$. The gray ellipse indicates a spin-singlet state across the wells. Parallel local spins correspond to a spin-triplet. (d) Same as in (c), but for totally antisymmetric states with $P = -1$.

and rewrite the Hamiltonian (1) as

$$\begin{aligned} H = & t (f_{\sigma}^{\dagger} f_{\sigma} - a_{\sigma}^{\dagger} a_{\sigma}) + \\ & + \frac{V}{2} (\mathbf{S}_1 + \mathbf{S}_2) \cdot (\mathbf{T}_{ff} + \mathbf{T}_{aa}) + \\ & + \frac{V}{2} (\mathbf{S}_1 - \mathbf{S}_2) \cdot (\mathbf{T}_{fa} + \mathbf{T}_{af}). \end{aligned} \quad (3)$$

Here $T_{ff} = \frac{1}{2} \sigma_{\alpha\beta} f_{\alpha}^{\dagger} f_{\beta}$, $T_{aa} = \frac{1}{2} \sigma_{\alpha\beta} a_{\alpha}^{\dagger} a_{\beta}$, $T_{fa} = \frac{1}{2} \sigma_{\alpha\beta} f_{\alpha}^{\dagger} a_{\beta}$ and $T_{af} = T_{fa}^{\dagger}$.

We will denote local spin states as $|\uparrow_i\rangle$ and $|\downarrow_i\rangle$ and use Clebsch-Gordan coefficients of build wavefunctions with a definite total spin out of the mobile $S = \frac{1}{2}$ and local two-spin states. When writing the three-particle states we will follow the naming scheme

$$|S_{\text{tot}}, S_{\text{tot}}^z; P; a\rangle,$$

where a is the index of a state with other quantum numbers fixed. Since the number of particles in the lowest levels is exactly one/level, the Hilbert space of Eq. (1) has dimension $2 \times 2^3 = 16$.

Sector $S_{\text{tot}} = S_{\text{tot}}^z = \frac{1}{2}$, $P = +$

There are two ways to have a state with total spin $S = \frac{1}{2}$: (i) local particles form a $S = 0$ singlet, so the spin comes from the mobile atom, and (ii) local spins are in the $S = 1$ triplet and we have to couple it to the

mobile spin using angular momentum addition theorems. Hence, there are two basis states:

$$\begin{aligned} \left| \frac{1}{2}, \frac{1}{2}; +; 1 \right\rangle &= \frac{1}{\sqrt{2}} |\uparrow_1 \downarrow_2 - \downarrow_1 \uparrow_2\rangle \otimes a_{\uparrow}^{\dagger} |\text{vac}\rangle, \\ \left| \frac{1}{2}, \frac{1}{2}; +; 2 \right\rangle &= \frac{1}{\sqrt{6}} [2|\uparrow_1 \uparrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle - \\ &\quad - |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle]. \end{aligned}$$

The first ket is built from antisymmetric local-spin singlet and an antisymmetric fermion combination, while the second one is constructed using all symmetric states. In this basis the Hamiltonian matrix has the form

$$H_{\frac{1}{2}, \frac{1}{2}; +} = \begin{pmatrix} -t & -\frac{\sqrt{3}V}{4} \\ -\frac{\sqrt{3}V}{4} & t - \frac{1}{2}V \end{pmatrix}.$$

On the diagonal, t originates from the kinetic energy, and $-\frac{V}{2}$ from the interaction of the mobile atom with a mean local magnetization in the 2nd state. The eigenvalues and corresponding eigenfunctions are:

$$E_{1,2} = -\frac{V}{4} \mp R$$

with $R = \sqrt{(t - V/4)^2 + (\sqrt{3}V/4)^2}$, and

$$\begin{aligned} |\psi_1\rangle &= \cos \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; +; 1 \right\rangle + \sin \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; +; 2 \right\rangle, \\ |\psi_2\rangle &= -\sin \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; +; 1 \right\rangle + \cos \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; +; 2 \right\rangle, \end{aligned}$$

where $\cos \vartheta = (t - V/4)/R$ and $\sin \vartheta = \sqrt{3}V/4Rt$.

