

The Pennsylvania State University
The Graduate School
Eberly College of Science

**NEUTRAL ATOM QUANTUM COMPUTING : QUANTUM GATES
AND MAXWELL'S DEMON**

A Dissertation in
Physics
by
Aishwarya Kumar

© 2019 Aishwarya Kumar

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

May 2019

The dissertation of Aishwarya Kumar was reviewed and approved* by the following:

David S. Weiss
Professor of Physics
Dissertation Advisor, Chair of Committee

Nitin Samarth
Professor of Physics
Department Head and Chair of Graduate Program

Kenneth O'Hara
Associate Professor of Physics

Benjamin J. Lear
Associate Professor of Chemistry

*Signatures are on file in the Graduate School.

Abstract

Quantum computers might solve several problems thought to be intractable on classical computers. These include problems like the prime factorization of large numbers, searching over a solution space and simulations of complex quantum systems, out of which the latter is the most exciting for physicists. Over the past few years, the experimental technologies have matured enough that there has been a strong effort to build the first 50 – 100 qubit computers in several candidate systems, one of which is neutral atoms. Neutral atoms in light traps have been a leading technology for quantum simulations for many years, but when trapped in optical lattices, they are also promising qubit candidates for a scalable quantum computer. Over the past few years, exquisite control over positions and the internal states of individual atoms has been developed towards this goal. This thesis describes development of such control in our three dimensional array of Cesium atoms.

We have developed an addressing scheme to target single atoms and execute high fidelity quantum gates without affecting the stored quantum information in neighboring atoms, a challenging feat in a 3D geometry. As we will see, atoms cannot be deterministically loaded in to these arrays, and typically only about half the sites have a single atom. Later in the thesis we use our single site control and combine it with state dependent lattices to move individual atoms and re-arrange them to yield fully-filled sub-lattices, creating nearly perfect registers of neutral atoms and initializing our quantum computer. We show that this atom "sorting" is an implementation of the famous "demon" thought experiment proposed by James Clerk Maxwell in 1872. Our experiment, like the demon, decreases the entropy of the system in the process of creating an ordered state from a random disordered one. Finally, we use our state dependent lattice to implement a loss-less state detection protocol with 20 times lower error than any other methods.

Table of Contents

List of Figures	vii
List of Tables	x
Acknowledgments	xi
Chapter 1	
Introduction	1
1.1 The state of the art in quantum computation	3
1.2 The state of the art in neutral atom quantum computation	5
1.3 Atomic physics preliminaries	7
1.4 Organization of this thesis	10
Chapter 2	
Overview of the experiment	12
2.1 Introduction	13
2.2 Three dimensional optical lattice	14
2.2.1 Loading the lattice	15
2.2.2 Imaging the atoms	15
2.2.3 Updates	18
2.3 Manipulating atomic states	19
2.3.1 Microwave system	19
2.3.2 Addressing system	21
2.3.3 Updates	23
2.4 Qubit quality	24
2.4.1 Projection sideband cooling	24
2.4.2 Coherence time	27
Chapter 3	
Single qubit gates	31

3.1	Introduction	32
3.2	Targeted addressing	33
3.3	A targeted gate	36
	3.3.1 Phase shifts	37
	3.3.2 Phase shift cancellation	38
	3.3.3 Results	40
3.4	A targeted phase gate	42
	3.4.1 Robust dynamical decoupling	45
	3.4.2 Experimental implementation	46
	3.4.3 Randomized benchmarking	48
	3.4.4 Sources of errors	52
3.5	Conclusion	53

Chapter 4

	Sorting atoms in 3D: an omniscient Maxwell's demon	55
4.1	Introduction	56
4.2	Moving atoms	58
	4.2.1 State dependent lattices	59
	4.2.2 Site selective state flips	61
	4.2.3 Filling a vacancy	64
	4.2.4 Motion fidelities	67
4.3	Sorting Algorithm	69
	4.3.1 Balancing	71
	4.3.2 Modifications to the algorithm	73
4.4	Results	74
	4.4.1 Entropy and Maxwell's demon	77
	4.4.2 Errors and scaling	79
	4.4.3 Improving the errors: throw and catch	81
4.5	Conclusion and outlook	86

Chapter 5

	Lossless state detection	88
5.1	Introduction	88
5.2	The X_s lattice	92
5.3	Imaging in the X_sYZ lattice	96
	5.3.1 PGC phase scrambling	97
	5.3.2 Fitting in the X_sYZ lattice	98
5.4	Final occupancy map generation in the X_sYZ lattice : a third imaging	99
	5.4.1 Re-loading atoms into the XYZ lattice	101
	5.4.2 Using information from the third set of images	103

5.5	Detection of $ 4, -4\rangle$ and $ 3, -3\rangle$	106
5.6	Detection from $ 4, 0\rangle$ and $ 3, 0\rangle$	110
5.7	Conclusion	113
Chapter 6		
	Future work	114
6.1	Introduction	114
6.2	Rydberg gates	115
6.2.1	Fidelity considerations	117
6.2.2	UV+microwave excitation	119
6.2.3	Single site addressing	123
6.3	Cluster states	124
	Bibliography	126
	Appendix A	
	Real time experimental control	138
	Appendix B	
	Notes on addressing beam alignment and calibration	142
B.1	Calibration of the addressing beam positions and powers	143
	Appendix C	
	Implementation of atom sorting in C++	145
C.1	class SortingClass	145
C.1.1	Key data structures in SortingClass	146
C.1.2	Glossary of key functions in SortingClass	148
C.1.3	A commented sorting example	149
C.2	Data flow in Panel	153
C.3	Some thoughts on code expansion	154
	Appendix D	
	Mathematica file for throw and catch	156

List of Figures

1.1	Cesium hyperfine structure	8
1.2	Bloch sphere representation	9
2.1	Experiment cell and lattice beam geometry	14
2.2	Imaging of different planes of atoms	16
2.3	Imaging histograms for four cases	18
2.4	Spectrum of AFP and Blackman pulses	21
2.5	Optical transfer system for the addressing beams	22
2.6	Results of projection sideband cooling	26
2.7	Dynamical decoupling pulse sequence	27
2.8	Interference fringes for coherence time measurement	28
2.9	Coherence time	29
3.1	Addressing scheme and Stark shifts	34
3.2	Spectrum with addressing and illustrative images of atoms	35
3.3	Targeted gates in a dynamical decoupling sequence	37

3.4	Crosstalk cancellation and results for the first gate.	40
3.5	Targeted phase gate	43
3.6	Robust dynamical decoupling	45
3.7	$R_z(\pi/2)$ gates	47
3.8	$R_z(\pi)$ gates	49
3.9	Randomized benchmarking microwaves	50
3.10	Randomized benchmarking data for non-target atoms	51
3.11	Randomized benchmarking data for target atoms	52
4.1	State dependent lattices	59
4.2	Depth and displacement of state dependent lattices	61
4.3	Setup for rotating the polarization.	62
4.4	AFP spectrum for addressed sites.	63
4.5	Filling a vacancy	65
4.6	Timing sequence for a motion step	66
4.7	Experiment cell and lattice beam geometry	68
4.8	Balancing and compaction	70
4.9	The balancing tree.	72
4.10	Sorted $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sub-lattices.	76
4.11	Filling fraction and entropy.	77
4.12	Throw and catch	82

5.1	Steps in lossless state detection	91
5.2	X_s lattice beam geometry.	93
5.3	X_s trap frequency measurement	94
5.4	In-situ displacement distribution	95
5.5	Typical state detection images.	97
5.6	Histograms with the updated fitting procedure.	100
5.7	Information from the third imaging.	102
5.8	Occupancy assignment flow chart for state detection.	104
5.9	Displacement distributions and state assignment.	107
5.10	Displacement asymmetry.	108
5.11	Re-initialization of the 3D qubit array.	109
5.12	State-selective detection from the clock states.	111
6.1	A Rydberg blockade based C_Z gate.	116
6.2	UV+microwave excitation.	119
6.3	Parameters for the two photon transition.	121
A.1	Schematic for real time control	139

List of Tables

3.1	Fidelities for the first kind of gate	42
3.2	Errors and contribution from various sources for the second kind of gate in units of 10^{-4}	52
C.1	Motion direction convention in <code>SortingClass</code>	146

Acknowledgments

‘No man is an island’. While Donne used this phrase in a different context, I think it is entirely applicable to the long journey of getting a Ph.D., especially mine. There are many people without whom this dissertation would not have been possible, and to whom I am deeply grateful.

First and foremost, I want to thank my thesis advisor, Prof. David S. Weiss, who has diligently guided me every step of the way. Over the years, he has not only taught me the intricacies of atomic physics, but also, scientific writing and effective communication. Most importantly though, he taught me how to be a good scientist and a physicist. He showed me the importance of thinking about physics by drawing pictures and effectively using “ball-and-stick” models to get to the heart of the problems. He was all that I could have hoped for in a Ph.D. advisor and for that I consider myself extremely fortunate.

While one might learn all the theory behind any experiment on their own, being able to run an experiment is an entirely different ball game. Prof. Ted Corcovilos, Dr. Yang Wang, and Dr. Xianli Zhang were some of the first people that I worked with on this experiment and they taught me the details of running the experiment. For that, I owe them a debt of gratitude. I was later joined by Tsung-Yao Wu and Felipe Giraldo Mejia, who were instrumental in the demonstration of atom sorting and lossless state detection. Tsung-Yao, in addition to being a great scientist, has some of the greatest organizational skills I have ever seen and I was lucky that he was there to supplement the relatively abysmal organizational skills that I have. Felipe has a keen eye for detail and I would also like to thank him for proofing this thesis. I grateful to other members of the Weiss Lab: Josh Wilson, Neel Malvania, Teng Zhang, Cheng Tang, Laura Zundel, Lin Xia, and Yuan Le, for all their help and many fruitful discussions over the years.

In the same vein, I would like to thank all the members of the AMO community here at Penn State, in particular: Prof. Kurt Gibble, Prof. Kenneth O’Hara, Prof. Marcos Rigol, Prof. Mikael Rechtsman, Dr. Andrew Marcum, Dr. Craig Price, Daniel Schussheim, and Francisco Fonta, who were always generous with their time

an resources. I am especially thankful to Prof. O'Hara for all his help with the Direct Digital Synthesizer and the designs for the tapered amplifier mounts. I am also grateful to the other members of my comittee, Prof. Nitin Samarth and Prof. Ben Lear for their invaluable feedback.

Getting a Ph.D. thousands of miles away from your home and family, in a foreign land, is a challenging task. I would not have been able to do it without the help of some great friends I was lucky enough to make in my time here. I would like to thank Dr. Nicholas Benjamin Senno and Justin Lanfranchi for showing me the ropes when I first got here. I am also thankful to Dr. Chris Dilks and Andre Harner. Making and playing music with them was some of the best time I had here and grad school would have been nigh impossible without that outlet. I am grateful to Dr. Thomas Foley, Dr. Amber McCreary and Dr. Jorge Zanudo for their friendship.

I was extremely lucky to have found all-weather friends in Justin Lanfranchi, Liz McCarthy and Dr. Nicole Benevento. There were there to help whenever I needed it the most. Thank you for everything.

Finally, I would like to thank my parents and my brother, who made me who I am and have always supported me in whatever I wanted to do. This is their work as much as it is mine.

The work described in this dissertation was supported by the DARPA QuEST program and National Science Foundation grant numbers PHY-1520976 and PHY-1820849.

For my parents

Chapter 1 |

Introduction

"But I am not afraid to consider the final question as to whether, ultimately – in the great future – we can arrange the atoms the way we want; the very atoms, all the way down! What would happen if we could arrange the atoms one by one the way we want them....."

– Richard Feynman, *There's Plenty of Room at the Bottom*

It is almost a convention to start a text on quantum computing with Richard Feynman's famous quote about how efficient simulations of quantum phenomena must themselves be quantum mechanical [1]. I have chosen a different quote, more appropriate for this thesis, in which we describe the kind of control over individual atoms that Feynman might have imagined, with the hope of one day realizing Feynman's (and many others') dream of a building quantum computer using these atoms.

Simply speaking, a quantum computer is a device that utilizes quantum superpositions, entanglement and many-body interference to solve problems that have no known efficient¹ solutions on classical computers. While the details of how quantum algorithms achieve this speed-up are out of the scope of this thesis, their development laid the foundations of the field as it stands today and had a deep impact on the theory of computation. Models of computation which used quantum mechanics were first explored by Benioff [2] and Manin [3] in 1980. In 1982, Feyn-

¹For the purposes of this thesis, efficient means that the running time scales polynomially with the size of the problem.

man [1] proposed the idea of a universal quantum simulator which would "exactly" simulate a quantum mechanical system. In 1985, David Deutsch [4] described the quantum Turing machine and proposed the idea of a universal quantum computer, providing a solid theoretical foundation for the nascent field of quantum information. The first quantum algorithms that demonstrated a significant speed-up over their classical counterparts were discovered in the 1990s, starting with the Deutsch-Josza algorithm [5]. In 1994, Peter Shor [6] proposed his famous algorithm to factorize integers in polynomial time on a quantum computer. Although this problem has not been shown to be inefficient on a classical machine, it is widely believed that it is the case, so much so that many of the most powerful encryption schemes are based on the hardness of factoring large numbers.

Even after these developments, quantum computers remained little more than intellectual curiosities. Quantum bits or qubits, the quantum analogues of classical bits, can interact with the environment and rapidly de-cohere, losing any quantum information stored in them. Quantum gates, which are physical implementations of unitary transformations on qubits and are required to execute a quantum algorithm, always have errors that can propagate through the course of execution. A quantum computer that could implement Shor's algorithm on an integer of reasonable size would require error rates many orders of magnitudes smaller than thought possible [7], relegating quantum computers to the same status as the analog classical computers of yore.

The first breakthroughs came in 1995 and 1996, when Shor [8] and Steane [9] described the first quantum error correction codes and showed that a *logical* qubit could be encoded in several *physical* qubits and the errors introduced by the environment and quantum gates could be corrected to execute quantum algorithms with exponentially more gates than previously anticipated. These results were further refined by Aharonov and Ben-Or [10], who proved the threshold theorem, which essentially says that errors can be corrected in a quantum computer faster than they accumulate as long as the error per gate is smaller than some constant threshold. This was called *fault-tolerant* quantum computation and sets quantum computation apart from analog classical computers, where errors cannot be corrected. The generally accepted threshold for fractional error rate for fault tolerance is about 10^{-4} [11], although more recent error correction codes have pushed the thresholds to as large as a percent [12]! It is worthwhile to mention that even with 0.1% error

per gate, several thousand physical qubits would be required to realize a single fault-tolerant logical qubit, in addition to the fact that creating many qubits and gates with such a small error is still a daunting task. Therein lies the difficulty of quantum computation; the extreme fragility of a quantum system due to unwanted couplings with the environment. Thus, the requirement is to build quantum systems which are exquisitely isolated from the environment and, at the same time, can be manipulated by using external means and made to controllably interact with each other. These systems also should be able to scale to large sizes without impacting the gate fidelities [13].

Soon after these results, qubits were implemented in nuclear magnetic resonance (NMR) systems [14], with proof of principle demonstrations of quantum gates and even Shor's algorithm with a few qubits [15]. Although these NMR systems could not be scaled to large sizes [16], their demonstration along with the inherent appeal of quantum computation, spurred on the proposals for and development of many other systems. This thesis explores one such system : neutral atoms trapped in a 3D optical lattices. It is enlightening to begin with a brief review of the current state of the art in experimental quantum computing followed by a look at the current standing of neutral atom technology in the sections that follow. This chapter then concludes with a brief overview of this thesis.

1.1 The state of the art in quantum computation

The field of quantum computation is arguably having a watershed moment, with researchers at universities, large corporations such as Google, IBM, Microsoft, and Intel, and a multitude of start-ups racing to build the first 50 – 100 qubit quantum computers and usher in what John Preskill at Caltech has called the NISQ (Noisy Intermediate-Scale Quantum) era [17]. This number of qubits is special because it is widely thought that even a slightly noisy quantum computer of this size might be powerful enough to surpass current classical computers in some specially designed tasks, a milestone that has often been dubbed *quantum supremacy* [18]. Even though these computational tasks might not be of any practical importance and there are many subtleties associated with demonstration of quantum supremacy, it is an important goal. It will serve as a proof of principle demonstration that quantum computation is a viable technology and will also provide platforms to use

as springboards to scale up further.

To this end, many experimental systems are being explored as possible candidates for quantum computation, including but not limited to, trapped ions [19], superconducting Josephson junctions [20], neutral atoms [21], NV centers in diamond [22], photons [23], quantum dots [24] and NMR [14]. While all these technologies have their own advantages and drawbacks, from the perspective of high fidelity quantum gates, the field has been led by trapped ions and superconducting qubits.

Ions trapped in 1D chains by radio frequency quadrupole electric fields (Paul traps) were some of the first systems explored as qubit candidates. The qubits are typically stored in the long lived internal states of the ions. These states can be coupled by lasers or microwaves to execute quantum gates. Following seminal proposals by Cirac, Zoller, and others [25, 26], high fidelity single and two-qubit gates were demonstrated by several groups, with fidelities exceeding 99% for few ion systems [27–29]. Subsequently, error correction and detection has been demonstrated in small systems by the groups of Wineland [30], Blatt [31], and Monroe [32]. A major challenge in these systems has been scaling these fidelities to a large number of qubits. For ions in Paul traps, the strong Coulomb interaction leads to problems in trapping and executing gates for long chains, making it difficult to scale to chains of more than ~ 10 ions. Other trap geometries have been explored, like the Penning Trap, which uses static electric and magnetic fields to trap ions in large 2D arrays. High fidelity control of individual qubits has not yet been demonstrated in these systems owing to the complications introduced by the cyclotron motion of the ions. Most of the recent proposals in scaling trapped ion systems have focused on building modular systems, for example, systems where the small 1D chains of ions are connected by shuttling ions between them, like in the so called Quantum Charged Coupled Device (QCCD). Eventually, this scheme will be limited by the speed of moving ions over large distances and further scaling would require a photonic link, where the entanglement is distributed between two distant qubits using single photons [33].

Superconducting qubits consist of various qubit architectures with Josephson junctions as the core ingredient. Since their inception, the coherence times for superconducting qubits have been dramatically improved from a few nano seconds to $> 100 \mu s$. Subsequently, they have followed a similar trajectory as trapped ions,

with initial demonstrations of high fidelity gates in small systems of 5–10 qubits and a more recent focus on scaling to 50 – 100 qubits [34]. Although superconducting systems can leverage the existing fabrication technology for scaling, control of large systems has been difficult because of the inherent inhomogeneities in these qubits. This is in contrast to trapped ions, where every ion is the same, leading to much longer coherence times. Neutral atoms, the focus of this thesis, retain this advantage of ions and lend themselves particularly well to scaling to large numbers in optical traps.

The vision for neutral atom quantum computation has been very similar to that of ions [21]. The qubits are stored in a long lived hyperfine spin states of an atom, typically an alkali metal like Cesium or Rubidium. The atoms can be isolated from the environment by laser cooling and then trapping them in either optical or magnetic traps. They can be cooled down to the trap ground state using various side band cooling techniques. They can be initialized with very high fidelity using optical pumping, and quantum gates can be executed using lasers and microwaves. Single atoms can then be fluorescence imaged for state detection. Unlike ions, the lack of strong interaction between neutral atoms allows for creation for large and scalable arrays of trapped and isolated atoms. On the other hand, this feature of neutral atoms also makes it harder to engineer the controllable interactions that are required for generating entanglement and implementing high fidelity two-qubit gates.

1.2 The state of the art in neutral atom quantum computation

The workhorse for trapping neutral atoms has been optical lattices, where individual atoms can be localized by standing waves created by interfering lasers in one, two or three dimensions. Thousands of atoms can be loaded into these traps to create large arrays. Since the creation of the first degenerate quantum gases of neutral atoms, optical lattices have been used extensively to study bosons and fermions in periodic potentials with the ultimate goal of effectively simulating condensed matter systems of interest. [35].

The appeal of such a system lies in its versatility. Various trap geometries can be

created by changing the wavelength, polarization and the arrangement of the laser beams. The tunneling and the on-site interaction between atoms can be controlled by changing the lattice depth. Lattices with relatively smaller spacings (hundreds of nano-metres) and lower depths have been used to explore an impressive range of many-body physics, starting with the demonstration of the super-fluid to Mott insulator transition [36]. On the other hand, deeper lattices with larger spacings (a few microns) can be used to isolate the atoms from the environment and each other, yielding pristine quantum systems for interrogation. In such a configuration these atoms can be used to build atomic clocks or be used as qubits with long coherence times. For instance in our 3D lattice of 50 qubits, the coherence time now exceeds 12 seconds [37], long enough to execute thousands of gates. On a related note, another widely used trapping technology in recent years has been reconfigurable arrays of individual dipole traps [38–40]. Such traps have been generated using acousto-optic modulators and spatial light modulators in one, two and even three dimensions. These traps allow for a greater degree of control over atom positions, but it is harder to ensure the homogeneous trapping for all the atoms, compromising the ground state cooling and coherence time for large arrays. As we will see in chapter 4, although atoms cannot be deterministically loaded to all the sites in an optical lattices or arrays of dipole traps, the control over positions of individual atoms that has been developed in recent years, allows one to "sort" the atoms by moving and re-arranging them in any desired geometry [37–42].

Atoms in such optical traps can also be imaged with high fidelity and single site resolution by collecting the fluorescence from the cooling light. For a large system of atoms trapped in a large spacing optical lattice, these techniques were first demonstrated by Nelson et. al. [43], and are also described in Dr. Xiao Li's dissertation [44], which was the first thesis based on the experimental setup described in this thesis. Development of such imaging techniques eventually culminated in the creation of the quantum gas microscope, where atoms can be imaged with single site resolutions in small spacing optical lattices, allowing one to directly measure many-body correlations [45].

While high fidelity transitions can be driven by microwaves or Raman lasers in an ensemble of neutral atom qubits, a major challenge had been demonstration high fidelity single qubit gates on individual qubits in a large array with minimal crosstalk, preserving the quantum information stored in the non-target atoms. The

proposals for achieving site addressability in neutral atom systems consisted of using addressing beams to shift the resonances for target atoms, allowing microwaves to selectively manipulate a chosen qubit [46] or using gradient magnetic fields to similarly select qubits [47]. These ideas have been implemented in various systems with varying degrees of gate fidelities and mostly uncontrolled crosstalk [48, 49]. The first, and so far only, implementation of site selective single qubit gates in a large system with gate and crosstalk fidelities approaching fault tolerance has been in our 3D optical lattice [50, 51]. This has been described in detail in Dr. Yang Wang’s dissertation [52], as well as chapter 3 of this thesis.

Two qubit gates have proved to be even more challenging. Although the required entanglement has been generated by various means such as inducing a strong interaction by exciting the atoms to a Rydberg level or using local interactions based on controlled collisions or spin exchange collisions in motional qubits, the fidelities have so far been limited to $< \sim 80\%$, well below the fault tolerance threshold [53–57]. In spite of this, analysis of these schemes suggests that the challenges are mostly technical and fidelities exceeding 99.99% are possible [58]. Moreover, combining atom sorting techniques with Rydberg excitations has recently led to impressive quantum simulations of Ising like models in one and two dimensional systems of up to 51 atoms, opening up another exciting avenue for exploration [59, 60].

Finally, there has been a recent surge in the development of cooling and trapping technologies for alkaline earth and alkaline-earth-like atoms such as Strontium and Ytterbium. Strontium atoms trapped in a 3D optical lattice have been used to create the most accurate atomic clock yet [61] and are also being explored by various groups as viable candidates for quantum information processing when trapped in optical dipole traps [62, 63]. A major advantage of these atoms is the existence of a metastable excited electronic state, connected to the qubit states by an optical transition. This can be used to shelve one of the qubit states, which can allow site selective state detection protocols, necessary for error detection and correction, that are simpler than those proposed for alkali atoms.

1.3 Atomic physics preliminaries

This thesis describes our recent progress in building a quantum computer using Cesium atoms trapped in a 3D optical lattice as qubits. Specifically, it describes




Figure 1.1. Cesium hyperfine structure. (a) shows the hyperfine structure of lowest three fine structure states of the Cesium atom. The various lasers for cooling (molasses), magneto optical trap (MOT), optical pumping (OP) and the lattice are also shown. (b) shows the positive m_F magnetic sub-levels of $F = 3$ and $F = 4$ states of the $6S_{1/2}$ manifold. the degeneracy of these levels is broken in a magnetic field. Image source : Wikipedia

the implementation of high fidelity and low crosstalk single qubit gates, creation of nearly perfect registers of up to 50 qubits using an implementation of Maxwell's demon and high fidelity lossless detection of qubit states. Why a 3D lattice one might ask? As we will see, although using a 3D geometry makes manipulating individual atoms more difficult, it is significantly more advantageous from the point of view of scaling to a large number of qubits, offering a higher density of atoms and more near neighbors for creating and distributing entanglement. Moreover, techniques that are implemented in a 3D system can be easily adapted to lower dimensions.

Since this is the third dissertation based on this experimental setup, I will not describe the theory of atomic structure, and cooling and trapping atoms. I instead refer the reader to the previous dissertations by Dr. Xiao Li [44] and Dr. Yang Wang [52] as well as a multitude of books written on this subject [64]. I will, however, re-iterate some facts about our experiment and two level systems interacting with electro-magnetic fields, for quick reference while reading this thesis.

Figure 1a shows the hyperfine structure of the Cesium atom relevant to our experiment, along with the various lasers we use for cooling and trapping. Here F




Figure 1.2. Bloch sphere representation. A general qubit state, $|\psi\rangle$, can be represented as a vector of unit length or equivalently, a point on a unit sphere. Image source : Wikipedia

is the total angular momentum of the atom. Each F level is further divided into $2F + 1$ degenerate sub-levels with a fixed z component of the angular momentum, m_F , with values ranging from F to $-F$, as shown in Figure 1b for the $6^2S_{1/2}$ (abbreviated as $6S_{1/2}$ here on) ground state. The degeneracy of these sub-levels is broken by applying a constant magnetic field. Often in this thesis, these states will be written as $|F, m_F\rangle$ and this always refers to the ground $6S_{1/2}$ manifold unless otherwise stated. Our qubits are encoded in the $|3, 0\rangle$ and $|4, 0\rangle$ states, also called the clock states, with an energy difference of exactly $9, 192, 631, 770 \text{ Hz}$ in frequency units. We will also make use the stretched states $|F, \pm F\rangle$, to initialize the atoms by optical pumping and move atoms in state-dependent traps in chapters 4 and 5.

A concept that is used extensively in this thesis is that of the Bloch sphere representation of a qubit. A qubit is a quantum system with two states, say $|0\rangle$ and $|1\rangle$, with an energy difference $\hbar\omega_0$. A general normalized qubit state can be written as :

$$|\psi\rangle = \cos\theta |0\rangle + e^{i\varphi} \sin\theta |1\rangle \quad (1.1)$$

This can be graphically represented as a vector of unit length, called the Bloch vector, that can lie anywhere on a unit sphere, called the Bloch sphere, with a of polar angle θ and an azimuthal angle of φ as shown in Figure 2. Single qubit gates

are now simply rotations of this "Bloch vector" on the Bloch sphere. In atomic qubits, this rotation is induced by interaction with an electro-magnetic field with some frequency, say ω . When the two states can be coupled by the electric dipole moment, the interaction hamiltonian is given by²:

$$H_I = -\vec{d} \cdot \vec{E}, \quad (1.2)$$

and the interaction is often parameterized by the Rabi frequency,

$$\Omega = \langle 1 | \frac{-\vec{d} \cdot \vec{E}}{\hbar} | 0 \rangle, \quad (1.3)$$

where \vec{d} is the electric dipole moment operator and \vec{E} is the electric field. This interaction can be visualized on the Bloch sphere simply as a torque vector given by $\Omega \hat{x} - \delta \hat{z}$, where $\delta = \omega - \omega_0$ is the detuning. The Bloch vector representing the qubit state precesses around this torque vector, analogous to precession of a spin in a magnetic field.

1.4 Organization of this thesis

The organization of the rest of this thesis is as follows. Chapter 2 provides a review of the experimental setup along with the description of our optical lattice, imaging system and the addressing system required to execute single qubit gates. It also describes various updates that the system has gone through over the years. A measurement of the qubit coherence time in our system, which characterizes the quality of qubits, is also described. Additionally, the projection sideband cooling technique that cools the atoms to the trap ground state is described. The ground state cooling is crucial to improving the coherence time and we have improved the cooling significantly over the past two years.

Chapter 3 is devoted to the implementation of site resolved single qubit gates. This is enabled by our development of single site addressing in three dimensions. Two different kinds of gates are described. The first scheme works by selectively transferring the atoms out of our qubit basis to execute the gate. The second scheme

²The treatment for magnetic dipole coupled states (like our qubits) is analogous, with the interaction hamiltonian given by $-\vec{\mu} \cdot \vec{B}$ where $\vec{\mu}$ is the magnetic dipole moment operator and \vec{B} is the magnetic field.

is based on targeted phase shifts while keeping the atom in the qubit states. We show that these phase shifts are highly insensitive to the alignment of the addressing beams and can be used to realize a general rotation on the Bloch sphere when combined with global microwave pulses. This results in a robust and scalable gate that has a high fidelity (0.9962) and low crosstalk (0.9979).

In chapter 4 we combine the single site addressing developed in chapter 3 with state dependent optical lattices to move individual atoms, re-arranging them in a half filled lattice to nearly perfectly filled sub-lattices. This procedure initializes our quantum computer. It required development of real-time control and decision making along with the implementation of an efficient atom sorting algorithm. Furthermore, we show that, since our atoms are sufficiently cold, this sorting procedure decreases the entropy of our system by a factor of 2.44, in essence, realizing Maxwell's famous demon thought experiment in a relatively large system.

The standard way to measure the qubit state in neutral atom systems has been to push atoms in one of the states out of the lattice followed by imaging of what is left. This is inefficient and not conducive to error correction. In chapter 5, the state dependent lattices are leveraged to implement a lossless detection scheme which coherently splits the atomic wave-function and maps the internal states of the atoms to different spatial locations. The atoms can then be imaged to yield a state detection fidelity of 0.9994, 20 times less error than any other lossless detection scheme in neutral atom systems and comparable to detection fidelities in ion traps.

Finally, chapter 6 discusses a scheme to implement high fidelity, site-selective two qubit gates in our system using the Rydberg blockade. We propose a new two-photon excitation scheme with an optical and a microwave photon which has significant advantages over the usual two optical photon excitation scheme. I also comment on the possibility of generating large scale entanglement in our system using cold-collision gates. Such an entangled state can eventually be used to implement one-way quantum computation.

Chapter 2 |

Overview of the experiment

This chapter describes our experimental apparatus for trapping atoms in a 3D optical lattice, imaging the atoms with single site resolution and manipulating the atomic states with microwaves and addressing single sites. Most of the material covered here has been described in the previous two dissertations on this experiment [44,52]. This chapter serves, in part, as a brief review for the sake of completeness. Additionally, there have been several substantial updates which are described in sub-sections of this chapter, unimaginatively titled "Updates". They only apply to experiments described Chapter 4 onwards and some of the experiment specific additions like polarization control to move atoms and a short spacing lattice are not described until later chapters where they are required. In the interest of brevity, I have omitted many details, including a description of the vacuum system, the magneto optical trap (MOT) which first cools and traps a cloud of atoms, and the polarization gradient cooling (PGC) of atoms in the lattice. These details can be readily found in Dr. Xiao Li's dissertation [44] and Dr. Yang Wang's dissertation [52]. I have also avoided details of data acquisition, software and real time control of the experiment. They are discussed in Appendix A. A section is devoted to the quality of qubits, the primary measure of which is the coherence time. We use projection sideband cooling [65] to cool the atoms to the trap ground state, which is critical to improving qubit quality. Since its inception, we have improved the projection cooling to $> 90\%$ atoms in the 3D ground state [37]. A description of projection cooling as well as the improvements is provided.

2.1 Introduction

The ground zero for the experiment is a cuboid fused silica glass cell (see Figure 2.1a) with good optical access on all sides, under ultra high vacuum. An experiment starts with loading atoms into a magneto optical trap (MOT) from a Cesium vapor. The atoms are laser cooled by the MOT and loaded with 55% occupancy (at best) into a three dimensional optical lattice (Section 2.2). The typical lifetime of trapped atoms is about 90 s, limited by background gas collisions. This implies a pressure of a few 10^{-11} torr [44]¹. The atoms are imaged by collecting the scattered polarization gradient cooling (PGC) light, allowing us to create an occupancy map for the lattice. Three pairs of Helmholtz coils provide a bias magnetic field along the $\hat{x} + \hat{y}$ direction (a plane diagonal of the lattice) to lift the degeneracy of the m_F sublevels in the Cesium ground state. The atoms are optically pumped to either $|F, m_F\rangle = |4, 4\rangle$ or $|4, -4\rangle$ using a circularly polarized beam, nearly resonant with the $6S_{1/2}, F = 4$ to $6P_{3/2}, F' = 4$ transition in the Cesium D2 line, with the aid of re-pumping beam on the $6S_{1/2}, F = 3$ to $6P_{3/2}, F' = 4$ transition². The atoms are then cooled to the vibrational ground state using projection sideband cooling (Section 2.4). At this stage atoms can be sorted to yield a fully occupied sub-lattice (Chapter 4) and cooled again after that. This concludes the initialization stage of our system.

A quantum computation can start at this point. Various gates can be executed using microwaves and addressing beams (Section 2.3) or other experiments to characterize the system can be performed. We finally detect the F state of the atoms by applying a *clearing* beam, which is resonant with atoms in the $F = 4$ states on the $6S_{1/2}, F = 4$ to $6P_{3/2}, F' = 5$ transition. This pushes the $F = 4$ atoms out of the lattice in about 2 ms. We then image the atoms again to see what is left, giving us the population in the $F = 3$ states. This was the primary detection method until the implementation of the lossless detection scheme described in Chapter 5.

¹The lifetime had decreased to ~ 20 s for the experiments described in the first half of chapter 3 due to not running the titanium sublimation pump. We were able to improve it to the nominal value for subsequent experiments.

²In principle it doesn't matter which side of the stretched states we optically pump to, but this has flip flopped a few times for various reasons. Most recently, we have been optically pumping to $|4, -4\rangle$ because we can get better σ^- polarization, which is better for projection cooling.

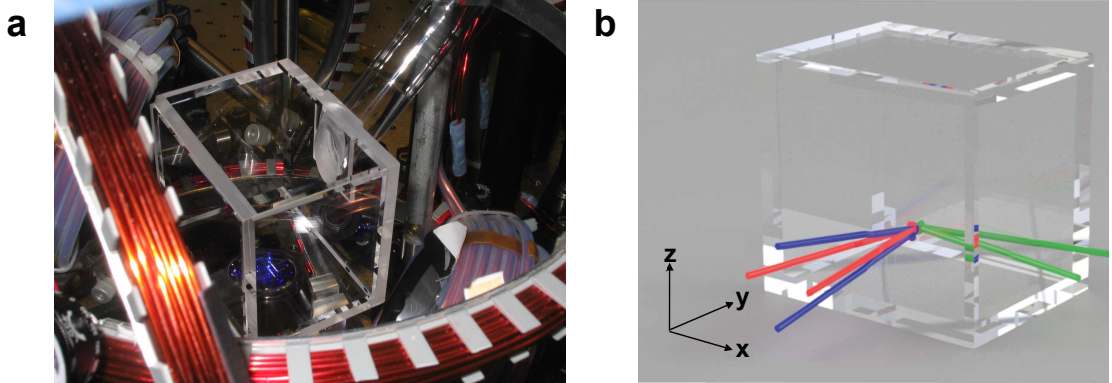


Figure 2.1. Experiment cell and lattice beam geometry. (a) shows the fused silica cell from the top is shown. Also visible are the imaging objective below the cell, the bias field coils and the MOT coils. Source: reference [44] (b) shows the geometry of the lattice beams relative to the cell. The beam pairs for different directions are shown in different colors, red for x , green for y and blue for z . Beams in each pair interfere at a 10° angle leading to a standing wave with a $\sim 4.9 \mu m$ spacing. Angles and distances have been exaggerated for clarity.

2.2 Three dimensional optical lattice

Our 3D lattice is formed by 3 pairs of beams focused to a $\sim 60 \mu m$ waist, with each pair interfering at a shallow angle ($\sim 10^\circ$), see Figure 2.1b. The source laser had a wavelength of $847.780 nm$ (old value that only applies to experiments described in Chapter 3, see "Updates"), blue detuned from the Cesium D2 line ($852.3 nm$). This results in a lattice spacing of nearly $4.9 \mu m$ in each of the three directions. Each pair is offset in frequency from the other two by several tens of MHz to avoid mutual interference. The source laser is a 1 W Ti:Sapphire, giving 70 mW (old value, see Updates) in each beam, resulting in a lattice depth of nearly $190 \mu K$ or equivalently, $4 MHz$ in each direction. Since the lattice is blue-detuned, the dipole force on the atoms is repulsive and they are trapped near the intensity minima of the standing wave. This results in a significant suppression of the spontaneous emission of lattice photons and consequently, long coherence times. Near the trap minimum, the potential is well approximated as a harmonic function, yielding a trap frequency of $\sim 15 kHz$ in each direction.

The linear polarizations of the lattice beams ensure nearly equal Stark shifts as well as spatially aligned potentials for the $F = 3$ and $F = 4$ states. The purity

of the linear polarization is critical in projection sideband cooling to ensure high fidelity sideband transitions (see reference [44]).

2.2.1 Loading the lattice

The atoms can be collected in a MOT after either desorbing them from the walls of the cell using UV light, or passing current through a Cesium dispenser. The desorption, although not as effective as the dispenser, allows us to reuse the atoms before they are pumped out by the vacuum system. The MOT is then compressed and the atoms are cooled further due to polarization gradients. The lattice is then turned on, overlapped with the MOT, leading to 5-6 atoms being loaded at each lattice site. The MOT is turned off and PGC beams are turned on to further cool the atoms, resulting in pairwise light-assisted collisions at each site [43, 66–68]. Sites that started out with an odd number of atoms are left with one atom and sites with even number of atoms are finally left with no atoms. Although this preparation ensures single atoms at occupied sites, it also limits the loading to 50% occupancy. We can, in practice, get a slightly better loading of up to 55%. Although, higher loading rates are possible by adding a blue detuned beam to protect against pair-wise loss (see Chapter 4), we are not sure of the exact mechanism at play in our experiment. After polarization gradient cooling, the atom temperature is about $4 - 5 \mu K$, although the distribution might not be entirely thermal [44]. At the best vacuum we have seen, the lifetime of the atoms in the lattice is about 90 s, limited by loss due to collisions with the background gas particles.

2.2.2 Imaging the atoms

We image the atoms on a CCD with single site resolution by collecting the scattered PGC light with a high numerical aperture lens ($N.A. = 0.55$) [43]. The depth of field is about $3.5 \mu m$, allowing us to resolve different $x - y$ planes by moving the lens in the z direction. We collect light for 150 ms from each plane, with 20 ms to move the lens to focus on a different plane and typically image five planes in this way in 850 ms. Illustrative images of atoms from different planes are shown in Figure 2.2a.

To determine which sites were loaded with atoms, an 11×3 pixels section is chosen around each site (Figure 2.2b), the values are averaged over the y direction

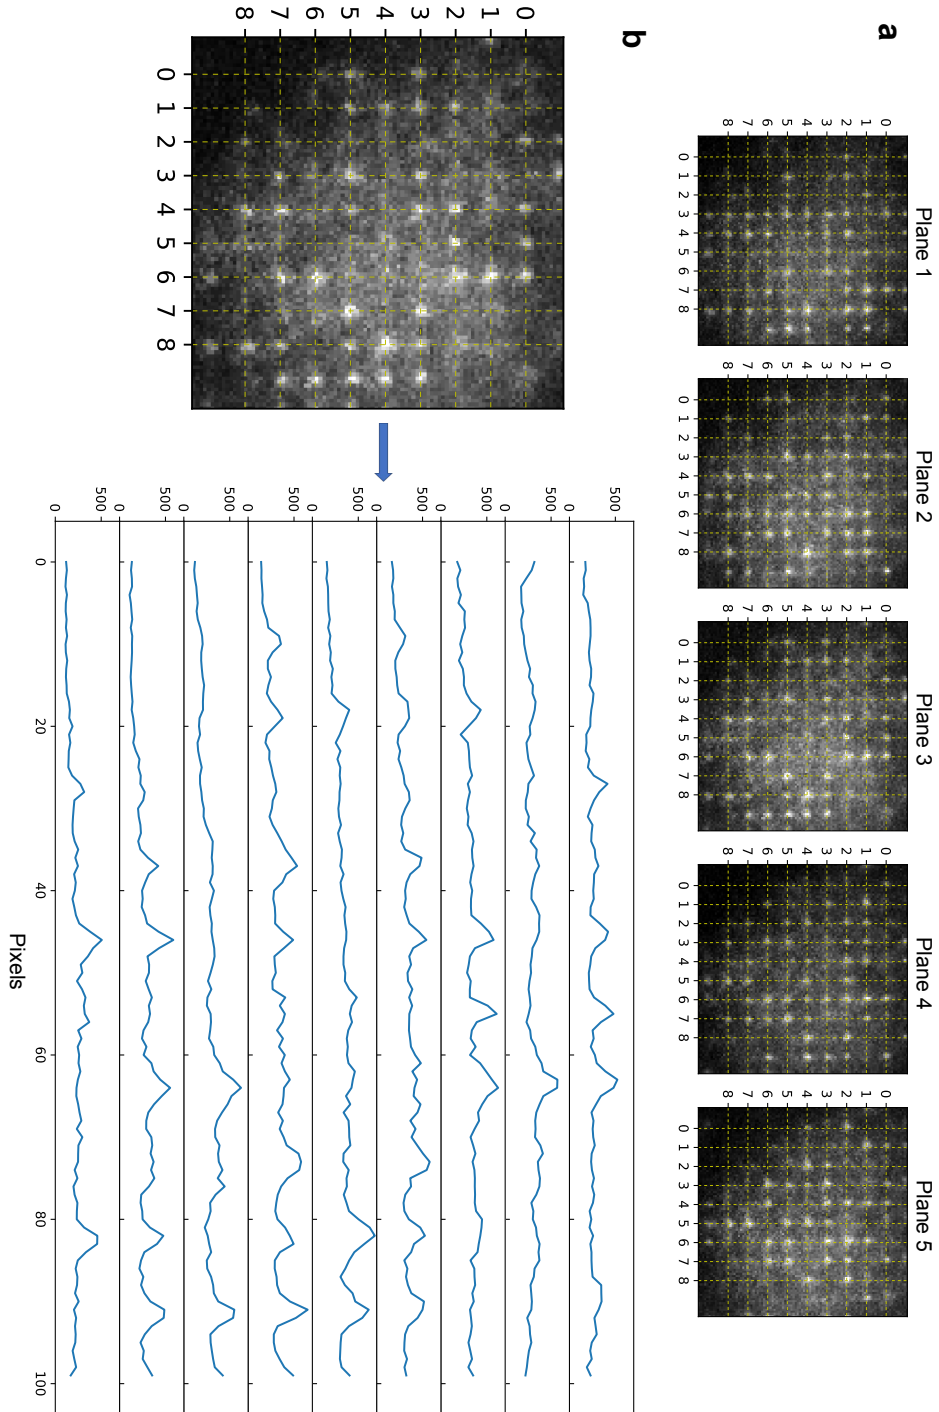


Figure 2.2. Imaging of different planes of atoms. (a) shows images of 5 different planes in the lattice. We typically capture and analyze a $9 \text{ sites} \times 9 \text{ sites}$ (100×100 pixels) section of the lattice in each plane. The grid lines show the locations of the lattice site. Each bright spot is a single atom, while the haze is the out of focus light from atoms in the surrounding planes. (b) shows the data from plane 3 in the set of images in (a). The panel on the right shows the data in arbitrary units for each of the 9 rows averaged over 3 pixels in the vertical direction. Clear peaks can be seen at the occupied sites. We pick an 11 pixel section around each site and fit it to a Gaussian with a fixed center and variance, but a variable amplitude and offset. The amplitude is then used to determine occupancy.

