

ULTRACOLD ATOMS IN A DISORDERED OPTICAL LATTICE

BY

MATTHEW ROBERT WHITE

B.S., University of California, Santa Barbara, 2003

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2009

Urbana, Illinois

Doctoral Committee:

Professor Paul Kwiat, Chair
Assistant Professor Brian DeMarco, Director of Research
Assistant Professor Raffi Budakian
Professor David Ceperley

Acknowledgments

This work would not have been possible without the support of my advisor Brian DeMarco and coworkers Matt Pasienski, David McKay, Hong Gao, Stanimir Kondov, David Chen, William McGehee, Matt Brinkley, Lauren Aycock, Cecilia Borries, Soheil Baharian, Sarah Gossett, Minsu Kim, and Yutaka Miyagawa. Generous funding was provided by the University of Illinois, NSF, ARO, and ONR.

Table of Contents

List of Figures	v
Chapter 1 Introduction	1
1.1 Strongly-interacting Boson Systems	1
1.2 Bose-Hubbard Model	2
1.3 Disordered Bose-Hubbard Model with Cold Atoms	3
1.4 Observations	4
1.5 Future Work on the DBH Model	5
1.6 Quantum Simulation	5
1.7 Outline	6
Chapter 2 BEC Apparatus	8
2.1 Introduction	8
2.2 Atomic Properties	8
2.3 Overview of BEC Apparatus	10
2.4 Laser System	12
2.5 Vacuum System	17
2.6 Quadrupole Coils	19
2.7 MOT	21
2.8 Capture	22
2.9 Ioffe-Pritchard Magnetic Trap	25
2.10 Evaporative Cooling	30
2.11 Imaging	33
2.12 Dipole Trap	37
2.13 Plugged Quadrupole Trap	42
Chapter 3 Control System	45
3.1 Control Software	45
3.2 Control Hardware	49
3.3 Isolation	50
3.4 Experiment Monitor	52
3.5 Image Processing	54
3.6 FPGA Sequencer	54

Chapter 4	Optical Lattice	59
4.1	Optical Lattices	59
4.2	Band Structure	60
4.3	Bose-Hubbard Hamiltonian	65
4.4	Expansion from the Lattice	72
4.5	Technical Overview	76
4.6	Light Source	76
4.7	Intensity Stabilization	79
4.8	Lattice Beam Path	79
4.9	Alignment and Calibration	81
4.10	Loading into Optical Lattice	83
4.11	Release from Optical Lattice	85
4.12	Superfluid/Mott insulator Phase Transition	86
Chapter 5	Disordered Lattice	88
5.1	Introduction	88
5.2	Overview of Experiment	88
5.3	Optics	89
5.4	Speckle Characterization	91
5.5	Alignment	93
5.6	Density-Density Correlations with Disorder	96
5.7	Disordered Bose-Hubbard Hamiltonian	97
5.8	Results	99
5.9	Discussion	104
5.10	Future Directions	106
Chapter 6	Simulating Quantum Magnetism	108
6.1	Overview	108
6.2	Multilevel Atom-Light Interaction	108
6.3	Optical Lattices for Multilevel Atoms	110
6.4	Stroboscopic Quantum Simulation	111
6.5	Systems of Interest	114
6.6	Implementation	114
6.7	Microwave System	122
References		127
Vita		136

List of Figures

2.1	Electronic structure of ^{87}Rb .	9
2.2	Schematic of diode laser system.	11
2.3	CAD model and photograph of external cavity diode laser design. A thermoelectric element between the enclosure and the plate to which the ECDL components are mounted is used to adjust the temperature of the plate.	12
2.4	Schematic of improved ECDL design.	12
2.5	Polarization spectroscopy setup for laser stabilization.	14
2.6	Phase-locked loop for stabilizing repump ECDL relative to master ECDL.	15
2.7	CAD design for ultra-high-vacuum system.	17
2.8	Isometric view from above of science cell.	18
2.9	Schematic of quadrupole coil circuit.	20
2.10	Ringling of closed-loop Hall probe current sensor in response to a sudden change in current.	20
2.11	Sequence for transferring atoms into quadrupole trap from the MOT.	22
2.12	Fraction of atoms recaptured in MOT after transport to given position in transfer tube and return.	24
2.13	Lifetime of trapped atoms in the science cell.	25
2.14	Photographs of core portion of Ioffe-Pritchard magnetic trap and assembled trap showing bias coils and connections for current and cooling water.	27
2.15	End-on view of Ioffe bars during construction.	27
2.16	Sudden jumps in current through the Ioffe bars during and after a slow ramp down of the setpoint current.	28
2.17	Schematic of IP trap circuit.	29
2.18	Adiabatic potentials for the $F = 1$ state in the presence of gravity.	30
2.19	Displacement from the IP trap center of atoms trapped in dressed potential vs. RF frequency.	31
2.20	Formation of a BEC during evaporative cooling, viewed after several ms time-of-flight.	32
2.21	Optical system used for absorption imaging of atoms in science cell.	35
2.22	Sample images taken with and without the quarter-wave plate installed in the imaging system.	35
2.23	Layout of CCD in kinetics mode for a Princeton Instruments VersArray XP camera.	36
2.24	Auxiliary imaging system.	37
2.25	Schematic of optical dipole trap system.	39