$$\text{Sector } S_{\text{tot}} = S_{\text{tot}}^z = \frac{1}{2}, P = -$$

In complete analogy with the previous subsection, we have the basis:

$$\begin{aligned} \left| \frac{1}{2}, \frac{1}{2}; -; 1 \right\rangle &= \frac{1}{\sqrt{2}} |\uparrow_1 \downarrow_2 - \downarrow_1 \uparrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle, \\ \left| \frac{1}{2}, \frac{1}{2}; -; 2 \right\rangle &= \frac{1}{\sqrt{6}} [2|\uparrow_1 \uparrow_2\rangle \otimes a_{\downarrow}^{\dagger} |\text{vac}\rangle - \\ &\quad - |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes a_{\uparrow}^{\dagger} |\text{vac}\rangle]. \end{aligned}$$

The Hamiltonian matrix can be obtained from that in the previous subsection by replacing $t \rightarrow -t$:

$$H_{\frac{1}{2}, \frac{1}{2}; -} = \begin{pmatrix} t & -\frac{\sqrt{3}V}{4} \\ -\frac{\sqrt{3}V}{4} & -t - \frac{1}{2}V \end{pmatrix}.$$

The spectrum and eigenstates can be written as before, but with a small change in labeling,

$$E_{1,2} = -\frac{V}{4} \pm R$$

with $R = \sqrt{(t + V/4)^2 + (\sqrt{3}V/4)^2}$, and

$$\begin{aligned} |\psi_1\rangle &= \cos \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; -; 1 \right\rangle + \sin \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; -; 2 \right\rangle, \\ |\psi_2\rangle &= -\sin \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; -; 1 \right\rangle + \cos \frac{\vartheta}{2} \left| \frac{1}{2}, \frac{1}{2}; -; 2 \right\rangle, \end{aligned}$$

where $\cos \vartheta = (t + V/4)/R$ and $\sin \vartheta = -\sqrt{3}V/4Rt$.

$$\text{Sector } S_{\text{tot}} = \frac{1}{2}, S_{\text{tot}}^z = -\frac{1}{2}, P = \pm$$

The basis states can be chosen as

$$\begin{aligned} \left| \frac{1}{2}, -\frac{1}{2}; +; 1 \right\rangle &= \frac{1}{\sqrt{2}} |\uparrow_1 \downarrow_2 - \downarrow_1 \uparrow_2\rangle \otimes a_{\downarrow}^{\dagger} |\text{vac}\rangle, \\ \left| \frac{1}{2}, -\frac{1}{2}; +; 2 \right\rangle &= -\frac{1}{\sqrt{6}} [2|\downarrow_1 \downarrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle - \\ &\quad - |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle], \\ \left| \frac{1}{2}, -\frac{1}{2}; -; 1 \right\rangle &= \frac{1}{\sqrt{2}} |\uparrow_1 \downarrow_2 - \downarrow_1 \uparrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle, \\ \left| \frac{1}{2}, -\frac{1}{2}; -; 2 \right\rangle &= -\frac{1}{\sqrt{6}} [2|\downarrow_1 \downarrow_2\rangle \otimes a_{\uparrow}^{\dagger} |\text{vac}\rangle - \\ &\quad - |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes a_{\downarrow}^{\dagger} |\text{vac}\rangle]. \end{aligned}$$

The Hamiltonian matrices and their eigenstates are the same as in corresponding parity sectors with $S_{\text{tot}}^z = \frac{1}{2}$.

$$\text{Sector } S_{\text{tot}} = \frac{3}{2}, \text{ all } S_{\text{tot}}^z, P = +$$

There are $2S_{\text{tot}} + 1 = 4$ states with different S_{tot}^z :

$$\begin{aligned} \left| \frac{3}{2}, -\frac{3}{2}; + \right\rangle &= |\downarrow_1 \downarrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle, \\ \left| \frac{3}{2}, -\frac{1}{2}; + \right\rangle &= \frac{1}{\sqrt{3}} [|\downarrow_1 \downarrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle + \\ &\quad + |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle], \\ \left| \frac{3}{2}, \frac{1}{2}; + \right\rangle &= \frac{1}{\sqrt{3}} [|\uparrow_1 \uparrow_2\rangle \otimes f_{\downarrow}^{\dagger} |\text{vac}\rangle + \\ &\quad + |\uparrow_1 \downarrow_2 + \downarrow_1 \uparrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle], \\ \left| \frac{3}{2}, \frac{3}{2}; + \right\rangle &= |\uparrow_1 \uparrow_2\rangle \otimes f_{\uparrow}^{\dagger} |\text{vac}\rangle. \end{aligned}$$

Of course, these states are not mixed by the Hamiltonian, hence they are already eigenstates with an energy

$$E = t + \frac{1}{4}V.$$

$$\text{Sector } S_{\text{tot}} = \frac{3}{2}, \text{ all } S_{\text{tot}}^z, P = -$$

This sector can be obtained from the previous one by replacing $f \rightarrow a$ and $t \rightarrow -t$.