(3 pixels) and the data is fit to a 1D Gaussian with a fixed center and waist to obtain an amplitude and offset. The amplitude for each site is compared to a threshold to create a preliminary occupancy map of the whole lattice. Each site is binned into one of four classes based on this occupancy map. The four classes are, for a given site at location $\{x, y, z\}$: no atoms at $\{x, y, z - 1\}$ and $\{x, y, z + 1\}$ denoted by 0_0, an atom at $\{x, y, z - 1\}$ denoted by 1_0, an atom at $\{x, y, z + 1\}$ denoted by 0_1 and atoms at both $\{x, y, z - 1\}$ and $\{x, y, z + 1\}$ denoted by 1_1. Different thresholds are applied for each of these classes and the occupancy map is updated. This process is iterated up to 5 times. Figure 2.3 shows typical histograms of the fitted amplitudes for the four classes. The histograms show a clear separation between the distribution for unoccupied sites and the distribution for the sites occupied by a single atom in each of the four cases. The mis-identification rate for a single imaging is $< 3 \times 10^{-4}$ [43]. The error is limited by spatial intensity variations of the molasses beams and the out of focus light from atoms in the surrounding planes.

The averaged data from 9×9 pixels sections around all the sites that belong to the 0_0 subset is also fit to a 2D Gaussian with floating centers to determine the alignment of the lattice relative to the CCD. The lattice position is actively stabilized by feeding back to Brewster angle plates in the path of one lattice beam from each pair to adjust the position of the lattice [52] [51]. The x and y positions are directly obtained from the fits, while the z position is inferred from the already well-characterized astigmatism in the optical system. This active stabilization is critical in creating the robust site addressing scheme described in Chapter 3.

The foregoing analysis is managed by a single piece of homemade software written in C++, colloquially called *CameraGUI*. It is triggered by an external, fixed timing sequence to interface with the camera and take images, analyze them in real time and manage the feedback to the lattice stabilization system. The timing is controlled by another, Labview-based homemade timing software called *Supertime*, operating on a different computer. Supertime also manages all the other required digital triggers and analog signals to various devices with a $2 \mu s$ resolution.

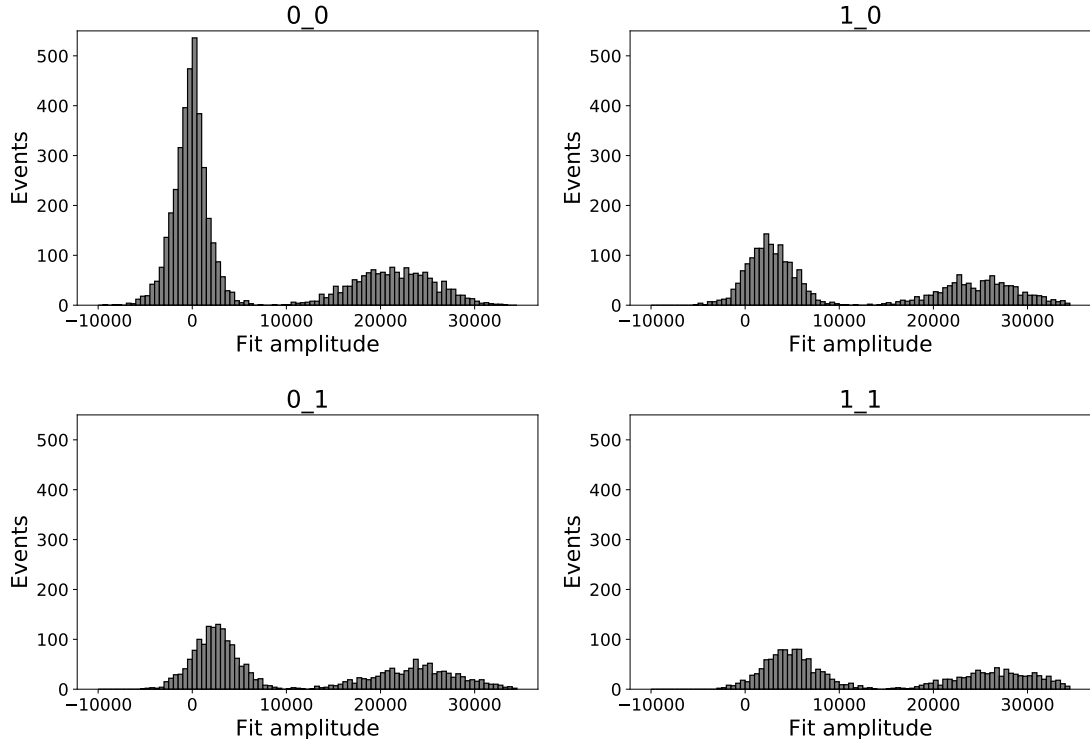


Figure 2.3. Imaging histograms for four cases. The Figure shows histograms for Gaussian fit amplitudes for sites in the central three planes. The data is divided into four cases depending on the presence of an atom in the nearest neighbor sites in the adjacent planes. 0_0 means no atoms in the nearest neighbor sites in the adjacent planes, 0_1 and 1_0 imply an atom in one of the adjacent planes, and 1_1 implies atoms in both the adjacent planes. In each of the cases, there is a clear separation between empty and occupied sites. The amplitude threshold for assigning occupancy to any site is decided based on these histograms. The amplitude units are arbitrary.

2.2.3 Updates

After the implementation of the single qubit gates described in chapter 3, the lattice wavelength was changed to 838.9 nm , which is a threefold increase in the detuning from the Cesium D2 line. To compensate for the lower lattice depth, three tapered amplifier systems were designed and built, with a 2 W output each³. The amplifiers are seeded by the previously mentioned Ti:Sapphire laser and provide $\sim 250\text{ mW}$ of power in each lattice beam. The new trap frequencies are 15 kHz , 14 kHz , 13 kHz in the x, y and z directions respectively. A critical requirement in using a tapered

³The initial design of the tapered amplifier setup was provided by Prof. Kenneth O'Hara. It was then modified to suit our purposes.

amplifier is to filter the amplified spontaneous emission, which is done by using short pass interference filters (Thorlabs FESH 0850) in our setup.

Increasing the detuning and the power threefold results in a threefold decrease in the spontaneous emission rate from the lattice. This was necessary for the high fidelity motion of atoms and the implementation of atom sorting described in chapter 4. Other significant updates to data analysis and fitting of light from imaged atoms for our lossless detection scheme are the subject of chapter 5.

2.3 Manipulating atomic states

The hyperfine sub-levels, in the same fine structure state of an atom, can be coupled by the magnetic dipole moment induced by an oscillating electromagnetic field. These couplings are analogous to those mediated by the electric dipole moment, only much weaker. In the case of the Cesium $6S_{1/2}$ ground state, the role of the exciting electromagnetic field can be played by microwaves around the Cesium clock frequency of 9.19 GHz . These microwaves allow us to drive transitions between different hyperfine sub-levels with very high fidelity; but the microwaves can't be focused to address a single atoms. To achieve single site control, we combine the microwaves with off resonant "addressing" beams which induce ac Stark shifts to preferentially change the transition frequencies for the targeted atoms.

2.3.1 Microwave system

Our microwave system starts with a homemade direct digital synthesizer (DDS)⁴ based on the Analog Devices AD9852 DDS chip. The DDS typically produces a frequency around 60 MHz and the amplitude and frequency of a pulse can be shaped with a 1 MHz resolution. The DDS also allows us to change the phase of the pulses when control over atomic coherence is necessary. The output of the DDS is filtered and mixed with that of a fixed frequency 9.19 GHz synthesizer. The mixed signal is amplified and fed to a horn antenna which bathes the lattice in microwaves. The polarization of the microwaves at the atoms is a complicated function of the orientation of the horn, the reflections from cell and components

⁴The design of the DDS board is due to Prof. Florian Schreck and Todd Meyrath. It can be found here: <http://www.strontiumbec.com/indexControl.html>

around the cell. We have enough power in all three polarization components and their relative strengths can be adjusted to some degree by rotating the horn⁵.

The DDS allows us to apply shaped pulses, which are useful when high fidelity state transfers and suppression of *sinc*² sidebands (typical of a square pulse) is required. The commonly used kinds of pulses are adiabatic fast passage (AFP) [69, 70] and Blackman [52, 71]. In an AFP pulse the frequency and power in the pulse are changed slowly such that the state vector adiabatically follows the torque vector on the bloch sphere. As the frequency is swept across the resonance, the population is flipped. A Blackman pulse, on the other hand, is resonant, with its amplitude shaped in time to suppress the sidebands. This allows us to suppress residual motional sideband transitions and discriminate transitions better when addressing single sites. We use an AFP pulse when the best possible transfer fidelity is required and Blackman pulses when control of the atomic coherence is necessary along with high fidelity. In the experiment, the AFP pulses are used to drive high fidelity sideband transitions in projection sideband cooling. After optical pumping to the stretched state, a series of AFP is also used to transfer the atoms to the clock states, our qubit basis, with high fidelity (for example: $|4, 4\rangle \rightarrow |3, 3\rangle \rightarrow |4, 2\rangle \rightarrow |3, 1\rangle \rightarrow |4, 0\rangle$). Figure 2.4a shows transfer of atoms from $|4, 4\rangle$ to $|3, 3\rangle$ as the central frequency of an AFP pulse is scanned. The typical fidelity of the AFP pulses exceeds 0.999. Once in the clock states, we almost exclusively use Blackman pulses to execute gates or probe atomic coherence. Figure 2.4b shows transfer of population from $|4, 0\rangle$ to $|3, 0\rangle$ as the frequency of a Blackman pulse is scanned around the resonance. The fidelity of a typical Blackman pulse is around 0.994 [52], limited by inhomogenous broadening and microwave power fluctuations.

The DDS is programmed using a National Instruments PCI-6534 interfaced with a computer and the pulse sequence is typically pre-programmed before the start of an experimental "shot". The pulses are then triggered externally at appropriate points in the sequence. For details of the programming and control of the DDS, see reference [52].

⁵This has been critical to the implementation of single qubit gates and lossless state detection.

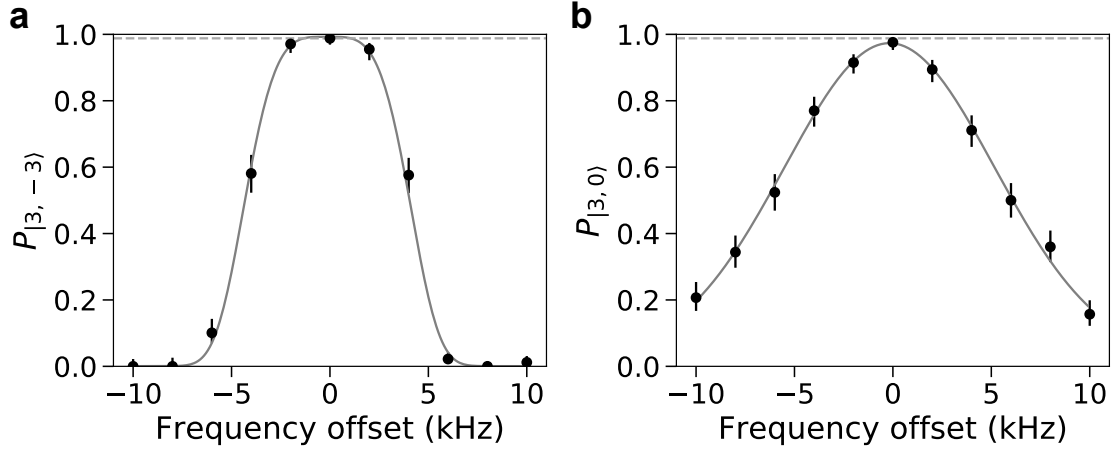


Figure 2.4. Spectrum of AFP and Blackman pulses. (a) shows the population transferred from $|4, -4\rangle$ to $|3, -3\rangle$ as the central frequency of an AFP pulse is scanned around the previously known resonance. The pulse was 3 ms long with a frequency sweep of 8 kHz . The gray curve is a fit to a super-Gaussian function of order 4. (b) shows the population transferred from $|4, 0\rangle$ to $|3, 0\rangle$ as the frequency of a $120\text{ }\mu\text{s}$ Blackman pulse is scanned around the clock states' resonance. The gray curve is a Gaussian fit. The dashed gray lines in both (a) and (b) are the vacuum limited populations.

2.3.2 Addressing system

We use two far detuned ($\lambda = 880.250\text{ nm}$), orthogonal and focused beams along the x and y lattice directions to address single sites in the lattice. The beams cross at a single site to pick out a target atom. The physics of this addressing and its use in the construction of a single qubit gate is described in chapter 3. The addressing system is responsible for translating each addressing beam in two directions to target different sites and maintaining the power stability of the beams. Spatial translation of each beam is achieved using a pair of MEMS mirrors (one for each direction) [72] [73], which tilt by up to 0.9° when a high voltage is applied. This angle is amplified to $\sim 1.7^\circ$ by double bouncing the beam on each mirror⁶. This angular change is converted to a translation on the lattice by the optical transfer system shown in Figure 5. The requirement is to generate parallel translations on the lattice while focusing the beams on the atoms. Diffraction is required to meet this requirement, which would be impossible using ray optics. This system results in a beam focused to $2.7\text{ }\mu\text{m}$ at the atoms. The pair of beams can address any site

⁶For details, see Dr. Yang Wang's dissertation.

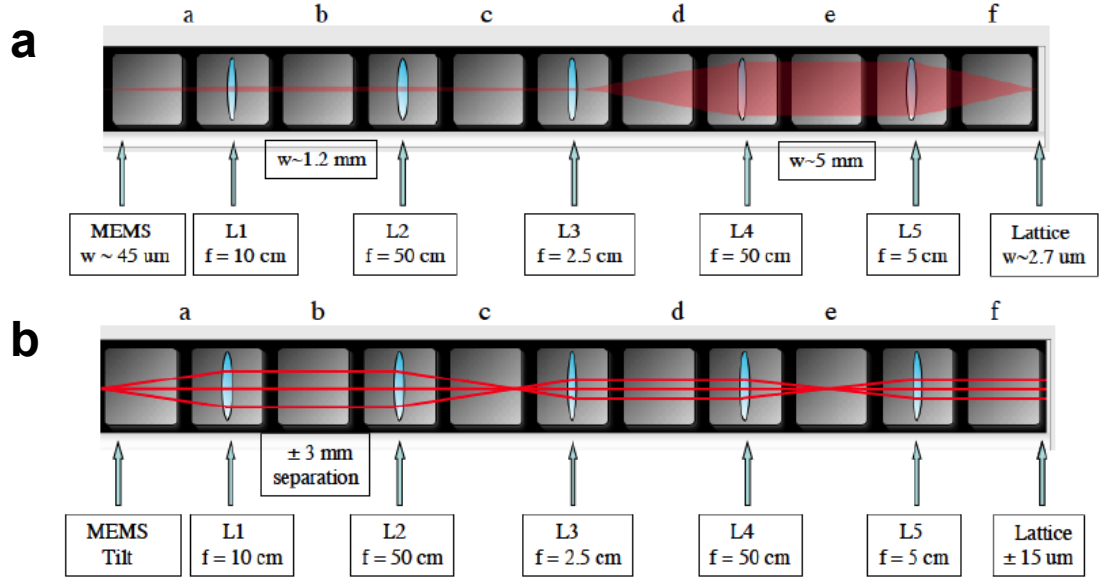


Figure 2.5. Optical transfer system for the addressing beams. (a) shows the ABCD matrix propagation calculation of a Gaussian beam focused to $45 \mu\text{m}$ on the MEMS mirrors. (b) shows the ray optics propagation of a beam as the angle of the MEMS mirrors is changed. This system results in a beam focused to $2.7 \mu\text{m}$ on the atoms which can be translated by $30 \mu\text{m}$ by changing the angle of the MEMS mirrors. The beams are sampled between L4 and L5 and imaged on a PSD and a photodiode. There is also a mirror on a pico-motor based mirror mount between L2 and L3 at the crossing point for day-to-day fine adjustment of the alignment. Source: reference [52]

in a $5 \times 5 \times 5$ section of the lattice. The beams can be repositioned to any site in that section in $< 5 \mu\text{s}$.

High fidelity addressing imposes stringent requirements on the intensity noise of the addressing beams on the target atom. The ac Stark shifts on the atoms are directly proportional to the addressing beam intensity, so any intensity fluctuations shift the transition resonances linearly. This requires maintaining both the position and power stability of the addressing beams. To this end, the beam are sampled before the final lens and imaged on a position sensitive detector (PSD) at the conjugate point after a magnifying telescope. A small part of the beam is sampled just before the PSD and diverted to a photodiode to measure the power. Before every addressing, the signal from the photodiode is used to lock the power to a pre-determined value which typically depends on the ac Stark shift required and the site being addressed. Active stabilization of the position was also tried based on the beam position on the PSD without much success. No useful, repeatable correlations

were found between the beam position on the PSD and the beam position on the atoms. Thus, we rely on the passive stability of the addressing beam positions.

It is also critical to carefully align the beams to each lattice site. This ensures robust performance across the lattice, as well as maximal insensitivity to the pointing fluctuations of the beam. The power of the addressing beams also has to be calibrated for every site to ensure equal shifts, as there are variations in the beam shape when different sites are targeted due to diffraction at the MEMS mirrors. This alignment and calibration process is described in Appendix B. The intensity lock and the targeting for the addressing beams are managed by a National Instruments 7841R FPGA board. Parameters for sequential targeting are calculated by a Python program and communicated to the board using a Labview program. Each targeting in the sequence is then triggered externally.

2.3.3 Updates

Initially, the microwave pulses and the movement of the addressing beams to target different sites in a given sequence was pre-calculated and loaded to their respective cards before the sequence was executed. They were then triggered by a fixed timing sequence, with microwave and addressing triggers at pre-decided points in time. The imaging, microwaves and addressing were all managed by different pieces of software. Experiments described in chapters 4 and 5 require real-time decision making for site selective state flips using addressing beams and microwaves. To facilitate this, the management of both microwaves and addressing has been integrated into CameraGUI. The general idea is based around another FPGA board, also managed by CameraGUI, which can pause Supertime and take control of the timing sequence and some of the electronic channels. This FPGA then triggers the DDS and the addressing FPGA based on image analysis by CameraGUI⁷. The data for microwaves is still pre-calculated by a different program, which creates a repository of different, possibly required, microwave pulses. CameraGUI then chooses and arranges different microwave pulses in the right order before communicating them to the DDS to avoid any latency related to calculation during the sequence.

⁷See appendices A and C for more details

2.4 Qubit quality

Any quantum computation is ultimately limited by the decoherence of the constituent qubits and detailed analysis and control of all the decoherence mechanisms in a system is essential to ascertain the quality of qubits. Decoherence mechanisms arise from the interactions of the qubits with their environment, which for atoms in our optical lattice are dominated by interaction with the lattice light⁸. For an ensemble of qubits, the homogeneity of qubit properties is also important for robust control.

The essential measure of the quality of qubits is the coherence time [74], which is typically obtained by preparing a qubit in an equal superposition state and measuring the probability of finding the qubit in the same superposition after some time. There are two time scales that affect the overall coherence time. The time T_1 measures the rate of processes which change the populations in the two qubit states, represented by the polar angle on the bloch sphere and the length of the bloch vector. They can couple the qubit states to each other or states outside the qubit basis. For our system, that process is the spontaneous emission of the lattice light. Even though the atoms are trapped at the intensity minima, the non-zero spread of the vibrational wavefunction samples some of the intensity of the lattice light, leading to spontaneous emission. The time T_2 measures the rate of processes which affect the atomic 'coherence' i.e. the phase of the qubit state represented by the azimuthal angle on the bloch sphere. These processes represent the perturbations to the transition frequency of the qubits as well as variations in this transition frequency across a sample of qubits. Factors that affect T_2 are: temporal fluctuations of the lattice power, inhomogeneous broadening from different trap depths for different qubits and the vibrational distribution of qubits. It is typically much smaller than T_1 and will be our primary concern.

2.4.1 Projection sideband cooling

An important factor that affects both T_1 and T_2 is the vibrational state of the qubits. Atoms in higher vibrational states spend less time closer to the trap center, thus

⁸The clock states are insensitive to changes in the magnetic field up to the first order. It is therefore a much smaller effect. It will become important for us when we work in the stretched states for sorting. The ultimate limitation is, of course, our vacuum lifetime of 90 s.

sampling more of the lattice light, and experiencing more spontaneous emission, thus reducing T_1 . A small difference in the trap depth of the $F = 3$ and $F = 4$ states⁹ means that atoms in different vibrational states have slightly different transition frequencies, which contributes to inhomogeneous broadening, reducing T_2 . This inhomogeneous broadening also contributes to reducing the fidelity of the microwave pulses. Thus, cooling atoms to the vibrational ground state to ensure best possible performance. To this end, we use projection sideband cooling [65], which has also been described previously in Dr. Xiao Li's dissertation [44].

Projection sideband cooling is based on creating a non-zero matrix element for vibrational state changing microwave transitions, $\Delta\nu \neq 0$, by slightly translating the lattice experienced by atoms in $F = 4$ state, relative to the lattice experienced by atoms in $F = 3$ state ($m_F \neq 0$). This creates a non-zero projection between the spatial wave-functions of different vibrational levels in different F states. This translation of the lattices is implemented by slightly rotating the linear polarization of one of the beams in a given lattice beam pair. This adds a circular polarization component to the overall polarization with displaced intensity minima. One of the states is primarily affected by the σ^+ component, while the other by σ^- , leading to traps displaced in opposite directions (relative to the original lattice) for the two states¹⁰. Briefly, projection cooling proceeds as follows: the atoms are optically pumped to $|4, -4\rangle$, the lattice is displaced for one of the three directions, a $\Delta\nu = -2$ transition from $|4, -4\rangle$ to $|3, -3\rangle$ is driven by an AFP microwave pulse for all the atoms, a $\Delta\nu = 1$ transition is driven by another AFP pulse from $|3, -3\rangle$ to $|4, -4\rangle$, and finally, an optical pumping pulse resets the population to $|4, -4\rangle$ for atoms that failed to make one of these transitions. These steps reduce the vibrational quantum number (in that particular direction) by one for atoms that successfully made both the transitions. The steps are then repeated for the other two directions. Typically, we repeat the whole sequence 30-50 times to reach the 3D vibrational ground state. This sequence also ensures that the ground state is reached with the smallest possible number of optical pumping steps. For instance, if all the microwave pulses worked perfectly, only one optical pumping step would be required¹¹.

The projection cooling was first reported in reference [65], with a final 3D ground

⁹This is due to the 9.19 GHz frequency splitting.

¹⁰See Chapter 4 for details of state dependent lattices

¹¹The optical pumping can lead to substantial heating because we are only weakly in the Lamb-Dicke limit.

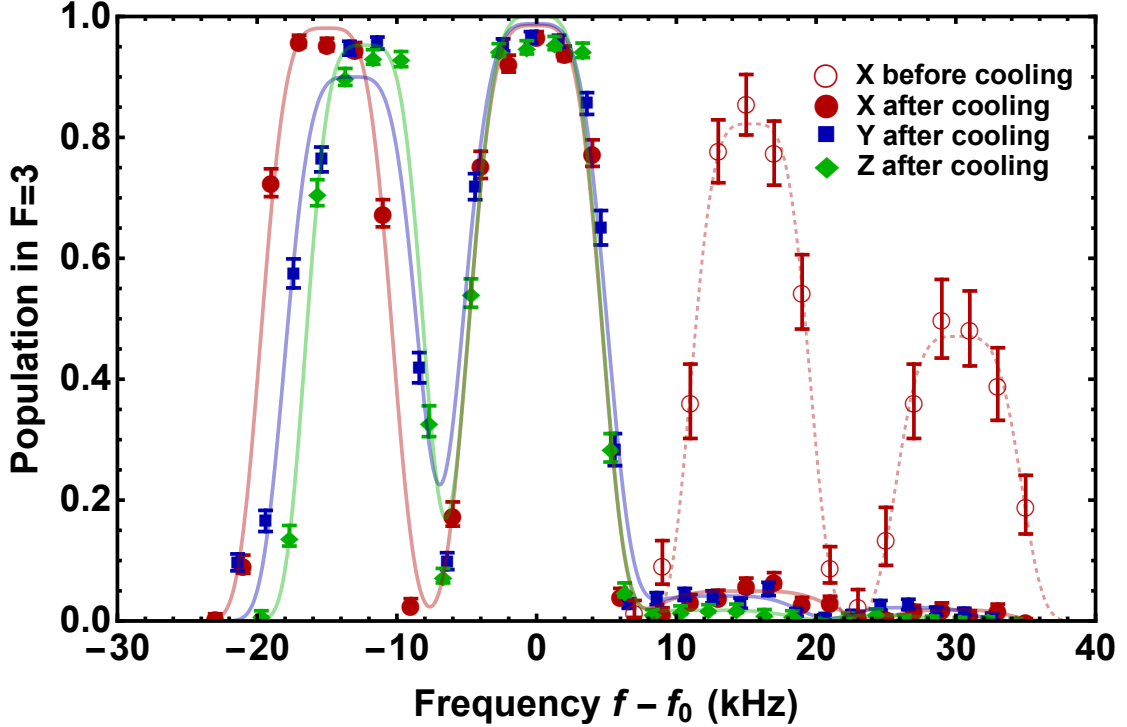


Figure 2.6. Results of projection sideband cooling. The Figure shows the spectrum of the vibrational transitions for x (solid maroon circles), y (solid green diamonds), and z (solid blue squares) lattice directions as the frequency of an AFP pulse is varied across the $|4, -4\rangle$ to $|3, -3\rangle$ $\Delta\nu = 0$ resonance frequency, f_0 , after projection sideband cooling. The spectrum from the x lattice direction before the projection cooling (hollow maroon circles) is also shown for comparison. The carrier, $\Delta\nu = 0$ transitions are at $f - f_0 = 0$, with the $\Delta\nu = 1$ sidebands on the left of the carrier. The $\Delta\nu = -1$ and the $\Delta\nu = -2$ sidebands can be seen on the right of the carrier. It can be seen that the $\Delta\nu = -1$ and the $\Delta\nu = -2$ sidebands are significantly suppressed after cooling, because for atoms in the ground state, there are no lower vibrational levels. The ratio of the heights of the $\Delta\nu = -1$ and the $\Delta\nu = 1$ sidebands give the ground state population for each direction.

state population of $\sim 76\%$. We have, since then, significantly improved this cooling. Figure 6 shows a recent measurement of the vibrational spectra of the $|4, -4\rangle$ to $|3, -3\rangle$ transition after 50 projection cooling steps for the three lattice directions. The ground state populations are obtained by fitting to the data and normalizing the $\Delta\nu = -1$ sideband to $\Delta\nu = 1$ sideband. The ground state populations for the x, y and z lattices are: 94.5%, 95.5% and 98.2% respectively, which implies that $\sim 89 \pm 0.9\%$ of the population is in the 3D ground state. Two changes were important for improving the efficiency of the cooling. First, we have increased the efficiency of the $\Delta\nu = -2$ sideband transition significantly by using separately




Figure 2.7. Dynamical decoupling pulse sequence. This Figure shows the pulse sequence used to measure the coherence time of the atoms. The atoms are prepared in an equal superposition of the qubit states using a $\pi/2$ pulse. They are protected against inhomogeneous broadening by repeated π pulses which re-phase the various Bloch vectors. The time, t , between the π pulses is fixed to 20 ms and the number of π pulses is varied to adjust the total time, T . The final superposition is determined by changing the phase of the final $\pi/2$ pulse and obtaining an interference fringe. Sequences like these will form the basis for executing gates in our system while preserving the qubits.

optimized lattice translations for the two required sideband transitions. Second, we have improved the purity optical pumping circular polarization by almost a factor of five¹². More recently, we have been able to further increase the efficiency of the cooling by eliminating some power imbalances in the molasses beams, thus starting from a possibly colder initial state. We can now achieve the same ground state fraction in 30 instead of 50 steps, and a ground state population of $\sim 91\%$ after 50 steps.

2.4.2 Coherence time

The standard method for determining T_2 is *Ramsey interferometry* [64, 75]. We start by preparing atoms in one of the qubit states, say $|0\rangle$. A $\pi/2$ microwave pulse prepares an equal superposition: $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$. After time, T , another $\pi/2$ pulse is applied and the state of the atoms is measured. As the phase of the final $\pi/2$ pulse is varied, an interference fringe is obtained for final population in either of the qubit states. The decay of the fringe contrast as a function of T yields T_2 . For our system, this measurement yields a T_2 of $\sim 20\text{ ms}$ [52] for the central $5 \times 5 \times 5$ section of the lattice. This time can be improved significantly to a new value T'_2 by using dynamical decoupling techniques first pioneered in nuclear magnetic resonance

¹²The polarization of the optical pumping beam is critical in ensuring correct initial state preparation and efficient optical pumping to minimize heating. As a stroke of luck we discovered that we were able to get better σ^- polarization than σ^+ . We are not sure why that is the case, but it is presumably related to the effect of residual birefringence due to the vacuum cell.




Figure 2.8. Interference fringes for coherence time measurement. This Figure shows the interference fringes in the population of the $|4, 0\rangle$ state, obtained for various sequences shown in Figure 7, as the total sequence time, T , is varied. The curves are sinusoidal fits to the data.

(NMR) systems starting with the Hahn echo experiment [76–78]. The idea is to rephase the inhomogeneously broadened evolution of the different Bloch vectors by using repeated, judiciously positioned π pulses, as shown in Figure 7. In the presence of inhomogeneous broadening different qubits evolve at different rates in the equatorial plane of the Bloch sphere after a superposition has been created. A π pulse inverts the Bloch vectors and they rephase at the half way points between the π pulses. In a quantum computing context, this echoing structure can form the "backbone" of any algorithm, and quantum gates can be executed around the rephasing points¹³.

Sequences like the one shown in Figure 7 now allow us to measure the overall coherence time. Similar to Ramsey interferometry, as the phase of the final $\pi/2$

¹³A more detailed discussion of constructing robust dynamical decoupling sequences can be found in chapter 3.




Figure 2.9. Coherence time. This Figure shows a semi-log plot of the fringe contrast obtained by fitting the data shown in Figure 7, against the sequence time, T . The contrast is fit to an exponential decay to obtain the coherence time T_2' .

pulse is varied, an interference fringe is obtained. A decay of the fringe contrast is obtained when the time T between the sandwiching $\pi/2$ pulses is varied. The time between the π pulses is kept fixed to be around 20 ms in our measurement and the number of π pulses is increased to change T . The choice of timing between π pulses for a sequence like this is a balancing act between the decoherence of the qubit and imperfections in the π pulses to optimize the final performance. Figure 8 shows the fringes for various times and Figure 9 shows shows the decay of fringe contrast with T on a semi-log plot. The coherence time, T_2' is obtained to be $\sim 12.6\text{ s}$ by fitting the decay to an exponential function. This data was obtained at a quarter lattice depth of 1 MHz , which is deep enough for us to execute high fidelity quantum gates. I should note that the this measurement was done with the lossless detection developed in chapter 5 as a demonstration of that scheme, and will be revisited.

This coherence time is an improvement of a factor of three over our previously measured value before the lattice wavelengths were changed to a further detuning,

reducing the spontaneous emission rate by a factor of three for the same lattice depth. This is to be expected as the coherence time is eventually limited by T_1 . The calculated value of the spontaneous emission rate of 0.108 s^{-1} at the full lattice depth of 4 MHz and the updated wavelength of 838.9 nm matches well with the value 0.12 s^{-1} , determined by measuring the depump rate of atoms prepared in $F = 4$ to $F = 3$. This implies a T_1 of 8.28 seconds for a full lattice and, because scattering rate scales as the square root of lattice depth, a T_1 of 16.56 seconds for quarter lattice depth. The measured coherence time of 12.6 seconds for the quarter depth lattice is slightly smaller than the scattering rate would suggest. This is due to the imperfections in the π pulses arising from the fluctuations of the microwave power, which influence the measurement of coherence times. Since this dynamical decoupling sequence will form the backbone of our implementation of quantum gates, the measured value of the coherence time is more indicative of the overall performance rather than the scattering rate.

Chapter 3 |

Single qubit gates

This chapter describes the development and implementation of high fidelity and low crosstalk site-selective single qubit gates in our three dimensional optical lattice. This subject has been covered in detail in Dr. Yang Wang’s dissertation [52]. Here I provide a somewhat brief review.

The site selectivity of the gates is based on targeting single atoms using two orthogonal addressing beams. A scheme is developed to cancel the phase shifts on the non-target atoms due to addressing beams and microwaves, thereby reducing the crosstalk. Two qualitatively different types of gates are implemented. The first one is a targeted qubit rotation about any axis in the $x - y$ plane of the Bloch sphere which requires mapping the target qubit out of the clock state basis. Transferring the the target atom out of the clock states also compromises the fidelity. The second gate is a targeted phase shift (rotation about the z of the Bloch sphere), which is highly insensitive to the pointing fluctuations of the addressing beams. It doesn’t require the target atom to be transferred out of the clock states, thus yielding a high fidelity. We further show that an arbitrary targeted single qubit unitary can be realized by combining the targeted phase gate with global $\pi/2$ rotations which leave the non-addressed atoms unaffected. We implement a Randomized Benchmarking (RB) protocol with the second kind of gate to show a gate fidelity of 0.9962(16) and a crosstalk fidelity of 0.9979(2). The work described in this chapter was published in references [51] and [50].

3.1 Introduction

Gate-based quantum computation requires action of arbitrary unitary operators on a collection of qubits [74]. A common approach to constructing arbitrary unitaries is through separate implementations of arbitrary single qubit rotations and a two qubit gate such as Controlled Phase (CPHASE), Controlled NOT (CNOT) or SWAP. It is also possible to arbitrarily approximate any unitary by using a finite "universal" set such as the Hadamard gate (a π rotation about the $x + z$ axis), $R_z(\pi/4)$ (a $\pi/4$ rotation about the z axis on the Bloch sphere), and the CNOT gate¹ [79]. In addition to being important in constructing arbitrary unitaries, single qubit gates are also required to change the measurement basis of a qubit in error correction schemes [74, 80], as well as implementations of measurement-based quantum computation [81, 82]. In general, a complete set of single qubit unitaries requires arbitrary rotations about at least 2 axes on the Bloch sphere.

Several approaches to implementing single qubit gates are possible in alkali atom hyperfine qubits. Microwaves can be applied to ensemble of trapped atoms to drive transitions. Excellent power and phase control of the microwaves has enabled high fidelity, non-targeted, arbitrary (to within instrumental resolution) gates [83, 84]. Transitions can also be driven by two phase locked Raman lasers [85]. This approach offers built in addressability but suffers from high sensitivity to the pointing fluctuations of the Raman beams. Addressing in microwave based gates can be implemented through targeted frequency shifts by using DC Zeeman shifts from magnetic field gradients or ac Stark shifts from tightly focused addressing beams (this is the approach in our experiment as well) [46]. The frequency of the microwaves then preferentially selects the addressed qubits. A plane of Rb atoms was addressed in reference [47] using magnetic field gradients and a fidelity of 0.95(3) was achieved. Another experiment used magnetic field gradients to address one of two qubits in an ensemble of two qubit registers in a double well optical lattice [49]. Single atoms were addressed in a 2D trap array using an addressing beam and direct microwave transitions between the qubit states, with a fidelity of 0.9923(7) [48]. In that work the state flip crosstalk error was 0.002(9) but the

¹This is guaranteed by the Solovay-Kitaev theorem to be efficient. The number of gates required for such a construction scales as $O(\log^{3.97}(1/\epsilon))$, where ϵ is the desired error in the unitary being constructed.

phase shift errors at non-target atoms were completely uncontrolled and were not measured.

Section 3.2 outlines our 3D addressing scheme, where two perpendicular beams are used to address a single atom in a 3D lattice at their intersection. Eliminating crosstalk in 3D is challenging, since it is impossible to avoid having some atoms in the direct path of one of the addressing beams. Section 3.3 is devoted to construction and implementation of a simple targeted gate with a scheme that cancels the crosstalk on non-target qubits. Section 3.4 describes implementation and characterization of a different kind of single qubit gate, with the same crosstalk cancellation scheme but with the much desirable insensitivity to addressing beam intensity fluctuations. This insensitivity results in a high fidelity gate with errors approaching the generally accepted fault tolerant threshold of 10^{-4} and a clear path to improvement.

3.2 Targeted addressing

As mentioned earlier, single qubits in our experiment are addressed using two orthogonal tightly focused beam (see Figure 3.1a). The waist of each of the beams is $\sim 2.7 \mu m$ and the Rayleigh range is around $27 \mu m$. The beams are circularly polarized, leading to a vector light shift on the atoms. The beams are at the "tune out" wavelength of $880.250 nm$ [86], which leads to a minimal absolute Stark shift on the clock states, $|3, 0\rangle$ and $|4, 0\rangle$ ². The vector light shift behaves like a fictitious magnetic field, resulting in a Stark shift proportional to the m_F level (Figure 3.1b) leading to differential shifts between all hyperfine sub-levels (except the clock state) large enough to be discriminated by a microwave. Each of the beams can be steered in 2D with a pair of MEMS mirrors to target any site in a $5 \times 5 \times 5$ lattice. This repositioning of the beams can be done in $< 5 \mu s$. The optical transfer system to convert the angular changes of the MEMS mirrors to beam translations on the lattice has been previously described in reference [52].

It is useful to divide atoms into three classes: the *cross* atom at the intersection of the two beams is shown in orange, *line* atoms in the path of only one of the addressing beams are shown in blue and experience nearly half the Stark shift

²This is the wavelength at which the scalar polarizability of the $6S_{1/2}$ state vanishes

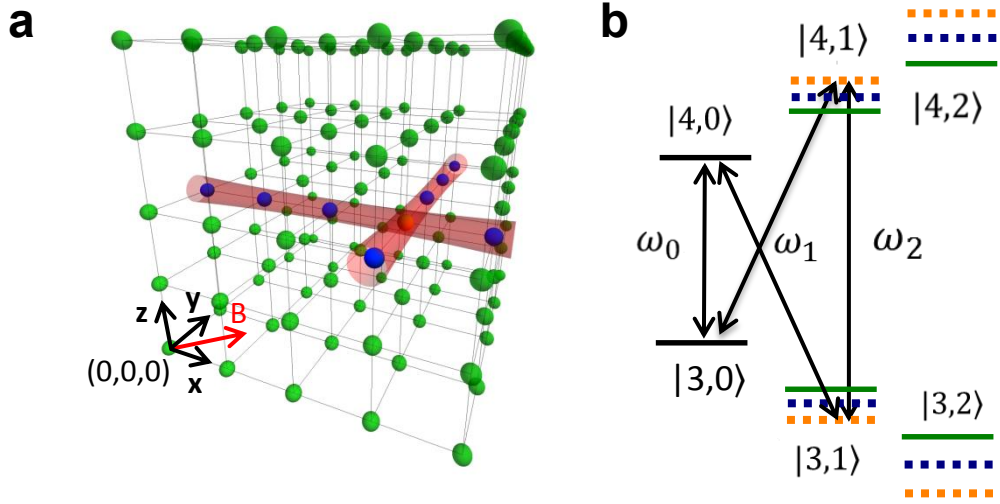


Figure 3.1. Addressing scheme and Stark shifts. (a) Two orthogonal beams address an atom at their crossing point in a 3D lattice. The beams are σ^+ polarized and partition the qubits in two 3 classes: *spectators* in green, *line* atoms in blue and the orange *cross* atom. (b) shows the Stark shifts for the 3 different classes of atoms. The Stark shifts are proportional to the m_F value and the addressing intensity on the atom. The frequencies for three transitions of concern are also shown as ω_0 , ω_1 , and ω_2 . In a targeted gate, pulse at ω_1 maps the state from $m_F = 0$ basis to $m_F = 1$ basis for the target atom. Another pulse at ω_2 performs the qubit rotation for the transferred atom. A final ω_1 pulse maps the superposition back.

as cross atoms, and *spectators* which are not in the direct path of either of the addressing beams are shown in green and experience no Stark shifts. To measure these shifts for different classes of atoms, we start by optically pumping the atoms to the $|4,4\rangle$ state and transferring them to the $|4,0\rangle$ state with Adiabatic Fast Passage (AFP) pulses. We then apply a $240\ \mu\text{s}$ blackman π pulse to transfer atoms from $|4,0\rangle$ to $|3,1\rangle$, the frequency of which is varied. The atoms in the $F = 4$ manifold are then cleared with a resonant beam, and the remaining atoms are imaged again. Figure 3.2a shows the populations in $F = 3$ for different classes of atoms as a function of microwave detuning from the *spectator* resonance. Figure 3.2b shows illustrative average figures for targeting four lines of atoms with either of the addressing beams and Figure 3.2c shows averaged images targeting four single atoms with both the addressing beams. Typical alignment of the addressing beams involves looking at the "line" signals from one addressing beam at a time while varying the position of the addressing beams (see Appendix B). The typical shifts in




Figure 3.2. Spectrum with addressing and illustrative images of atoms. (a) Spectrum for the different classes of atoms as the microwave frequency is scanned around the $|4,0\rangle$ to $|3,1\rangle$ transition. The detuning from the spectator resonance at 0 is $f = 50 \text{ kHz}$ for the line atoms and $kf = 91 \text{ kHz}$ for the cross atoms. (b) shows averaged images of three planes with four rows in planes 1 and 3 sequentially addressed by the vertical (V) beam in the top row and horizontal (H) beam in the bottom row. (c) shows averaged images for four atoms sequentially addressed by both beams. The light in plane 2 is from out of focus atoms in both (b) and (c).

our experiments for the $m_F = 1$ sub-levels range from $20\text{ kHz} - 50\text{ kHz}$ due to each addressing beam, which corresponds to a beam power of about $120\text{ }\mu\text{W} - 300\text{ }\mu\text{W}$. The shift from two addressing beams is slightly less than double of that due to just one beam. This is due to a comparable Zeeman shift of $\sim 200\text{ kHz}$ from the bias magnetic field, which is at 45° to the addressing beams. There is a significant drop in contrast for line and target atom signals due to inhomogeneous broadening and a relatively long microwave pulse. The imperfect contrast for the spectators is due to vacuum loss and imperfect AFP transfer to the clock states. The inhomogeneous broadening is mostly due to the different addressing intensities on different atoms and addressing beam pointing fluctuations. There are two effects that lead to different addressing intensities for different atoms: the tight focusing of the addressing beams and the changes in beam intensity profile at different locations due to diffraction from the MEMS mirrors. We deal with this problem by carefully measuring the Stark shifts for each site from each beam and calibrating the intensity accordingly (Appendix B). The beam pointing fluctuations are harder to fix due to the long path and several approaches including active stabilization have been tried, as described in reference [52]. This further underscores the need for a gate that is insensitive to pointing fluctuations of the addressing beams.