2.26	Evaporation sequence for producing nearly pure BEC in the hybrid optical/magnetic trap.	41
2.27	Temperature, in-trap density, and phase space density vs. total atom number during evaporation sequence.	42
2.28	Lifetime of cold atoms (after RF evaporation) in the plugged quadrupole trap vs. plug beam power.	43
3.1	Sample of sequencer code, generated C code, and corresponding timeline.	46
3.2	Another sample of sequencer code, showing the ability to interleave C code.	47
3.3	Schematic of a direct digital synthesizer.	49
3.4	Isolation circuit for digital signals.	51
3.5	Screenshot of experiment monitor software.	53
3.6	Structure of FPGA-based digital sequencer.	56
3.7	Data format generated by <code>out_fpga</code> and used by <code>plb_fastdio</code> core.	57
3.8	Schematic of control system software and FPGA sequencer.	57
4.1	Bloch bands for a red-detuned 1D optical lattice potential.	61
4.2	Mean-field theory calculation of ground state condensate fraction for the Bose-Hubbard model in 3D.	69
4.3	Variance of phase between two lattice sites calculated using mean-field ground state.	71
4.4	Single particle density (without envelope and normalized) after expansion from 1D lattice with 10 sites for superfluid regime.	73
4.5	Distribution of light for the optical lattice.	77
4.6	Optical system for one lattice beam.	80
4.7	Fluorescence from one lattice beam on resonance passing through hot atom cloud.	81
4.8	Diffraction atoms using a single lattice beam.	83
4.9	Lifetime of BEC in the optical lattice, for two different lattice depths.	84
4.10	Condensate fraction recovered vs. (linear) ramp down time of $20 E_R$ lattice potential.	85
4.11	Observation of superfluid to Mott insulator phase transition.	87
5.1	Schematic view of science cell showing speckle beam and optics and optical lattice beams and creation of disordered lattice.	89
5.2	Speckle beam optics.	90
5.3	Autocorrelation function of the optical speckle field determined from ex situ measurement.	91
5.4	Probability distribution of site energy offsets for several different disorder potentials which have been realized experimentally.	92
5.5	CAD design of flexure mount for speckle lens and diffuser and image of finished mount.	94
5.6	Displacement of atom cloud due to speckle beam.	94
5.7	Displacement of a 500 nK thermal cloud vs. speckle intensity.	95
5.8	Density-density correlation image for the disordered lattice.	96
5.9	Probability density of DBH parameters for several different lattice depths.	98
5.10	Correlation between nearest-neighbor energy offset and tunneling matrix element.	98

5.11	Correlation between site energy and tunneling matrix elements averaged over all nearest neighbors.	99
5.12	Lattice and speckle intensity vs. time for disorder experiments.	99
5.13	Recovery of condensate fraction vs. ramp down time of disordered lattice.	100
5.14	Effect of strong, fine-grain disorder on condensate fraction across a range of lattice depths, comparing adiabatic and bandmap turn off of disordered and clean lattices.	101
5.15	Change in condensate fraction for lattice depths of 6, 12, and 14 E_R vs. speckle intensity.	102
5.16	Regions of the Bose-Hubbard phase diagram explored by the data in figure 5.15.	103
5.17	Reversible change in condensate fraction induced by disorder.	103
5.18	Sampling of proposed phase diagrams for the disordered Bose-Hubbard model.	105
6.1	Lattice beam and magnetic field geometry for state-dependent optical lattice calculations.	110
6.2	State dependent potentials for a 785 nm lattice vs. polarization of the retro-reflected beam.	116
6.3	Potentials for $ F = 1, m_F = -1\rangle$ state for the two other lattice beams in our system with retro-reflection polarizations of 30° and 90°	117
6.4	Optical system for implementing shifting lattice technique.	117
6.5	Interaction energies vs. retro beam polarization angle for lattice beam 1.	118
6.6	Effective spin-spin interactions using the shifting lattice technique.	121
6.7	Rabi oscillations between $ F = 1, m_F = -1\rangle$ and $ F = 2, m_F = -2\rangle$ driven by a microwave signal.	123
6.8	Phase-locked loop for YIG oscillator.	123
6.9	Phase noise calculated for PLL shown in figure 6.8.	124
6.10	Microwave coils as fabricated.	125