RKKY REGIME $V \ll t$

It is instructive to build a low-energy effective Hamiltonian for the case of small exchange coupling (compared to the bandwidth). In this regime, the model (1) can be seen as a frustrated triangular plaquette with two local and one mobile spin being at the vertices. The sign of the effective exchange coupling can be controlled by populating either symmetric (f -type) or antisymmetric (a -type) band.

When $V = 0$, there are two 8-fold degenerate sectors corresponding to the mobile atom being either in a symmetric or antisymmetric state (the degeneracy comes from spins). For small but finite V , this manifold is split by the exchange interaction. Hence, we start from Eq. (3) whose first two terms are diagonal and do not mix the hybridized states in the upper band, and the last line is an off-diagonal perturbation. We will call the diagonal terms in Eq. (3) as H_0 and the last off-diagonal exchange correction – as H_1 . The latter can be removed in each order on V/t by a Schrieffer-Wolff transformation $H \rightarrow e^\Lambda H e^{-\Lambda} \approx H_0 + \frac{1}{2}[\Lambda, H_1]$ with Λ chosen so that $[H_0, \Lambda] = H_1$. To the lowest order in V we have

$$\Lambda = \frac{V}{2} \sum_{E'_0, E_0} \frac{\mathbf{S}_{12} \cdot P'_0 \mathbf{T}_{fa} P_0}{E'_0 - E_0} - \text{h.c.} = U - U^\dagger,$$

where $\mathbf{S}_{12} = \mathbf{S}_1 - \mathbf{S}_2$. In the summation, E_0 (and E'_0) denotes the full set of quantum numbers that characterize a degenerate state of H_0 (here only the kinetic energy).

P_0 (and P'_0) are projectors on this subspace.

If the mobile particle is in the lowest-energy antisymmetric state [see Fig. 1(b)], the 2nd order correction $[\Lambda, H_1] = [U, H_1] + \text{h.c.}$ is

$$\begin{aligned} & \frac{1}{2} \langle \text{vac} | a_{\sigma'} [\Lambda, H_1] a_\sigma^\dagger | \text{vac} \rangle = \\ & = -\frac{V^2}{4} \sum_{E'_0, E_0} \frac{S_{12}^i S_{12}^j}{E'_0 - E_0} \langle \text{vac} | a_{\sigma'} T_{af}^i P'_0 T_{fa}^j P_0 a_\sigma^\dagger | \text{vac} \rangle = \\ & = -\frac{V^2}{32t} S_{12}^i S_{12}^j (\sigma^i \sigma^j)_{\sigma' \sigma} = \\ & = -\frac{3V^2}{64t} + \frac{V^2}{16t} [\mathbf{S}_1 \cdot \mathbf{S}_2 + \boldsymbol{\tau} \cdot (\mathbf{S}_1 + \mathbf{S}_2)]. \end{aligned}$$

Adding the linear contribution [second line in Eq. (3)], we obtain an effective spin Hamiltonian (up to a constant energy shift)

$$H_{\text{ef}} = V \boldsymbol{\tau} \cdot (\mathbf{S}_1 + \mathbf{S}_2) + \frac{V^2}{16t} [\mathbf{S}_1 \cdot \mathbf{S}_2 + \boldsymbol{\tau} \cdot (\mathbf{S}_1 + \mathbf{S}_2)],$$

where $\tau_{\alpha\beta} = \langle \text{vac} | a_\alpha \mathbf{T}_{aa} a_\beta^\dagger | \text{vac} \rangle$ is the spin- $\frac{1}{2}$ operator of the antisymmetric mode. The quadratic term in this expression is an antiferromagnetic Heisenberg model on a geometrically frustrated triangle.

The case when the mobile atom is in the symmetric state can be treated in exactly the same way. The resulting effective Hamiltonian can be obtained from H_{ef} by replacing \mathbf{T}_{aa} with \mathbf{T}_{ff} and $t \rightarrow -t$. In this case the effective Heisenberg model is ferromagnetic and the triangle is non-frustrated.