3.3 A targeted gate

The targeted addressing of the atoms can now be used to construct a simple site selective gate, the basic idea of which is outlined in Figure 3.1b. A resonant microwave pulse at frequency ω_1 along with the addressing beams is used to map the quantum state from the $m_F = 0$ (called the storage basis) states to the $m_F = 1$ states (called the computational basis) for a single atom:

$$a|3, 0\rangle + b|4, 0\rangle \leftrightarrow b|3, 1\rangle + a|4, 1\rangle,$$

where a and b are complex numbers that represent some normalized qubit state. A second microwave pulse at frequency ω_2 , resonant with the $|3, 1\rangle \leftrightarrow |4, 1\rangle$ transition (and thus only with the previously targeted atom), can be used to perform an arbitrary rotation about any axis in the $x - y$ plane of the Bloch sphere. The third microwave pulse with addressing beams then maps the state back to the




Figure 3.3. Targeted gates in a dynamical decoupling sequence. The black pulses are at frequency ω_0 , resonant with the clock transition. The red pulses are the ω_1 mapping pulses. The purple pulse is the ω_2 gate pulse. The black pulses form a spin echo like infrastructure to protect the qubits against inhomogeneous broadening. The targeted gate, comprised of red-purple-red pulse sequence can be applied around the rephasing points of the spin echo sequence. Different sites can be addressed sequentially.

storage basis. The coherent mapping of the quantum state between the two bases by the same microwave can be thought of as two decoupled two-level systems if the interaction cross terms are neglected³. This step is further complicated due the different transition matrix elements for the $|3, 0\rangle \leftrightarrow |4, 1\rangle$ and $|4, 0\rangle \leftrightarrow |3, 1\rangle$ transitions. Several solutions to this problem were explored, including composite π pulses but the simplest solution was found to be the equalization of the two Rabi frequencies by simply rotating the microwave horn⁴ [52]. The gates can be further embedded in a spin echo type dynamical decoupling sequence to preserve the qubits against inhomogeneous broadening⁵ as shown in Figure 3.3. The echo π pulses are at frequency ω_0 , resonant with the clock transition.

3.3.1 Phase shifts

Hitherto, we have neglected the effects of the addressing light and the off-resonant microwaves on the storage basis. These effects turn out to be important for both

³This primarily includes the off-resonant microwave coupling between the same m_F states. This is a good approximation if the Rabi frequency is much smaller than the total shift of the $m_F = 1$ states relative to the $m_F = 0$ states due to the addressing beams and the bias magnetic field.

⁴Rotating the horn changes the balance of the different polarization components in the microwaves. It was possibly a stroke of luck that there exists an orientation of the horn which gives the right ratio of intensities between σ^+ and σ^- polarizations with reasonable power in all three polarization components.

⁵This also includes the temporal fluctuations in transition frequencies due to fluctuations of the lattice powers.

crosstalk and gate fidelity. Atoms in the path of the addressing beams can scatter the addressing photons and the microwaves can cause off-resonant state transfer. This leads to state flip and qubit leakage errors. These effects, while small, are ultimately important considerations for fidelities but a much larger effect is the phase shift due to addressing beams and the microwaves.

Addressing beams induce a non-zero differential Stark shift on clock states due to their slightly different resonances owing to the hyperfine splitting of the ground state, leading to an unwanted phase shift on the line atoms. We measured this phase shift using a spin echo type sequence to probe the coherence of 4 different rows of atoms, targeted sequentially by an addressing beam. The atoms were prepared in a superposition of $|3, 0\rangle$ and $|4, 0\rangle$ states and the lines were targeted around the re-phasing points of the spin echo sequence by an addressing beam for $700 \mu s$. The addressing beam power was such that it induced a vector light shift of $30 kHz$ on the $m_F = 1$ states. Successively targeted lines get opposite sign phase shifts due to the spin echo π pulse. For these parameters, the addressing beams lead to a phase shift of about 0.14π which implies a differential Stark shift of $\sim 250 Hz$.

The microwave pulses, which are off-resonant for all but the cross atoms, also induce an ac Zeeman shift (analogous to ac Stark shifts for magnetic dipole transitions) on the non-target atoms. This also leads to significant phase shifts and thus crosstalk. These phase shifts are given by:

$$\Delta\phi \approx \frac{\Omega^2 T}{4\delta}$$

,where Ω is the microwave Rabi frequency, T is the time of the pulse and δ is the microwave detuning from the clock state. These phase shift affect all the atoms, unlike the phase shifts from the addressing light, which are only experienced by the line and the cross atoms.

3.3.2 Phase shift cancellation

In principle, it might be possible to model and calibrate the phase shifts for all the atoms and incorporate them in the quantum computation. This approach is significantly complicated by the inhomogeneity of these phase shifts over the sample

of qubits as the shifts depend on the intensity of the addressing light experienced by the qubits. The slow pointing fluctuations of the addressing beams also add a temporal fluctuation. Moreover, the extra computational overhead required to account for these shifts might be undesirable. Instead we pursue a more integrated approach to deal with this issue, where active cancellation of these shifts is built into the structure of the gate.

This cancellation stems from the observation that the sign of the accumulated phase shift can be flipped by applying a π pulse before phase shift inducing process. The existing π pulse infrastructure of the spin echo sequence can be utilized to perform this cancellation as illustrated in Figure 3.4a. Two atoms in different planes of the lattice can be addressed⁶, interleaved with global π pulses on the clock transition. Targeted gates are performed in stages 1 and 3 of the sequence, while in stages 2 and 4 no atoms are addressed by both the beams. The two dummy stages (2 and 4) work to cancel the phase shifts already accumulated or about to be accumulated during the gate operations. The microwaves for the dummy stages are exactly the same as the addressing stages, but no atoms are resonant as there is at most one addressing beam in any plane of the lattice. For example, in stage 2, the phase shifts accumulated by the horizontal line atoms in plane A and the spectators are cancelled. Phase shifts are added for some of the atoms, but are similarly cancelled in following stages. Another way to think about this is the following: when an atom is a line (spectator) atom in some stage, another stage can always be found where the atom is again line (spectator) with its superposition flipped by an odd number of π pulses. This ensures that after every four stages (two gates), the phase shifts for the non-target atoms are perfectly cancelled.

The situation is more complicated for the target atoms. None of the phase shifts are cancelled as the target atoms switch between cross, line and spectator status in an unsymmetrical fashion throughout the sequence. Moreover, changing the qubit basis compromises the effectiveness of the spin echo sequence. In the actual implementation of the gate, these phase shifts are accounted for by empirically adjusting the phase of the ω_2 gate pulse to achieve the desired result.

⁶It is also possible to design this sequence for just one atom. Addressing two atoms is just more efficient.

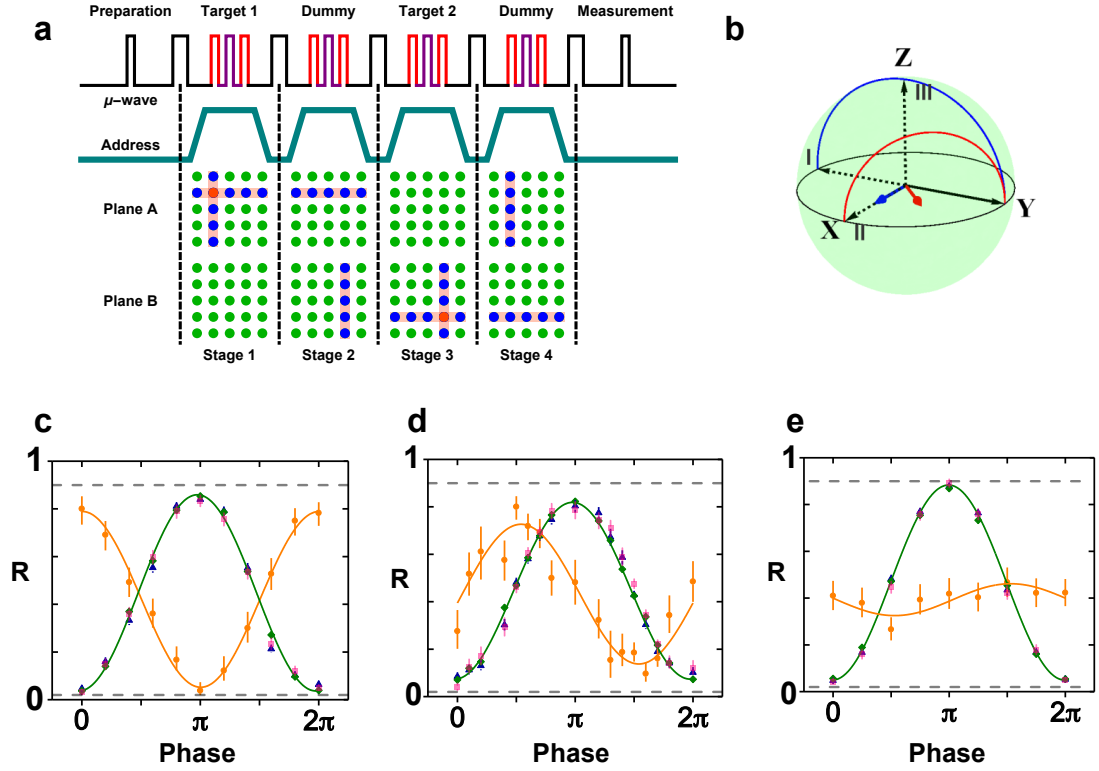


Figure 3.4. Crosstalk cancellation and results for the first gate. (a) Gates are executed in pairs of atoms in planes A and B. We start by preparing an equal superposition of clock states with a $\pi/2$ pulse. Two targeted gates (in stages 1 and 3) are interleaved with two dummy stages (2 and 4). There are no cross atoms in the dummy stages and the addressing beams target rows or columns in different planes. The microwaves are the same for all stages. Phase shifts on non-target atoms are cancelled by the end. A final $\pi/2$ pulse with a variable phase analyzes the quantum state. (b) The paths taken by the Bloch vector for the 3 different gates are shown. The qubits are prepared along the y axis (solid black arrow). The final state after each kind of gate is labelled by I, II and III (dotted black arrows). The blue and red arrows show the directions of the torque vector for the gates. The first gate, I, is a π rotation about the \hat{x} axis. The second gate, II, is a π rotation about the $\hat{x} + \hat{y}$ axis. The third gate, III, is a $\pi/2$ rotation about the \hat{x} axis. (c) The ratio, R, the fraction of the population in the $F = 3$ state, is plotted against the phase of the final $\pi/2$ pulse showing an interference fringe for gate I. The orange points are for the target atoms, green for the spectators, blue for the line atoms and pink for the nearest neighbors. The curves are sinusoidal fits to reconstruct the Bloch vector. The gray dashed lines are experimental limits due to vacuum loss, imperfect state preparation and detection. (d) Interference fringe for gate II. (e) Interference fringe for gate III.

3.3.3 Results

To demonstrate this single qubit gate scheme and assess its fidelity, we executed three different gates on two qubits in our three dimensional array. The gates were:

π rotation about the \hat{x} axis, $\pi/2$ rotation about the \hat{x} axis and π rotation about the $\hat{x} + \hat{y}$ axis. We started by preparing an equal superposition (along the \hat{y} axis) of the clock states with a $\pi/2$ pulse. The Bloch vector paths, as well as the final states for different gates are shown in Figure 3.4b. The microwave pulses and the addressing beam targeting was as shown in Figure 3.4a. We measured the ratio, R , which is the fraction of population in the $F = 3$ states (by clearing $F = 4$ atoms with a resonant beam) after a final $\pi/2$ pulse with a varying phase which probes the qubit state, yielding an interference fringe. The fringes are plotted in figures 3.4c-e for four different classes of atoms, the aforementioned target, line and spectators, as well as six nearest neighbours.

It is evident from the fringes of the non-target atoms that crosstalk is minimal and the phase shift cancellation works well. There is some loss of contrast from various sources of state preparation and measurement (SPAM) errors. For example, the dashed line shows the vacuum loss and the Adiabatic Fast Passage (AFP) transfer of the atom from the stretched state to the storage basis is also imperfect due to imperfect microwave pulses. The vibrational ground state fraction for these experiments was about 60%, which also affects the AFP pulse fidelity. The spontaneous emission of the lattice light also adversely affects the contrast. To deduce the fidelity of these gates, we use process tomography. We first normalize the data to account for the 10% loss in contrast from SPAM errors⁷. The density matrix, ρ , is then reconstructed by fitting the Bloch vector parameters (n, θ, ϕ) to the interference fringes. The fidelity is then given by:

$$\mathcal{F} = Tr(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}),$$

where σ is the expected density matrix for a perfect gate. For pure states, this expression reduces to overlap between two quantum states.

The fidelities are summarized in table 3.1 for different classes of atoms. For the non-target atoms, these fidelities characterize the crosstalk. While the crosstalk fidelities are promising, the gate fidelities of $\sim 93\%$ fall well short of the commonly accepted error correction threshold of 10^{-4} . The gate fidelities are primarily

⁷Most of the SPAM error here is from the atom loss due to background gas collisions which was significantly higher due to degraded vacuum for these experiments. The trap lifetime had reduced to 20 s. The vacuum was fixed for all subsequent experiments bringing the trap lifetime back to its nominal value of 90 s.

Table 3.1. Fidelities for the first kind of gate

	π rotation around \hat{x}	π rotation around $\hat{x} + \hat{y}$	π rotation around \hat{x}
Spectator	0.988 ± 0.002	0.978 ± 0.002	0.990 ± 0.001
Line	0.984 ± 0.007	0.977 ± 0.005	0.992 ± 0.003
Target	0.946 ± 0.008	0.913 ± 0.023	0.925 ± 0.047
Nearest neighbors	0.983 ± 0.004	0.985 ± 0.007	0.993 ± 0.004

compromised because of the extraneous phase shifts which are not cancelled for the target atoms. Moreover, transferring atoms out of the storage basis adds further complications. The $m_F = 1$ states are sensitive to the magnetic field fluctuations, an additional source of phase shifts. The spin echo pulses on the clock states don't work to undo the effects inhomogeneous broadening once the atom has been transferred to the computational basis. A partial suppression can be achieved depending on the gate, for example with π rotation about the \hat{x} axis, as evident in the differing fidelities among the gates. The fidelity of the spin-echo π pulses, which is about 0.994, also affects the gate fidelity, limiting it to $0.994^4 = 0.976$. While it might be possible to engineer a better cancellation of the phase shifts on the target atoms, a robust and scalable solution would be hard to implement. Ultimately, a gate which does not transfer the atoms out of the clock states is more desirable, which is described in the following sections.

3.4 A targeted phase gate

The phase shifts incurred by atoms suggest a way to implement a phase gate while keeping atoms in the clock state. The basic requirement of such a gate is to cause a finite phase shift on the target atom while still cancelling the phase shifts on non-target atoms. Figures 3.5 a and b illustrate such a scheme. The basic structure and cancellation scheme remains the same, but the resonant gate pulses are replaced by a single off-resonant microwave pulse with frequency between the line and cross $m_F = 0 \leftrightarrow m_F = 1$ transitions. The off-resonant microwaves do not cause appreciable population transfer but induce an ac Zeeman shift on all the atoms, including the cross atoms, unlike the previous gate. The phase shifts for the non-target atoms are cancelled exactly as before. The target atoms, on the other hand, are addressed by a different number of addressing beams in different stages without the same underlying symmetry and accumulate a phase shift. For example,

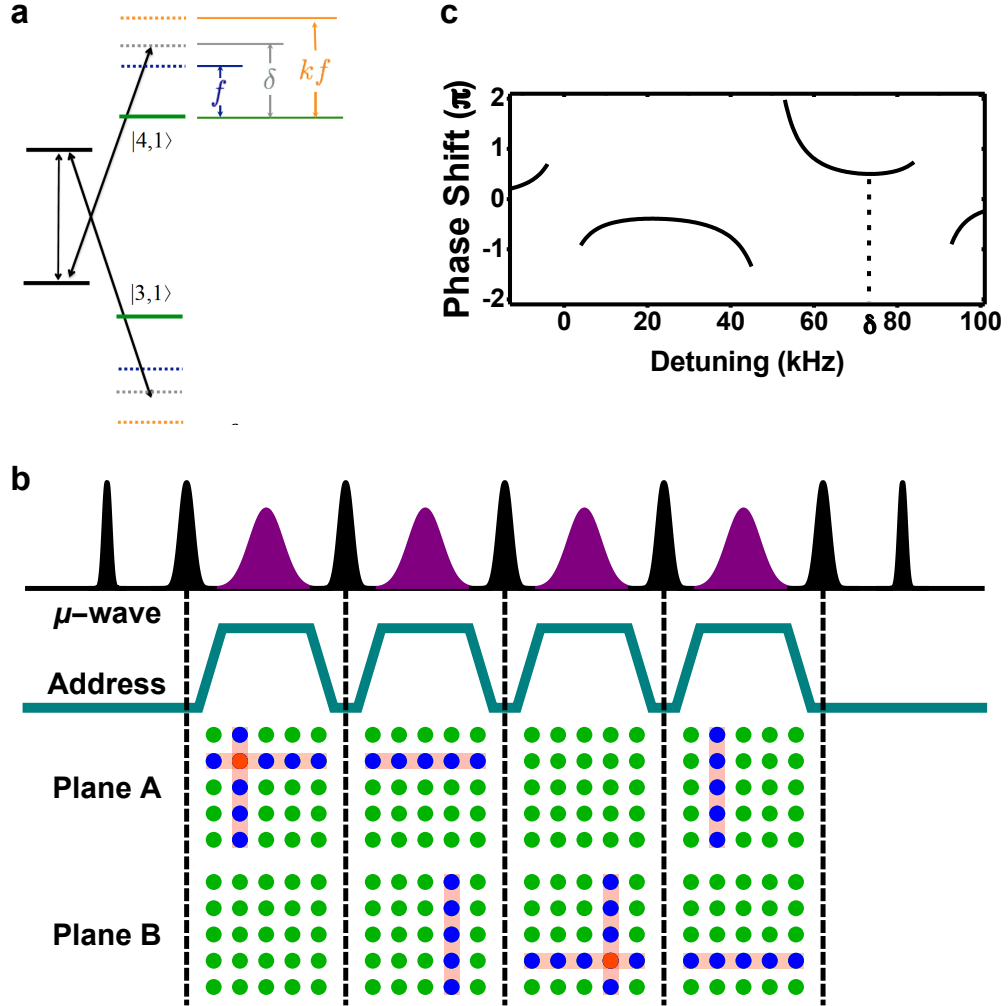


Figure 3.5. Targeted phase gate. (a) The addressing shifts are similar to before. The phase gate microwave is at the detuning δ from the spectator resonance, between the line and cross atom resonances. These microwaves are off-resonant from any transition. (b) shows the pulse and addressing sequence for this gate. The crosstalk cancellation is the same as before. The black pulses form the spin echo infrastructure and the purple pulses are phase shifting, off-resonant microwaves. (c) The calculated phase shift for the target atom as a function of the detuning of the gate microwaves (purple pulses) from the spectator resonance. δ is our typical operating point.

for the target atom in plane A in Figure 3.5b, the phase shift is given by

$$\Delta\phi_{\text{Target}} \approx \frac{1}{2} \left(\underbrace{\frac{\Omega^2 T}{kf - \delta}}_{\text{Stage 1}} + \underbrace{\frac{\Omega^2 T}{\delta - f}}_{\text{Stage 2}} - \underbrace{\frac{\Omega^2 T}{\delta}}_{\text{Stage 3}} + \underbrace{\frac{\Omega^2 T}{\delta - f}}_{\text{Stage 4}} \right),$$

where Ω is the microwave Rabi frequency, f is the frequency shift for the line atoms, k is the ratio of frequency shifts for cross and line atoms⁸, δ is the detuning of the off-resonant microwave pulses from the spectator resonance and T is the time. The phase shifts for the target atom when it is a cross atom (stage 1) is almost equal and opposite to the phase shift when it is a line atom (stages 2 and 4). But because of the spin echo π pulses these shifts end up adding. There is also a smaller negative shift in stage 3 when it's a spectator. The magnitude of the phase shift can be easily controlled by adjusting the power and the time of the off-resonant microwave pulse, yielding an arbitrary phase gate. This phase shift is plotted in Figure 3.5c as a function of the microwave detuning δ . The absolute value of the phase shift has minima with extended flat regions, where it is highly insensitive to the microwave detuning or nearly equivalently⁹, the ac Stark shift from the addressing beam. The phase shift is only sensitive to the second and higher order fluctuations of the Stark shift at the minima, and thus only to the fourth and higher order fluctuations of addressing beam pointing. The extended flatness results from a small second order coefficient in the dependence of phase shift on δ or f . A more physical understanding of this insensitivity is as follows: when the Stark shift fluctuates, say higher, the phase shift for the target atom during the *cross* stage (1) decreases, while the phase shifts during the *line* stages (2 and 4) increase, resulting in a small overall change.

A phase gate by itself is not enough to construct arbitrary single qubit unitaries. An arbitrary rotation about an orthogonal axis (to z) on the Bloch sphere is also required. Such a rotation can be constructed by sandwiching the phase gate between two $\pi/2$ rotations in the following way:

$$R_n(\theta) = R_{n'}\left(\frac{\pi}{2}\right).R_z(\theta).R_{n'}\left(-\frac{\pi}{2}\right),$$

where n , n' and z are mutually orthogonal axes obeying the right hand rule. Such a sequence can be implemented in our system with global $\pi/2$ and $-\pi/2$ microwave pulses sandwiching a targeted phase gate. The non target atoms don't experience the phase gate, so the action of this sequence on those is just $R_{n'}\left(\frac{\pi}{2}\right).R_{n'}\left(-\frac{\pi}{2}\right) = I$,

⁸For our typical values of magnetic field (560 mG) and 50 kHz addressing shift, $k \approx 1.8$

⁹For a given f it is not optimal to have the microwave frequency at the phase shift minimizing δ when insensitivity to the fluctuations in f is desired. This is because of terms in the phase shift which do not depend on just $\delta - f$.

leaving their state unchanged to within crosstalk errors. This allows us to implement an arbitrary rotation about any axis in the $x - y$ plane of the Bloch sphere. It is note worthy that any of these rotations (about any axis in the $x - y$ plane or the z axis) requires only a single targeted phase shift¹⁰.

3.4.1 Robust dynamical decoupling




Figure 3.6. Robust dynamical decoupling. The contrast of interference fringes obtained from a 100 pulse dynamical decoupling sequence with the $\{x, x, -x, -x\}$ phase scheme. The contrast is plotted against the Rabi frequency of each pulse in the units of the π pulse Rabi frequency. The gray dashed lines is the limit due to vacuum loss and imperfect state preparation.

As mentioned before, imperfect π pulses in the spin-echo infrastructure are a major source of error for these gates. These errors accumulate exponentially with the number of gates. These errors are especially deleterious, because they accumulate on the whole ensemble of qubits and not just the target atoms, as the π pulses are global. To combat this issue we turn to some of the dynamical decoupling techniques first pioneered in nuclear magnetic resonance (NMR) experiments [76–78, 87]. The

¹⁰The Hadamard gate still requires two phase gates.

essence of one of those techniques is in adjusting the phases of the echo π pulses so as to minimize the accumulated errors by the end of the sequence. A simple example of such a sequence is to have the phase of the pulses alternating between 0 and π i.e to have the direction of torque vector alternating between the x and $-x$ axes on the Bloch sphere. In this case, for qubits that starts out along the y axis of the Bloch sphere and suffer some inhomogeneous broadening, errors caused by imperfect Rabi frequency of a π pulse are always cancelled by the following pulse assuming that the π pulse imperfection does not vary over the time scale two pulses. But for qubits starting along the x axis, this phase scheme doesn't work at all. For these qubits, a phase scheme like $\{y, -y\dots\}$ or simply $\{x, x, x\dots\}$ would work similarly to cancel the errors. In a general quantum computation, the qubits can be in different states when a gate is applied, so a scheme that works for all initial states is required. The desired result is obtained by combining the two different phase schemes into a cycle of four pulses: $\{x, x, -x, -x\}$ ¹¹.

Figure 3.6 shows the result of implementing such a phase scheme on a spin echo sequence with 100 pulses. The contrast of the interference fringe is plotted against the Rabi frequency as a fraction of a π pulse. Almost no decrease in contrast is observed, even with a 20% pulse error. Moreover, simulations show that such a scheme is also robust against frequency errors.¹².

3.4.2 Experimental implementation

The experimental implementation of this gate is similar to the previous one. We start by preparing an equal superposition of the clock states with a $\pi/2$ pulse as shown in Figure 3.5b. The spin echo π pulses are $80 \mu s$ long with a blackman profile. For the phase gate off-resonant pulse, the addressing beams are first ramped up in $50 \mu s$ with an additional $20 \mu s$ for the intensity lock to settle. The off-resonant microwave pulse with a blackman shape is then applied centered around the re-

¹¹Several such schemes have been explored in NMR experiments. An important scheme is the XYn class of pulses, where n is the length of a cycle ($n = 4, 8, 16$ etc.) [88], with longer cycles performing better. For example XY4 is $\{x, y, x, y\}$ and XY8 is $\{x, y, x, y, y, x, y, x\}$. It is desirable to have the cycle length coincide with the number of pulses required for each gate implementation. Experimentally, we observed that the scheme $\{x, x, -x, -x\}$ performs at least as well as XY16, so we chose that.

¹²See reference [52] for a more detailed discussion of inner workings of these phase schemes, as they apply to our system.

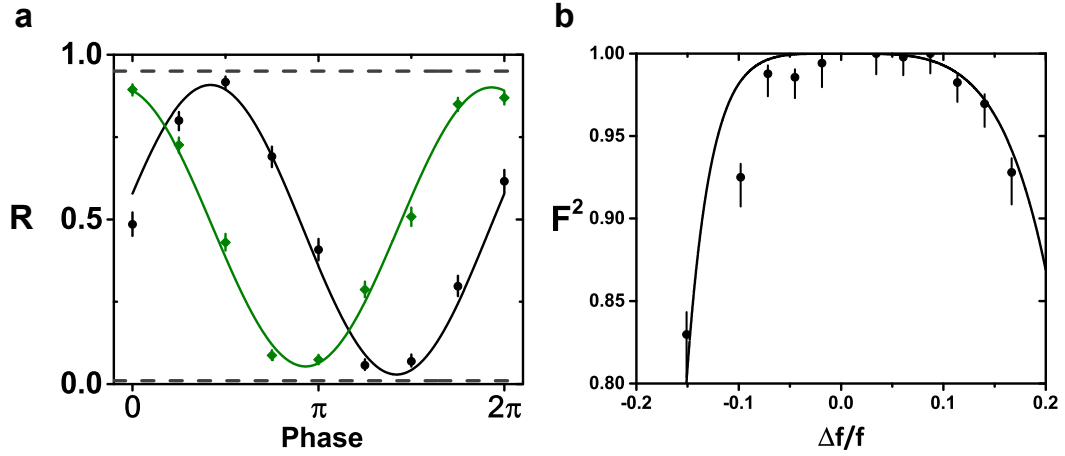


Figure 3.7. $R_z(\pi/2)$ gates. (a) shows averaged interference fringes for spectators (green) and target (black) atoms after sequential execution of $R_z(\pi/2)$ gates on 48 arbitrarily chosen sites. The curves are sinusoidal fits. (b) shows the measured fidelity, \mathcal{F}^2 , for an $R_z(\pi/2)$ gate as a function of fractional changes in the Stark shifts due to the addressing beams. The solid black curve is the theoretical prediction.

phasing points of the spin echo sequence for $120 \mu s$. A final $\pi/2$ pulse with a variable phase is used to characterize the final state of the qubits. The gates are applied regardless of whether there is an atom present at a given site or not and data is post-selected based on which sites had an atom initially.

Figure 3.7a shows the resulting interference fringes after sequential application of an $R_z(\pi/2)$ gate on 48 different sites. The sites were chosen and paired arbitrarily throughout a $5 \times 5 \times 5$ addressable section of the lattice. The black points show the fringe averaged over the 48 target atoms, while the green points show the fringe average over the remaining atoms. A clear $\pi/2$ phase shift for the target atoms can be seen in the data. Good contrast for the target atoms suggest the robustness of this gate against the inhomogeneities in the lattice and addressing. The comparable contrast of the target and non-target fringes also suggests that the primary sources of errors are common to both the classes of atoms and thus unconnected to the addressing. The obvious culprits are: fluctuations of the microwave power, inhomogeneous broadening and spontaneous emission of the lattice light. The fringes are fit to a sinusoid (solid curves) to estimate the quantum state. The error per gate, \mathcal{E} , can then be obtained by calculating the overlap with the expected state. This results in an \mathcal{E} of about 13×10^{-4} , which is similar for

both targets and non-targets.

Figure 3.7b shows the fidelity of a single $R_z(\pi/2)$ gate measured by process tomography, as a function of the fractional change in the addressing beam Stark shift. The curve is the theoretical calculation and matches well with the data. This demonstrates the expected robustness of this gate to addressing beam intensity fluctuations.

To further demonstrate the versatility and scalability of this gate, we applied an $R_z(\pi)$ gate on 32 purposefully chosen sites after preparing the qubits in an equal superposition of $|3,0\rangle$ and $|4,0\rangle$ starting from the $|4,0\rangle$ state. The phase of the final $\pi/2$ pulse is fixed at π . This sequence flips the states of only the target atoms to $|3,0\rangle$. Figure 3.8 shows the average images for each plane after 50 implementations of this sequence.

On an operational note, the typical value of the Rabi frequency is $\sim 4\text{ kHz}$ for a $\pi/2$ phase shift. This is a balance between causing minimal off-resonant transfer and getting enough of a phase shift. This results in a typical gate time of $\sim 1.4\text{ ms}$ for a pair of gates. The gate can be shortened by increasing Ω and the addressing shift (and thus δ) proportionately. This leaves the errors from off-resonant transfer and spontaneous emission of the addressing photons unchanged. We are presently limited by the microwave power and the bandwidth of the addressing beam intensity lock.

3.4.3 Randomized benchmarking

Process tomography, which we employed to characterize the first gate, is useful in estimating performance of specific gates. Often the performance of different gates (like different amounts of rotations or rotations about different axes) is affected differently by errors in the system. A useful standard technique to characterize average gate performance in a system is randomized benchmarking (RB) [48, 83, 89–91]. Unlike process tomography, RB can also decouple the SPAM errors from gate errors. In RB, gate sequences of increasing lengths are randomly sampled from a pre-defined set and applied to the system. The sequences are such that the final state is always one of the qubit states, say $|1\rangle$, for perfect gates. The decay of the population in the measured state as the length of the sequence increases, measures the error per gate, \mathcal{E} . The gate set commonly used is the Clifford group,




Figure 3.8. $R_z(\pi)$ gates. The states of 32 atoms in planes 1, 3 and 5 are flipped from $|4, 0\rangle$ to $|3, 0\rangle$ with the help of π phase gate. Atoms in $F = 4$ are then cleared and the remaining atoms imaged. The data is averaged over 50 implementations. The intensity variations at target sites are due to atom loading shot noise. The light in planes 2 and 3 is from out of focus atoms.




Figure 3.9. Randomized benchmarking microwaves. The black pulses form the spin echo infrastructure. The red pulse is a global Pauli randomization pulse. The targeted phase gate is sandwiched by two $\pi/2$ pulses shown in blue to implement rotations about x and y .

which is essentially all the permutations of the Bloch sphere axes. As shown in Figure 3.9, an RB sequence consists of alternating computational gates and Pauli gates. Computational gates (CG) are chosen from the set $\{R_x(\pm\pi/2), R_y(\pm\pi/2)\}$ and Pauli gates (PG) are chosen from the set $\{I, R_x(\pm\pi), R_y(\pm\pi), R_z(\pm\pi)\}$. Each PG-CG pair adds one to the length of the sequence. The CG set is enough to generate the whole Clifford group, while the PGs work to randomize the "initial" state before each CG. The implementation of a PG-CG pair in our experiment is shown in Figure 3.9. The CGs are targeted phase gates sandwiched by two global $\pi/2$ pulses, while for PGs, rotations about x and y are implemented using global π pulses, rotations about z are implemented by shifting the phase of all the following pulses, and the identity operation is implemented by simply waiting. Finally, the population in $|3, 0\rangle$ (also called $|1\rangle$) is measured after a conditional π pulse (in case the state was supposed to be $|4, 0\rangle$ after the random sequence). Several random sequences of the same length are implemented to appropriately sample the various paths the Bloch vector can take. The average population in $|1\rangle$ as a function of sequence length, represented by $\overline{\mathcal{F}^2}$, is fitted to:

$$\overline{\mathcal{F}^2} = \frac{1}{2} + \frac{1}{2}(1 - d_{if})(1 - 2\varepsilon)^l$$

Here $\overline{\mathcal{F}^2}$ is a measure of fidelity that represents the average probability (not just the overlap) of finding the qubit in the right state, $d_{if}/2$ is the SPAM error, ε is the error per CG and l is the length of the sequence. For our implementation of CG on pairs of target atoms, ε is the error per gate pair, represented by \mathcal{E}_{2t} for the target atom. A similar calculation can also be done by analyzing the data for the non target atoms after applying similar sequences. The ε s for these measure

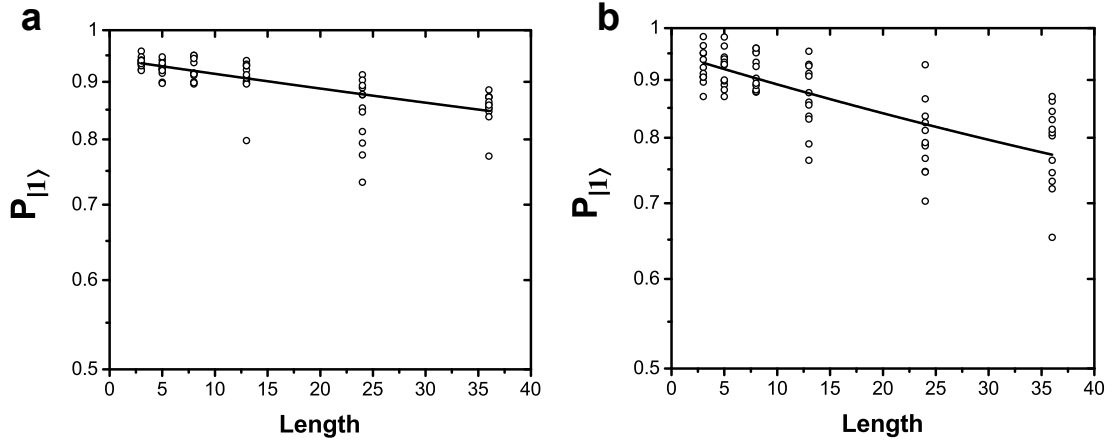


Figure 3.10. Randomized benchmarking data for non-target atoms. (a) shows the population in the $|1\rangle$ state on a semi-log plot as a function of sequence length for spectators and (b) for line atoms. For each length 4 random Pauli sequences are combined with 3 random computational gate sequences yield a total of 12 points. The average of the points is fit to the equation for $\overline{\mathcal{F}^2}$ and the result is plotted as the solid lines.

the crosstalk errors, yielding \mathcal{E}_{2s} and \mathcal{E}_{2l} for spectators and line atoms respectively. The errors per gate are then given by, $\mathcal{E}_s = \mathcal{E}_{2s}/2$ for the spectators, $\mathcal{E}_t = \mathcal{E}_{2t} - \mathcal{E}_s$ for the target atoms and $\mathcal{E}_l = \mathcal{E}_{2l} - \mathcal{E}_s$ for the line atoms. The RB data for the line and spectators is shown in Figure 3.10 on a semi-log plot of $\overline{\mathcal{F}^2}$ with length. The average crosstalk error per gate for the non-target atoms, defined as the weighted (by number of spectators and line atoms) mean of \mathcal{E}_s and \mathcal{E}_l , is $21(2) \times 10^{-4}$. RB data for the target atoms is shown in Figure 3.11. The data is fit with $d_{if} = 1.1 \times 10^{-2}$, obtained by fitting the non-target data, yielding $\mathcal{E}_t = 38(16) \times 10^{-4}$.

I should mention that the phase scheme mentioned in the previous section, that works well for a simple spin echo sequence, doesn't work optimally for randomized benchmarking because of the global $\pi/2$ pulses. We came up with and used the following empirical rules for generating the phase scheme of an RB sequence after extensive experimental testing :

1. The spin echo pulses around successive Pauli pulses should follow the cycle: $\{y, -y, -y, y\}$.
2. If the PG preceding a CG has a torque vector along the same axis as the CG's $\pi/2$ pulse, then the torque vector of the first π pulse inside the CG should be opposite that of PG. Otherwise, it can have the same phase as the $\pi/2$ pulse.




Figure 3.11. Randomized benchmarking data for target atoms. Semi-log plot of $|1\rangle$ state population as a function of sequence length. For each length 2 random Pauli sequences are combined with 3 random computational gate sequences yield a total of 9 points. The average of the points is fit to the equation for \mathcal{F}^2 and the result is plotted as the solid line.

3. The final four π pulses in a CG should follow the cycle: $\{y, y, -y, -y\}$.

Table 3.2. Errors and contribution from various sources for the second kind of gate in units of 10^{-4}

Index	Items	Spectators	Line	Target
	\mathcal{E}_2 , error per gate pair	34(4)	63(9)	55(16)
	\mathcal{E} , error per gate	17(2)	46(9)	38(16)
i	Spontaneous emission of the addressing light	-	8	16
ii	Phase error from fluctuations in f	-	-	0.7
iii	Phase error from microwave fluctuations	-	-	0.2
iv	Off-resonant microwave transitions	-	3.6	5.4
v	Other sources	17	34	16

3.4.4 Sources of errors

The errors per gate for different classes of the atoms are summarized in table 3.2. We have independently measured the SPAM errors, with vacuum loss contributing

around 3%¹³, and imperfect AFP transfer to the clock states contributing around 2%¹⁴. Various sources contribute to the gate errors: the spontaneous emission from each addressing beam per gate is calculated to be 8×10^{-4} , the spontaneous emission from the lattice photons is $< 10^{-4}$ per gate, the phase errors from addressing beam intensity and pointing fluctuations can be calculated to be $< 10^{-4}$ and the error from off-resonant transfer due to the gate microwaves is about 5.6×10^{-4} . The remaining error is inferred to be from the microwave power fluctuations which compromise the spin echo infrastructure. This error is common to all atoms and contributes to the crosstalk.

It should be possible to bring the microwave power errors to $< 10^{-4}$ by upgrading to state of the art microwave control. The spontaneous error from the addressing beam can be reduced eightfold by doubling the wavelength. There would be a significant Stark shift on the qubit states, but the effects of that can be mitigated by adiabatic turn on in a deeper lattice. Increasing the detuning of the addressing beam more has diminishing returns because we rely on the differential Stark shift of the $m_F = 1$ states. In fact, this gate would work even better in an atom where only one of the qubit states is shifted by the addressing and the microwave can simply be off-resonant with the qubit transition. Further twofold improvement is also possible with a faster intensity lock of the addressing beams.

3.5 Conclusion

In conclusion we have demonstrated two kinds site-selective single qubit gates in a three dimensional array of neutral atom qubits. The first gate sets the foundation for targeted coherent addressing of qubits in a challenging 3D geometry with exceedingly small crosstalk. The second gate is a qualitatively different targeted phase gate with a very high fidelity. The low crosstalk results from the active cancellation of phase shifting errors in a dynamical decoupling sequence. This scheme for crosstalk cancellation can also be adapted to 1D and 2D geometries and even other qubits like ions. The high fidelity of the second gate results from the

¹³The vacuum was improved considerably between the experimental implementations of the two gate schemes.

¹⁴This is now lower than 1% after we have improved the cooling and better optimized the AFP pulses

extreme insensitivity of the ac Zeeman shifts accumulated over several addressings to the addressing beam pointing and intensity fluctuations. We have demonstrated both these features with the standard technique of randomized benchmarking.

Chapter 4 |

Sorting atoms in 3D: an omniscient Maxwell's demon

This chapter describes our work on sorting atoms in the lattice by combining the single site addressing developed in the previous chapter with state selective lattices to move individual atoms, rearranging them in the lattice to yield fully filled sub-lattices. Starting with a $5 \times 5 \times 5$ randomly half-filled lattice and guided by an efficient algorithm, we selectively move atoms to generate fully-filled $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sub-lattices. We obtain a filling fraction of $\sim 96\%$ with a perfect filling rate of $\sim 30\%$. We show that the most significant error is the spontaneous emission of the lattice light, which is enhanced during a motion step. We use Monte-Carlo simulations to explore the scaling of this technique to larger systems and show that a threefold reduction in spontaneous emission should allow us to effectively sort 400 atoms starting from a $10 \times 10 \times 10$ half-filled lattice. Furthermore, we propose a "throw and catch" scheme for moving atoms which can significantly reduce the spontaneous emission, allowing us to scale to thousands of atoms.