Chapter 1

Introduction

Quantum degenerate atom gases have emerged as valuable tools for studying theoretical models of condensed matter systems. The ability to manipulate atoms with magnetic and electric fields, from DC to optical frequencies, coherently if desired, and tune the strength and sign of interatomic interactions means that a wide range of effective Hamiltonians can be realized. Moreover, the Hamiltonian can be known to very high accuracy, far better than in most condensed matter systems.

Within the past 15 years, Bose-Einstein condensate (BEC) [1, 2] and Fermi degeneracy [3] have been achieved in ultracold (henceforth, cold) atom gases, adding to the short list of systems which display macroscopic quantum behavior, and culminating decades of work on cooling atoms to μK temperatures. Early experiments with BEC primarily explored the nature of this quantum behavior in the weakly-interacting regime, studying, for example, interference between separate condensates and vortices [4, 5]. Cold atoms are also being employed for precision measurement, besides their longstanding use in timekeeping, even probing the standard model of particle physics and beyond [6–8].

Loading a BEC into an optical lattice, a periodic potential produced using light, opens a new realm of possibility. For shallow lattice potentials, the atoms remain condensed and weakly interacting. As the lattice is strengthened, atoms are increasingly confined to individual lattice sites and the strongly-interacting regime can be reached, a hallmark of which is the superfluid/Mott insulator phase transition. The observation of this transition by Greiner *et. al.* [9], as well as the advent of techniques such as Feshbach resonances [10] for controlling interatomic interactions, stimulated interest in the use of cold atoms to study unsolved problems in condensed matter physics. In this work, we discuss the realization of one of the many tools in the cold atom toolbox, a model of strongly-interacting, disordered systems, and progress toward the realization of a more general scheme for simulating model Hamiltonians.

1.1 Strongly-interacting Boson Systems

We will begin by enumerating several bosonic condensed matter systems in which strong interactions and disorder play important roles.

Numerous experiments [11] have studied the behavior of superfluid helium in porous media, e.g., Vycor glass. In the high density case, where the pores are filled with helium, the observed behavior is not strikingly different from that of bulk helium. However, in the low density case, where only a thin layer is adsorbed to the surface of the pores, the superfluid transition temperature can be reduced by a factor of several hundred and only occurs once a critical coverage is reached. These observations have been explained in terms of boson localization in a disordered potential; sufficient coverage is needed to screen the disordered potential created by the substrate before a layer of superfluid can form.

Non-classical rotational inertia in solid ^4He at sub-Kelvin temperatures, possible evidence of a so-called supersolid, was first observed in porous media (Vycor glass and porous gold), although it has since been observed in bulk ^4He [12, 13]. However, the strong dependence on ^3He fraction and annealing indicates disorder may still be relevant to the observed behavior [14, 15].

Experiments on granular superconductors have shown a zero-temperature superconductor/insulator transition depending only on the normal state conductivity, across a range of materials. The universality of this behavior prompted studies using micro-fabricated Josephson junction arrays to model systems of superconducting grains separated by thin insulating regions; see [16] for a review. The effect of disorder, unavoidable in granular superconductors, in Josephson junction arrays has also been examined [17].

At low temperatures, a global transition from a superconducting to an insulating state is observed in Josephson junction arrays (while each island in the array remains superconducting) as the ratio E_J/E_C is decreased. Here, E_J is the Josephson energy, which characterizes the tunneling between islands, and $E_C = e^2/2C$ is the electrostatic interaction energy on a island with capacitance C .

Scanning tunneling microscopy of high- T_c superconducting cuprates has revealed that, at low hole doping, hole-rich superconducting domains (on the order of a few nm) form, separated by hole-depleted non-superconducting regions [18]. It is unclear whether the separation is spontaneous or due to underlying disorder and defects in the material. This observation suggests that the long-range behavior of the cuprates is that of a granular superconductor, although it does not directly address the open question of the pairing mechanism in the cuprates.