We further show that this implementation of sorting is also a realization of Maxwell's "demon" thought experiment, which reduced the entropy of a gas by observing individual particles and sorting them by speed. Since we can cool our atoms down to the vibrational ground state, most of the entropy in our system is contained in the disorder associated with random loading. We show that the sorting process reduces the configurational entropy of the system by a factor of 8 and the total entropy by a factor of 2.44. Unlike Maxwell's thought experiment, our demon is "omniscient" as it observes the whole system at once, reducing its

entropy, to create the plan for sorting the atoms.

The work described in this chapter is based on reference [37].

4.1 Introduction

The standard technique of trapping neutral atoms in optical lattices and tweezer or micro-trap arrays involves loading from a magneto optical trap and then cooling the atoms with red-detuned optical molasses. The optical molasses photo-associates two atoms on the same site to an attractive molecular state in which the atoms gain enough kinetic energy to leave the lattice, leading to pairwise loss and at best, a half filling of the lattice [43, 66, 92]. This random 50% loading of atoms in our optical lattice due to light assisted collisions is a double edged sword. On the one hand, it ensures that every occupied site has only one atom, which is desirable, but on the other, the presence of empty sites distributed randomly is not conducive to large scale entanglement generation and distribution among qubits.

Several approaches have been tried to overcome this problem. A blue detuned beam can be applied to protect atoms against this loss [67, 68, 93]. The blue detuned beam couples pairs of atoms to repulsive molecular states, with energy large enough to induce a single atom loss but small enough to not cause double loss. Single microtraps were loaded with up to 90% efficiency using this technique, but it requires traps that are several milli-Kelvins deep, which is a significant challenge for large arrays of traps. More recently, atoms were loaded in 10×10 array of 1 mK deep traps with 80% efficiency by replacing the blue-detuned light with gray molasses [94]. Much higher efficiencies (on the order of 99%) can be obtained with a superfluid to Mott insulator transition [36, 95–97], but this is untenable in a large spacing optical lattice where the tunneling is heavily suppressed. Moreover, the non-zero temperature of the superfluid leads to randomly distributed residual defects in the Mott insulator, which would be a significant source of error in a quantum computation without the knowledge of defect positions in the array.

Here we use the approach proposed by Weiss et. al. in reference [46] to solve this problem. The basic idea is to start with a randomly half filled lattice in which single atoms can be imaged to ascertain their position. The atoms are then moved individually in state dependent lattices to create a fully-filled sub-lattice and yield a deterministic initialization for subsequent quantum computation. Since the atoms

can be moved in parallel, it can be shown that the number of motion steps required to "sort" N qubits scales as $O(N^{1/3})$ in a three dimensional array [98]. This is a very favorable dependence for scaling to a large system of qubits.

Furthermore, this selective motion of atoms to yield an ordered state from a disordered one can also be thought of as an implementation of Maxwell's demon. The "demon" thought experiment was proposed by Maxwell in 1872 [99], as a challenge to the second law of thermodynamics. In essence, Maxwell imagined a creature that could observe individual particles of a gas in a box with two compartments. The demon could selectively open door between the compartments, letting the faster particles accumulate in one side and the slower ones in the other, thus sorting the gas into hot and cold components. This would decrease the entropy of the gas and the temperature difference would allow work to be extracted from the system, in violation of the second law of thermodynamics. This apparent paradox was resolved by theoretical work in the twentieth century. It was shown that the entropy of the universe would increase in collecting information about the system with irreversible processes, which is often the case [100,101]. Although it is possible to collect that information in a reversible way, there is always an entropy increase as well as an energy cost associated with the demon's memory [101,102].

Several experiments, in atomic [103], solid state [104], and photonic systems [105] have previously been thought of as demon-like. In some of these experiments the system is thought of as an analogue for the demon, but there is no intermediate information storage and feedback based on collected information. In that sense, any experiment which involves lowering the entropy by selectively acting on particles based on their state can be thought of as demon-like. This is essentially true for all laser cooling methods. But, the resolution of the aforementioned paradox is intimately tied together with the intermediate information collection and storage step, and the demon's memory. Therefore, it is an essential feature of Maxwell's thought experiment.

Atoms have also been sorted in arrays of dipole traps in 1D, 2D and, concurrently with this work, in 3D [38–40, 42, 106, 107]. These experiments have essentially followed the ideas proposed in reference [46], by observing the the locations of atoms and moving the traps by changing the parameters of the dipole array generating acousto-optic modulators or spatial light modulators. The entropy decrease in these experiments has been limited to $< 10\%$, because most of the

entropy in these systems is contained in the vibrational excitation of the atoms in the trap.

A small system of four atoms was sorted in a 1D optical lattice, with the ground vibrational state population of $\sim 60\%$, resulting in an appreciable decrease in system entropy [108]. Here we sort 50 atoms in a 3D optical lattice, with a ground state population of $\sim 89\%$, significantly reducing the total entropy, and realizing the true essence of Maxwell’s demon in a large system. From a quantum computing perspective, on the order of 50-100 qubits are required to solve a classically intractable problem [60]. Moreover, colder atoms make better qubits in general, with higher gate fidelities and coherence times. Good vibrational cooling is especially advantageous in a blue detuned lattice like ours, where cold atoms have a significantly smaller spontaneous emission rate for lattice photons.

4.2 Moving atoms

Our experiment typically proceeds as follows. After loading the atoms into the lattice we image them and create an occupancy map. Based on the occupancy map, a sorting algorithm creates a plan for moving individual atoms to yield a desired fully-filled target sub-lattice. This plan consists of a list of motion steps and each motion step has a list of atoms, referenced by their location in the lattice, and the direction they need to be moved in by one lattice site. Essentially, each motion step is a parallel translation of some number of atoms in a given direction. We can move atoms in any of the six lattice directions. This plan is then communicated to our real-time control system (Appendix A), which then takes control of all the analog and digital channels required to execute the plan. After the sorting is finished, we can image the atoms again and if there are any errors, vacancies in the target sub-lattice, execute the sorting again to fill them.

Sorting requires selective motion of individual atoms in all the lattice directions. This is accomplished by combining two previously touched upon techniques in this thesis: *single site addressing*, which was used to construct single qubit gates, to flip the atom between two states and *state dependent optical lattices*, which were used to translate the atomic wave-function in projection sideband cooling, to move the atoms.

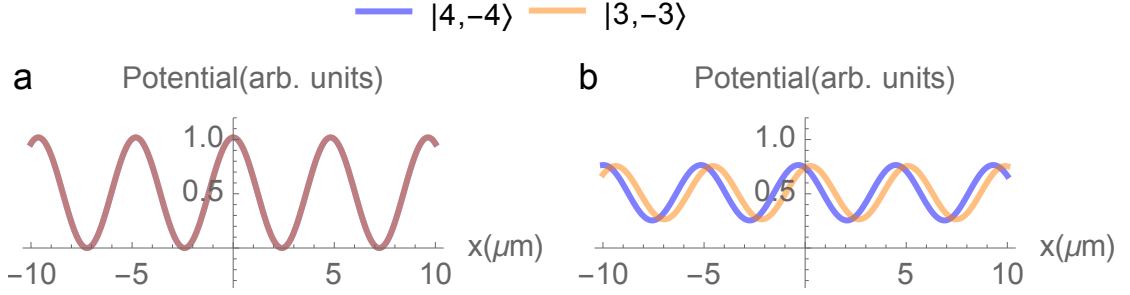


Figure 4.1. State dependent lattices. The potentials experienced by atoms in states $|4, -4\rangle$ and $|3, -3\rangle$ are shown. They are overlapped when (a) $\phi = 0$ and displaced in opposite directions when (b) $\phi = \pi/4$. Here ϕ is the angle between the polarizations of the two lattice beams. in a pair. When $\phi \neq 0, \pi$, the lattice depth is smaller and the intensity at the trap minima is not zero.

4.2.1 State dependent lattices

Consider a one dimensional optical lattice formed by two counter-propagating beams in the x direction, with their linear polarizations at some angle ϕ . The electric field for the resulting standing wave can be expressed as [98]:

$$E(x) = \sqrt{2}E_0(e^{i\phi} \cos(kx - \phi/2)\mathbf{e}^- - e^{-i\phi} \cos(kx + \phi/2)\mathbf{e}^+), \quad (4.1)$$

where E_0 is the peak electric field of one beam, k is the wave vector magnitude and \mathbf{e}^- and \mathbf{e}^+ are the unit vectors for σ^- and σ^+ polarization components, respectively. This looks like two standing waves with different polarizations shifted in space by $\phi/2k$. Now consider an unrealistically simple atom with total angular momentum $J = 1/2$ for ground and excited states. Due to polarization selection rules, the $m_J = 1/2$ ground state only couples to the excited state through the σ^- polarization while the $m_J = -1/2$ state only couples through the σ^+ polarization. As a result, the lattices experienced by atoms in these two states are shifted from each other by a distance proportional to ϕ .

In real atoms, there are many different magnetic sub-levels which complicate the situation. Atoms in any state are typically influenced by all the polarizations, and the resultant lattice differentials depend on variations in transition strengths for different sub-levels and polarizations. The ac stark shift in these atoms is typically a combination of a scalar shift, independent of the m_F sub-level and a vector light shift which is proportional to the m_F value of sub-level. Moreover, in our system,

there are three lattices, all at some non-zero angles to the quantization axis, and thus they can also have a π polarization component. For Cesium atoms in our system, we use the $|4, -4\rangle$ and $|3, -3\rangle$ of the $6S_{1/2}$ ground manifold as our sorting states. In our configuration, the lattice potentials for these two states are given by [44]:

$$U_{4,-4} = \frac{1}{6}U_{\sigma^+} + \frac{1}{2}U_{\sigma^-} + \frac{1}{3}U_{\pi} \quad (4.2)$$

$$U_{3,-3} = \frac{11}{24}U_{\sigma^+} + \frac{5}{24}U_{\sigma^-} + \frac{1}{3}U_{\pi}, \quad (4.3)$$

where U_i is proportional to the light intensity in the i th polarization component. It is clear that the two states are predominantly influenced by either σ^+ or σ^- polarizations, and equally influenced by the π polarization. Figure 4.1 shows a calculation of these potential for the two states for $\phi = 0$ and $\phi = \pi/4$. To move atoms, we rotate the polarization of one of the lattice beams in the beam pair forming the lattice in the direction in which motion is desired. A rotation by π moves atoms in the two states by half a lattice spacing in two opposite directions. It is noteworthy that the light intensity at trap minima is non-zero for $\phi \neq 0, \pi$. This leads to a significant enhancement of lattice photon spontaneous emission in such a configuration, and will turn out to be a major source of errors in moving the atoms. Moreover, as shown in Figure 4.2b, the displacements of the atoms are not linearly proportional to ϕ any more, due to the dependence of the potential on all polarization components. The lattice depths also drop sharply during the fastest part of the motion, increasing the time constant for adiabatically translating the atom (Figure 4.2b).

We rotate the polarization of one lattice beam in each pair by using two electro-optic modulators (Pockels cells) and a quarter wave plate in its path. The Pockels cells are Fastpulse model 1145SG, which can safely reach up to a half-wave voltage, $V_{\lambda/2}$. The setup is shown in Figure 4.3. The Pockels cells are aligned with their axes at 45° to the incoming linear polarization of the lattice beams. The Pockels cells cause a variable differential phase shift, θ , between the polarization components along their axes. This phase shift is proportional to the applied voltage and the resulting polarization is elliptical in general. The quarter wave plate, which is aligned with the incoming linear polarization before the Pockels cell, converts this elliptical polarization into a linear polarization rotated by $\theta/2$. Thus, using

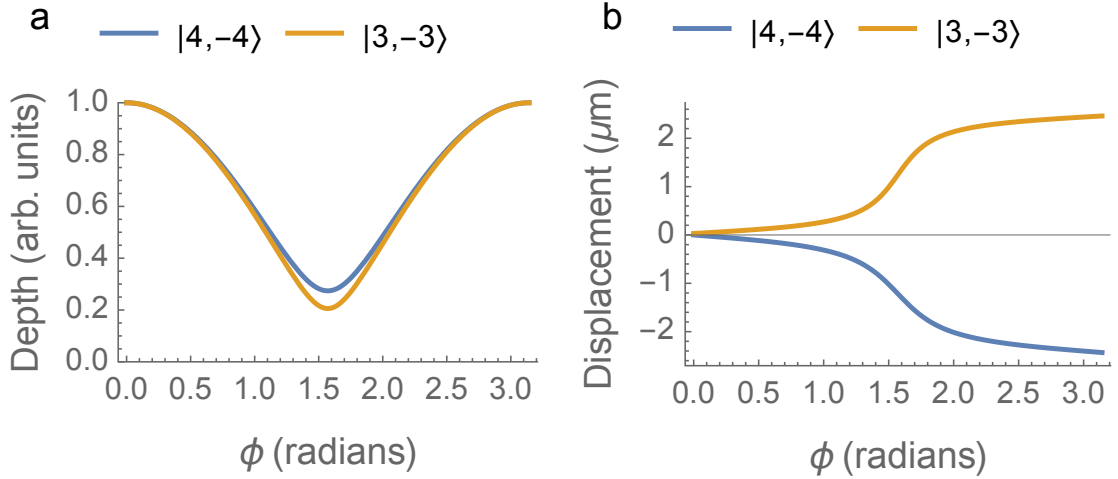


Figure 4.2. Depth and displacement of state dependent lattices. (a) shows the depth of the lattices experienced by atoms in $|4, -4\rangle$ and $|3, -3\rangle$ as a function of the angle between the polarization of the lattice beams, ϕ . (b) shows the displacement of the state dependent lattices as a function of ϕ . The lattices move in opposite directions in a highly non-linear fashion. After the polarization has been rotated by π the lattices converge again, having moved half a lattice spacing in opposite directions. These two parameters (depth and the rate of displacement) set the adiabatic timescale for changing ϕ to move the atoms.

two cells, we can rotate the polarization by π . The direction of the polarization rotation (clockwise or counter-clockwise, controlled by the sign of the voltage) determines the direction of motion (for example $+x$ or $-x$) for the atoms, thus allowing us to move atoms in any of the six directions in the lattice.

4.2.2 Site selective state flips

To create a site selective motion step, we use our single site addressing to flip the state of the atoms to be moved from $|4, -4\rangle$ to $|3, -3\rangle$. The requirements for this state flip are much less stringent than arbitrary single qubit gates because no coherence needs to be maintained. Moreover, the required addressing beam powers are a factor of seven smaller than to comparably shift, say, the $|4, -1\rangle$ to $|3, 0\rangle$ transition, because of the higher m_F values. This significantly reduces the spontaneous emission from the addressing beams and allows us to use the longer and more efficient Adiabatic Fast Passage pulses for this transfer.




Figure 4.3. Setup for rotating the polarization. Pictured is the y-lattice optical components for rotating the polarization of one of the lattice beams. The lattice beams are shown in red. One of them (left) passes through two Pockels cells and a quarter wave plate. The polarization of the beam is linear and nearly perpendicular to the plane formed by both the beams before the cells. The Pockels cells have their axes at 45° to the incoming polarization. The quarter wave plate has one of its axis aligned to the polarization before the Pockels cells. The rotation angle of the polarization, ϕ , is directly proportional to the voltage on the Pockels cells.

The performance of this site selective transfer depends on various factors. For fixed parameters of an AFP pulse, there is an optimum addressing shift that provides balance between frequency noise due to addressing beam pointing fluctuations and the AFP scan range. The AFP scan range, on the other hand, is limited by the requirement of not scanning over the vibrational sidebands. Thus, for robust performance, it becomes important to ensure roughly equal addressing shifts for all the sites in spite of the frequency insensitivity of the AFP pulse. To accomplish this, we calibrate the addressing powers for each site individually as described in Appendix B. The $|4, -4\rangle$ to $|3, -3\rangle$ transition is also sensitive to the periodic 60 Hz

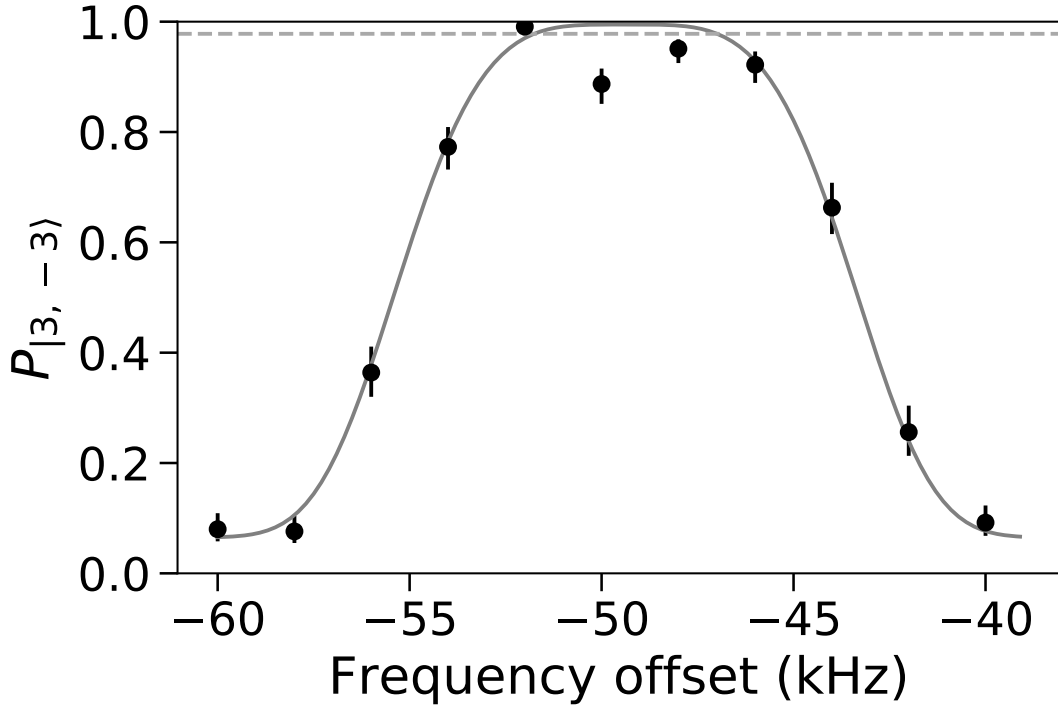


Figure 4.4. AFP spectrum for addressed sites. The Figure shows the spectrum of target atoms with a 3ms long AFP microwave pulse with a 12kHz frequency sweep. All the available atoms in a $5 \times 5 \times 5$ section of the lattice were targeted sequentially in each shot with the addressing beams and an AFP microwave pulse to drive the $|4, -4\rangle$ to $|3, -3\rangle$ transition. The horizontal axis is the offset from transition frequency of the un-addressed atoms. Each point in the plot is the signal from ~ 100 atoms. The gray curve is the fit to a super-gaussian of order 4. The dashed gray line is the limit set by atom loss due to background gas collisions.

(and harmonics) magnetic field fluctuations. Typically, these fluctuations have a peak to peak amplitude of $\sim 1\text{mG}$, which implies a transition frequency fluctuation of $\sim 2.5\text{kHz}$. Furthermore, we have observed that the nature of these fluctuations doesn't change drastically for weeks or even months. This allows us to calibrate the magnetic field related shift over a 16.67ms cycle and "re-sync"¹ the sequence to the line before the sorting starts. The calibrated frequencies can now be used depending on when a given pulse occurs relative to the re-sync point.

¹While implementing the first kind of single qubit gate, we discovered that there is significant phase noise in the line. Even though the sequence was synchronized to the line at the beginning, we observed shot to shot fluctuations due to the phase fluctuations of the 60Hz B-field. To alleviate this problem, we devised a simple digital circuit to re-sync the sequence to the line at critical points like before projection cooling, again before gates and before sorting.

Typically, we use a 3 ms long AFP pulse with a 12 kHz sweep range, along with an addressing shift of $\sim 50\text{ kHz}$ to drive these transitions. Figure 4.4 shows the transition spectrum for addressed atoms, with an addressing shift of $\sim 50\text{ kHz}$, as the central frequency of the AFP pulse is scanned. To obtain this data, we individually targeted all the atoms that are loaded in a $5 \times 5 \times 5$ section of the lattice in each shot, as we changed the microwave frequency through different shots. The data is thus the average performance over the entire $5 \times 5 \times 5$ section. To ascertain the addressing fidelity, we conducted a similar experiment, but this time we fixed the frequency at the resonance and repeated for 50 shots. This gives the average fidelity of our site selective state flips to be ~ 0.986 . It should be noted that atoms were projection cooled for this measurement, and the fidelity of AFP transfers is about a factor of two better for projection cooled atoms compared to atoms that are just polarization gradient cooled. The state flip fidelity is limited by pointing fluctuations of the beams over short time scales and the slow drift of the magnetic field and the alignment of the beams to the atoms over longer time scales. Typically, we observe that the addressing fidelity falls by 0.02 over the course 100 implementations of the experiment, mostly due to the slow drift of the addressing beams.

4.2.3 Filling a vacancy

The site selective state flips and state dependent lattices can now be combined to move individual atoms and fill vacancies in the lattice. The basic sequence of steps is outlined in Figure 4.5 for two atoms separated by a vacant lattice site. In general, all atoms are initially prepared in $|4, -4\rangle$, also called the *stationary* state, by optical pumping. The atoms to be moved are sequentially flipped to $|3, -3\rangle$, also called the *motion* state. As the polarization of a lattice beam for a given lattice direction is rotated, the atoms in stationary and motion states move in opposite directions by half a lattice spacing after the polarization has been rotated by π . At this stage, the atoms in the motion state ($|3, -3\rangle$) are optically pumped back to the stationary state ($|4, -4\rangle$) by a global pulse. As the polarization of the lattice beam is rotated back to zero, all the atoms move in the same direction, with atoms initially in the stationary state returning back to their original location and the atoms initially flipped to the motion state moving by one lattice site. This sequence




Figure 4.5. Filling a vacancy . The curves show the lattice potential as a function of position for the *motion* state (orange curve) and the *stationary* state (blue curve). Brown curves indicate overlapping potentials. The arrows denote the direction of a time series in which the angle (ϕ) between the polarizations of the two lattice beams is adiabatically ramped to π and back to 0. The atom to be moved is transferred to the motion state (orange circle) using targeted addressing at the beginning of the time series. As the polarization of one of the lattice beams is rotated, the atoms in the motion state and the stationary state (blue circle) move in opposite directions, settling half a lattice spacing away from their original positions when $\phi = \pi$. The atom in the motion state is then optically pumped to the stationary state (illustrated by the red arrow). As the polarization is rotated back, both atoms move in the same direction, with the atom that started in the stationary state returning to its original position and the atom that started in the motion state moving by one lattice site. 65




Figure 4.6. Timing sequence for a motion step. A motion step to move n atoms is shown. n atoms are sequentially targeted by the addressing beams and transferred from the *stationary* state to the *motion* state using microwaves. The Pockels cell voltages are ramped up to the half-wave voltage ($V_{\lambda/2}$) on each in order to move atoms by half of the lattice spacing. After motion, the atoms are optically pumped so that they all return to the stationary state. The voltages are then ramped back down. A final optical pumping ensures optimal preparation for the next motion step.

of steps moves a chosen set of atoms by one lattice spacing and resets the system back to the initial point, where the same sequence can be repeated with another set of atoms in another direction.

Figure 4.6 shows the pulse sequence schematic for one motion step. The sequence starts by flipping a given set of atoms one by one from the stationary state to the motion state with single site addressing and 3 ms long AFP pulses. The addressing laser powers are ramped up over $40\ \mu\text{s}$, after which we wait for another $110\ \mu\text{s}$ for our intensity lock to settle before every AFP pulse. The voltages on each of the two Pockels cells in a lattice beam are linearly ramped to the half-wave values in 3 ms . This results in a cumulative differential phase shift of 2π and a rotation of the polarization by π . The half-wave voltages for the Pockels cells are $\sim 5.5\text{ kV}$ and are provided by TREK 610 HV supplies with a maximum slew rate of $20\text{ V}/\mu\text{s}$, although here we are limited to slower ramps by the adiabaticity requirements to not heat the atoms during motion. After the ramp, we wait for $250\ \mu\text{s}$ for the voltages to settle and apply a $250\ \mu\text{s}$ long optical pumping pulse to flip all the atoms in the motion state back to the stationary state. The polarization is then rotated back to 0 radians in another 3 ms , moving the initially flipped atoms by one lattice site. A final optical pumping pulse ensures clean preparation for the

next motion step.

There are several technical considerations related to the execution of this motion step. First, throughout the motion step, both the Pockels cells have the same voltage, even though the half-wave voltage values for the two cells might be different. The small differences from the half-wave values at the peak voltage point in the sequence do not matter, as long as the cumulative phase shift is 2π . To precisely find this value of the peak voltage, we use the fact that when polarization is rotated by 0 or π there is no matrix element for the vibrational level changing microwave transitions from $|4, -4\rangle$ to $|3, -3\rangle$. Thus, we vary the voltage around the expected half-wave value for each cell and look for the suppression in $\Delta\nu = 1$ vibrational sideband transition from $|4, -4\rangle$ to $|3, -3\rangle$, where ν is the vibrational quantum number². Second, the optical pumping in the motion step heats about $\sim 7\%$ of the population to higher vibrational levels. This can be avoided by replacing the optical pumping with another set of site selective state flips. We chose to avoid the hassle of aligning addressing beams to half-way in between the lattice sites at the cost of heating. Finally, all the timing parameters were empirically optimized by measuring errors in motion steps as described in the next section.

4.2.4 Motion fidelities

The primary process that introduces errors in motion, in the limit that the motion is adiabatic, is spontaneous emission of the lattice photons. In a blue detuned lattice like ours, the spontaneous emission is heavily suppressed because the atoms are trapped at the intensity minima. But during a motion step, the atoms experience much more light intensity, as shown in Figure 4.4, leading to a 10 times higher rate of spontaneous emission. During the motion, an atom can spontaneously emit a lattice photon and scatter to the "wrong" state, in which it is anti-trapped at its location, leading to atom loss from the lattice. Sometimes it can scatter to the wrong state, stay trapped and start moving in the wrong direction. As a result, an atom that was supposed to move, would have not moved at all and vice versa. In the process of sorting, this can eventually lead to two atoms at the same site, which are lost due to pair-wise light assisted collisions when they are imaged, leading to

²This is the same technique we use to find the "zeroing" voltages for the cells. This is important for projection sideband cooling.

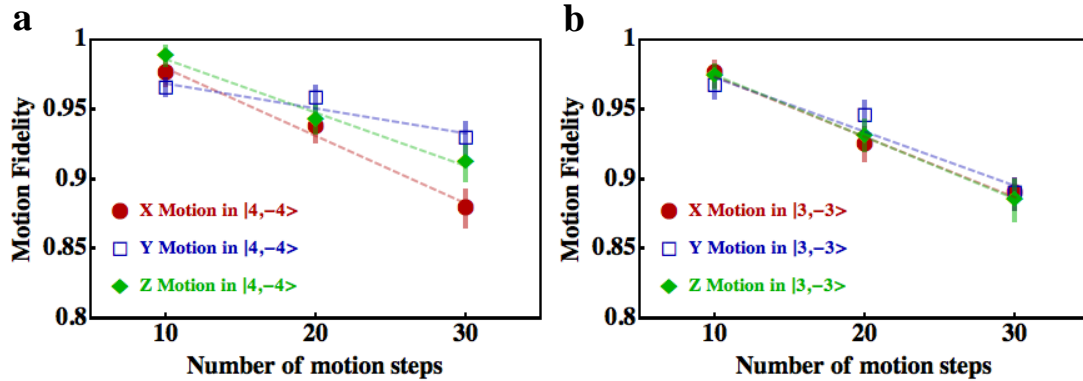


Figure 4.7. Measured motion fidelity as a function of the number of motion steps in (a) $|4, -4\rangle$ and (b) $|3, -3\rangle$ in the x (maroon circles), y (blue squares) and z (green diamonds) directions. The lines are fits to the data. The error bars represent one standard deviation. Each point corresponds to about 600 atoms.

"double" loss.

We measured the motion fidelities for both the states ($|4, -4\rangle$ and $|3, -3\rangle$) in each of the lattice directions. For $|4, -4\rangle$, we started by preparing all the atoms in $|4, -4\rangle$ by optical pumping. We then ramped each of the Pockels cells to the half-wave voltage, $V_{\lambda/2}$, followed by a ramp to $-V_{\lambda/2}$, and another final ramp to 0. This sequence was repeated several times in a shot leading to atoms moving back and forth around their initial lattice site by half a lattice spacing. Every N repeats of this sequence correspond to motion during $2N$ motion steps. After N repeats, the atoms were imaged again and the fraction of atoms left at their original site was treated as the fidelity of $2N$ motion steps. Figure 4.7a shows this fidelity as a function of number of the motion steps for each lattice direction. Each point is an average of 10 experiment implementations, corrected for loss due to background gas collisions. A linear fit to this data gave us the fidelity per motion step. For atoms in $|4, -4\rangle$, the fidelities are: $\{0.9951(6), 0.9982(6), 0.9962(4)\}$ for the $\{x, y, z\}$ directions. A similar experiment was done with all the atoms in the $|3, -3\rangle$ state by using an AFP pulse to transfer the atoms to $|3, -3\rangle$ after optical pumping to $|4, -4\rangle$. The motion fidelities for atoms in $|3, -3\rangle$ are: $\{0.9956(4), 0.9961(10), 0.9956(1)\}$ for the $\{x, y, z\}$ directions (Figure 4.7b).


4.3 Sorting Algorithm

The sorting algorithm to generate a sequence of motion steps that compacts a collection of atoms in a three dimensional lattice to a smaller sub-lattice was proposed in reference [98]. The implementation of that algorithm results in a roughly even distribution of atoms, similar to one that would result from starting at one corner of the lattice and filling it site by site iterating over the co-ordinates one by one. While this algorithm yields a compact distribution of all the atoms, it doesn't result in a regularly filled arbitrary sub-lattice. For our experiment we have modified the algorithm to do just that i.e. starting from a half filled $N \times N \times N$ lattice, to yield a fully filled $X \times Y \times Z$ sub-lattice given $XYZ < N^3/2$. This section provides a brief review of the sorting algorithm as well as our modifications to the algorithm. The details of the implementation of the algorithm in C++ are provided in appendix C. The sorting algorithm has two stages: *balancing* and *compaction*. Suppose the objective was to compact the atoms in the z direction, with fully filled $x - y$ planes as far as possible. In that case, balancing results in an equal number of atoms in all the z -rows (i.e. x and y co-ordinates are constants) of the atoms as far as possible. For example, in a $5 \times 5 \times 5$ lattice loaded with 52 atoms, balancing would result in two atoms in 23 of the z -rows and three atoms in the remaining 2 z -rows. A simplified example of this is shown in Figure 4.8 for a $3 \times 3 \times 3$ lattice, where balancing leads to two atoms in each z row. After balancing, atoms in each of the z -rows are moved towards one of the $x - y$ planes in parallel (compaction), without putting them on the same site. Say the objective was to compact the already balanced z -rows towards the lowest $x - y$ plane, as shown in the figure. Compacting requires following simple steps:

1. For each z -row start at $z = 0$ and find the first vacancy at coordinate z_V , say. There will always be vacancies for 50% loading and a balanced lattice.
2. Move all the atoms at $z > z_V$ by one lattice site towards $z = 0$. If there are no atoms beyond the first vacancy, the z -row is already compacted and is to be ignored for future iterations.
3. Repeat these steps until all the rows are compacted.

This fills the $x - y$ planes one by one leading to a compacted lattice. Balancing, on

Balancing



Compaction




Figure 4.8. Balancing and compaction. Simplified illustration of two parts of the sorting algorithm in a $3 \times 3 \times 3$ lattice. Orange and blue circles denote atoms in *motion* and *stationary* states respectively; empty circles denote empty sites. The first motion step 'balances' the array so that every z -row has exactly two atoms. The second motion step 'compacts' atoms into two planes.

the other hand, is more complicated and will be discussed in the next section.

4.3.1 Balancing

The basic idea of balancing is to divide the lattice into two halves and move atoms from one half to the other, such that each one has an equal number of atoms. This procedure is then repeated recursively for those halves. Consider an $N \times N \times N$ lattice, L with M atoms. After L is balanced, each z -row would have at least $k = \lfloor \frac{M}{N^2} \rfloor$ atoms. An informal description of the algorithm for balancing is as follows:

1. If this is the first iteration, choose a dividing plane, P , to be an $x - z$ plane. Otherwise, choose the dividing plane to be perpendicular (either $x - z$ or $y - z$) to the previous iteration. Choose P to divide the lattice into two sub-lattices, S_1 and S_2 , that are as similar as possible (that is, a difference of one plane between S_1 and S_2 is permitted if the lattice dimension is odd).
2. If the number of z -rows is p in S_1 and q in S_2 , the required number of atoms in S_1 is pk and the required number of atoms in S_2 is qk . Move atoms between the two sub-lattices until they each have at least the required number of atoms.
3. Repeat these steps for both S_1 and S_2 until each one of them is a single z -row.

There are several non-trivial considerations with regard to the implementation of the balancing algorithm. The first consideration relates to the order in which the sub-lattices are balanced. The balancing algorithm can be thought of as construction of a binary tree with a sub-lattice and a sequence of sorting steps associated with each node of the tree. As shown in Figure 4.9, it starts with the whole lattice, L at the root node, with steps for the balancing of its two halves. It splits into its two sub-lattices, which further split into their own sub-lattices and so on until a single z -row at each of the "leaves". All the nodes at the same depth of the tree are balanced across the same plane direction, either $x - z$ or $y - z$. Therefore, for an optimized implementation that minimizes the total number of motion steps, it is necessary to combine the motion steps at all nodes at the same depth as far as possible. A simple recursive implementation of the algorithm would result in a depth first construction and traversal of the tree, i.e. each branch

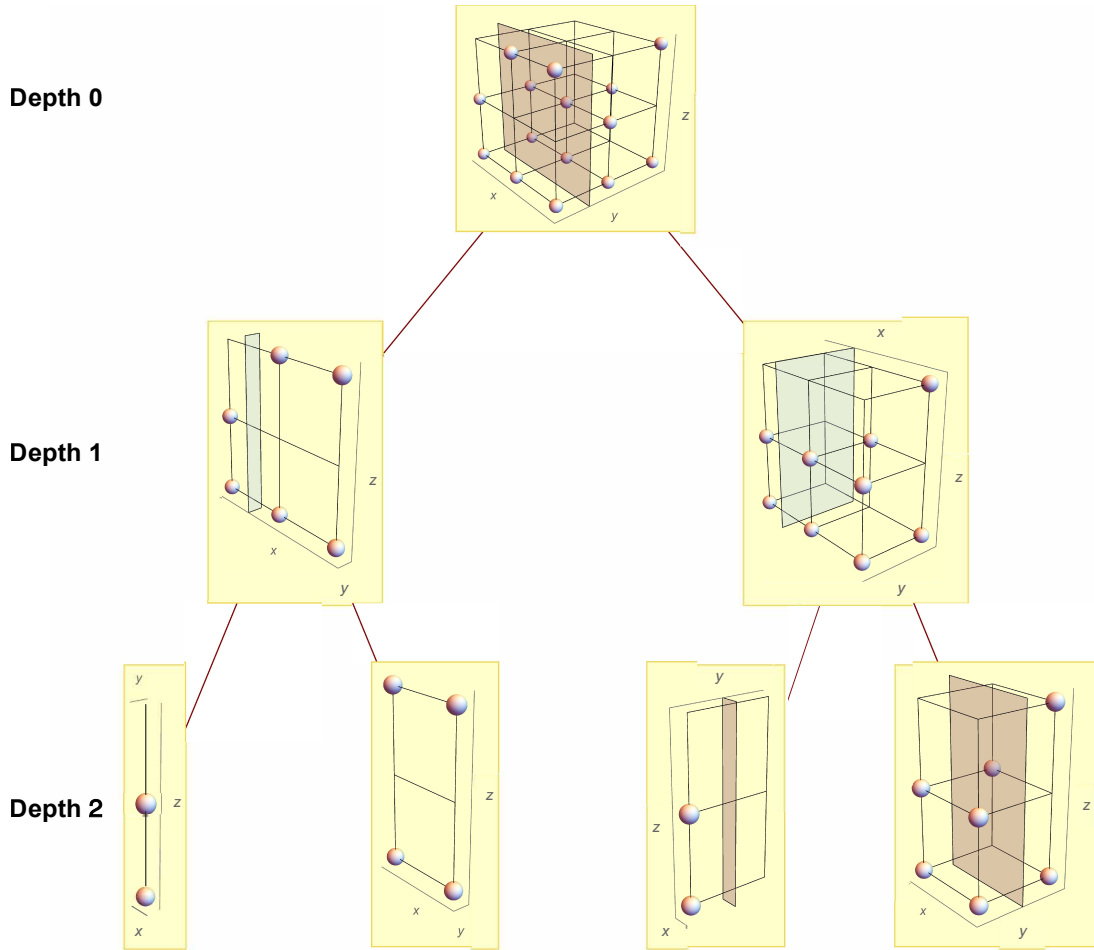


Figure 4.9. The balancing tree. The Figure shows the first three levels of the balancing tree for a $3 \times 3 \times 3$ lattice. The balancing is done level by level starting at the root node with the whole lattice. To start, the lattice is divided into two sub-lattices by an $x - z$ plane (brown). Atoms are moved across the plane to balance the sub-lattices. In the next level, the the two sub-lattices are divided by $y - z$ planes (blue) and balanced in parallel. The process is continued, exhausting the tree one level at a time until all the leaves (z -rows) are reached. The dividing planes alternate as the levels increase.

traversed wholly before moving onto the next branch. This would lead to an inefficient sequence of motion steps. Thus, a breadth first traversal of the tree is required, where each depth of the tree is exhausted before the next one. The details of this implementation are provided in appendix C.

Second, moving a given number of atoms, say n , from one sub-lattice to the other, say S_1 to S_2 , has many possibilities for which n atoms to choose in S_1 and where to move them in S_2 . Our algorithm chooses a sequence that minimizes the

total number of motion steps and addressings required to balance the two sublattices, with an additional constraint that atoms are only moved in the direction perpendicular to the dividing plane, P between S_1 and S_2 . Briefly, the steps are as follows :

1. If this is the first iteration, start with the plane, P' in S_1 , closest to P . If not chose P' to be the next plane in S_1 , from the plane chosen in previous iteration.
2. Find all the atoms in P' , and for every atom in P' find the closest vacancy in S_2 in the same "column". Keep track of the distances of the atoms to their closest vacancies.
3. If p total atoms have been found before this iteration, choose $n - p$ atoms found in this iteration with the lowest distances to be moved. If the the number of atoms found in this iteration are not enough, repeat steps 1-3, otherwise, terminate.

It will often be the case that an atom found in S_1 , to fill a vacancy in S_2 , will have several atoms "in the way" in S_2 . In that case, all those atoms, including the atom in S_1 , are moved until the vacancy is filled. The atom in S_1 is moved until it enters S_2 . It is also possible that there would be multiple atoms in the same "column" of S_1 to be moved to fill some vacancies. The previous sequence of steps then easily generalizes. These steps can be thought of as, for that particular column, compacting all atoms starting at the farthest atom found in S_1 towards the farthest vacancy found in S_2 .

4.3.2 Modifications to the algorithm

In quantum computation and quantum simulation experiments, the requirement is often to sort the atoms in a half filled larger $N \times N \times N$ to a fully filled cuboid $X \times Y \times Z$ lattice. While the previously discussed sorting algorithm would yield a compacted lattice for which a fully filled $N \times N \times Z$ lattice would always be a subset, it doesn't result in an arbitrary cuboid given enough atoms. Moreover, even for an $N \times N \times Z$ target sub-lattice the number of motion steps and state flips needed to implement the full sorting would always be more than just sorting to that target.

When the target sub-lattice is $N \times N \times Z$, there is a simple modification to the balancing algorithm that results in an optimum number of motion steps. Instead of doing a full balance i.e. ensuring close to equal number of atoms in each of the z -rows, it is sufficient to balance such that each z -row has at least Z atoms. For example, in sorting atoms in a $5 \times 5 \times 5$ lattice to a single plane, it is acceptable to not move any atoms into a sub-lattice if every z -row in that sub-lattice has at least one atom, even though its other half might have substantially more atoms. The full balancing, on the other hand, would lead to moving atoms into this sub-lattice until every z -row had at least two atoms, in case of half-filling. Similarly, the compaction has to be modified to move only Z atoms in each z -row to the target planes, instead of compacting all the atoms in each row towards a plane.

In the more general case of an $X \times Y \times Z$ target sub-lattice, atoms are first moved into an $X \times Y \times N$ section of the lattice. This is done by looking at each of the sub-lattices outside the $X \times Y \times N$ section in the larger lattice, and moving as many atoms as possible from each one of them with motion steps in only one direction, which is perpendicular to the dividing plane of the outer sub-lattice and the $X \times Y \times N$ section. This ensures that we have enough atoms in the $X \times Y \times N$ section, which can now be run through our modified balancing and compacting to yield a fully filled $X \times Y \times Z$ lattice. For example, if it was required to sort atoms in a $5 \times 5 \times 5$ lattice to a $4 \times 4 \times 3$ section, atoms are moved from the outer $1 \times 5 \times 5$ and $5 \times 1 \times 5$ planes to the remaining $4 \times 4 \times 5$ section with motion steps in x and y directions respectively. The $4 \times 4 \times 5$ section is then balanced and compacted to a $4 \times 4 \times 3$ lattice. While it may seem that this moving of atoms from outer sub-lattices does not guarantee enough atoms in the $X \times Y \times N$ section, monte-carlo simulations show that it always works with $> 50\%$ loading and sorting to a lattice with $XYZ < N^3/2$. This approach of bringing atoms from outer sub-lattices as a first step is probably sub-optimal from an algorithmic perspective, but works well nonetheless.

4.4 Results

Starting with a $5 \times 5 \times 5$ half-filled lattice, we implemented the sorting algorithm in our system to create a fully filled 5×5 plane of atoms, and fully filled $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sub-lattices. We started by cooling atoms to the ground state and

implemented the sorting algorithm based on the distribution of atoms and the required target sub-lattice. After sorting, we re-imaged the atoms to determine the vacancies in the target sub-lattice due to errors in the first "sort". The sorting algorithm was then applied again to fill these vacancies, this time without projection sideband cooling. The last two steps were repeated again for a total of three sorts.

After three sorts, we can fill a plane perfectly $> 70\%$ of the time, with a filling fraction $> 98.5\%$, essentially limited only by loss due to background gas collisions. Filling a plane is considerably simpler than sorting to other more complicated sub-lattices, as it almost never requires more than one balancing step, since the expected number of empty z -rows is 1.5 for 50% filling. The average number of motion steps required for filling a plane are 3.86 with 14.4 state flips in total, dominated by the compaction step, which can require atoms to be moved to the central plane from two planes away in either direction.

For $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sub-lattices, we achieved 97% and 95% filling fractions respectively, after three sorts. The perfect filling rate was 32% for the $5 \times 5 \times 2$ sub-lattice and 27% for the $4 \times 4 \times 3$ sub-lattice. $5 \times 5 \times 2$ required an average of 6.4 motion steps and 38 state flips in total, while $4 \times 4 \times 3$ required 5.6 motion steps with 62 state flips. The significantly higher number of state-flips for $4 \times 4 \times 3$ result from the first two motion steps which try concentrate all the atoms in the lattice to a $4 \times 4 \times 5$ section. The total time for each sort is ~ 190 ms. Figure 4.10 shows implementations of sorting to $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sections, which were successful after two sorts. The filling fractions for both the target sub-lattices as a function of number of sorts are shown in Figure 4.11 by the solid red lines. The filling fraction saturates to a value less than 1 when the rate at which errors are introduced in the already sorted atoms balances against the rate at which the vacancies are filled. The errors are due to the motion steps that are required to fill the remaining vacancies, leading to double loss and hopping of already sorted atoms and the atom loss because of background gas collisions. Thus, more sorts would not be helpful at this point and the filling fraction would eventually decrease when, due to loss, there aren't enough atoms available to fill the vacancies. The extra atoms left after sorting can either be cleared out of the lattice or left for later use to potentially correct atom loss errors.




Figure 4.10. Sorted $5 \times 5 \times 2$ and $4 \times 4 \times 3$ sub-lattices. (a)-(f) The five images in each row correspond to the five lattice planes (labelled 1-5). The colour map shows intensity. The associated grid patterns are real-time occupancy maps, generated by processing the five images. The two sets of images, (a)-(c) and (d)-(f), are from two different experimental implementations. (a) An initial, unsorted atom distribution. The occupancy maps are used as the basis of a series of site-selective state flips and state-selective translations that execute our sorting algorithm. (b) Result after one sorting sequence with the goal of filling a $4 \times 4 \times 3$ sub-lattice in planes 2-4. There are three errors after this sort (one in plane 3, two in plane 4). (c) Result after a second sorting sequence starting from the distribution in b. The sorting goal has been reached. Atoms outside the target sub-lattice can be kept as spares, or they can be selectively state-flipped and removed by a resonant clearing beam. (d) Another initial, unsorted atom distribution. (e) The result after one sorting sequence with the goal of filling a $5 \times 5 \times 2$ sub-lattice in planes 2 and 3. There are four errors after this sort (one in plane 2, three in plane 3). (f) The result after a second sorting sequence starting from the distribution in e. The sorting goal has been reached. The absence of spare atoms in f is coincidental.




Figure 4.11. Filling fraction and entropy. The empty red symbols show the filling fraction as a function of the number of sorts for $5 \times 5 \times 2$ (circles) and $4 \times 4 \times 3$ (squares) target sub-lattices. The circles (squares) show results based on 85 (48) experimental implementations. The red horizontal dashed line is the limit associated with loss from collisions with background gas atoms during the 1 s required to image and sort. The solid blue symbols show the configurational entropy as a function of the number of sorts for $5 \times 5 \times 2$ (circles) and $4 \times 4 \times 3$ (squares) target sub-lattices. The total entropy at the beginning and at the end is the sum of the vibrational entropy (blue horizontal dotted line) and the configurational entropy; sorting reduces it by a factor of 2.44. The $1 - \sigma$ error bars are smaller than the size of the symbols.

4.4.1 Entropy and Maxwell's demon

After sorting the atoms, we projection sideband cooled them to the ground state, measured the vibrational sidebands to ascertain the populations in various vibrational states and calculated the vibrational entropy. This data was shown in Figure 2.6 in Chapter 2. Projection sideband cooling leads to ground-state occupation probabilities of 0.949(7), 0.954(6) and 0.985(1) in the x , y and z directions, respectively, which implies 89% occupation of the 3D vibrational ground state. The state is not thermal, but population is overwhelmingly in the lowest three vibrational states. Say, the normalized sidebands for the three lattice directions are given by X_i , Y_i , and Z_i , where $i = \Delta\nu$ is the order of the sideband, with ν being the

vibrational quantum number and $X_0 = Y_0 = Z_0 = 1$. The probability of finding an atom in the n^{th} vibrational level of the X lattice is given by :

$$p_X(n) = X_n - X_{n+1}, \quad (4.4)$$

with similar expressions for Y and Z represented by $p_Y(n)$ and $p_Z(n)$. The probability of finding an atom in the vibrational state (n_1, n_2, n_3) is then given by:

$$p(n_1, n_2, n_3) = p_X(n_1)p_Y(n_2)p_Z(n_3) \quad (4.5)$$

The vibrational entropy is then simply given by the Boltzmann formula :

$$S_v = -k_B \sum_{n_1, n_2, n_3} p(n_1, n_2, n_3) \ln(p(n_1, n_2, n_3)), \quad (4.6)$$

where k_B is the Boltzmann constant. For our final state, we calculated the vibrational entropy to be $0.59k_B$.

Sorting the atoms decreases the configurational entropy of the system, which is the entropy associated with the distribution of atoms in the lattice. It can be shown that for a lattice with filling fraction \bar{n} , the configurational entropy is given by [109]:

$$S_c = \frac{1}{\bar{n}} \left(\bar{n} \ln \left(\frac{1}{\bar{n}} \right) + (1 - \bar{n}) \ln \left(\frac{1}{1 - \bar{n}} \right) \right) k_B \quad (4.7)$$

The configurational entropy for sorting implementations for $5 \times 5 \times 2$ and $4 \times 4 \times 3$ target sub-lattices as a function of number of sorts is shown in Figure 4.10 by the solid blue lines. Initially, the configurational entropy is significantly higher than vibrational entropy of cold atoms (shown by the dashed blue line), and thus it is the primary contributor to the total entropy. Sorting reduces the configurational entropy by a factor of 8 to $0.16k_B$ and the total entropy by a factor of 2.44 to $0.75k_B$.

Our experiment increases the entropy of the outside world, which is carried away by the scattered imaging photons, while observing the initial positions of the atoms. In that process, the configurational entropy of the system goes to zero, as we have complete information about the site occupancy of the lattice. The stored occupancy information is then used as a guide to the execution of reversible operations that leave the system in a manifestly low-entropy state, making our

experiment conceptually like a Maxwell’s demon. While the sorting of atoms may seem no different from sorting checkers on a board, the difference here is that most of the initial entropy of our system is in the initial configurational disorder, so that by measuring and sorting we considerably reduce the total system entropy. Maxwell’s demon collected information and acted on one particle at a time. In contrast, our demon is omniscient and obtains an occupancy map of the whole system as the first step, so that it can map out a plan to act on all the particles in parallel.

Unlike Maxwell’s demon, our experiment doesn’t create a temperature gradient from which work could be extracted by driving a heat engine. On the other hand, it might be possible to create a much colder gas with sorted cold atoms by carrying out the approximately inverse process of a superfluid to Mott insulator transition. It was shown in reference [109] that sorted atoms in an optical lattice with sufficiently low entropy can be mapped to a Bose Einstein Condensate by adiabatically transforming the lattice to a 3D box potential or a gaussian trap. For the 3D box potential, which is more challenging than a gaussian trap in entropy requirements, the entropy per particle threshold for such a transition is $1.24k_B$, which is comfortably surpassed in our experiment. Although the adiabatic time scale for such an experiment is too large in our optical lattice, it would be possible if this experiment were to be done in accordion lattices, which implement variable lattice spacings by dynamically changing the angle between the lattice beams [110]. That would present a third path to quantum degeneracy of cold atoms, joining many evaporative-cooling experiments and one laser-cooling experiment [111–113].

4.4.2 Errors and scaling

To study the effects of various errors in a motion step on the sorting scheme, as well as the scalability of sorting, we implemented monte-carlo simulations for systems of various sizes³. There are three sources of error of importance. Errors in site selective state flipping can lead an atom not moving from its site at all, resulting in failure of all subsequent motion steps involving that atom. Another adjacent atom in the direction of motion that was also moved can end up at the same site,

³These simulations also served as a test bed for algorithms before plugging them into the experiment.

eventually leading to double loss. As mentioned before, spontaneous emission of the lattice photon can lead to both hopping and loss for both motion and stationary atoms in a motion step. Atom loss from the background gas collisions sets an ultimate limit on the highest obtainable filling fraction. For our experiment, the background loss is 2.2% during the first sort (including time during final imaging and projection cooling) and 1.1% for all subsequent sorts, which are shorter by a factor of 2 as they lack projection cooling. This puts an upper limit of 98.9% on the filling fraction.

Monte Carlo simulations of this sorting algorithm start with a randomly half-filled 3D array. Errors are probabilistically applied at each motion step and atom loss is considered after the completion of a sort. We calculate a separate motion fidelity for each internal state as the average of the measured fidelities in the three directions. One thousand simulations were run for various lattice dimensions and various target sub-lattices. The simulations predict an average filling fraction of about 0.97 for $5 \times 5 \times 2$ and $4 \times 4 \times 3$ target sub-lattices after three sorts, in agreement with our measured filling factor to within the uncertainty associated with our measured errors. It is clear that a steady state is reached after three sorts, and this would eventually start decaying due to atom loss.

Insight into the effects of different kinds error on the sorting can be obtained by looking at the number of motion steps as well as the number of state flips in each sort. The number of state flips drops sharply to 10 from the second sort onwards, while the number of motion steps stays roughly constant. This is because there are only a few errors after the first sort and atoms often need to be moved from far away sections of the lattice to fill them. Moreover, state flip errors affect only a few atoms in each motion step, while spontaneous emission errors affect all the atoms. Thus, the final filling fraction obtained after multiple sorts depends much more strongly on errors due to spontaneous emission than those due to imperfect addressing. A factor of 5 increase in the addressing error, keeping all the other errors at the current values in the experiment, barely affects the filling fraction after three sorts, decreasing it to 95%. On the other hand, a factor of 5 increase in the spontaneous emission error reduces the filling fraction to $\sim 77\%$ after 3 sorts. In fact, multiple sorts do not help if the error is so high. The highest filling fraction of $\sim 83\%$ is obtained after the first sort. Similarly, if the spontaneous emission were to be decreased by a factor of 5, the filling fraction would be 98.5% after 3

sorts, essentially only limited by vacuum lifetime.

We also studied the scaling of this technique by simulating the filling of $7 \times 7 \times 7$ and $10 \times 10 \times 4$ sub-lattices starting from a half-filled $10 \times 10 \times 10$ lattice. With our current rates of errors, the highest filling fractions are obtained after four sorts, peaking at $\sim 93\%$. If the spontaneous emission were to be reduced by a factor of three, say by detuning that lattice further by a factor of 3, the filling fraction can be pushed to $> 97\%$. Filling these lattices requires an average of ~ 10 motion steps and ~ 740 state flips in the first sort, in line with the expected scaling.

4.4.3 Improving the errors: throw and catch

To move the atoms in our implementation of sorting, we slowly translate the lattice such that the atoms can follow the lattice potential. While this adiabatic procedure doesn't cause much heating by itself, the increased light intensity seen by the atoms for long times significantly increases the lattice light spontaneous emission and associated errors. A polar opposite, highly non-adiabatic approach is also possible which can significantly decrease the spontaneous emission. The idea is to kick or *throw* the atoms in the lattice and turn the lattice off. After the atoms have moved the required amount due to the throw, they are *caught* in a shifted lattice potential. At first glance, it may seem like this "throw-and-catch" might significantly heat the atoms in the lattice. It turns out though, that it is possible to structure the throwing and catching steps in such a way that the original motional wave-function before the throw is perfectly recovered in the caught position, causing no heating.

The general strategy is shown in Figure 4.12 and proceeds as follows. The atom is initially trapped in a lattice, with its harmonic trap characterized by "distance units" or the characteristic ground state wave-function size, $\alpha_1 = \sqrt{\frac{\hbar}{m\omega_1}}$, where ω_1 is the trap frequency. To throw an atom, the lattice is shifted suddenly (compared to the trap frequency) by some distance x_0 (by suddenly rotating the polarization by corresponding amount). In general the lattice depth is different after the polarization is rotated, and the trap at this stage can be characterized by distance units, α_2 . The atom accelerates towards and reaches the trap minima in a quarter cycle, at which point the lattice is turned off and the atom is released. The atom's wave-function translates and evolves in free space for some time, t_1 , at which point the lattice is turned on again with its minima at location, x_1 , compared

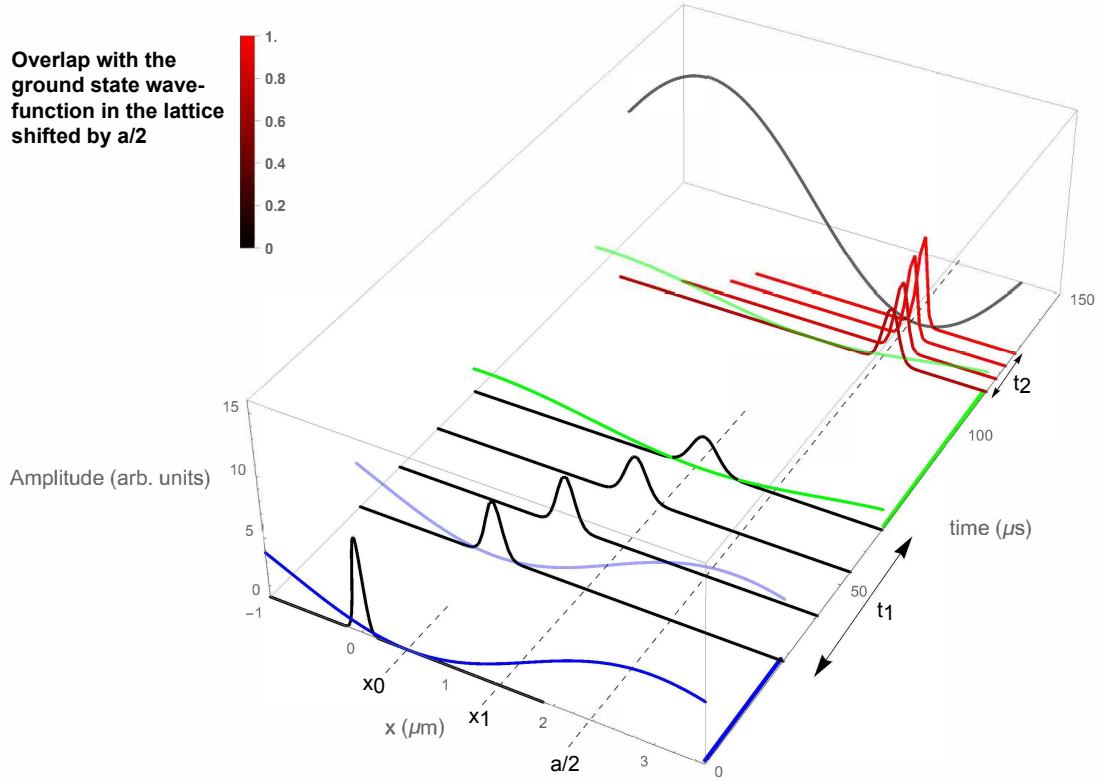


Figure 4.12. Throw and catch. The Figure shows the evolution of the probability distribution of the ground state atomic wave-function in a throw and catch scheme for moving atoms. The initial position of the wave-function is at $x = 0$. At time $t = 0$ the trapping lattice is suddenly shifted by distance x_0 by rotating the polarization (blue curve). This shift also decreases the lattice depth in accordance with Figure 4.2. The harmonic trap due to this lattice is parametrized by α_2 as described in the text. The wave-function evolves for a quarter cycle in this trap, denoted by blue line on the time axis. The atom is released by turning off the lattice and its wave-function is allowed to evolve for time t_1 , after which it reaches position x_1 . The lattice is turned on again with its minimum at x_1 (green curve), now parameterized by α_3 . During the quarter cycle evolution in this trap (green line on the time axis), the atom climbs the potential hill and comes to a stop at half the lattice spacing ($a/2$) away from its initial position. The lattice is turned off again and the atomic wave-function is allowed to evolve for time t_2 in free space. After time t_2 , the wave-function has evolved to the same as initial wave-function shifted by $a/2$. The color of the curve denotes the overlap-squared with the shifted ground-state wave-function. The original lattice, shifted by half the lattice spacing (gray curve), can now be turned on to re-trap the atom.

to the initial position of the atom. The trap here is characterized by α_3 . The wave-function climbs the potential hill in this trap for a quarter cycle and reaches it's final position. At this point, the lattice is turned off again and the wave-function is allowed to evolve again in free space for time, t_2 . The wave-function doesn't translate anymore because the average velocity is zero. If the parameters, t_1 , t_2 , x_1 , and the depth of the catching lattice (related to α_3) were chosen correctly, the atom's motional wave function would be just the original wave-function shifted by $a/2$ where a is the lattice spacing. Now the original lattice, shifted by $a/2$, can be turned on again (polarization rotated by π). If this turn on is sudden enough, the atomic wave-function would project to the same vibrational state. Now the state of the atom can be flipped and the preceding steps can be performed again with the opposite kick, leading to the atoms initially in the *motion* state translating by a lattice spacing and atoms initially in the *stationary* state returning to their original positions.

If the laser intensity is kept fixed, the α_i only depend on the rotation angle, θ , of the polarization. These can, in general, be controlled by controlling the laser intensity. We will typically be looking for solutions that start with an atom in the deepest possible lattice and only involve decreasing the laser intensity for subsequent stages, since that is typically a limiting factor. An analytical, albeit complicated solution to find these parameters is possible if the atom is initially in the ground vibrational state. Here I outline the steps to obtain the solution and present the final results. The full solution is calculated in a mathematica notebook in Appendix D.

We start with the original ground state atomic motional wave function centered at $x = 0$ in the trap characterized by α_1 . This wave-function is given by:

$$\psi_0(x) = \left(\frac{1}{\alpha_1^2 \pi}\right)^{\frac{1}{4}} \exp\left(-\frac{x^2}{2\alpha_1^2}\right) \quad (4.8)$$

The trap is now suddenly shifted by x_0 and is now characterized by α_2 . To find the wave-function of the atom released from this trap after a quarter cycle evolution, $\psi_{release}$, we will use the quarter cycle propagator of the shifted harmonic trap given by :

$$K_q(x, x') = \frac{\exp\left(-\frac{i(x-x_0)(x'-x_0)}{\alpha_2^2}\right)}{\sqrt{2\pi i\alpha_2}} \quad (4.9)$$

This gives,

$$\psi_{release}(x) = \int_{-\infty}^{\infty} K_q(x, x')\psi_0(x')dx' \quad (4.10)$$

This wave-function evolves in free space for time t_1 and the wave function after such a free space evolution, ψ_{free} , is given by:

$$\psi_{free}(x) = \mathcal{F}^{-1}\left(\exp\left(\frac{-i\hbar k^2 t_1}{2m}\right)\mathcal{F}, (\psi_{release}(x))\right) \quad (4.11)$$

where m is the mass of the atom and \mathcal{F} and \mathcal{F}^{-1} represent Fourier and inverse Fourier transforms, respectively. The wave-function is then caught in the lattice characterized by α_3 , with a trap minimum at x_1 . The wave-function after a quarter cycle evolution in this trap, ψ_{caught} , is obtained by using equations similar to (4.9) and (4.10), starting with ψ_{free} . This trap is then turned off and ψ_{caught} is evolved again in free space for time t_2 . The final wave-function is of the form:

$$\psi_{final}(x) = A \exp(i(c_0 + c_1x + c_2x^2)) \exp(-c_3(x - x_f)^2), \quad (4.12)$$

where c_i and x_f depend on various evolution times and lattice depths considered until now. For this wave-function to match the original wave-function, only shifted by some amount, the imaginary part of the argument of the exponential should be a constant, i.e. $c_1 = c_2 = 0$. Another condition is given by $c_3 = \frac{1}{\alpha_1^2}$. A final condition is given by ensuring that the total distance travelled by the atom be equal to half the lattice spacing, i.e. $x_f = a/2$. Simultaneous solution of these conditions yields the following four equations:

$$\boxed{t_2 = \frac{\alpha_1^4 \alpha_3^4 t_1}{\alpha_2^8 + 4\alpha_1^4 b^2 t_1^2}} \quad (4.13)$$

$$\boxed{x_1 = \frac{x_0 (\alpha_2^2 + 2bt_1)}{\alpha_2^2}} \quad (4.14)$$

$$\boxed{\alpha_3 = \frac{\sqrt[4]{\alpha_2^8 + 4\alpha_1^4 b^2 t_1^2}}{\alpha_2}} \quad (4.15)$$

$$\boxed{x_0 \left(\frac{\alpha_3^2 + 2bt_1}{\alpha_2^2} + 1 \right) = \frac{a}{2}} \quad (4.16)$$

Here $b = \frac{\hbar}{2m}$. Here the second equation (4.14) simply says that the catching lattice should be turned on with its minimum at the atom position at time t_1 after the release.

This is an overdetermined system with many solutions, and we will proceed by fixing α_1 , α_2 , and x_0 to experimentally reasonable values to obtain a particular solution. We will start with atoms in the deepest lattice available to us ($\alpha_1 = 0.071 \mu m$), at maximum laser intensity, and shift it by $x_0 = 0.5 \mu m$. This translation also decreases the trap depth by a factor of ~ 3 , giving $\alpha_2 = 0.092 \mu m$. Solving eqns (4.13 - 4.16) yields the values of other parameters, $\alpha_3 = 0.118 \mu m$, $t_1 = 39 \mu s$, $t_2 = 14 \mu s$, and $x_1 = 1.61 \mu m$. At maximum laser intensity, the lattice depth at $x_1 = 1.61 \mu m$ is larger by a factor of two than needed for this value of α_3 . So, the sequence requires the lattice laser power to be decreased to a half at the catch. The quarter cycle evolution time is $\sim 28 \mu s$ in the α_2 lattice and $\sim 46 \mu s$ in the α_3 lattice. This gives the total time for half a motion step to be $\sim 127 \mu s$. The lattice spacing here is taken to be $4.86 \mu m$. This is the solution shown in Figure 4.12.

In addition to picking experimentally reasonable values, one should try to find a solution such that the wave-function doesn't expand so much in free evolution that the harmonic trap approximation is violated. The atom shouldn't be caught too far from its final position for the same reason. The fact that this solution causes no vibrational excitations has also been verified numerically. Although, the solution was constructed by only considering the ground vibrational states, it has also been numerically verified to work for higher vibrational states ($\nu = 1$ and 2). It is not obvious why this should be the case, but not entirely surprising since the whole time evolution can be thought of as an action of a unitary operator, and

the simplest possibility for that unitary operator is just the translation operator as suggested by its action on the ground state.

Calculation shows that this scheme can reduce the spontaneous emission rate per motion step to $< 2 \times 10^{-4}$, improving the errors in motion step by a factor of > 30 . This would allow us to sort 400 atoms in a half-filled $10 \times 10 \times 10$ lattice to a $\sim 99\%$ filling fraction, with the current rates of other errors, limited almost entirely by the background loss rate. The implementation of this scheme requires sudden translation of the lattice potential (compared to the trap frequency) as well as fast turn on and turn-off of the lattice laser power. This requires a high voltage amplifier for the Pockels cells with a large slew rate to quickly rotate the polarization of the lattice beam. The amplifiers currently in use have a slew rate of $20 V/\mu s$, which is too slow for this throw and catch scheme. We plan to replace these amplifiers with the TREK PD05034-L, which has a maximum slew rate of $1000 V/\mu s$ and would allow us to shift the lattice by $0.5 \mu m$ in $\sim 2 \mu s$.

4.5 Conclusion and outlook

The experiment described in this chapter sort individual atoms in our randomly half-filled 3D optical lattice. In addition to being a realization of Maxwell's famous though experiment, it prepares a favorable initial state for a neutral-atom quantum computer with one atom at nearly every 3D lattice site, each cooled near its vibrational ground state. The cold array minimizes many potential errors in Rydberg-gate-based quantum computations. The 3D optical lattice allows entanglement with many near neighbors and provides favorable scaling, often as $N^{1/3}$ or $N^{2/3}$, to minimize computation time and the laser power requirements. If we can further improve cooling, for instance by temporarily transferring the atoms to a lattice with smaller detuning, where the atoms are trapped more deeply in the Lamb-Dicke limit, we might be able to create large-scale entanglement through cold-collision gates [114, 115] and thus ultimately implement one-way quantum computation [81, 82]⁴. Our sorted array could also be used for a variety of Rydberg-based quantum simulations with different geometries, dimensionalities and anisotropy of interactions. For instance, the simulations of Ising-like Hamiltonians

⁴See Chapter 6

that have recently been implemented in one- and two-dimensional (2D) tweezer and microtrap arrays [59,60] could be extended to three dimensions in our optical lattice. Our demonstrated coherent site-selective control allows the possible implementation of a universal quantum simulator, which might be used to implement Kitaev's toric code in 2D sublattices and in 2D or 3D lattice gauge theories [116].

Chapter 5 |

Lossless state detection

In this chapter, we extend our state-dependent motion of atoms to implement a lossless state detection of the atomic states in our 3D array. The state dependent lattice coherently separates the spatial wave-function for two internal states. Imaging the atoms in a new, short-spacing lattice, collapses the wave-function to a particular location. The measured position of an atom is related to the internal state of the atom before measurement. This measurement is reminiscent of the Stern-Gerlach separation [117], except here the wave-functions are separated using state dependent light traps instead of magnetic fields. We demonstrate a measurement fidelity of 0.9994 for ~ 160 qubits. The fidelity is essentially independent of the number of qubits measured. This fidelity is ~ 20 times better than any other lossless detection protocol in neutral atom arrays [118, 119]. To our knowledge, it also significantly exceeds the demonstrated measurement fidelity of other arrays with more than four qubits, including those with ion and superconducting qubits [27, 29, 120, 121]. Moreover, our measurement causes no observable loss. After the measurement we can replace any atoms lost due to background gas collisions and re-initialize the 3D qubit array.

5.1 Introduction

Gate based quantum computation requires a high fidelity measurement of qubit states in its final stage [74]. High fidelity read-out is also important in one way quantum computation, which consists of repeated local measurements in an entangled system to execute quantum gates [81]. Qubit measurement is even more

critical in error correction, where repeated measurements on "ancilla" qubits are used to feedback and correct the errors in a quantum computation [30,31,80,122,123]. In a typical quantum error correction scheme, the measurements are not as common as entangling gates. Nonetheless, the requirements on measurement fidelity are still stringent. For example, reference [124] sets a measurement error limit of 10^{-3} for a 300 qubit quantum computer with 10^9 logical operations.

A standard technique for high fidelity state detection in neutral atom qubits consists of pushing atoms in one of the qubit states out of the trap with a resonant "clearing" beam. The remaining atoms are then imaged and taken to be in the other qubit state. For instance, in our experiment, state detection consisted of shining a laser beam resonant with the $6S_{1/2}, F = 4$ to $6P_{3/2}, F' = 5$ transition at saturation intensity, which pushed out all the atoms in the $F = 4$ ground state in 2 ms . The detection fidelity of the remaining atoms is > 0.9997 for our experiment, limited by the occupancy determination errors. This technique suffers from two major drawbacks from a quantum computation perspective. First, atoms in the ejected qubit state are indistinguishable from loss due to background gas collisions. Thus the loss rate sets the ultimate error rate in measurement. Second, this technique leads to ejection of half the atoms in the lattice on average. Error correction would require a major chunk of the lattice to be re-filled and re-initialized with atoms, a significantly harder prospect than filling a few vacancies due to background gas collisions. Another technique to detect states of atoms relies on photoionization of one of the qubit states. State detection of a single ^{87}Rb atom in a dipole trap was achieved by two photon ionization of the $F = 2$ state and coincidence counting of the Rb^+ ion and electron on channel electron multipliers [125]. Although this technique allows for sub micro-second measurement time, the fidelity was limited to 98%. One could potentially distinguish loss from state detection in this method by sequentially photoionizing the qubit states or fluorescence imaging the non-photoionized state, but this technique would still suffer from the loss of the atoms (due to ionization).

Other approaches have been developed to achieve lossless state detection of neutral atom qubits [118, 119, 126, 127]. The approaches are typically based on resonance fluorescence on a closed, cycling transition, for example $6S_{1/2}, F = 4, m_F = 4$ to $6P_{3/2}, F' = 5, m_F = 5$ in Cesium, coupled by σ^+ polarized light. Such a scheme relies on existence of "bright" and "dark" states, where the dark states

are completely decoupled from any excitation. For example in Cesium, the qubit states could be mapped to $|4, -4\rangle$ and $|3, -3\rangle$ and they could be made to fluoresce on the cycling transition in turn. The fidelity in such a scheme is typically limited by the number of photons that can be scattered without heating the atoms out of the trap and the polarization imperfections in the probe. Experiments with small 1D ensembles of Rubidium atoms have led to a 0.987 fidelity and 2% heating loss in reference [119], and 0.98 fidelity and 1% heating loss in reference [118]. The photon collection efficiency can be improved by trapping the atoms in a high finesse cavity, although scaling this technique to large arrays of atoms is a significant challenge. A 0.994 state detection fidelity was achieved for a single Rubidium atom weakly coupled to a cavity [128]. Another experiment with single Rubidium atom strongly coupled to the cavity achieved a detection fidelity of 0.9992 [129]. An experiment in a quantum gas microscope separated atoms based on internal state by using a magnetic field gradient [130]. The separated atoms were then trapped in adjacent sites in a super-lattice potential and imaged again, leading to a 0.98 fidelity for ~ 50 atoms, limited by lattice phase fluctuations.

Our experiment also spatially separates the wave-functions for two internal states, but by using a state dependent optical lattice instead of a magnetic field. Since the internal states are mapped to different positions, the atoms can be cooled by the imaging light, avoiding the detection fidelity and heating trade-off in the other fluorescence experiments. Our state detection is illustrated in Figure 5.1. We start with atoms in our 3D lattice, denoted by XYZ , with a $\sim 4.9 \mu m$ spacing in all directions. For the detection, the qubit is mapped to field sensitive states for which a state dependent optical lattice is possible, say $|4, -4\rangle$ and $|3, -3\rangle$. We make the X lattice state dependent by rotating the polarization of one its generating laser beams. The potentials for the two states move in opposite directions and coherently split the atomic wave-function into two parts. These changes are adiabatic so that the atoms can follow the potential without heating out of the lattice, with the two states eventually separating by half the lattice spacing, enough to be resolved in our imaging system with high fidelity. At this stage another lattice with a 10 times smaller spacing, called X_s , is turned on in the x direction. The X lattice is then turned off leaving the spatial wave-function for each internal state typically distributed over a two X_s sites since the X and X_s lattices are not commensurate. We then image the atoms with the polarization gradient cooling (PGC) light, which

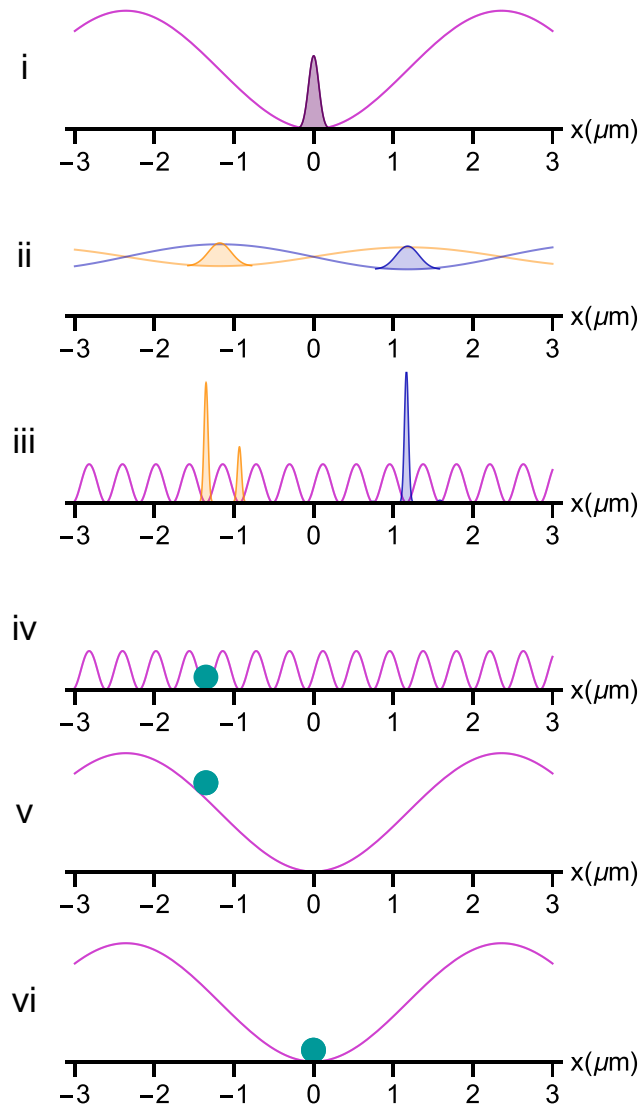


Figure 5.1. Steps in lossless state detection. The figure schematically shows the different stages in our state detection procedure. The potential energy curves have a consistent vertical scale and the atomic wavefunctions are normalized. i: Pre-measurement state. The purple curve denotes the state-independent potential energy in the x direction (X lattice). The purple shaded region represents the wavefunction of an atom in the vibrational ground state and an equal superposition of two states. ii: Displaced state-dependent lattices. The potential has been adiabatically transformed into two shallower state-dependent potentials, where the potential energy and wavefunctions for each of the two states are shown in orange and blue respectively. iii: Transfer atoms to X_s . For each internal state, the number of X_s lattice sites with significant wavefunction amplitude depends on the relative position of the X_s and the state-dependent X potential energy minima, which is not fixed in our experiment. iv: Image atoms in X_s . The wavefunction has been projected onto a single lattice site. The atom, now spread among many internal states and several vibrational levels, is denoted by the cyan ball. Its location is used for state assignment. v: Atom transferred back to X lattice. vi: Atom⁹¹ at the bottom of an X lattice site after one quarter oscillation period. At this point another image is taken.

projects the atom to a single X_sYZ lattice site, and measure its position. The change in the position of the atom compared to its initial position in the XYZ lattice constitutes a detection of the internal state. After the state detection, the X lattice is turned on again without any polarization rotation, the X_s lattice is turned off and the atoms are PGC cooled, re-trapping them in the XYZ lattice. The qubits can now be re-initialized for another quantum computation.

The outline of this chapter is as follows. Section 5.2 describes the physical setup of the short-wavelength X_s lattice as well as the sequence for loading atoms into this lattice from the original X lattice. The shorter spacing of the X_s lattice, the fact that it is not commensurate with X lattice and the lack of its phase stabilization lead to several complications in imaging and occupancy determination in this new lattice. Although the final data for state detection is shown starting in section 5.5, we use some of the same data to illustrate our resolution of these complications in sections 5.3 and 5.4, and show that we can bring the occupancy determination error below 10^{-4} . Section 5.4 then characterizes the performance of the state detection in the $|4, -4\rangle$ and $|3, -3\rangle$ states, and shows that the primary source of error is the spontaneous emission of the lattice light, same as in sorting atoms in the previous chapter. Section 5.5 shows how to adapt this technique for measuring atoms from the clock states, $|4, 0\rangle$ and $|3, 0\rangle$, our primary qubit states, even though atoms in these states don't experience state dependent lattices. We used our new state detection to measure the coherence time in the clock states and the result of this measurement was discussed in chapter 2.

5.2 The X_s lattice

The schematic of the apparatus for creating the X_s lattice is shown in Figure 5.2. The X_s lattice laser is derived from a 2W tapered amplifier seeded by an 838.9 nm source. The tapered amplifier output is detuned from every other lattice laser by at least tens of MHz by using an acousto-optic modulator, to avoid mutual interference. This finally results in a 700 mW beam which is retro-reflected in the x direction by an interference filter to form the X_s lattice, which has spacing of ~ 419.5 nm between its sites¹. The X_s beam is focused to 100 μ m at the

¹We have interference filters in the x and y directions, next to the cell, right after the addressing beam objectives. These are bandpass filters around 880 nm (the addressing beam wavelength) and




Figure 5.2. The X_s lattice is created by the beam shown in pink, retroreflected from an interference filter. The other lattice beams are also shown for reference. The X_s beam is focused to a waist of $\sim 100 \mu m$ at the interference filter, leading to a waist of $\sim 150 \mu m$ at the atoms for both the incoming and the retroreflected beams.

interference filter, leading to matched waists of $\sim 150 \mu m$ for both the incoming and the retro-reflected beams at the atoms. We measured the trap frequency in the X_s lattice by modulating the laser intensity by 10% at varying frequencies for 4 s. The modulation leads to parametric heating [131] of the atoms and would lead to an atom loss resonance at twice the trap frequency in a harmonic trap. We observe a loss resonance in our lattice at $196 kHz$ (Figure 5.3), which slightly underestimates twice the frequency spacing in our sinusoidal traps. Half this value, $98 kHz$, is roughly consistent with the expected lattice depth. To spatially split the wave-function we rotate the X-lattice beam polarization by $\sim \pi/2$ in $250 \mu s$, keeping the motion adiabatic and wait another $50 \mu s$ for the pockels cell voltages to settle. We then suddenly turn on the X_s lattice to a trap frequency of $43 kHz$, localizing the spatial wave-function for each internal state to within a couple of X_s lattice

highly reflective for the lattice wavelength. The primary purpose of these filters was to control the reflections of lattice beams while allowing the addressing beams to pass through. This is to prevent stray reflections from increasing the scattering rate at the atoms.




Figure 5.3. X_s trap frequency measurement. The Figure shows the result of a parametric heating experiment in the X_s lattice. The intensity of the X_s lattice beam is modulated by 10% at varying frequencies. This intensity modulation couples vibrational levels with $\Delta\nu = 2$ and heats the atoms, eventually leading to loss from the trap. The fraction of atoms remaining in the lattice, f , after 4 s of parametric heating is measured. The resonance in this signal at ~ 196 kHz gives twice the trap frequency.

sites. The trap frequency is increased from 43 kHz to 98 kHz in 78 μ s and the X lattice is turned off over the next 50 μ s. These parameters were experimentally optimized by minimizing the hopping in X_s , minimizing the displacement spread of the atoms in each internal state and minimizing the scattering to the other qubit state after preparing all the atoms in one of the states. Imaging uses the similar parameters for the PGC light powers and time as imaging in XYZ lattice. Since the X_s lattice is not phase stabilized and commensurate with the X lattice, the final position of the atoms can vary depending on the relative positions of the atomic wave-function and the nearest X_s minimum. Thus to ascertain the positions of the atoms, we fit the atom intensity to a gaussian with a floating center instead of the fixed center fitting which was used to determine the occupancy in XYZ lattice (Chapter 2).

We began by studying loading and imaging in the X_s lattice by loading the




Figure 5.4. In-situ displacement distribution. The displacement distribution obtained from floating center fits after the atoms are loaded into the X_s lattice with an optimized sequence and without state dependent translation. The x axis is the displacement from their initial location in the XYZ lattice. The gaussian fit (dashed curve) gives a root mean square width $167 \mu\text{m}$.

atoms "in-situ", i.e. from their original position without moving them in the state dependent X lattice. As a first step, the loading sequence here can simply be optimized by generating occupancy with our usual fixed center fits for both XYZ and X_sYZ lattices. If the atoms are too hot after loading into X_sYZ , they will hop in the X_s lattice leading to farther excursions from their initial positions and thus, they will not be captured by the fixed center fit. The fraction of atoms detected in the X_sYZ lattice at or near (near enough to be counted as atoms by the fixed center fit) their initial positions provides a useful signal for preliminary optimization and fine alignment of the X_s lattice. After this initial optimization step, we looked at the displacement distribution obtained from the floating center fit, a representative example of which is show in Figure 5.4. The sequence can be further optimized here by minimizing the root mean square width of this distribution as well as minimizing the non-gaussian tails. It should be noted though that this data is for atoms that were projection cooled, as they would be in a quantum computations,

and that the widths are significantly higher for atoms that were just polarization gradient cooled. The optimized sequence for in-situ loading will in general not be optimal for loading atoms from the state dependently displaced X lattice, which has a fourfold smaller depth and consequently, a larger characteristic size of the vibrational wave-functions, but it still provides a reasonable starting point.

5.3 Imaging in the X_sYZ lattice

Imaging of the atoms and the occupancy map generation in the X_sYZ lattice, although similar in principle to those in the XYZ lattice, are complicated by several factors. The imaging is done by scattering light from three pairs of PGC beams which create a standing wave pattern, but since the phases of the individual beams are not controlled, the intensity at any point in the 3D space is not stable. The PGC intensity at any given lattice site can vary from shot to shot and, to the extent that atoms are localized to less than half wavelength of the PGC standing wave, contribute to the broadening of the occupied peaks in the histograms shown in Figure 2.3. This effect is much more pronounced for atoms trapped in the X_sYZ lattice, where the wave function is much more localized in the x direction compared to when they were trapped in the XYZ lattice.

Figure 5.5 shows an example of images of five planes after the atoms were prepared in an equal superposition of $|4, -4\rangle$ and $|3, -3\rangle$, and loaded into the X_sYZ lattice after the state dependent translation as described in the previous section. One can clearly see the translations of individual atoms based on their internal states. Since the X_s lattice is not commensurate with the X lattice, when the atoms are loaded into it, they can be distributed over several lattice sites when the motional wave function of the translated atom overlaps considerably with a maximum of the X_s lattice. As will be described in sections 5.3.1 and 5.4, the non-regular distribution of atoms along with the previously mentioned floating center fit to get the atom positions leads to numerous complications, requiring us to significantly update our fitting and occupancy determination algorithm.




Figure 5.5. Exemplary state detection images. Images of five planes of the lattice before (top row, in the XYZ lattice) and after (bottom row, in the X_sYZ lattice) state detection. The atoms were prepared in an equal superposition of $|4, -4\rangle$ and $|3, -3\rangle$ after the initial imaging and then state detected using the previously described procedure. The atoms detected in $|4, -4\rangle$ ($|3, -3\rangle$) have moved to the left (right). The ellipses point out two potentially problematic cases for fitting atom intensities with a floating center gaussian. The red ellipse outlines two atoms which, being in different states, move towards each other and complicate the intensity signal in the region around their neighboring site. The yellow ellipse outlines two sites for which atoms were lost, but could be fit to a high amplitude because of the noise from the out of focus planes.

5.3.1 PGC phase scrambling

To address the site-to-site variation of the PGC intensity we modulate the phase of the PGC beams using liquid crystal variable retarders. The idea is to use the phase modulation to vary the intensity at each lattice site so the atoms experience an averaged and equalized PGC intensity over an imaging cycle. The electric field configuration for a 3D standing wave depends on two phases, as there are two variables required to specify a phase difference between every beam pair out of three beams. Thus, sampling every possible electric field configuration requires sampling the full 2D space generated by these phase differences. This is generally not possible, and even a dense sampling of this 2D space is prohibited by the slow speed of the liquid crystal retarders and the requirement that this modulation does not significantly affect the performance of PGC cooling.

We installed Meadowlark LV-100 retarders at the input port of the two of the three retro-reflected PGC beam pairs. The phase is varied from 0 to 2π with a 33 ms

period for one of the beams and from 0 to π with a 30 ms period for another. These two periods are incommensurate during the 150 ms imaging time per plane and lead to a broad sampling of the phases in the 2D space. The liquid crystal retarders require a 2 kHz ac square wave to drive them, where the peak to peak voltage of up to 10 V determines the retardation, with maximum retardation of $\sim 1.5\pi$ at 0 V . The voltage response of the liquid crystals is highly non-linear, which is further complicated by long (several ms) response times. To obtain a linear phase modulation with sharp turning points (triangle wave) we used customized driving waveforms which were empirically optimized using a Michelson interferometer setup and generated by a homemade Labview code. For these devices, the fastest we could scan over a 0 to π range was 30 ms , which is slow enough to not affect the PGC cooling of the atoms i.e., it must change more slowly than the characteristic PGC cooling time.

The incorporation of this phase scrambling results in a $\sim 20\%$ narrower width for the occupied site peaks in the histogram compared to those without phase scrambling. We have some evidence that this improvement is not limited by the amount of phase space that can be explored by this modulation as doubling the modulation range from π to 2π for the second beam does not affect the amount of improvement.

5.3.2 Fitting in the X_sYZ lattice

A simple approach for occupancy and state determination in the X_sYZ lattice would be to start with some data around each XYZ lattice site, which is then averaged over the y direction and fit to a 1D gaussian with a variable center. Whereas the fixed center fit leads to natural rejection of the noise from the out of focus planes due to the inherent astigmatism in the optical system, the floating center fit can easily chase this noise appearing between the lattice sites of the XYZ lattice. For a given site, it can also chase the signal from the neighboring occupied sites in both in and out-of-focus planes when the atom moves towards the site in question. This effect is especially pronounced when fitting empty sites, where there is no in-plane signal to guide the fit. It also complicates the fitting for the occupied sites as the intensity signal can deviate significantly from a gaussian due to the increased noise. Another complication arises when two neighboring atoms move towards each other

and intrude into each other's 11 pixel section ($1 \text{ pixel} = 0.52 \mu\text{m}$). Examples of these effects are shown in Figure 5.5.

To combat these issues we have significantly modified the fitting procedure for the X_sYZ lattice. We use the occupancy information from the initial image in the XYZ lattice and use the floating center fit only for the sites that were initially occupied by an atom. Around initially occupied sites, we integrate a 3-pixel-wide region in the y direction, search for the highest intensity among 8 pixels in the x direction, and select 7 out of these 8 pixels to perform a least-squares fit to a Gaussian with variable amplitude, offset, and center, and a fixed width. If the site is still occupied, we use the fit center to determine its state. This requires the maximum translation of the atoms to not be more than $\sim 3 - 4 \text{ pixels}$ to be adequately captured by the fits. Ideally, the states still need to be separated enough to be resolvable with a $< 10^{-4}$ error. This gives us a very tight window to operate and requires the hopping in the X_s lattice to be minimized by optimizing the X_s alignment and the loading procedure.

5.4 Final occupancy map generation in the X_sYZ lattice : a third imaging

After the previously described fitting procedure in the X_sYZ lattice, it is possible to follow our usual procedure for occupancy determination, but as we will show in this section, that procedure is not good enough to bring the occupancy determination errors below the 10^{-4} level. Usually, we would start by creating a preliminary occupancy map based on some threshold. For every site that is deemed to be occupied, this occupancy map also contains a preliminary state assignment based on the location of the fitted center of the gaussian fit. We then divide the sites into *eight* subsets, same as the previously described four subsets but for each of the two states to be detected. These subsets are assigned different thresholds for occupancy determination. This procedure is then iterated several times until the occupancy for each site converges.

Figure 5.6a show the 0_0 subset histogram for atoms that were prepared in $|4, -4\rangle$, translated and then loaded into the X_sYZ lattice. It is evident that all the previously described effects lead to a significant broadening of the histogram peaks

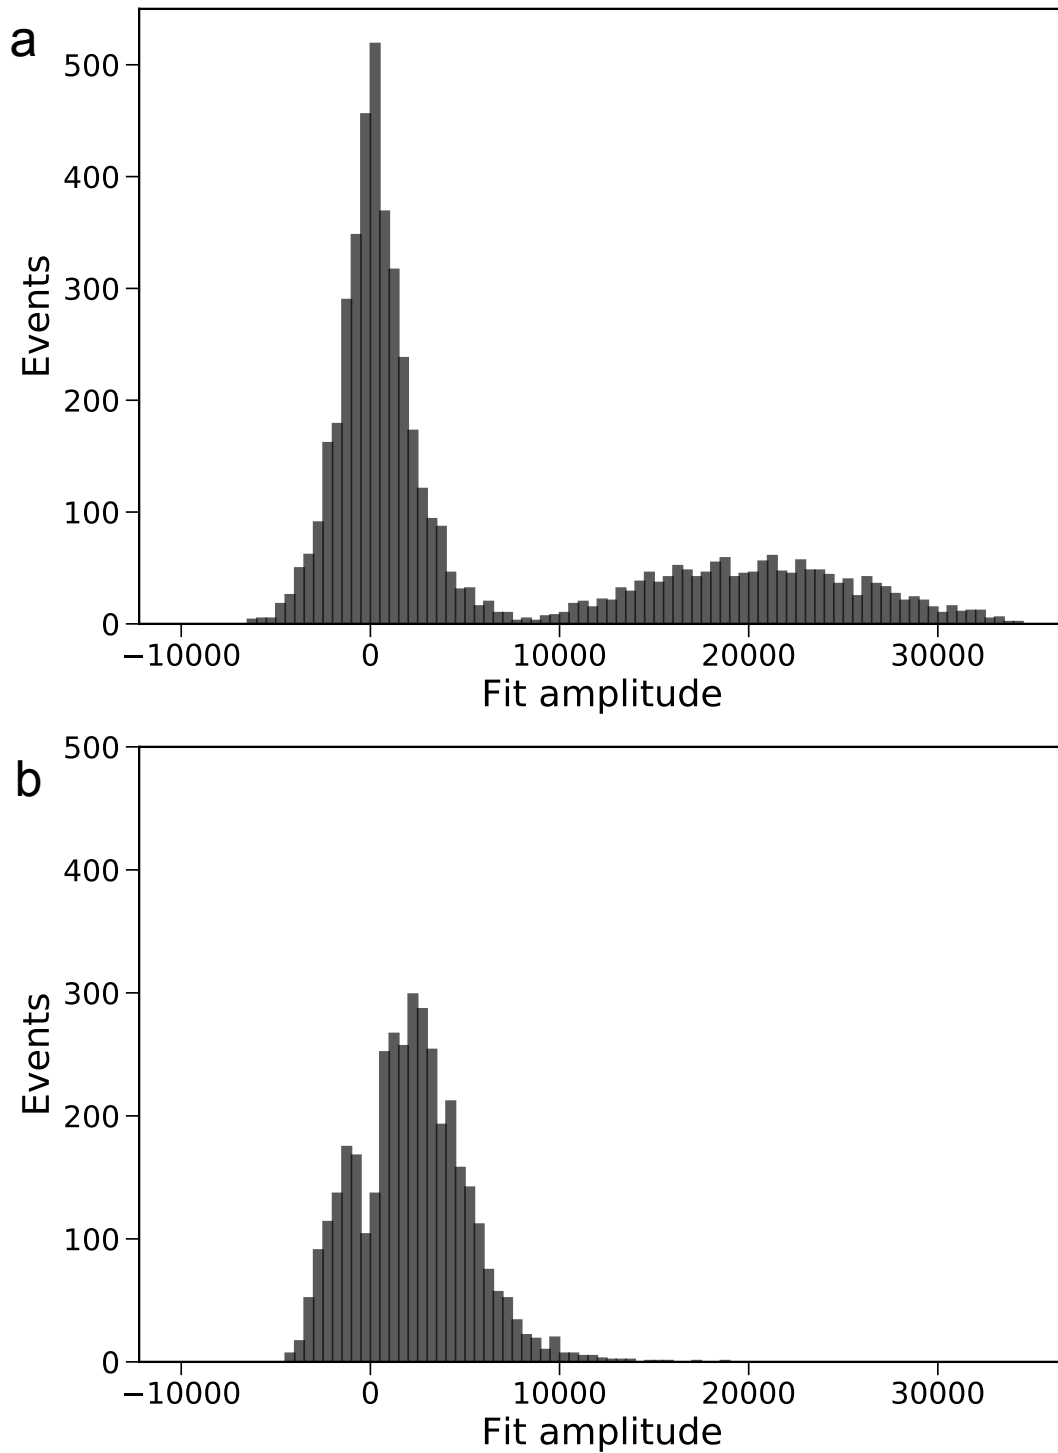


Figure 5.6. Histograms with the updated fitting procedure. (a) shows the histogram for the gaussian fit amplitues in the X_sYZ images for the 0_0 case with the fixed center fit for empty sites as determined by the initial images and a floating center fit for the sites that were initially occupied. (b) shows the histogram for the initially empty sites fitted with a floating center fit. This shows how errors are magnified for the sites for which the atom was lost between initial and state detection imaging. The amplitude units are arbitrary.

for both the occupied and un-occupied sites. The fixed center fit for the empty sites significantly reduces the variance for the unoccupied site peak compared to if those sites were fit with the floating center fit (Figure 5.6b). Though, it still has a longer tail than the similar distribution in the XYZ lattice (Figure 2.3a) due to the increased noise at the empty sites from the atoms moving to between lattice sites. The occupied site fit is broader because of the same reasons, in addition to the higher variance inherent in the floating center fit.

Even after this updated fitting procedure, there is still a lack of a clear gap between the unoccupied and occupied site peaks in Figure 5.6a. Still, the larger problem here is the 2% of sites for which the atoms were lost between initial imaging and state detection imaging in the X_sYZ lattice (due to background gas collisions) and were thus fit with a floating center fit in spite of being empty. As can be seen in Figure 5.6b, empty sites fitted with a floating center fit can have significantly high fit amplitudes and a much higher probability of being counted as occupied, leading to a random state assignment based on the local noise. To address these problems, we reload the atoms into the XYZ lattice and take another set of images which help us to verify which sites had an atom in the state detection images.

5.4.1 Re-loading atoms into the XYZ lattice

The re-loading of the atoms in the XYZ lattice from the X_sYZ lattice is schematically depicted in figures 1-v and 1-vi. After the state detection, we turn on the X lattice with a 2ms ramp without any state dependent translation. The atoms are still pinned to their positions by the shorter spacing X_s lattice. Next we ramp down the X_s lattice from a trap frequency of 98kHz to 43kHz in 5ms and then suddenly shut it off. This ensures that the atoms are kept cold until they are released from the X_s trap, typically half-way up in the X lattice potential around each site. The released atoms then fall to the site centers in the X lattice, gaining energy, in a quarter oscillation cycle ($\sim 17\mu\text{s}$). After a quarter cycle, the PGC beams are turned on to cool the atoms near the site centers, re-trapping them. The atoms can now be imaged again with the PGC light yielding a third set of images.

The re-loading sequence was optimized by minimizing the total observed loss between the initial imaging in the XYZ lattice and this third and final set of images. After the optimization the loss is consistent with the expected loss from background




Figure 5.7. Information from the third imaging. The data shown in Figure 5.6 is replotted with information from both the initial and the third set of images. The red bars result from floating center fitting of sites which are known to be empty based on the initial set of images (same as Figure 5.6b). The blue bars result from floating center fitting of sites which are known to be occupied based on the third set of images. Sites for which an atom was lost between initial and state detection imaging would be fit with a floating center gaussian and would have a distribution same as the red distribution. A choice of threshold that counts all the blue atoms also counts a fraction of sites for which the atoms were lost as occupied. This fraction is given by the overlap of the red and blue distributions (6.7%). Given 2% atom loss, this would amount to 0.13% of the sites being given a random state assignment based on the local noise. Therefore, following the updated logic (Figure 5.8) we pick the threshold high enough (dashed gray line) such that none of the empty sites are counted as occupied. Most (98.9%) of the 19% of the blue sites below the threshold are counted by the third set of images. The remaining, which make up 0.2% of the total number of atoms, are counted as the "pseudo-loss".

gas collisions in 3 s (2 s between the initial and the state detection imaging and 1 s between the state detection and final imaging), to within a measurement uncertainty of 0.1%. This also suggests that there is no observable loss due to the state detection itself. In addition to allowing for a third set of images in the original lattice, the re-loading is also necessary for us to be able to re-use the atoms after state detection.

5.4.2 Using information from the third set of images

This third set of images is no longer plagued by the issues in imaging the translated atoms in the X_sYZ lattice and an occupancy map can be generated for these images using exactly the same procedure as for the initial set of images. The occupancy in the third set of images provides additional information about atoms that were definitively there during the state-detection imaging. To understand the issues we are trying to resolve, it is useful to plot the amplitudes shown in Figure 5.6 with the information from the first and third set of images. This is shown in Figure 5.7, where the red peaks are floating center fit amplitudes for sites that were empty during the initial set of images and the blue peaks are floating center fit amplitudes for sites that were occupied during the final set of images. While it is clear that the sites that had an atom in the third set of images also had an atom during state detection, we cannot always distinguish between sites for which the atom was lost between the first set and state detection images (2% of the atoms) from those for which the atom was lost between the state detection images and the third set (1.1% of the atoms). The extent of this problem is illustrated by the overlap between the tails of the red and blue peaks. This overlap amounts to $\sim 6.7\%$ of the empty sites being fitted with an amplitude high enough to overlap with the lowest amplitudes in the blue distribution. Since this is only a problem for 2% of the sites that suffered an atom loss between the first set and the state detection², the problem we are trying to avoid relates to $0.13\% = 2\% \times 6.7\%$ of the atoms.

We can substantially minimize the impact of this effect on our measurement by using the logic summarized in the flowchart in Figure 5.8. If the atom is in the final image, it is counted as being there for the state detection images. If it is not in the final image, then it was lost either before the state detection or after the state detection. We assign a set of thresholds (one for each of the eight subsets) to the state detection histograms, placing them at the upper edge of the tail of the histogram of empty sites (the dashed vertical line in Figure 5.7). If the amplitude is above the threshold in the state detection image, we infer that there was an atom during state detection that was lost afterwards, and its position is counted in the state measurement. If it was below the threshold, we place it in the category of lost before state detection.

²Because the sites that were already empty were fit with a fixed center fit.

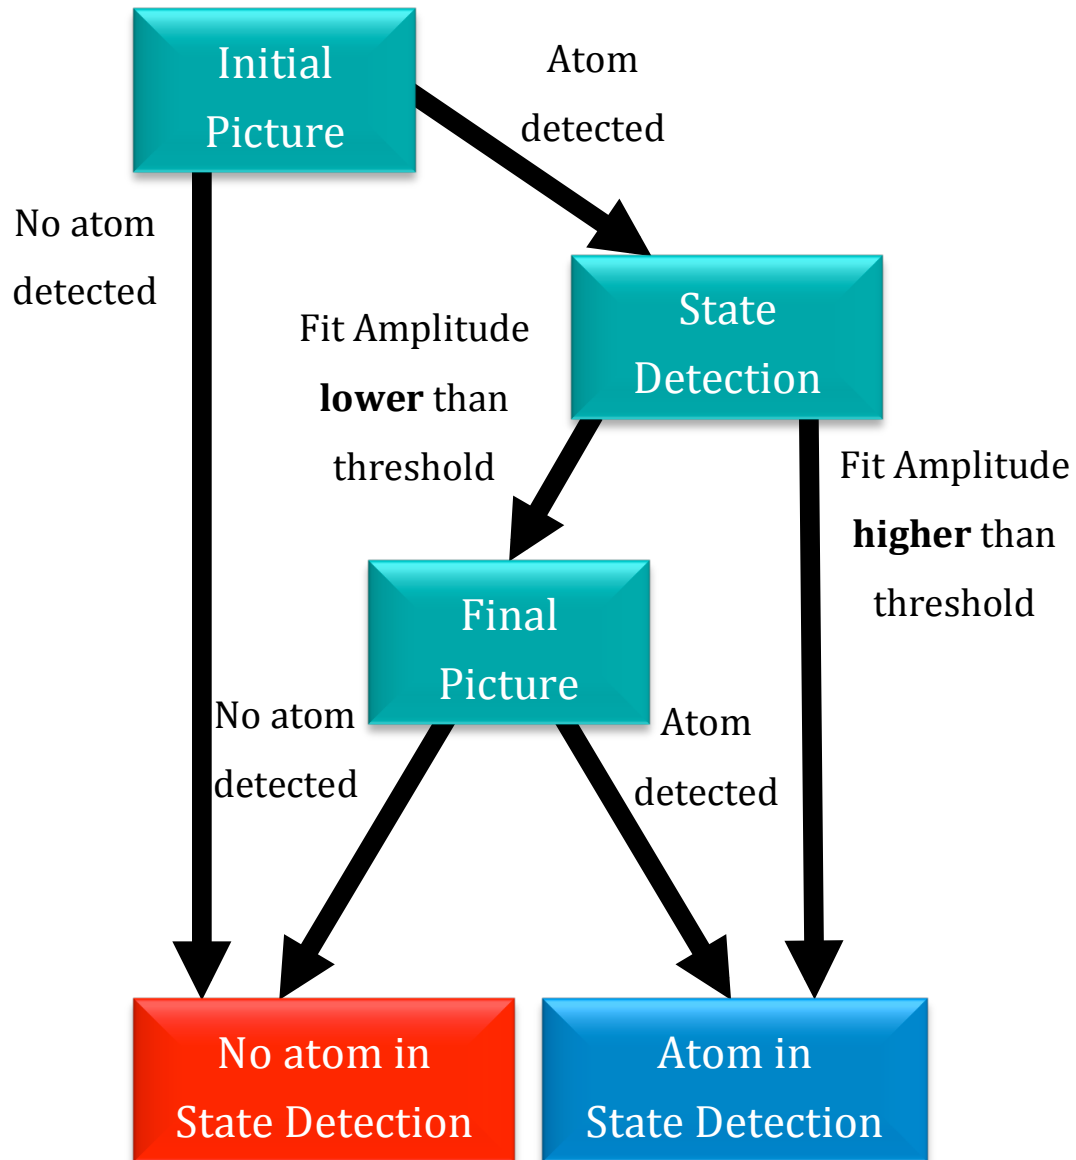


Figure 5.8. Occupancy assignment flow chart for state detection. There is a similar flow chart with different thresholds for each occupancy case. The end result is a contribution of $< 7 \times 10^{-5}$ to the state detection error, and a $\sim 10\%$ virtual increase in the number of atoms lost to background gas collisions, which amounts to $\sim 0.2\%$ of the atoms given the quality of our vacuum.

This logic ensures that the aforementioned 0.13% empty sites are always counted as unoccupied because of the high enough threshold. This advantage though, comes at the cost of counting some of the atoms that would have been there during the state detection, as lost. Approximately 19% of atoms that we know are present at state detection, are detected with an amplitude below the threshold (the blue bars to the left of the dashed vertical line in Figure 5.7). The vast majority of these, 98.9%, are still there when we make the final occupancy map in the third set of images, and they contribute to the state detection measurement. But because of background gas collisions, $\sim 1.1\%$ of them are lost between the state detection imaging and the third imaging. Thus, 1.1% of the 19% of the atoms, i.e. 0.2% of all the atoms, were lost between the state detection imaging and the third imaging *and* had too low a floating center fit amplitude to be treated as atoms in logic dictated by Figure 5.8. These atoms were thus counted as lost between the initial imaging and the state detection imaging.

This procedure is equivalent to there being a 10% worse background gas pressure (causing 2.2% instead of 2% loss between the initial imaging and the state detection), and indeed, it is simply a secondary consequence of background gas collisions. This “pseudo-loss” is not inherent to the detection scheme and if the loss due to background collisions is reduced, the pseudo-loss would be proportionally reduced. If there were no background loss, there would be no pseudo-loss, and no need for a third image either. Note also that were this state detection technique to be applied to only a sparse subset of atoms (as would be the case during error correction) or to 2D arrays, the red and blue distributions in Figure 5.7 would be disjoint as usual, despite the floating center fits, and there would be no need for a pseudo-loss. If X_s were made commensurate and phase stabilized, there would be no need for floating center fits, and there would also be no pseudo-loss.

If the objective is to yield a faithful state detection, the pseudo-loss is preferable to assigning a random state to 0.13% of the empty sites. After this logic is applied, the errors in the state detection happen when there are false positive occupancy assignments based on the first and third set of images. This, just as before, forces a floating center fit on an empty site for the the state detection images, leading to a state assignment based on the noise around the site. Given our choices of thresholds, we infer that this error occurs 3.6×10^{-5} of the time. There is an additional contribution of 2.9×10^{-5} from the far end of the red distribution

in Figure 5.7. Thus, the contribution to state detection error from occupancy misidentification is only 6.5×10^{-5} .

5.5 Detection of $|4, -4\rangle$ and $|3, -3\rangle$

We now turn our attention to the performance of the state detection itself. After optimizing the occupancy determination, we first executed state detection after preparing starting from the stretched states, $|4, -4\rangle$ and $|3, -3\rangle$. Figure 5.9 shows the state detection of the atoms that were prepared in only $|4, -4\rangle$ (a), $|3, -3\rangle$ (b), and an equal super-position of these states (c). For Figure 5.9a the atoms were state detected after optical pumping them to the $|4, -4\rangle$. For Figure 5.9b the atoms were transferred to from $|4, -4\rangle$ to $|3, -3\rangle$ using an AFP pulse after optical pumping, followed by a clearing beam to remove any $F = 4$ atoms and ensure a clean preparation. An equal super-position of these states was created using a blackman pulse to yield the data in Figure 5.9c. The data shown here is for the entire $9 \times 9 \times 5$ section of the lattice that we image, with 30 – 40% loading in each shot, totaling to about 8000 atoms for each case.

The data was analyzed with the occupancy determination logic described in the previous sections. In fact, the histograms presented in the previous sections correspond to the data set that yields Figure 5.9. The translation of the atoms relative to their initial positions indicated by the dashed yellow grid lines is clearly visible in the representative images of single planes on the right. The displacement distributions in the left most column are obtained from the floating center fit and quantify this translation. The typical uncertainty in the center obtained from the fit is $0.15 \mu m$. The Gaussian root mean square widths of the displacement distributions are $\sim 203 nm$, which is about half an X_s lattice spacing ($210 nm$).

There is a clear asymmetry in the translations of the two states, probably due to the imperfect polarization of the lattice beams. While the atoms in the $|4, -4\rangle$ state move by $\sim 1.2 \mu m$, the atoms in the $|3, -3\rangle$ state move only by $\sim 0.8 \mu m$. Thus, for the state assignment we choose the dividing line to be at $\sim -0.2 \mu m$ to minimize the average errors. The insets in the Figure 5.9 are the zoomed in images of the indicated sections of the histograms and show the fraction of atoms that were detected in the wrong state. From this data, we ascertain that $7 \binom{+4}{-3} \times 10^{-4}$ of the atoms nominally prepared in $|4, -4\rangle$ are measured to be in $|3, -3\rangle$ and $5 \binom{+3}{-2} \times 10^{-4}$

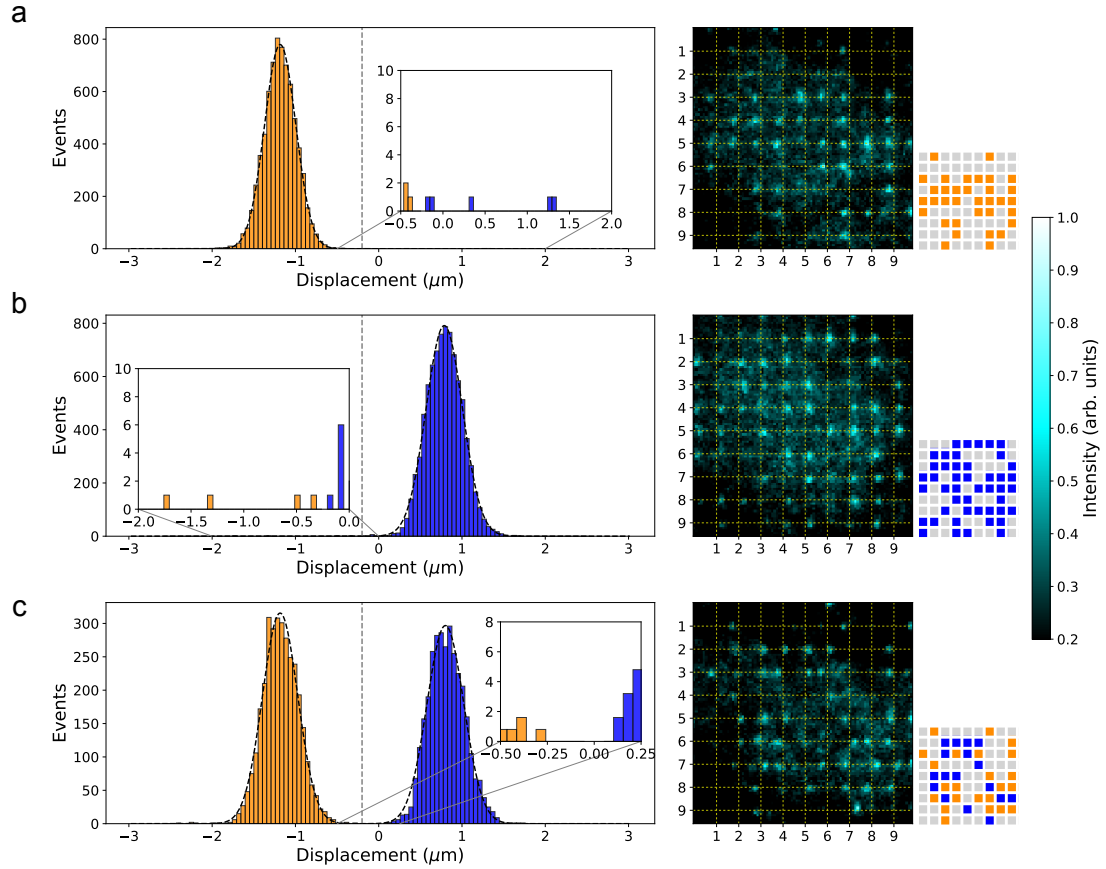


Figure 5.9. Displacement distributions and state assignment. (a) shows the displacement distribution during state detection for atoms prepared in $|4, -4\rangle$, (b) for atoms prepared in $|3, -3\rangle$, and (c) for atoms prepared in an equal superposition of these states. For each histogram we perform 50 implementations starting with a 30 – 40% randomly loaded $9 \times 9 \times 5$ lattice. The black dashed lines are Gaussian fits to the distributions. The vertical grey dashed line at $-0.2 \mu\text{m}$ sets the dividing line for state assignment. Atoms to the left (right) are ascribed to $|4, -4\rangle$ ($|3, -3\rangle$) and colored orange (blue). The state detection errors, those events that are nominally prepared in one state but are detected in the other state, are highlighted in the insets of a and b. An example image of one lattice plane is shown next to each histogram. The yellow grid marks the positions of the lattice sites in the XYZ lattice, and thus the possible initial positions of atoms. The occupancy map for each example image is shown in the associated square pattern. Orange or blue denotes an occupied site in one of the two states, while grey represents an empty site.

of the atoms nominally prepared in $|3, -3\rangle$ are measured to be in $|4, -4\rangle$ (see Fig. 5.9a and 5.9b insets), yielding an average state detection fidelity of $0.9994(3)$. Out of the 9 errors we see out of 16000 state detections, one seems to be due to a site occupation error, which seems consistent with occupancy errors described in the




Figure 5.10. Displacement asymmetry. The orange (blue) points show the measured displacement for $|4, -4\rangle$ ($|3, -3\rangle$) as a function of the rotation angle of the lattice beam polarization. The error bars are smaller than the size of the symbols. The dashed line is the calculation with perfect polarization of the lattice beams. The solid line is the calculation with a 1% polarization imperfection in the X lattice beam whose polarization is not rotated. It is noteworthy that the motion of $|3, -3\rangle$ is significantly faster near the center and limits how fast we can translate atoms without heating them.

previous section.

The observed asymmetry for the displacement peaks is significantly higher than that predicted by a theoretical calculation. We further studied this asymmetry in translations by measuring the displacement centers for both stretched states using a range of applied voltages on the pockels cells, as shown in Figure 5.10. If the polarization of the lattice light was perfect, we would expect the motion of atoms in the two states to follow the dashed lines, which have similar trajectories for the two states. The intensity of non-linearly polarized light for the X beams before the vacuum cell is smaller than 0.0007 of the total intensity. However, it is possible that the linear polarizations of the two lattice beams are not exactly perpendicular to

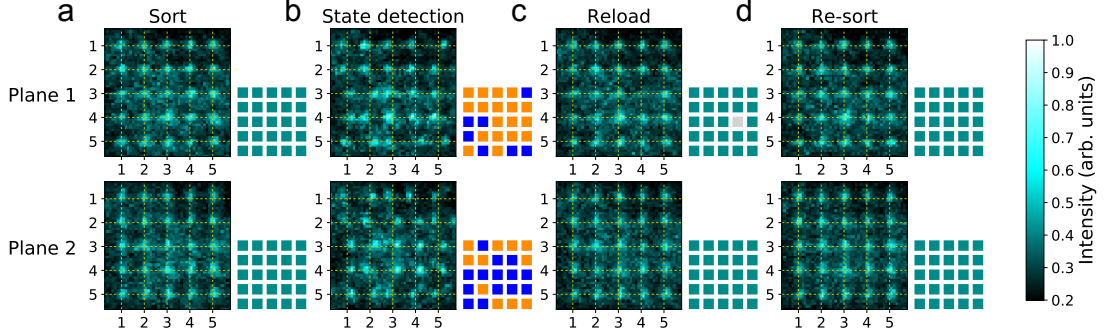


Figure 5.11. Re-initialization of the 3D qubit array. (a) Start with a $5 \times 5 \times 2$ sorted lattice. (b) shows state detection of an equal superposition of the two stretched states using our new method. (c) Atoms are reloaded into the XYZ lattice. An atom loss has occurred during this in Plane 1. (d) The atoms are re-sorted to replace the lost atom.

the propagation plane of the beams, and it is also possible that the cell birefringence compromises the polarization. To get an idea of how polarization imperfections might affect these displacements, we calculated the expected motion with a fixed 1% relative intensity of perpendicular $\pi/2$ out-of-phase linear polarization in the X beam whose polarization is not rotated (the solid lines). While the imperfect match to the data makes it clear that this is not the specific polarization flaw in X beams, the improved agreement with the observations support the idea that polarization imperfections underlie this discrepancy.

Another notable feature here is that this polarization imperfection makes the motion for $|3, -3\rangle$ a lot steeper near the half-wave voltage, limiting our ability to move atoms fast. When atoms spontaneously emit lattice light before they are captured in X_s , they can change states and move in the wrong direction. We calculate a scattering rate of 3×10^{-4} during the motion in X , which is consistent with this being the dominant source of error. If the X polarization were to be improved, we could move the atoms faster and thus reduce the scattering and the associated state detection error. We could also employ the “throw and catch” method described in chapter 4. For this purpose we would not need the elaborate catching scheme that doesn’t change the vibrational state. Here we could simply catch the atoms by turning on the X_s lattice some time after the throw. We estimate that this could reduce our error from spontaneous emission by a factor of five.

With essentially lossless detection, we can re-initialize our qubits after state

measurement. This capability is demonstrated in Figure 5.11. We first sort to get a perfectly filled $5 \times 5 \times 2$ pattern (Figure 5.11a). We then prepare an equal superposition of the two stretched states, execute state detection in X_sYZ lattice in (b), and re-load into the XYZ lattice in (c) using the previously described protocol. In this particular implementation, an atom was lost between state detection imaging and the third imaging after re-loading into the XYZ lattice. We re-sort, filling the vacancy with an atom from the reservoir region (within $5 \times 5 \times 5$ but outside $5 \times 5 \times 2$) in (d).

5.6 Detection from $|4, 0\rangle$ and $|3, 0\rangle$

As a final test of our lossless state detection, we adapted this technique to detect from the clock states, $|4, 0\rangle$ and $|3, 0\rangle$. While it is possible to use the stretched states as qubits, it is preferable to use the clock states because of their insensitivity to magnetic field noise. However, atoms in the clock states don't experience state dependent lattices and thus, don't move when the polarization is rotated. Thus, we use two AFP pulses to map the clock states to $|4, -2\rangle$ and $|3, -2\rangle$, similar to the mapping in the first kind of single qubit gate in chapter 3, except here we don't care about preserving the coherence. We perform the state detection by moving these two states now. In order to account for the shallower potentials, we increase the voltage ramping time from the $250 \mu s$ we use for the stretched states to $500 \mu s$. While it is preferable to move atoms in higher $|m_F|$ states because of deeper state dependent lattices, we are limited to transferring only to $|m_F| = 2$ states due to our lack of control over the microwave polarization. We rotated our microwave horn to equalize the rabi frequencies for the $|3, 0\rangle \leftrightarrow |4, -1\rangle$ and the $|4, 0\rangle \leftrightarrow |3, -1\rangle$ transitions. This means that the rabi frequency of the $|3, -1\rangle \leftrightarrow |4, -2\rangle$ is 0.7 times that of the $|4, -1\rangle \leftrightarrow |3, -2\rangle$ transition. Furthermore, the $|3, -2\rangle \leftrightarrow |4, -3\rangle$ has a Rabi frequency 0.29 times that of the $|4, -2\rangle \leftrightarrow |3, -3\rangle$ transition. Even though the AFP pulse is relatively insensitive to the rabi frequency, it is unlikely to ensure a mapping fidelity of better than 10^{-4} with such a large difference between rabi frequencies of two transitions being driven at the same time.

To characterize the performance of the state detection from the clock states, we performed the coherence time measurement that was described in Chapter 2. This kind of measurement allows us to study the state detection for a wide




Figure 5.12. State-selective detection from the clock states. (a) shows fringes from the spin-echo measurement sequence described in chapter 2. The evolution time is $480 \mu\text{s}$ and the phase of the final $\pi/2$ pulse is scanned. For state detection, atoms are transferred from $|4, 0\rangle$ to $|4, -2\rangle$ (orange points) and from $|3, 0\rangle$ to $|3, -2\rangle$ (blue points) using two AFP microwave pulses. At each point, the populations of both states are determined in the same state measurement. The data points at phase 2π are the same as those at phase 0 . The solid curves are fits to a sine function, and the fit amplitude gives the fringe contrast. (b) shows magnified regions of (a). At phases π and 0 , perfect state preparation and detection would lead to all atoms being in the same state. Error bars in (a) and (b) represent one standard deviation and are due to counting statistics. The error bars for the data points at phase π are smaller than the size of the symbols

range of superpositions in a systematic way. To recap, we started by preparing an equal superposition of the clock states using a $\pi/2$ pulse. This was followed by a dynamically decoupled sequence of π pulses to repeatedly spin echo the superposition. We then applied a final $\pi/2$ pulse with a variable phase. Finally, we measured the population in two clock states after using the aforementioned protocol. This yields an interference fringe for the populations in the two states as a function of the phase of the final $\pi/2$ pulse. The decay of the contrast of the interference fringe as a function of the time between the initial and the final $\pi/2$ pulses measures the coherence time.

While the coherence time is measured at a quarter lattice depth, we adiabatically raise the lattice to full power before the state detection. The full data for this measurement was shown in Figure 2.8. To ascertain the performance of state detection we focus on the shortest time ($480 \mu s$) data here. The interference fringe obtained here should essentially be perfect given the $12.6 s$ coherence time. This data has been re-plotted in Figure 5.12 for both the states along with the zoomed in sections near phases 0 and π . The fraction of atoms found in the wrong state at phase π is $0({}_{-0}^{+5}) \times 10^{-4}$, which is to say we found no errors in 2200 atom detections, consistent with as good a fidelity as we measured starting from the stretched states. The error at phase 0 is larger at $22({}_{-8}^{+12}) \times 10^{-4}$. We attribute the worse performance when measuring from the $|4, 0\rangle$ state to the factor of 0.7 rabi frequency difference in the final AFP mapping step. Some atoms are left behind in the $|3, -1\rangle$ state, from which they are detected to be in the wrong state. Using a dual-polarized horn antenna would enable the microwave polarization to be dynamically changed, which would allow the two rabi frequencies driven by each AFP pulse to be equalized and avoid this issue completely.

There are certain subtleties when measuring the atomic states using this method after a significant amount of decoherence. The scattering of the lattice light not only scrambles the phase of atoms in the qubit states, but also transfers atoms out of the qubit basis into other magnetic sublevels. Most of these will show up in one of the state detection peaks. Those that end up in the $|3, -1\rangle$ and $|4, -1\rangle$ levels are transferred back to the $|3, 0\rangle$ and $|4, 0\rangle$ states respectively by the first of the final two AFP pulses. These two states do not move in the otherwise state-dependent lattice, something we have checked by performing the state detection method without the final AFP pulses, which results in a single peak centered near $0 \mu m$. As

lattice scattering starts to accrue, the empty region between the two state-detection peaks in the displacement distributions (Figure 5.9c) starts to fill. These atoms, like all atoms that undergo spontaneous emission, degrade the contrast. Since they preferentially lie on the $F = 3$ side of the state dividing line at $-0.2 \mu m$, they lead to a uniform upward shift of the fringes (Figure 2.8), but this is inconsequential from the point of view of the coherence measurement as it doesn't affect the contrast of the fringes.

5.7 Conclusion

In summary, we have demonstrated a new state detection method, with no observable loss to within an uncertainty of $< 1 \times 10^{-3}$. Our method relies on coherently separating atomic states in state dependent lattices and trapping them in a shorter spacing lattice so they can be imaged without heating them out of the trap. The state detection error approaches the commonly accepted threshold of 10^{-4} for fault-tolerant quantum computation, and it already comfortably surpasses the thresholds of 10^{-3} for some surface codes [132, 133].

The concept of our state measurement method, where internal states are mapped onto atom positions, can be generalized to atoms in reconfigurable dipole trap arrays [38, 40, 106]. Each dipole trap can be formed by overlapping two traps with opposite circular polarizations. For state detection, the two traps can be spatially separated and if needed, dipole heating could be minimized by making the two traps linearly polarized before the atoms are imaged.

We have further demonstrated that we can re-initialize the array and replace lost atoms after state detection, which will ultimately be important for error correction. Error correction would also require selective detection of qubits. A possible approach is as follows. The atoms to be detected would be selectively transferred to larger $|mF\rangle$ states. During the motion, the clock state qubits will remain weakly trapped at their original position and transferring them to X_s will require matching the nodes of X lattice to those of the X_s lattice. It might also be possible to use phase contrast imaging [134] for fast detection of atoms with minimal heating. Preventing the undetected qubits from interacting with detection light will require shifting qubits out of resonance, transferring them to dark states or using a second atomic species for measurement [135].

Chapter 6 |

Future work

An outstanding challenge in neutral atom based quantum computation experiments has been the demonstration of high fidelity two-qubit gates to generate entanglement. They are a necessary milestone for universal quantum computation. This chapter describes our proposal to implement Rydberg excitation based targeted controlled phase gate in our 3D array of atoms. I also briefly touch upon the possibility of generating large scale entanglement in this system using cold controlled collisions [114] to yield a so called cluster state. It has been shown that cluster states along with site selective measurements can be used to implement one-way quantum computation [81, 82].

6.1 Introduction

Essentially, two approaches have been pursued to engineer interactions between neutral atoms. The first one involves exciting the atoms to a high n (principle quantum number) Rydberg state where the interaction between two atoms is dominated by the electric dipole-dipole interaction, which is many orders of magnitude stronger than the magnetic dipole-dipole interaction dominant in the ground state [135]. For instance, the interaction shifts for S states can be several tens of MHz in frequency units at a distance of $5\ \mu m$ for n values near 75 in Cesium. This is a direct result of the strong dipole transition matrix elements between nearby Rydberg states which leads to a large polarizability, scaling as n^7 . This also leads to a large electric field sensitivity for the Rydberg states, which on one hand is problematic for high fidelity excitation to a Rydberg state in the presence of field fluctuations, but on the other

hand, it can also be used to control the nature and strength of the interaction. The interaction typically also has a dependence on the angle between the quantization axis and the interatomic axis, with the extent of this dependence determined by the Rydberg levels in question.

The second approach involves bringing two atoms close together such that there is a significant overlap between their motional wave-functions. The interaction in this regime is mediated by two body-collisions with a delta function potential, characterized by the s-wave scattering length. High fidelity gates with such a scheme require pristine control over the positions and the motional state of the atoms.

Although two-qubit gates have been implemented using both these approaches [53–57], no such experiment has been done in a large scalable system with qubits cooled down to the motional ground states like our 3D lattice. For various, and as of yet uncertain reasons, the fidelities in these experiments have been limited to $\sim 80\%$, far below the commonly accepted fault-tolerant thresholds. Section 5.2 describes a commonly used Rydberg excitation based gate protocol and gives an overview of the error sources and limitations. A novel Rydberg excitation scheme is also discussed along with a proposal to expand our current single site addressing technology to implement site-selective two qubit gates in 3D. This, along with our cold and well-localized atoms, makes our system a promising candidate for demonstration of high fidelity Rydberg gates.

The control of atom positions in our system, developed for sorting, also offers the possibility of implementing cold-collision gates. Specifically, such an approach promises the creation of large scale entanglement in a few global motion steps to create cluster states which can be used for one-way quantum computation. A high fidelity realization of such a scheme would require significantly better ground state cooling. A possible path to such an implementation is discussed in section 5.3

6.2 Rydberg gates

A Rydberg blockade based controlled phase (also called controlled-Z or C_Z) gate for trapped neutral atoms was proposed by Jaksch et. al. in 2000 [136]. The basic idea of the gate, as shown in Figure 6.1, relies on the fact that an atom that is excited to a high lying Rydberg state can "blockade" the same excitation for a




Figure 6.1. A Rydberg blockade based C_Z gate. Three pulses (labelled 1,2 and 3 based on their timing) are applied to two different atoms. The control atom is transferred to a high lying Rydberg state, $|r\rangle$, if it was in qubit state $|1\rangle$ by the first π pulse. A second, 2π pulse is applied to the target atom. If the control atom was in state $|0\rangle$ and therefore failed to make the transition to $|r\rangle$, the target atom is on resonance with the second pulse and is excited to the Rydberg state and back, accumulating a π phase shift. If the control atom was in $|1\rangle$ and made the transition, the target atom is shifted out of resonance due to the blockade shift, B , and the second pulse fails. The third π pulse transfers the control atom back if it was excited by the first pulse.

nearby atom by shifting the transition frequency due to the strong interaction. The gate involves driving transitions from one of the qubit states (denoted by $|0\rangle$ or $|1\rangle$) to a high lying Rydberg level (denoted by $|r\rangle$). For instance, the $|1\rangle \rightarrow |r\rangle$ transition can be driven by a laser pulse, while nothing happens to atoms in the $|0\rangle$ states that are off-resonant. The gate sequence proceeds as follows. A π pulse (pulse 1) is applied to the control atom, which makes the transition to the Rydberg state if it was in the $|1\rangle$ state. This is followed by a 2π pulse (pulse 2) on the target atom, which will fail if the control atom made the transition to the Rydberg state and shifted the resonance for the target atom by a blockade shift, B . If it succeeds, the wave-function for the target atom gets an $e^{i\pi} = -1$ phase factor due to the 2π

pulse¹. A final π pulse (pulse 3) on the control atom transfers it back to the $|1\rangle$ state if it was excited by the first pulse. If the control atom started out in the $|1\rangle$ state, its wave-function also gets a minus sign due to the two π pulses².

In the two atom basis given by $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ where the first and the second numbers denote the states of the control and target atoms respectively, this sequence of pulses results in the following transformation:

$$\begin{aligned}
 |00\rangle &\rightarrow |00\rangle, \\
 |01\rangle &\rightarrow -|01\rangle, \\
 |10\rangle &\rightarrow -|10\rangle, \\
 |11\rangle &\rightarrow -|11\rangle
 \end{aligned}
 \tag{6.1}$$

A controlled not (CNOT) gate can now be constructed simply by sandwiching the C_Z gate between two Hadamard gates on the the target atom [74].

6.2.1 Fidelity considerations

The fidelity of such Rydberg interaction based gates depends on various factors. The Rydberg states have finite lifetimes, limited by the blackbody radiation at room temperature. As an example, the lifetime of the $75S_{1/2}$ state in Cesium is $\sim 100 \mu s$. This requires the π pulse Rabi frequencies to be on the order of several MHz to keep spontaneous emission from the Rydberg state at a tractable level. Since the atoms in the Rydberg state typically don't experience the same trap as the qubit states³, the excitation to the Rydberg state and back can lead to heating as well as imperfect transfer due to the evolution of the motional wave-function in the Rydberg state. Another problem arises because Rydberg atoms are prone to photo-ionization from the trapping light, the excitation light or the site addressing beams in case of targeted gates. Minimizing all these effects requires the gates to be fast with large Rabi frequencies for the π pulse. On the other hand, the value of the Rabi frequency has to be small enough to not cause off-resonant transitions which

¹This is the Berry phase in an $SU(2)$ system.

²Other gate schemes are also possible. Both the atoms can be transferred to the Rydberg state and allowed to interact. This allows for a much stronger interaction but buys into a lot of problems associated with atoms spending time the Rydberg states. A Rydberg dressing scheme has also been demonstrated where neither of the atoms has appreciable population in the Rydberg state [57], although this comes at the cost of a much slower gate.

³Unless a magic trapping scheme is used [137].

essentially degrade the blockade. This requires the blockade shift to be at least several times the Rabi frequency along with the use of shaped pulses to suppress off resonant transitions. Although higher n levels generally have larger blockade shifts, they are more sensitive to the electric fields. Thus picking the right n level for this gate is a balancing act to minimize the total error. It has been theoretically shown that gate fidelities > 0.9999 are possible in a cryogenic environment for Cesium and Rubidium [58].

Another set of errors arise depending on the way the atom is excited to the Rydberg state. The most straightforward way is to use a single laser pulse to drive a qubit to a high lying P -state. For Cesium, the wavelength for a such a laser is in the ultraviolet (UV) regime around 319 nm , which although not untenable, is difficult for obtaining high powers for a fast gate. More importantly, the excitation to a P state is not preferable because the blockade shift is strongly anisotropic compared to that for the S states and would significantly decrease the number of nearest or next-nearest neighbors that can be entangled, especially in a 3D lattice. Moreover, the P states are significantly more sensitive to external electric fields and generally have a higher photo-ionization cross-section than S states.

For these reasons, although gates with excitation to nP Rydberg states have been demonstrated, a more widely used technique is two-photon transition to nS states, with $6P$ or $7P$ states typically used as intermediate states in Cesium. In this scheme, the limiting factors are the spontaneous emission from the intermediate state and eventually, the photoionization from the exciting lasers. The spontaneous emission can be mitigated by increasing the detuning from the intermediate state, requiring increased laser powers to get the same two-photon Rabi frequency; but this also increases the photo-ionization rate. Another issue that arises is that of unwanted ac Stark shifts on the qubit and Rydberg states due to the exciting beams. The shifts that arise from coupling to far off resonant states increase with laser power and can lead to a substantial differential shifts between the qubit and Rydberg states and thus unwanted phase shifts. In case of focused beams for single site addressing, these can also lead to significant changes to the local trap potential, thereby heating the atoms.




Figure 6.2. UV+microwave excitation. (a) shows the proposed excitation scheme from $|4,0\rangle$ to $75S_{1/2}$ Rydberg state using a UV photon at 319 nm and a microwave photon near 8.7 GHz . The Rabi frequency is Ω_1 for the UV photon and Ω_2 for the microwave photon. The detuning from the $75P_{3/2}$ state is Δ . (b) shows the polarizations required to excite the atom to a single m_J level (in this case $m_J = 1/2$ in the $75S_{1/2}$ state. The σ^+ polarization of the UV light requires the beam to be along the magnetic field. Generating good π polarized microwaves will require at least a dual polarized horn.

6.2.2 UV+microwave excitation

Another two photon excitation scheme is possible using a UV and a microwave photon, instead of two optical photons as in the previous schemes. A possible implementation is shown in Figure 2a, with the $75S_{1/2}$ as the target state and the $75P_{3/2}$ state as the intermediate level⁴. The $75P_{1/2}$ level is not coupled to the ground state by the UV light because the transition matrix elements between

⁴The final choice for n would eventually depend on the local electric fields in our system. Here we use $n = 75$ to demonstrate the advantages and disadvantages of this scheme.

$6S_{1/2}$ and $nP_{1/2}$ are exceedingly small for large n [138]⁵. Although this approach has many of the same tradeoffs as in previous approaches and is more technically challenging, there are certain advantages that make it more attractive.

A major advantage of this scheme is that it almost completely avoids the spontaneous emission from the intermediate state due to the long lifetimes of Rydberg states compared to $6P$ or $7P$ states. The limiting factors for this scheme are: photoionization from the UV and addressing light (described in the following section), the ac Stark shifts of the various levels involved, and the power requirements for the UV light. To demonstrate the advantages and disadvantages of this scheme, it is useful to start with some concrete numbers for the various parameters. The two photon Rabi frequency is given by $\frac{\Omega_1\Omega_2}{2\Delta}$, where Ω_1 is the Rabi frequency for the UV coupling, Ω_2 is the Rabi frequency for the microwave coupling and Δ is the detuning from the $75P_{3/2}$ state. A two photon Rabi frequency of $2\pi \times 5 \text{ MHz}$ would allow for 100 ns π pulses with a total gate time of 400 ns , fast enough to cause minimal heating (around 10^{-4} fractional excitation out of the vibrational ground state). The blockade shift in these states is about 80 MHz at $5 \mu\text{m}$ and 25 MHz at $7 \mu\text{m}$, large enough to significantly suppress any off resonant transitions and effectively entangle up to next nearest neighbors⁶.

After fixing the two photon Rabi frequency, it is useful to plot the power requirements for the UV light and the Stark shifts as a function of Ω_2/Δ as shown in Figure 3. As can be seen from Figure 3a, this approach requires relatively large laser powers at 319 nm but a system with greater than 2 W of output power and a few kHz line-width has recently been demonstrated at this wavelength [140]. A reasonable starting point is to consider a 150 mW UV laser beam focused to $3 \mu\text{m}$ at the atoms. This gives the Rabi frequency for the $6S_{1/2} \rightarrow 75P_{3/2}$ transition, Ω_1 , to be $\sim 2\pi \times 10 \text{ MHz}$. Figure 2b shows the Stark shifts for the ground and Rydberg levels due to the UV light. For a fixed Ω_2/Δ , the Stark shift of the $75S_{1/2}$ state due to the microwaves can be quite large, scaling as Ω_2 . This large dynamical change of the Rydberg state energies can be avoided by simply keeping the microwave on throughout the gate, instead of pulsing it with the UV light. The small changes of the resonance frequency due to the UV related Stark shifts can be compensated by

⁵This is a well known effect in Cesium due to spin-orbit effects

⁶The calculation for the blockade shift and the other calculations that follow were done by Tsung-Yao Wu and Felipe Giraldo Mejia using the Atomic Rydberg Calculator package [139] for Python.

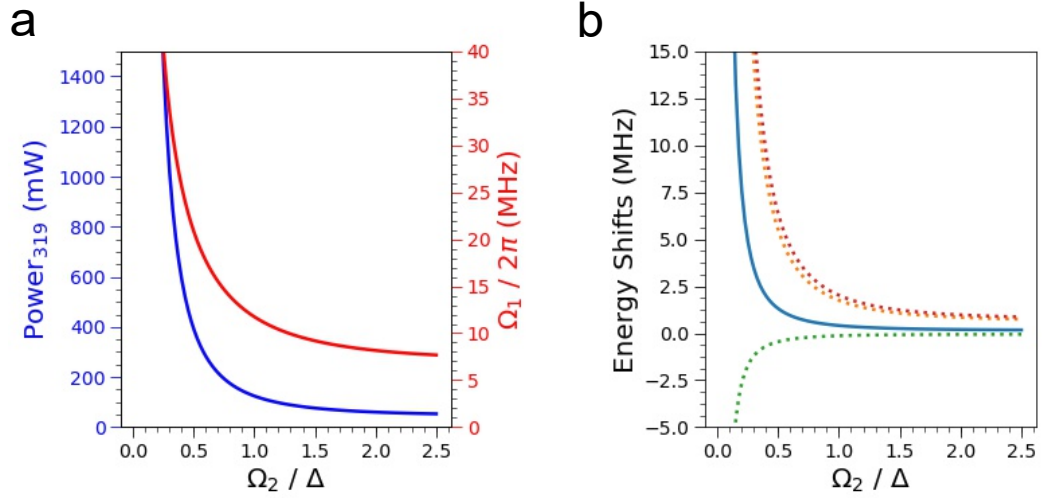


Figure 6.3. Parameters for the two photon transition. The calculations in this Figure are done after the fixing the two photon Rabi frequency to $2\pi \times 5 \text{ MHz}$ and as a function of Ω_2/Δ (a) shows the power requirements for a $3 \mu\text{m}$ waist 319 nm beam (blue curve) and the corresponding Rabi frequency, Ω_1 (red curve). (b) shows the various AC Stark shifts due to the UV light. The dotted green curve is the shift of the $6S_{1/2}$ state due to its coupling to the $75P_{3/2}$ state. The dotted orange curve is the shift of the $6S_{1/2}$ state due to its coupling to all the other states. Similarly, the dotted red curve is the shift of the Rydberg states due to the UV light. This is dominated by the free electron shift. The solid curve is the differential shift between the $6S_{1/2}$ state and the microwave dressed Rydberg state.

dynamic changes to the microwave frequency. This leads to the functional Rydberg state being an admixture of the $75S_{1/2}$, $75P_{1/2}$, and $75P_{3/2}$. It is important to minimize the p character in this dressed state to due the significantly higher electric field sensitivity and photoionization cross-sections of the $75P$ states. This can be done by using a large enough Δ .

As can be seen from Figure 3, these power considerations along with the requirement to minimize the differential shift of the ground and Rydberg levels require that the Ω_2/Δ be > 1 . Along with the requirement that the Δ be large, this implies that the Ω_2 be large as well. Although, the absolute microwave power requirements are very modest because of very high transition matrix elements, the limitation here comes from coupling of the $75S_{1/2}$ state to higher angular momentum $n = 71, l = 3, 4 \dots$ states, which are energetically very close to the $75S_{1/2}$ state. This is because the difference in quantum defects for the S states and the higher angular momentum states is almost exactly 4 (4.016 for the $l = 3$ states)

for Cesium [141]. The electric field of the microwaves, which is quasi-static for the Rydberg electron, can eventually lead to level crossings and a strong mixing of the s state and these higher angular momentum states and eventually population leakage out of the $75S_{1/2}$ state. Avoiding this requires that the total electric field, including the as of yet unknown DC electric field in our system, be small enough to not mix the s state with other states.

A reasonable value for Ω_2 and Δ is $\sim 2\pi \times 200 \text{ MHz}$. This keeps the microwave electric field well below the threshold set by the first level crossing and the admixture of the p states in the excited Rydberg state below 20%. The differential Stark shift of the ground and Rydberg states is almost negligible under these conditions, while the differential Stark shift between the ground states is about $\sim 200 \text{ kHz}$. The small phase shifts due to this as well as the ac Stark shifts on the non-target atoms can be spin echoed out like our single qubit gate schemes. One drawback of using the $75P_{3/2}$ as the intermediate state is that it is hard to excite to a single m_J sub-level in the $75S_{1/2}$ state⁷ and exciting to multiple m_J states can lead to de-phasing due to magnetic field noise. Coupling to only one of the m_J levels requires the UV light to be circularly polarized and the microwave to be linearly polarized (Figure 3b). The UV light can be directed along our current optical pumping axis ($\hat{x} + \hat{y}$) with a σ^+ polarization, but ensuring the a good linear polarization for the microwaves is challenging due to the reflections from the glass cell and components around the cell. It will at least require a dual polarization horn, which should get us close, but it is not enough for arbitrary control over the microwave polarization. In the end, a larger magnetic field (on the order of a $10G$) might be required to make the coupling to the $m_J = -1/2$ level off-resonant.

Finally, the photo-ionization due to the UV light is almost negligible because the ionization cross-sections generally decrease with decreasing wavelength. This is another advantage of using this scheme instead of two optical photons. As we will see in the next section the majority of the error comes from the photoionization due to the additional addressing beam required in a 3D system. If this scheme were to be used in a 1D or 2D system where no other beams are required, the only limiting factor might be the external electric field.

⁷J is the total electronic angular momentum. The total angular momentum, F, is no longer a good quantum number for the Rydberg states.

6.2.3 Single site addressing

For an off-resonant oscillating field, the ac polarizability of a Rydberg state is well approximated by that of a free electron, given by:

$$\alpha_e = -\frac{e^2}{m_e\omega^2}, \quad (6.2)$$

where e is the electron charge, m_e the mass and ω is the frequency of the field. This implies that our existing addressing beams at 880 nm , which are off-resonant with any transitions out of the Rydberg states, can be used to impose a size-able ac Stark shift and thus site-selectively drive the Rydberg excitation for an atom at the intersection of the UV beam and one of the existing addressing beams. This would also require site to site translation of the UV beam parallel to the $x - y$ plane diagonal of the lattice. This can be achieved using fast acousto-optic deflectors. We can still use our existing MEMS system for the addressing beam, with each of the current two addressing aligned to each of the atoms in the two-qubit gate prior to the actual execution of the gate. This avoids any latency associated with re-directing the beams.

Since the ac polarizability in equation 6.2 is negative, this leads to a positive sign ac Stark shift, the same sign as the blockade shift. Therefore, it can so happen, that after one atom has been excited to the Rydberg state, another atom in the path of the UV light could be blockade shifted to the target resonance. Although, it is possible to find an operating point for the addressing Stark shifts such that this doesn't happen for any atom in the whole lattice. With the previously mentioned values of blockade shifts, this requires an ac Stark shift of about 45 MHz . This would lead to the non-target atoms being off-resonant by at least 20 MHz , large enough to effectively suppress any off-resonant excitation. This means that we will have to increase the power in each addressing beam to about 300 mW , easily manageable with tapered amplifiers. At these powers, the addressing beams lead to spontaneous emission at a rate of 3×10^{-4} events per gate and a photo-ionization rate of 2×10^{-3} per gate. The photo-ionization rate can be reduced significantly if we could have higher UV power (than the assumed 150 mW for a $3\text{ }\mu\text{m}$ waist beam), allowing us to decrease the microwave power and the associated admixture of the high photo-ionization cross-section p states.

6.3 Cluster states

Another possibility in our experiment is generation of large cluster states using collisional gates. A cluster state is created when qubits in a lattice are prepared in an equal superposition of the qubit states followed by a C_Z gate for all the nearest neighbors. In neutral atoms, they have been realized by using state dependent lattices in one dimension which can coherently split the wave-function and lead to controlled collision for parts of the wave-functions of the atoms [114, 115]. For example, consider two atoms at neighboring lattice sites (called left and right), each prepared in the $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ states, where for instance in our system, $|0\rangle$ and $|1\rangle$ can be $|3, -3\rangle$ and $|4, -4\rangle$ states respectively. As the polarization of one of the lattice beams is rotated in such a way that the $|0\rangle$ state moves to the left and $|1\rangle$ to the right, the $|1\rangle$ part of the left atom wave-function and the $|0\rangle$ part of the right atom wave-function end up on the same site after a π rotation of the polarization, leading to a controlled collision, with the interaction energy determined by the two body s -wave scattering length and the motional wave-function overlap. The atoms are allowed to interact for some time, which leads to a phase shift, φ , for the part of the wave-function that collides. The polarization is then rotated back to bring the atoms back to their original positions. The final two-atom state is given by:

$$|\psi\rangle = |00\rangle + |01\rangle + e^{-i\varphi} |10\rangle + |11\rangle \quad (6.3)$$

The phase-shift φ can be easily controlled by changing the collision time. When $\varphi = \pi$ and this procedure is generalized to a full lattice with this operation occurring in parallel for all the nearest neighbors, it yields a cluster state. In our system we can create a three-dimensional cluster state after sorting the atoms and three additional motion steps, one in each direction. Such a cluster states when combined with site selective projective measurements along arbitrary directions, can be used to execute quantum gates and implement one-way quantum computation, which has been shown to be formally equivalent to circuit based quantum computation [81, 82].

A major source of error in generating neutral atom cluster states is the vibrational excitation of the atoms. Since the entangling phase shift is directly proportional to the motional wave-function overlap, it is of paramount importance that all atom be in the same vibrational state of the trap. This is required to ensure the same phase

shift for all the nearest neighbor pairs across the array. Our current projection sideband cooling with $\sim 90\%$ ground state population is not good enough for generating 50 atom cluster states with high enough fidelity. Atoms in a Mott-insulating state in an optical lattice, like the ones described in reference [115], are sufficiently cold but such systems lack high fidelity single qubit gates required for arbitrary projective measurements. The cooling in our system is primarily limited by our low trap depth, which means that our atoms are only weakly in the Lamb-Dicke limit (Lamb-Dicke parameter, $\eta=0.37$). Furthermore the low vibrational sideband spacing limits the frequency sweep range Adiabatic Fast Passage (AFP) pulses that drive the sideband transitions, compromising the darkness of the ground state.

Higher powers can be obtained by increasing the lattice powers, but it might be too technically challenging to get a lot more than the $2W$ provided by each of our three lattice tapered amplifiers. Another possibility is to change the lattice wavelength (currently 839 nm) closer to the D2 line, say 846 nm , while keeping the same power, increasing the trap depth by a factor of ~ 2 . Since we know that the scattering rate closer to the line is not good enough to sort atoms (see chapters 2 and 4), this would require dynamically changing the lattice wavelength for cooling. This could possibly be done by combining two different wavelengths on a beam splitter before seeding the tapered amplifiers and adiabatically changing from one seed to the other using a Pockels cell [142]. This wavelength change is small enough such that two wavelengths can be easily chosen that the individual traps are overlapped for a large sections of the lattice, allowing this transition to be made in a few ms with negligible heating. We estimate that this would allow the projection sideband cooling good enough to cool 99% of the atoms to the vibrational ground state.

Bibliography

- [1] R. P. Feynman, “Simulating physics with computers,” *Int. J. Theor. Phys.*, vol. 21, no. 6-7, pp. 467–488, 1982.
- [2] P. Benioff, “The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines,” *Journal of statistical physics*, vol. 22, no. 5, pp. 563–591, 1980.
- [3] I. Manin Yu, “Vychislimoe i nevychislimoe.”
- [4] D. Deutsch, “Quantum theory, the church–turing principle and the universal quantum computer,” *Proc. R. Soc. Lond. A*, vol. 400, no. 1818, pp. 97–117, 1985.
- [5] D. Deutsch and R. Jozsa, “Rapid solution of problems by quantum computation,” *Proc. R. Soc. Lond. A*, vol. 439, no. 1907, pp. 553–558, 1992.
- [6] P. Shor, “Polynomial-time algorithm for prime factorization and discrete logarithms on a quantum computer: proc,” in *35th Annual Symp. on the Foundations of Computer Science*, vol. 124, 1994.
- [7] S. J. Devitt, A. G. Fowler, and L. C. L. Hollenberg, “Robustness of shor’s algorithm,” *Quantum Info. Comput.*, vol. 6, pp. 616–629, Nov. 2006.
- [8] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory,” *Physical review A*, vol. 52, no. 4, p. R2493, 1995.
- [9] A. M. Steane, “Error correcting codes in quantum theory,” *Physical Review Letters*, vol. 77, no. 5, p. 793, 1996.
- [10] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error rate,” *arXiv preprint quant-ph/9906129*, 1999.
- [11] P. Aliferis and A. W. Cross, “Subsystem fault tolerance with the bacon-shor code,” *Physical review letters*, vol. 98, no. 22, p. 220502, 2007.

- [12] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg, “Surface code quantum computing with error rates over 1%,” *Physical Review A*, vol. 83, p. 020302, Feb. 2011.
- [13] D. P. Divincenzo, “The physical implementation of quantum computation,” *Fortschritte der Physik*, vol. 48, pp. 771–783, 2000.
- [14] Gershenfeld and Chuang, “Bulk spin-resonance quantum computation,” *Science (New York, N.Y.)*, vol. 275, pp. 350–356, Jan. 1997.
- [15] L. M. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, “Experimental realization of shor’s quantum factoring algorithm using nuclear magnetic resonance.,” *Nature*, vol. 414, pp. 883–887, 2001.
- [16] S. L. Braunstein, C. M. Caves, R. Jozsa, N. Linden, S. Popescu, and R. Schack, “Separability of very noisy mixed states and implications for nmr quantum computing,” *Physical Review Letters*, vol. 83, pp. 1054–1057, Aug. 1999.
- [17] J. Preskill, “Quantum computing in the nisq era and beyond,” *arXiv e-prints*, Jan. 2018.
- [18] J. Preskill, “Quantum computing and the entanglement frontier,” *arXiv e-prints*, Mar. 2012.
- [19] H. Häffner, C. F. Roos, and R. Blatt, “Quantum computing with trapped ions,” *Physics Reports*, vol. 469, pp. 155–203, Dec. 2008.
- [20] M. H. Devoret and R. J. Schoelkopf, “Superconducting circuits for quantum information: an outlook.,” *Science (New York, N.Y.)*, vol. 339, pp. 1169–1174, Mar. 2013.
- [21] D. Weiss and M. Saffman, “Quantum computing with neutral atoms,” *Physics Today*, vol. 70, no. 7, pp. 45–50, 2017.
- [22] L. Childress and R. Hanson, “Diamond nv centers for quantum computing and quantum networks,” *MRS Bulletin*, vol. 38, no. 2, pp. 134–138, 2013.
- [23] E. Knill, R. Laflamme, and G. J. Milburn, “A scheme for efficient quantum computation with linear optics,” *Nature*, vol. 409, pp. 46–52, Jan. 2001.
- [24] D. Loss and D. P. DiVincenzo, “Quantum computation with quantum dots,” *Phys. Rev. A*, vol. 57, pp. 120–126, Jan 1998.
- [25] K. Mølmer and A. Sørensen, “Multiparticle entanglement of hot trapped ions,” *Phys. Rev. Lett.*, vol. 82, pp. 1835–1838, Mar 1999.

- [26] Cirac and Zoller, “Quantum computations with cold trapped ions,” *Physical review letters*, vol. 74, pp. 4091–4094, May 1995.
- [27] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, “Demonstration of a small programmable quantum computer with atomic qubits,” *Nature*, vol. 536, no. 7614, pp. 63–66, 2016.
- [28] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, “High-fidelity universal gate set for be-9(+) ion qubits,” *Physical Review Letters*, vol. 117, no. 6, 2016.
- [29] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, “High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit,” *Physical Review Letters*, vol. 113, no. 22, 2014.
- [30] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, “Realization of quantum error correction,” *Nature*, vol. 432, pp. 602–605, December 2004.
- [31] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt, “Experimental repetitive quantum error correction,” *Science*, vol. 332, no. 6033, pp. 1059–1061, 2011.
- [32] N. M. Linke, M. Gutierrez, K. A. Landsman, C. Figgatt, S. Debnath, K. R. Brown, and C. Monroe, “Fault-tolerant quantum error detection,” *Science Advances*, vol. 3, no. 10, p. e1701074, 2017.
- [33] K. R. Brown, J. Kim, and C. Monroe, “Co-designing a scalable quantum computer with trapped atomic ions,” *Npj Quantum Information*, vol. 2, p. 16034, Nov. 2016.
- [34] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, pp. 500–503, 2014.
- [35] I. Bloch, “Ultracold quantum gases in optical lattices,” *Nature Physics*, vol. 1, p. 23, Oct. 2005.
- [36] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch, “Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms,” *Nature*, vol. 415, no. 6867, pp. 39–44, 2002.

- [37] A. Kumar, T.-Y. Wu, F. Giraldo, and D. S. Weiss, “Sorting ultracold atoms in a three-dimensional optical lattice in a realization of maxwell’s demon,” *Nature*, vol. 561, pp. 83–87, Sep 2018.
- [38] D. Barredo, V. Lienhard, S. De Leseleuc, T. Lahaye, and A. Browaeys, “Synthetic three-dimensional atomic structures assembled atom by atom,” *Nature*, vol. 561, pp. 79–82, September 2018.
- [39] D. Barredo, S. de Leseleuc, V. Lienhard, T. Lahaye, and A. Browaeys, “An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays,” *Science*, vol. 354, no. 6315, pp. 1021–1023, 2016.
- [40] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, “Atom-by-atom assembly of defect-free one-dimensional cold atom arrays,” *Science*, vol. 354, no. 6315, pp. 1024–1027, 2016.
- [41] H. Kim, W. Lee, H. G. Lee, H. Jo, Y. Song, and J. Ahn, “In situ single-atom array synthesis using dynamic holographic optical tweezers,” *Nature Communications*, vol. 7, p. 13317, 2016.
- [42] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, “Rapid production of uniformly filled arrays of neutral atoms,” *Physical Review Letters*, vol. 115, no. 7, p. 073003, 2015.
- [43] K. D. Nelson, X. Li, and D. S. Weiss, “Imaging single atoms in a three-dimensional array,” *Nature Physics*, vol. 3, no. 8, pp. 556–560, 2007.
- [44] X. Li, *Toward a neutral atom quantum computer*. PhD thesis, The Pennsylvania State University, 2009.
- [45] W. S. Bakr, J. I. Gillen, A. Peng, S. Folling, and M. Greiner, “A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice,” *Nature*, vol. 462, no. 7269, pp. 74–U80, 2009.
- [46] 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,” *Physical Review A*, vol. 70, no. 4, p. 040302, 2004.
- [47] J. H. Lee, E. Montano, I. H. Deutsch, and P. S. Jessen, “Robust site-resolvable quantum gates in an optical lattice via inhomogeneous control,” *Nat. Commun.*, vol. 4, June 2013.
- [48] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman, “Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits,” *Phys. Rev. Lett.*, vol. 114, March 2015.

- [49] N. Lundblad, J. M. Obrecht, I. B. Spielman, and J. V. Porto, “Field-sensitive addressing and control of field-insensitive neutral-atom qubits,” *Nat. Phys.*, vol. 5, pp. 575–580, August 2009.
- [50] Y. Wang, A. Kumar, T. Y. Wu, and D. S. Weiss, “Single-qubit gates based on targeted phase shifts in a 3d neutral atom array,” *Science*, vol. 352, no. 6293, pp. 1562–1565, 2016.
- [51] Y. Wang, X. L. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss, “Coherent addressing of individual neutral atoms in a 3d optical lattice,” *Physical Review Letters*, vol. 115, no. 4, p. 043003, 2015.
- [52] Y. Wang, *Single qubit addressing*. PhD thesis, The Pennsylvania State University, 2016.
- [53] K. Maller, M. Lichtman, T. Xia, Y. Sun, M. Piotrowicz, A. Carr, L. Isenhower, and M. Saffman, “Rydberg-blockade controlled-not gate and entanglement in a two-dimensional array of neutral-atom qubits,” *Physical Review A*, vol. 92, no. 2, p. 022336, 2015.
- [54] 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*, vol. 527, no. 7577, pp. 208–211, 2015.
- [55] 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,” *Physical review letters*, vol. 104, p. 010503, Jan. 2010.
- [56] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, “Entanglement of two individual neutral atoms using rydberg blockade,” *Physical review letters*, vol. 104, p. 010502, Jan. 2010.
- [57] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, “Entangling atomic spins with a rydberg-dressed spin-flip blockade,” *Nature Physics*, vol. 12, pp. 71–74, Jan. 2016.
- [58] L. Theis, F. Motzoi, F. Wilhelm, and M. Saffman, “High-fidelity rydberg-blockade entangling gate using shaped, analytic pulses,” *Physical Review A*, vol. 94, no. 3, p. 032306, 2016.
- [59] 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*, vol. 534, pp. 667–+, June 2016.

- [60] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. W. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletic, and M. D. Lukin, “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature*, vol. 551, no. 7682, pp. 579–584, 2017.
- [61] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye, “A fermi-degenerate three-dimensional optical lattice clock,” *Science*, vol. 358, no. 6359, pp. 90–94, 2017.
- [62] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, “Alkaline-earth atoms in optical tweezers,” *Phys. Rev. X*, vol. 8, p. 041055, Dec 2018.
- [63] M. A. Norcia, A. W. Young, and A. M. Kaufman, “Microscopic control and detection of ultracold strontium in optical-tweezer arrays,” *Phys. Rev. X*, vol. 8, p. 041054, Dec 2018.
- [64] C. Cohen-Tannoudji and D. Guéry-Odelin, *Advances in atomic physics: an overview*. Singapore: World Scientific, 2011.
- [65] X. Li, T. A. Corcovilos, Y. Wang, and D. S. Weiss, “3d projection sideband cooling,” *Physical Review Letters*, vol. 108, no. 10, p. 103001, 2012.
- [66] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, “Sub-poissonian loading of single atoms in a microscopic dipole trap,” *Nature*, vol. 411, pp. 1024–1027, June 2001.
- [67] N. Schlosser, G. Reymond, and P. Grangier, “Collisional blockade in microscopic optical dipole traps,” *Phys. Rev. Lett.*, vol. 89, July 2002.
- [68] T. Gruenzweig, A. Hilliard, M. McGovern, and M. F. Andersen, “Near-deterministic preparation of a single atom in an optical microtrap,” *Nat. Phys.*, vol. 6, pp. 951–954, December 2010.
- [69] F. Bloch, W. W. Hansen, and M. Packard, “The nuclear induction experiment,” *Phys. Rev.*, vol. 70, no. 7-8, pp. 474–485, 1946.
- [70] M. Khudaverdyan, W. Alt, I. Dotsenko, L. Forster, S. Kuhr, D. Meschede, Y. Miroshnychenko, D. Schrader, and A. Rauschenbeutel, “Adiabatic quantum state manipulation of single trapped atoms,” *Phys. Rev. A*, vol. 71, March 2005.
- [71] R. B. Blackman and J. W. Tukey, “The measurement of power spectra from the point of view of communications engineering .1.,” *Bell System Technical Journal*, vol. 37, no. 1, pp. 185–282, 1958.

- [72] C. Kim, C. Knoernschild, B. Liu, and J. Kim, “Design and characterization of mems micromirrors for ion-trap quantum computation,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 13, pp. 322–329, Mar-Apr 2007.
- [73] C. Knoernschild, C. Kim, F. P. Lu, and J. Kim, “Multiplexed broadband beam steering system utilizing high speed mems mirrors,” *Opt. Express*, vol. 17, pp. 7233–7244, April 2009.
- [74] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge: Cambridge University Press, 10th anniversary ed ed., 2010.
- [75] N. F. Ramsey, “A molecular beam resonance method with separated oscillating fields,” *Phys. Rev.*, vol. 78, no. 6, pp. 695–699, 1950.
- [76] E. L. Hahn, “Spin echoes,” *Phys. Rev.*, vol. 77, no. 5, pp. 746–746, 1950.
- [77] H. Y. Carr and E. M. Purcell, “Effects of diffusion on free precession in nuclear magnetic resonance experiments,” *Phys. Rev.*, vol. 94, no. 3, pp. 630–638, 1954.
- [78] S. Meiboom and D. Gill, “Modified spin-echo method for measuring nuclear relaxation times,” *Rev. Sci. Instrum.*, vol. 29, no. 8, pp. 688–691, 1958.
- [79] C. M. Dawson and M. A. Nielsen, “The solovay-kitaev algorithm,” *Quantum Inform. Comput.*, vol. 6, pp. 81–95, January 2006.
- [80] A. M. Steane, “Introduction to quantum error correction,” *Philos. Trans. R. Soc. A-Math. Phys. Eng. Sci.*, vol. 356, pp. 1739–1757, August 1998.
- [81] R. Raussendorf and H. J. Briegel, “A one-way quantum computer,” *Physical Review Letters*, vol. 86, no. 22, pp. 5188–5191, 2001.
- [82] R. Raussendorf, D. E. Browne, and H. J. Briegel, “Measurement-based quantum computation on cluster states,” *Phys. Rev. A*, vol. 68, August 2003.
- [83] S. Olmschenk, R. Chicireanu, K. D. Nelson, and J. V. Porto, “Randomized benchmarking of atomic qubits in an optical lattice,” *New J. Phys.*, vol. 12, November 2010.
- [84] W. Rakreungdet, J. H. Lee, K. F. Lee, B. E. Mischuck, E. Montano, and P. S. Jessen, “Accurate microwave control and real-time diagnostics of neutral-atom qubits,” *Phys. Rev. A*, vol. 79, February 2009.

- [85] C. Knoernschild, X. L. Zhang, L. Isenhower, A. T. Gill, F. P. Lu, M. Saffman, and J. Kim, “Independent individual addressing of multiple neutral atom qubits with a micromirror-based beam steering system,” *Appl. Phys. Lett.*, vol. 97, September 2010.
- [86] B. Arora, M. S. Safronova, and C. W. Clark, “Tune-out wavelengths of alkali-metal atoms and their applications,” *Phys. Rev. A*, vol. 84, October 2011.
- [87] A. M. Souza, G. A. Álvarez, and D. Suter, “Robust dynamical decoupling for quantum computing and quantum memory,” *Phys. Rev. Lett.*, vol. 106, p. 240501, Jun 2011.
- [88] M. A. A. Ahmed, G. A. Alvarez, and D. Suter, “Robustness of dynamical decoupling sequences,” *Phys. Rev. A*, vol. 87, April 2013.
- [89] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, “Randomized benchmarking of quantum gates,” *Phys. Rev. A*, vol. 77, January 2008.
- [90] C. A. Ryan, M. Laforest, and R. Laflamme, “Randomized benchmarking of single- and multi-qubit control in liquid-state nmr quantum information processing,” *New J. Phys.*, vol. 11, January 2009.
- [91] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland, “Single-qubit-gate error below 10^{-4} in a trapped ion,” *Phys. Rev. A*, vol. 84, September 2011.
- [92] 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.*, vol. 82, pp. 2262–2265, March 1999.
- [93] 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.*, vol. 10, December 2013.
- [94] M. O. Brown, T. Thiele, C. Kiehl, T. W. Hsu, and C. A. Regal, “Scaling atom array assembly with grey molasses,” *arXiv e-prints*, p. arXiv:1811.01448, Nov. 2018.
- [95] R. Joerdens, N. Strohmaier, K. Guenther, H. Moritz, and T. Esslinger, “A mott insulator of fermionic atoms in an optical lattice,” *Nature*, vol. 455, pp. 204–U34, September 2008.

- [96] U. Schneider, L. Hackermueller, S. Will, T. Best, I. Bloch, T. A. Costi, R. W. Helmes, D. Rasch, and A. Rosch, “Metallic and insulating phases of repulsively interacting fermions in a 3d optical lattice,” *Science*, vol. 322, pp. 1520–1525, December 2008.
- [97] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Folling, L. Pollet, and M. Greiner, “Probing the superfluid-to-mott insulator transition at the single-atom level,” *Science*, vol. 329, no. 5991, pp. 547–550, 2010.
- [98] J. Vala, A. Thapliyal, S. Myrgren, U. Vazirani, D. Weiss, and K. Whaley, “Perfect pattern formation of neutral atoms in an addressable optical lattice,” *Physical Review a*, vol. 71, no. 3, p. 032324, 2005.
- [99] J. C. Maxwell, *Theory of Heat*. London: Longmans, Green and co., 1871.
- [100] L. Brillouin, “Maxwell demon cannot operate - information and entropy .1,” *Journal of Applied Physics*, vol. 22, no. 3, pp. 334–337, 1951.
- [101] H. S. Leff and A. F. Rex, *Maxwell’s Demon : Entropy, Information, Computing*. Princeton series in physics, Princeton, N.J.: Princeton University Press, 1990.
- [102] R. Landauer, “Irreversibility and heat generation in the computing process,” *Ibm Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, 1961.
- [103] G. N. Price, S. T. Bannerman, K. Viering, E. Narevicius, and M. G. Raizen, “Single-photon atomic cooling,” *Physical Review Letters*, vol. 100, no. 9, p. 093004, 2008.
- [104] W.-B. Wang, X.-Y. Chang, F. Wang, P.-Y. Hou, Y.-Y. Huang, W.-G. Zhang, X.-L. Ouyang, X.-Z. Huang, Z.-Y. Zhang, H.-Y. Wang, L. He, and L.-M. Duan, “Realization of quantum maxwell’s demon with solid-state spins,” *Chin. Phys. Lett.*, vol. 35, April 2018.
- [105] M. D. Vidrighin, O. Dahlsten, M. Barbieri, M. S. Kim, V. Vedral, and I. A. Walmsley, “Photonic maxwell’s demon,” *Phys. Rev. Lett.*, vol. 116, February 2016.
- [106] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Beguin, A. Vernier, T. Lahaye, and A. Browaeys, “Single-atom trapping in holographic 2d arrays of microtraps with arbitrary geometries,” *Physical Review X*, vol. 4, no. 2, 2014.
- [107] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, “Cold matter assembled atom-by-atom,” *Science*, vol. 354, p. 1024, 2016.

- [108] C. Robens, J. Zopes, W. Alt, S. Brakhane, D. Meschede, and A. Alberti, “Low-entropy states of neutral atoms in polarization-synthesized optical lattices,” *Physical Review Letters*, vol. 118, no. 6, p. 065302, 2017.
- [109] M. Olshanii and D. Weiss, “Producing bose-einstein condensates using optical lattices,” *Physical Review Letters*, vol. 89, no. 9, p. 090404, 2002.
- [110] R. A. Williams, J. D. Pillet, S. Al-Assam, B. Fletcher, M. Shotton, and C. J. Foot, “Dynamic optical lattices: two-dimensional rotating and accordion lattices for ultracold atoms,” *Optics Express*, vol. 16, no. 21, pp. 16977–16983, 2008.
- [111] 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*, vol. 269, no. 5221, pp. 198–201, 1995.
- [112] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. Vandrueten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose-einstein condensation in a gas of sodium atoms,” *Phys. Rev. Lett.*, vol. 75, pp. 3969–3973, November 1995.
- [113] J. Z. Hu, A. Urvoy, Z. Vendeiro, V. Crepel, W. Chen, and V. Vuletic, “Creation of a bose-condensed gas of rb-87 by laser cooling,” *Science*, vol. 358, no. 6366, pp. 1078–1080, 2017.
- [114] D. Jaksch, H. J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller, “Entanglement of atoms via cold controlled collisions,” *Physical Review Letters*, vol. 82, no. 9, pp. 1975–1978, 1999.
- [115] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hansch, and I. Bloch, “Controlled collisions for multi-particle entanglement of optically trapped atoms,” *Nature*, vol. 425, pp. 937–940, October 2003.
- [116] H. Weimer, M. Mueller, I. Lesanovsky, P. Zoller, and H. P. Buechler, “A rydberg quantum simulator,” *Nat. Phys.*, vol. 6, pp. 382–388, May 2010.
- [117] W. Gerlach and O. Stern, “Der experimentelle nachweis der richtungsquantelung im magnetfeld,” *Zeitschrift für Physik*, vol. 9, no. 1, pp. 349–352, 1922.
- [118] M. Martinez-Dorantes, W. Alt, J. Gallego, S. Ghosh, L. Ratschbacher, Y. Volzke, and D. Meschede, “Fast nondestructive parallel readout of neutral atom registers in optical potentials,” *Physical Review Letters*, vol. 119, no. 18, p. 180503, 2017.
- [119] M. Kwon, M. F. Ebert, T. G. Walker, and M. Saffman, “Parallel low-loss measurement of multiple atomic qubits,” *Physical Review Letters*, vol. 119, no. 18, p. 180504, 2017.

- [120] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, “Experimental comparison of two quantum computing architectures,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 114, no. 13, pp. 3305–3310, 2017.
- [121] E. Jeffrey, D. Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O’Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, “Fast accurate state measurement with superconducting qubits,” *Physical Review Letters*, vol. 112, no. 19, p. 190504, 2014.
- [122] A. D. Corcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, “Demonstration of a quantum error detection code using a square lattice of four superconducting qubits,” *Nat. Commun.*, vol. 6, April 2015.
- [123] S. Bravyi and A. Kitaev, “Universal quantum computation with ideal clifford gates and noisy ancillas,” *Phys. Rev. A*, vol. 71, February 2005.
- [124] A. M. Steane, “How to build a 300 bit, 1 giga-operation quantum computer,” *Quantum Info. Comput.*, vol. 7, pp. 171–183, Mar. 2007.
- [125] F. Henkel, M. Krug, J. Hofmann, W. Rosenfeld, M. Weber, and H. Weinfurter, “Highly efficient state-selective submicrosecond photoionization detection of single atoms,” *Physical review letters*, vol. 105, no. 25, p. 253001, 2010.
- [126] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais, and A. Browaeys, “Free-space lossless state detection of a single trapped atom,” *Physical Review Letters*, vol. 106, no. 13, p. 133003, 2011.
- [127] M. J. Gibbons, C. D. Hamley, C. Y. Shih, and M. S. Chapman, “Nondestructive fluorescent state detection of single neutral atom qubits,” *Physical Review Letters*, vol. 106, no. 13, p. 133002, 2011.
- [128] J. Bochmann, M. Mucke, C. Guhl, S. Ritter, G. Rempe, and D. L. Moehring, “Lossless state detection of single neutral atoms,” *Physical Review Letters*, vol. 104, no. 20, p. 203601, 2010.
- [129] R. Gehr, J. Volz, G. Dubois, T. Steinmetz, Y. Colombe, B. L. Lev, R. Long, J. Esteve, and J. Reichel, “Cavity-based single atom preparation and high-fidelity hyperfine state readout,” *Physical Review Letters*, vol. 104, no. 20, p. 203602, 2010.
- [130] M. Boll, T. A. Hilker, G. Salomon, A. Omran, J. Nespolo, L. Pollet, I. Bloch, and C. Gross, “Spin- and density-resolved microscopy of antiferromagnetic

- correlations in fermi-hubbard chains,” *Science*, vol. 353, no. 6305, pp. 1257–1260, 2016.
- [131] T. Savard, K. O’hara, and J. Thomas, “Laser-noise-induced heating in far-off resonance optical traps,” *Physical Review A*, vol. 56, no. 2, p. R1095, 1997.
- [132] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.
- [133] R. Raussendorf and J. Harrington, “Fault-tolerant quantum computation with high threshold in two dimensions,” *Physical Review Letters*, vol. 98, no. 19, p. 190504, 2007.
- [134] R. Yamamoto, J. Kobayashi, K. Kato, T. Kuno, Y. Sakura, and Y. Takahashi, “Site-resolved imaging of single atoms with a faraday quantum gas microscope,” *Physical Review A*, vol. 96, no. 3, p. 033610, 2017.
- [135] M. Saffman, “Quantum computing with atomic qubits and rydberg interactions: progress and challenges,” *Journal of Physics B-Atomic Molecular and Optical Physics*, vol. 49, no. 20, p. 202001, 2016.
- [136] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, “Fast quantum gates for neutral atoms,” *Physical Review Letters*, vol. 85, pp. 2208–2211, Sept. 2000.
- [137] S. Zhang, F. Robicheaux, and M. Saffman, “Magic-wavelength optical traps for rydberg atoms,” *Phys. Rev. A*, vol. 84, p. 043408, Oct. 2011.
- [138] M. Fabry and J. R. Cussenot, “Theoretical and experimental determination of oscillator strengths of cesium atom transitions,” *Canadian Journal of Physics*, vol. 54, pp. 836–844, Apr. 1976.
- [139] “Atomic rydberg calculator, <https://arc-alkali-rydberg-calculator.readthedocs.io/en/latest/>.”
- [140] J. Wang, J. Bai, J. He, and J. Wang, “Development and characterization of a 2.2 w narrow-linewidth 318.6 nm ultraviolet laser,” *J. Opt. Soc. Am. B*, vol. 33, pp. 2020–2025, Oct 2016.
- [141] T. F. Gallagher, *Rydberg atoms*, vol. 3. Cambridge University Press, 2005.
- [142] M. Chi, O. B. Jensen, and P. M. Petersen, “High-power dual-wavelength external-cavity diode laser based on tapered amplifier with tunable terahertz frequency difference,” *Opt. Lett.*, vol. 36, pp. 2626–2628, Jul 2011.

Appendix A | Real time experimental control

This appendix briefly describes the real time control system for the quantum computing lab which is used for atom sorting. Historically, the timing sequence for the experiment is controlled by a homemade LabView-based software called Supertime. Supertime is rigid and can only execute a pre-decided timing sequence, with the values for all digital and analog channels pre-programmed at every timing edge. Atom Sorting requires decision making on the fly and the timing sequence needs to be changed based on the motion of the atoms required for a given loading of the lattice, which is random.

To achieve this, we have supplemented the system with field programmable gate arrays (FPGAs) from National Instruments, one of which is a master. When the atom sorting is supposed to start, Supertime transfers timing control to this FPGA (called “FPGA1” from here on) by pausing its own clock. FPGA1 has all the information for sorting atoms programmed into it by another computer (called the “Camera Computer” from here on) running a program called CameraGUI. At this point FPGA1 also takes control of all the channels required for sorting and executes the sequence of site flips and motion steps programmed into it by the Camera Computer. There is also a slave FPGA, FPGA2, which controls the targeting and the power of addressing beams. The information for targeting and addressing beam powers is programmed into FPGA2 by the Camera Computer as well. FPGA2 receives triggers from FPGA1 to output this information synchronously, to control the addressing beams. Other than these, there is also a DDS combined with a

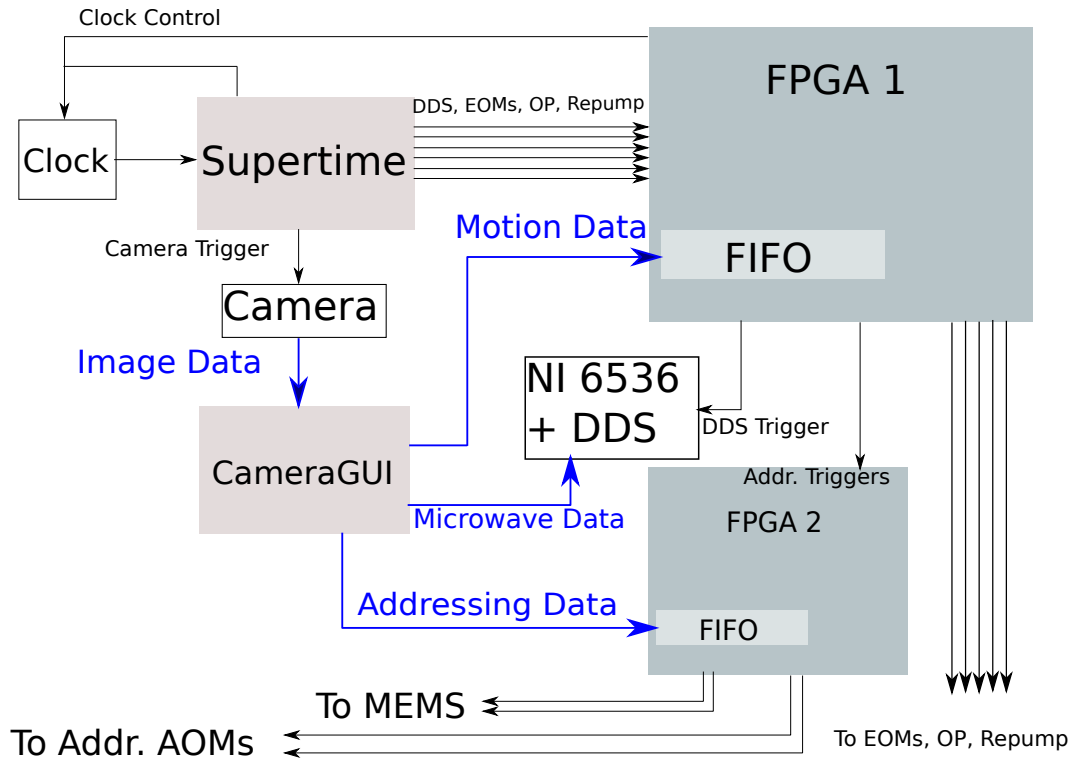


Figure A.1. Schematic for real time control. Supertime creates the timing backbone of the experiment through a fixed timing sequence. The channels required for sorting (DDS trigger, Pockels cell voltages, Repump and Optical pumping powers, all pass through FPGA1. At some point in the sequence Supertime triggers the Camera to take the images of the atoms. The camera passes this data to CameraGUI which analyzes the images and creates an occupancy map. It uses this map to create a sequence of motion steps. It writes the motion data to FPGA1, the addressing data to FPGA2 and the microwave pulse data to the DDS. At this point FPGA1 pauses the Supertime clock source and takes control of the channels required for sorting. It also manages the triggers for addressing and microwaves. After executing the sorting, it transfers the control back to Supertime by releasing the clock source.

9.19 GHz source that generates the microwaves required for the experiment. The DDS is programmed by the Camera Computer in real-time with the sequence of pulses required for cooling and/or sorting. It outputs those pulses based on triggers from Supertime or FPGA1. (see Figure A.1)

A brief description of the components and their function:

- **Supertime** : The timing backbone of the experiment. It is a homemade

software that can create and execute detailed timing sequences for all the analog and digital channels in the lab being used. These sequences cannot be modified on the fly. This software sits on a computer called the “Supertime Computer”. It also triggers the camera that takes pictures of the atoms. At some point in the sequence when the atom sorting is supposed to start, Supertime pauses its own clock and sends a signal to FPGA1 to start sorting.

- **CameraGUI:** This software, programmed in C++, sits on the “Camera Computer”. The Camera is triggered by Supertime and the images are sent to this software for analysis in real time. It creates the occupancy maps of the atoms and calculates the sequence of motion steps required to sort the atoms. It translates that sequence of motion-steps into a series of addressings (site-flips) and directions. It can then write that data into the FIFOs of FPGA1 and FPGA2. FPGA1 only needs the data for timing the addressings and motion steps as well as the directions of the motion in each motion step. FPGA2 needs the details of targeting for each addressing. At this point this program also writes the information required for generating microwave pulses to a PCI card that interfaces with the DDS.

- **FPGA1 :** National Instruments 7845R. This FPGA controls the master timing for sorting. The channels that are required for sorting : DDS trigger, EOMs for motion, Optical Pumping/Repump etc. originate at Supertime but pass through this FPGA. When Supertime is in control, this FPGA just passes the Supertime values through. When it gets the signal from Supertime to start sorting, it waits for the data to be written into its FIFO memory by CameraGUI. Once it has that data, it starts generating the signals for sorting. When done, it restarts the Supertime clock, transferring the control back. Since this FPGA generates all the triggers required to do different things, synchronous control is maintained. All ramps of the EO voltages as well as the relative timings of different triggers are internally controlled by this FPGA. The parameters of those can be set in the CameraGUI.

- **FPGA2 :** National Instruments 7841. This FPGA is a slave to FPGA1. CameraGUI writes the sequence of sites that need to be addressed for sorting to the FIFO memory of this FPGA. This data contains the information about the voltages to MEMS system used for targeting the addressing beams, as well as the

addressing beam powers for each site that needs to be addressed. These values are outputted sequentially as it receives triggers from FPGA1. All ramps of the addressing beam powers are internally controlled by this FPGA. The parameters of those can be set in the CameraGUI.

Appendix B |

Notes on addressing beam alignment and calibration

This appendix describes the procedures for the alignment and calibration of the addressing beam, for the benefit of the people who will use this system in the future.

The data day-to-day alignment of the addressing beams is done using pico-motor driven mirrors at the ray crossing points between lenses L2 and L3 shown in Figure 2.5. This corrects small drifts in the addressing beam positions associated with slight changes in the beam paths after the MEMS mirrors. The general procedure for the alignment is as follows. Assuming that the beams haven't drifted by a large amount, we target the central atom in each row or column (according to the beam being aligned) of the $5 \times 5 \times 5$ lattice. We try to drive the $|4, -4\rangle$ to $|3, -3\rangle$ transition at a microwave frequency slightly higher than the optimum expected Stark shift from the addressing beams. We then scan beam position using the mirrors one by one in two directions and try to maximize the transfer for the central sites. This procedure is iterated a couple of times because of slight coupling between the two directions. This procedure will typically bring back the alignment for the all the sites in the lattice.

Once in a while the MEMS mirror targeting data will have to be re-calibrated. This is especially the case if the system hasn't been used for a while, the final focusing lenses have been moved or the beams had been mis-aligned on the MEMS mirrors. If this calibration is needed, the first sign is that the aforementioned day-to-day alignment will not work for all the sites. The prudent first step is

to check that the beams are still centered on the MEMS mirrors and that the alignment of the MEMS system is un-affected. After this a coarse alignment needs to be done if there is no addressing beam signal on the atoms. The procedures for this have been mentioned in Dr. Yang Wang's dissertation.

B.1 Calibration of the addressing beam positions and powers

If there is addressing signal from all the central 5 rows/columns in the lattice, the calibration can be started. There are two parts to this calibration. First is the calibration for the absolute voltages on the MEMS mirrors to address each site in the $5 \times 5 \times 5$ section of the lattice. The objective is to find the position of the addressing beams to address every site in this section of the lattice. The nominal values of these positions for both beams are stored in files called "HBeam.csv" and "VBeam.csv" in a folder called "Addressing Corrections" on the Penn State cloud (Box) account for this lab, which is also synchronized to the Camera computer. These positions are converted to a voltage by FPGA2 for the MEMS mirror high voltage drivers.

We start by using previously known positions for the central sites in each row/column and scan around them, one beam and one direction for a beam at a time, while parking the microwaves slightly higher than the expected frequency shift. Once enough data is collected, it can be analyzed separately for each site by a Python program called "HV-analyze". This program can also fit the scan for every site to a gaussian and it will automatically point out the fits that are not good enough. The fitted center for these fits give the relative positions for every site compared to the previously known positions of the central sites. It is always a good idea to look at these fits and keep track of ones that are chasing noise or don't look reasonable. For other fits, this program can update the new positions in the two files. Once this is done in one direction for a given addressing beam, it can be repeated in the other direction for the same beam. Typically this procedure would need to be iterated several times for each beam.

A second site-by-site calibration for addressing powers needs to be done to make the shifts due to the addressing beams same for every site. They would typically

be different because of the tight focusing of the beams as well as diffraction at the MEMS mirrors. The information for this calibration is also stored in the two files mentioned before. The first thing that need to be ensured is that the shifts of the central sites are same for each of the beams. After that, we can execute a sequence in which we target every loaded atom with both beams and scan the microwave frequency. This yields a transition spectrum for each site, that can be fitted. The beam powers are then scaled for every site according to give some target frequency shift. This can also be handled by "HV-analyze". Typical we equally divide this scaling between the two beams, although care must be taken if it is suspected that the shifts due to the beams are imbalanced for some site.

This concludes the calibration of the addressing beams.

Appendix C | Implementation of atom sorting in C++

This appendix describes the key software side details for implementing atom sorting in our system, mostly for the benefit of people who will use this system in the future. The sorting is managed entirely by the homemade C++ CameraGUI software. The top level class in this software is called *Panel*. This class handles all the UI as well as other top level functions like controlling the data flow and communication with hardware like the FPGAs. The class that implements the sorting algorithm and creates the sequence of motion steps and single site addressings is called *SortingClass*. After the occupancy map has been created, the Panel class can create and instance of *SortingClass* and call the appropriate sorting function in *SortingClass* according to what the user has selected in the GUI. Based on the output of the algorithm, the Panel class passes this information FPGAs 1 and 2, which can then start sorting atoms.

C.1 class *SortingClass*

SortingClass relies heavily on the data structures in the C++ standard template library (STL) and it is highly recommended that the reader familiarize themselves with that before tinkering with the code. A reference can be found here : <http://www.cplusplus.com/reference/stl/>

SortingClass is initialized by *Panel* with a $5 \times 5 \times 5$ integer array of the occupancy map. There are also several user defined data structures in *SortingClass* that I will

Table C.1. Motion direction convention in `SortingClass`.

direction	real space interpretation
0	no motion
1	positive x direction (right on the screen)
2	negative x direction (left on the screen)
3	positive y direction (down on the screen)
4	negative y direction (up on the screen)
5	positive z direction
6	negative z direction

review now, from simplest to more complex.

C.1.1 Key data structures in `SortingClass`

```
struct sortstep
{
    int direction;
    vector<int> sites;
};
```

This struct represents a sorting step. It contains a direction for the motion step and a list of sites that need to be moved by one lattice spacing. The direction is represented by an integer which can range from 0 to 6. The direction convention is shown in table C.1. Each site is represented as an integer "site", $site = 100x + 10y + z$, where x , y , and z are array indices that represent a lattice site in the $5 \times 5 \times 5$ integer array that is used for occupancy map. `SortingClass` maintains two copies of this array, one is the original and another one that is updated as it figures out individual sorting steps¹.

```
struct sublatticez
{
    int i;
```

¹In `SortingClass` the x , y and z co-ordinates run from 0 to 4 for the $5 \times 5 \times 5$ lattice while in Panel they run from 2 to 7 for x and y directions. This is because Panel typically works with the full $9 \times 9 \times 5$ section that we image. This conversion is handled by Panel.

```

        int j;
        int k;
        int l;
    /*
    * All sites (x,y,z) such that
    *  $i \leq x \leq j$ ,  $k \leq y \leq l$ ,  $1 \leq z \leq n$ .
    * where n (=5) is the total size in z.
    */
};

struct balancetree
{
    balancetree ()
    {
        left = NULL;
        right = NULL;
    }

    ~balancetree ()
    {
        delete left;
        delete right;
    }

    //A queue of sorting steps for balancing this node
    queue<sortstep> steps;

    //children
    balancetree* left;
    balancetree* right;
};

```

This is a recursive implementation of a node of the balancing tree. Each node

has a queue of sortsteps to sort the sublattice associated with it, as well as pointers to the two "children" that result from the division of the sublattice at this node.

```
struct treetuple
{
    balancetree* sorttree;
    bool dir; //0 for x-z and 1 for y-z
    sublatticez s;
};
```

This data structure is only used for the breadth first creation of the balancing tree. It contains a node of the balancetree and the direction (dir) of the plane that divides the sublatticez.

C.1.2 Glossary of key functions in SortingClass

```
queue<sortstep> moveAtoms(sublatticez s1, sublatticez s2,
                          int n, int direction);
```

This function moves n atoms from sublattice s1 to s2 or vice-versa depending on the direction. This function is NOT user error proof. It is assumed that the user has provided s1, s2 and direction that make sense i.e. s1 and s2 have adjacent planes with the same size and that the direction corresponds to motion perpendicular to those planes. Also, it assumes the s1 is "before" s2 in the lattice, i.e. all the indices of s2 are larger than or equal to those of s1. If n is larger than the number of atoms that can be moved by just moving atoms in their own columns across the plane, the function will just move as many atoms as possible.

```
void BalanceZ3D(queue<treetuple> q, int num);
```

This is the main balancing function that creates the balancing tree breadth first using a queue (q) such that each z-row would have at least num atoms. If num is greater than the maximum possible number of atoms in each z-row a full balanced lattice it will just fully balance the lattice. This function dequeues and element from q, which is a treetuple containing a balancetree node, sublattice information and the direction of the plane that needs to divide this sublattice. It then figures out the motion steps to balance this sublattice and stores the sequence of steps

in the `balancetree` nodes. Finally, it pushes the two children sublattices into `q` and calls itself. The execution stops when `q` is empty. This is the standard way of doing a breadth first traversal. To use this function `q` is first initialized with the whole lattice as the only element, with no sorting steps, with a pointer to a `balancetree` object. By the end of the execution the `balancetree` pointer points to a fully constructed tree.

```
queue<sortstep> SortingClass::generateBalanceQueue
(balancetree* sorttree)
```

After the balancing tree has been created by `BalanceZ3D`, it is passed to this function as `sorttree`. This function does a breadth first traversal of `sorttree` to collect all the sorting steps at the same tree level and efficiently combine them. A final sequence of sorting steps for the balancing is created in a queue and returned.

```
queue<sortstep> fillPlanesGeneral(sublatticez s,
                                int numplanes);
```

This function creates a list of functions for compacting a balanced lattice to "numplanes" number of planes in the center of the lattice and returns that list as a queue of sortsteps.

C.1.3 A commented sorting example

I will now show how to put these functions together to sort atoms to a $4 \times 4 \times 3$ target sub-lattice.

```
queue<sortstep> SortingClass::fourFourThree(bool outer)
{
    /* For sorting to a 4x4x3 sub-lattice ,
     * we will first need t move atoms from
     * the outer planes.
     * bool outer controls this behaviour
     * Let's start by defining the various sublattices
     * we would need.
     */
```

```

sublatticez s1,s2,s3,s4,s5;

// First the target sub-lattice
s1.i = 1;
s1.j = 4;
s1.k = 1;
s1.l = 4;

// Outer y-z plane
s2.i = 5;
s2.j = 5;
s2.k = 1;
s2.l = 5;

// Outer x-z plane
s3.i = 1;
s3.j = 5;
s3.k = 5;
s3.l = 5;

// Complement sub-lattice of s2
s4.i = 1;
s4.j = 4;
s4.k = 1;
s4.l = 5;

// Complement sub-lattice of s3
s5.i = 1;
s5.j = 5;
s5.k = 1;
s5.l = 4;

queue<sortstep> move1; // sortsteps for s2 to s4

```

```

queue<sortstep> move2; // sortsteps for s3 to s5
if(outer)
    {
        /* Generate these motion steps.
           * 200 is just a ridiculously large number
           * that ensures that all the atoms that can
           * be moved, will be.
           */
        move1 = moveAtoms(s4 ,s2 ,200 ,2);

        move2 = moveAtoms(s5 ,s3 ,200 ,4);
    }

// Now we will balance s1

// This is the prescription to be
// followed for any balancing

balancetree *sorttree;
sorttree = new balancetree;
sorttree->left = NULL;
sorttree->right = NULL;

queue<treetuple> q;
treetuple t;

t.dir = 0; // begin by dividing with the x-y plane
t.s = s1;
t.sorttree = sorttree;
q.push(t);

//Balance with 3 atoms in each $$-row

BalanceZ3D(q,3);

```

```

// generate the queue containing sorting steps
// for balancing
queue<sortstep> balancesteps =
generateBalanceQueue(sorttree);

// Now generate the queue containing sorting steps
// for filling 3 planes in s1
queue<sortstep> compactsteps
= fillPlanesGeneral(s1,3);

// Now just combine all the different
// steps into one queue
queue<sortstep> steps;
while (!move1.empty())
{
    steps.push(move1.front());
    move1.pop();
}
while (!move2.empty())
{
    steps.push(move2.front());
    move2.pop();
}
while (!balancesteps.empty())
{
    steps.push(balancesteps.front());
    balancesteps.pop();
}
while (!compactsteps.empty())
{

```

```

        steps.push(compactsteps.front());
        compactsteps.pop();

    }
    delete sorttree;
    return steps;
}

```

C.2 Data flow in Panel

To see how Panel uses SortingClass to generate motion steps, we will again look at some example code.

```

void Panel::fourFourThree(short index)
{
    // Create an instance of sorting class
    // Initialize it with the occupancy map
    SortingClass *sclass = new SortingClass(sortarray3D);

    //Call the appropriate sorting function
    queue<sortstep> steps = sclass->fourFourThree();

    // This function creates the data structures
    // that can be written to the FPGA
    generateMotionFromSteps(steps);

    // This just keeps are record of the expected
    // array after sorting for data analysis later.
    fillExpectedArray(sclass, index);

    delete sclass;

}

```

This function starts by creating an instance of `SortingClass` using `"sortarray3D"`, which is the $5 \times 5 \times 5$ array containing the occupancy map.

For sorting, `FPGA1` expects a one dimensional array with information about number of and directions of different motion steps. In `Panel`, this array is called `"motion"`. The first element of this array is the total number of motion steps, followed by a list of number of sites to be addressed in each motion step, which is followed by a list of directions, one for each motion step. For n sorting steps, $n+1$ motion steps are implemented. The last motion step is just a dummy with an addressing and microwave trigger to reset everything. To ensure a clean initialization, dummy triggers are also added to the first motion step by adding an extra site to be addressed. The structure of the motion array is thus : $\{n + 1, N_1 + 1, N_2, \dots, N_n, 1, D_1, D_2, \dots, D_n, 0\}$, where N_i is the number of state-flips in the i -th sorting step and D_i is the direction of the i -th sorting step. Note that number of motion steps is $n + 1$ and the number of state-flips for the last step is 1 and the direction is 0 (i.e. no motion). This is for the final dummy triggers. For the initial dummy triggers, 1 is added to the number of site-flips for the first sorting step.

This `"motion"` array is created by the `"generateMotionFromSteps(steps)"` function shown above. Additionally, it also takes all the sites to be addressed in the sequence of sorting steps, puts them in order and stores them in a global array called `"sitelist"`². Other functions of `Panel` class convert `"sitelist"` array into targeting data for `FPGA2`. Using all this information, `Panel` also creates a timing map of the required microwave pulses. This facilitates incorporating the 60 Hz magnetic field calibration for the microwave frequencies.

Eventually, the motion array and the targeting data are written to `FPGAs 1` and `2` respectively by `Panel`. A required number of microwave pulses is written to the `DDS`.

C.3 Some thoughts on code expansion

The latest version of the `SortingClass` code is good for up to a lattice size of $10 \times 10 \times 10$, where algorithms have been implemented for filling a maximum possible number of z -planes, as well as filling the largest possible cubic lattices.

²During this process it also adds 220 to each site in `"sitelist"`, for the reasons mentioned in the first footnote of this appendix.

If one wants to use larger lattice sites, the only thing that needs to be changed is the handling of the "site" variables. The conversion $100x + 10y + z$ wouldn't work when indices exceed 9. This is not as hard as one may think. Once this is fixed, everything else should just work automatically.

Currently, the only algorithms that are implemented are the ones that balance in x and y directions and compact in the z direction. One might want to exchange these roles for these directions in the future. I think the simplest way to do that might be to rotate the initializing occupancy map for the `SortingClass` and rotate back the output of the `SortingClass` functions.

Appendix D | Mathematica file for throw and catch

This appendix is the mathematica notebook for finding the analytical solution for the atom throw and catch scheme without vibrational excitation when starting in the vibrational ground state described in Chapter 4.

This notebook contains an analytical solution for throwing and catching atoms in the state dependent optical lattice for $n = 0$. This solution causes no heating and preserves the wavefunction. This solution should probably work for any arbitrary n in the harmonic limit. It has been verified for $n = 1$. Some definitions and steps:

1. The various $\alpha_i = \sqrt{\frac{\hbar}{m\omega_i}}$ are the distance units that characterize the various harmonic potentials generated by the lattice during the time evolution.
2. The initial lattice displacement to kick the atoms is x_0 . The initial lattice has the distance units α_1 . The lattice is suddenly shifted by some distance, x_0 , and the atom is allowed to evolve for a quarter cycle. The shifted lattice has the distance units α_2 (determined by the rotation angle and lattice power). The wave function gets a net velocity during that time. The atom is released by turning off the lattice when atom is at the trap minimum (after a quarter cycle).
3. After the kick, the wavefunction travels and evolves for time t_1 .
4. The wavefunction is caught in a lattice with distance unit α_3 . The trap minima for this lattice is at x_1 . We will see that x_1 turns out to be exactly the distance travelled by the atom in t_1 .
5. After a quarter cycle evolution in α_3 , the lattice is turned off again for t_2 , for another cycle of free space evolution.
6. After t_2 the atoms should have moved by half the lattice spacing and the wave function should just have a spatial shift.

The general strategy for finding the solution is as follows. The unknowns are α_2 , α_3 , t_1 , x_1 , t_2 . The final wavefunction is calculated step by step using green's function propagators for quarter cycle evolution in the harmonic trap and the free space propagator. The final wavefunction should match to just the initial wavefunction shifted by half the lattice spacing. The final wave function is quite complicated and we will just look at the argument of the exponential, denoted by "expr". For the wave functions to match the, the coefficients of x and x^2 in $\text{Im}(\text{expr})$ should be 0, only allowing for a constant phase shift. This leads to expressions for t_2 and x_1 . It can be seen that the condition for x_1 is just that the wavefunction is caught at the minima. These can be plugged in to the real part and then another equation can be formulated to require the real part to match up, leading to another equation for α_3 . The final equation is given by ensuring that the total distance travelled by the atom is half the lattice spacing.

The whole system is still over determined with many solutions. Choosing a particular α_1 and α_2 is enough to fix all the parameters. These give a good handle on varying t_1 and t_2 . Care must be taken to ensure that during the first free evolution, the wave function doesn't expand so much that the harmonic approximation fails.

Needs["FourierSeries`"]

\$Assumptions = _Symbol ∈ Reals && _Symbol > 0;

(*Define the wavefunction in a harmonic trap with distance units α_1 , vibrational level n and location given by $y=\text{disp}$ *)

$$\text{psior}[n_, \text{disp}_] := \frac{1}{\sqrt{2^n n!}} * \sqrt{\frac{1}{\sqrt{\pi} * \alpha_1}} * e^{-\frac{(y-\text{disp})^2}{2\alpha_1^2}} * \text{HermiteH}[n, \frac{y-\text{disp}}{\alpha_1}]$$

(*find ψ_{release} by using the quarter cycle propagator for the α_2 lattice displaced by x_0 and $\text{psior}[0,0]$ *)

$$\psi_{\text{release}} = \text{FullSimplify}\left[\frac{1}{\alpha_2 * \sqrt{2 * \pi * \text{I}}} * \text{Integrate}\left[e^{-\frac{\text{I} * (x-x_0) * (y-x_0)}{\alpha_2^2}} * \text{psior}[0, 0], \{y, -\text{Infinity}, \text{Infinity}\}, \text{GenerateConditions} \rightarrow \text{False}\right]\right]$$

$$\text{Out}[4]= -\frac{(-1)^{3/4} e^{-\frac{(x-x_0) ((x-x_0) \alpha_1^2 - 2 i x_0 \alpha_2^2)}{2 \alpha_2^4}} \sqrt{\alpha_1}}{\pi^{1/4} \alpha_2}$$

\$Assumptions = True;

(*Turn the lattice off and find ψ_{free} after time t_1 of free space evolution using Fourier and inverse Fourier transforms. Here $b=\hbar/2m$ *)

$$\psi_{\text{free}} = \text{FullSimplify}\left[\text{InverseFourierTransform}\left[\text{FourierTransform}[\psi_{\text{release}}, x, k] * \text{Exp}[-\text{I} * b * k^2 * t_1], k, x\right]\right]$$

$$\text{Out}[6]= -\frac{(-1)^{3/4} e^{-\frac{-2 b t_1 x_0^2 + (x-x_0) (i (x-x_0) \alpha_1^2 - 2 x_0 \alpha_2^2)}{4 b t_1 \alpha_1^2 - 2 i \alpha_2^4}} \sqrt{\frac{\alpha_1^2}{\alpha_2^4}} \alpha_2^3}{\pi^{1/4} \alpha_1^{3/2} \sqrt{2 i b t_1 + \frac{\alpha_2^4}{\alpha_1^2}}}$$

(*Now catch the atom in the α_3 lattice and find the caught wave function, ψ_{caught} , after a quarter cycle*)

$$\text{In}[7]= \psi_{\text{caught}} = \frac{1}{\alpha_3 \sqrt{2 * \pi * \text{I}}} * \text{Integrate}\left[e^{-\frac{\text{I} * (x-x_1) * (y-x_1)}{\alpha_3^2}} * (\psi_{\text{free}} /. \{x \rightarrow y\}), \{y, -\text{Infinity}, \text{Infinity}\}, \text{GenerateConditions} \rightarrow \text{False}\right]$$

$$\text{Out}[7]= -\left(\left(\left(1-i\right) (-1)^{3/4} e^{-\frac{2 i b t_1 (x-x_1)^2 \alpha_1^2 + x^2 \alpha_2^4 - x_1^2 \alpha_2^4 - 2 i x_0 x_1 \alpha_1^2 \alpha_3^2 - 2 i x_1^2 \alpha_1^2 \alpha_3^2 + 2 x_0 x_1 \alpha_2^2 \alpha_3^2 + x_0^2 \alpha_3^4 - 2 x (x_0 (-i \alpha_1^2 \alpha_2^2) \alpha_3^2 + x_1 (\alpha_2^4 + i \alpha_1^2 \alpha_3^2))}{2 \alpha_1^2 \alpha_3^4}}\right) \sqrt{\frac{\alpha_1^2}{\alpha_2^4}} \alpha_2^3\right) / \left(\sqrt{2} \pi^{1/4} \alpha_1^{3/2} \sqrt{-\frac{i \alpha_1^2}{2 b t_1 \alpha_1^2 - i \alpha_2^4}} \sqrt{2 i b t_1 + \frac{\alpha_2^4}{\alpha_1^2}} \alpha_3\right)$$

```
In[8]:= (*At this point the expected value of the momentum of the atom is
0. Turn off the lattice again and evolve in free space for t2 *)
ψevolve = InverseFourierTransform[
FourierTransform[ψcaught, x, k] * Exp[-I * b * k^2 * t2], k, x]
```

$$\text{Out[8]} = - \left(\left((1 - i) (-1)^{3/4} \right. \right. \\ \left. \left. \frac{-4 b^2 t_1 t_2 x_0^2 - x^2 \alpha^2 + x_1^2 \alpha^2 + 2 i b (t_1 (x - x_1)^2 \alpha^2 + t_2 (x_1^2 \alpha^2 - 2 x_0 x_1 (\alpha^2 + i \alpha^2)) + x_0^2 (\alpha^2 + 2 i \alpha^2))}{8 b^2 t_1 t_2 \alpha^2 - 4 i b t_2 \alpha^2 - 2 \alpha^2 \alpha^3} \right)^{-2 i x_0 x_1 \alpha^2 \alpha^3 - 2 i x_1^2 \alpha^2 \alpha^3 - 2 x_0 x_1 \alpha^2 \alpha^3 - x_0^2 \alpha^3 - 2 x (x_0 (-i c} \right. \\ \left. \frac{\sqrt{\frac{\alpha^2}{\alpha^2}} \alpha^2}{\alpha^2} \right) / \left(\sqrt{2} \pi^{1/4} \alpha^{3/2} \sqrt{-\frac{i \alpha^2}{2 b t_1 \alpha^2 - i \alpha^2}} \sqrt{2 i b t_1 + \frac{\alpha^2}{\alpha^2}} \right. \\ \left. \sqrt{\frac{2 i b t_1 \alpha^2 + \alpha^2}{\alpha^2 \alpha^3}} \alpha^3 \sqrt{\frac{4 i b^2 t_1 t_2 \alpha^2 + 2 b t_2 \alpha^2 - i \alpha^2 \alpha^3}{2 b t_1 \alpha^2 - i \alpha^2}} \right) \right)$$

```
In[11]:= (*Woah!!Simplify this expression*)
$Assumptions = _Symbol ∈ Reals && _Symbol > 0;
FullSimplify[ψevolve]
```

$$\text{Out[12]} = - \left(i e^{\frac{-4 b^2 t_1 t_2 x_0^2 + 2 i b (t_1 (x - x_1)^2 + t_2 (x_0 - x_1)^2) \alpha^2 + 4 b t_2 x_0 (-x_0 - x_1) \alpha^2 + (x - x_1)^2 \alpha^2 + 2 i (x - x_1) ((x_0 - x_1) \alpha^2 + i x_0 \alpha^2) \alpha^3 + x_0^2 \alpha^3}{8 b^2 t_1 t_2 \alpha^2 - 4 i b t_2 \alpha^2 - 2 \alpha^2 \alpha^3}} \sqrt{\alpha^2} \alpha^2 \alpha^3 \right) / \\ \left(\pi^{1/4} \sqrt{2 i b t_1 \alpha^2 + \alpha^2} \sqrt{2 i b t_2 - \frac{i \alpha^2 \alpha^3}{2 b t_1 \alpha^2 - i \alpha^2}} \right)$$

```
In[13]:= (*Lets just look at the argument of the exponential
and separate it into real and imaginary parts*)
expr = (-4 b^2 t1 t2 x0^2 + 2 i b (t1 (x - x1)^2 + t2 (x0 - x1)^2) α^2 + 4 b t2 x0 (-x0 + x1) α^2 +
(x - x1)^2 α^2 + 2 i (x - x1) ((x0 - x1) α^2 + i x0 α^2) α^3 + x0^2 α^3) /
(8 b^2 t1 t2 α^2 - 4 i b t2 α^2 - 2 α^2 α^3)
```

$$\text{Out[13]} = \frac{-4 b^2 t_1 t_2 x_0^2 + 2 i b (t_1 (x - x_1)^2 + t_2 (x_0 - x_1)^2) \alpha^2 + 4 b t_2 x_0 (-x_0 + x_1) \alpha^2 + (x - x_1)^2 \alpha^2 + 2 i (x - x_1) ((x_0 - x_1) \alpha^2 + i x_0 \alpha^2) \alpha^3 + x_0^2 \alpha^3}{(8 b^2 t_1 t_2 \alpha^2 - 4 i b t_2 \alpha^2 - 2 \alpha^2 \alpha^3)}$$

```
In[14]:= rexp = FullSimplify[ComplexExpand[expr + Conjugate[expr]] / 2]
cexp = FullSimplify[ComplexExpand[expr - Conjugate[expr]] / 2 / I]
```

$$\text{Out[14]} = - \left(\left(\alpha^2 (-2 b t_2 (2 b t_1 x_0 + (x_0 - x_1) \alpha^2) + (-x + x_1) \alpha^2 \alpha^3 + x_0 \alpha^3)^2 \right) / \right. \\ \left. (32 b^4 t_1^2 t_2^2 \alpha^4 + 2 \alpha^4 \alpha^3 + 8 b^2 t_2 (t_2 \alpha^2 - 2 t_1 \alpha^4 \alpha^3)) \right)$$

$$\text{Out[15]} = \left(b t_2 (4 b t_2 x_0 (-x_0 + x_1) \alpha^2 + (x - x_1)^2 \alpha^2 + \right. \\ \left. 4 b^2 t_1 ((t_1 (x - x_1)^2 + t_2 (x_0 - x_1)^2) \alpha^4 - t_2 x_0^2 \alpha^2) \right) - \\ 2 b t_2 (x - x_1) (2 b t_1 (-x_0 + x_1) \alpha^4 + x_0 \alpha^2) \alpha^3 - \\ b ((t_1 (x - x_1)^2 + t_2 (x_0 - x_1)^2) \alpha^4 - t_2 x_0^2 \alpha^2) \alpha^3 + (x - x_1) (-x_0 + x_1) \alpha^4 \alpha^3 / \\ (16 b^4 t_1^2 t_2^2 \alpha^4 + \alpha^4 \alpha^3 + 4 b^2 t_2 (t_2 \alpha^2 - 2 t_1 \alpha^4 \alpha^3))$$

(*Now make impose the conditions that the complex part can only be a constant*)
 Solve[Coefficient[cexp, x] == 0 && Coefficient[cexp, x, 2] == 0, {t2, x1}]

$$\text{Out[16]= } \left\{ \left\{ t2 \rightarrow \frac{t1 \alpha 1^4 \alpha 3^4}{4 b^2 t1^2 \alpha 1^4 + \alpha 2^8}, x1 \rightarrow \frac{x0 (2 b t1 + \alpha 2^2)}{\alpha 2^2} \right\} \right\}$$

In[18]:= (*Use these results to simplify the real part*)

$$\text{FullSimplify[rexp /. \{t2 \rightarrow \frac{t1 \alpha 1^4 \alpha 3^4}{4 b^2 t1^2 \alpha 1^4 + \alpha 2^8}, x1 \rightarrow \frac{x0 (2 b t1 + \alpha 2^2)}{\alpha 2^2}\}]}$$

$$\text{Out[18]= } - \frac{(4 b^2 t1^2 \alpha 1^4 + \alpha 2^8) (x \alpha 2^2 - x0 (2 b t1 + \alpha 2^2 + \alpha 3^2))^2}{2 \alpha 1^2 \alpha 2^8 \alpha 3^4}$$

In[19]:= (*Now make sure that the coefficient of (x-something)² in the expression above is the same as the original wavefunction i.e. 1/α1², to find α3.

The final condition is that "something = half the lattice spacing"*)
 Solve[(α2⁸ + 1⁴ b² t1²) / α2⁴ / α3⁴ == 1, {α3}]

$$\text{Out[19]= } \left\{ \left\{ \alpha 3 \rightarrow - \frac{(b^2 t1^2 + \alpha 2^8)^{1/4}}{\alpha 2} \right\}, \left\{ \alpha 3 \rightarrow - \frac{i (b^2 t1^2 + \alpha 2^8)^{1/4}}{\alpha 2} \right\}, \right. \\ \left. \left\{ \alpha 3 \rightarrow \frac{i (b^2 t1^2 + \alpha 2^8)^{1/4}}{\alpha 2} \right\}, \left\{ \alpha 3 \rightarrow \frac{(b^2 t1^2 + \alpha 2^8)^{1/4}}{\alpha 2} \right\} \right\}$$

(*Really only one of these makes sense.*)

Vita

Aishwarya Kumar

Aishwarya Kumar grew up in a small town called 'Kotdwara' in India, in the foothills of the Himalayas. His interest in physics was ignited by his high school teacher, Mr. B. Tripathi, who introduced him to the Feynman Lectures in Physics. He got his bachelors degree from The Indian Institute of Technology, Delhi in 2012 and joined Penn State as a graduate student in the same year. He started working for Prof. David Weiss as a graduate student on the Quantum Computation experiment in 2013. During his Ph.D. the group was able to demonstrate Single qubit gates, atom sorting and lossless state detection in a neutral atom based experiment.