1.2 Bose-Hubbard Model

A first attempt at a model of the strongly-interacting boson systems described above might reduce the system to a set of discrete sites, parameterize the on-site interaction and nearest neighbor tunneling with single numbers, and ignore coupling between non-neighbor sites and off-site interactions. Such a model, based on the Hubbard model familiar in Fermi systems, was first popularized by Fisher *et. al.* [19]. Denoting the operator which destroys

a particle on site i by \hat{b}_i and the number operator by $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$, the Hamiltonian is

$$\hat{H} = - \sum_{\langle i,j \rangle} J_{ij} \hat{b}_i^\dagger \hat{b}_j + \sum_i \varepsilon_i \hat{n}_i + \frac{1}{2} \sum_i U_i \hat{n}_i (\hat{n}_i - 1), \quad (1.1)$$

with site (kinetic and potential) energies ε_i , tunneling energies between nearest neighbors J_{ij} , and on-site interaction energies U_i . In the absence of disorder (and any other non-uniform potential), ε , J , and U are site-independent; this is the (clean) Bose-Hubbard (BH) model. The opposite limit, where one or more the parameters are distributed randomly, is called the disordered Bose-Hubbard (DBH) model.

Even in a simple mean-field theory [19, 20], the BH model predicts a zero-temperature superfluid–insulator phase transition as the ratio J/U is decreased. This result is evidence for the applicability of the BH model to the physical systems described in the previous section. In analogy with the fermionic metal-insulator transition, the insulating state in the BH model is often called a Mott insulator (MI).

In realistic condensed matter systems, the DBH model is more applicable than the BH model and theoretical work has focused as much, if not more, on the DBH model. The nature of the phase diagram for the DBH model is more contentious. Theoretical studies predict the appearance of a Bose glass (BG) phase, but disagree on how it enters the phase diagram. The BG phase has zero superfluid fraction but with a gapless excitation spectrum and finite compressibility. One picture is that the BG phase consists of many small, isolated “puddles” of superfluid. In contrast, the MI state, which also has zero superfluid fraction, has a gapped excitation spectrum and is incompressible. Fisher *et al.* argue in [19] that a BG phase intervenes between the MI and SF phases, implying that disorder transforms the MI phase into a BG phase. However, recent results using, e.g., stochastic mean-field [21] and replica theory [22] indicate that disorder can change the MI phase directly into a SF.

The problem has also been studied using various mean-field methods [23–25], renormalization group [26], replica theory [22], and computationally with quantum Monte Carlo (QMC) algorithms [27–32], but no consensus on the DBH model phase diagram has emerged.

1.3 Disordered Bose-Hubbard Model with Cold Atoms

Starting with a BEC allows atoms to be loaded into the ground state of an optical lattice. This, together with the nature of the interatomic interactions, make cold bosonic atoms in an optical lattice an ideal tool for studying the BH model on a periodic lattice, as was first pointed out by Jaksch *et al.* [33]. Because the interaction strength, parameterized by U/J , can be tuned over a wide range with a simple adjustment of the lattice strength, most of the parameter space can be covered.

In this work, we use cold atoms in a disordered optical lattice, created by combining a fine-grained optical speckle field and a 3D optical lattice, to study the DBH model. The

disorder strength is continuously tunable by controlling the intensity of the speckle field, and the statistical properties of the disorder are known, allowing the distribution of the DBH parameters ε_i , J_{ij} , and U_i to be calculated. This level of knowledge is seldom, if ever, available in disordered condensed matter systems, making comparison of experimental observations with results from theoretical models difficult. We access the previously unexplored combination of strong interactions, enabled by the optical lattice, and true disorder, enabled by the fine-grained speckle field, realizing the DBH model in a cold atom system.

A number of previous studies have experimentally investigated cold atoms in disordered potentials, mostly in 1D. One set of studies [34–36] has created quasi-disorder in 1D by overlapping two optical lattices with incommensurate periods (415 nm and 538 nm, for example). These potentials, which show perfect correlations in measures such as the joint-probability distribution of the energy shift for neighboring lattice sites, realize the so-called Aubry-André Hamiltonian. Another set of studies [37–40] has used optical speckle fields to create non-periodic disorder. However, in most of these studies, the characteristic speckle size is large, on the order of 10 μm . The only previous study [41] to combine true disorder with an optical lattice used a 1D optical lattice and a coarse-grain speckle field with speckle size nearly an order of magnitude larger than the lattice period. Theories of the DBH model typically assume uncorrelated disorder between neighboring sites, which requires speckle sizes on the order of the lattice period. An entirely different route to disorder, which may or may not probe similar physics, is the addition of impurity atoms to an optical lattice [42]. The results of these studies have been challenging to interpret.

Universal behavior, such as Anderson localization [43], in disordered non-interacting systems is well understood. Indeed, Anderson localization has been studied in a host of systems, from light in photonic crystals [44] to sound waves [45], and, recently, cold atom systems using both true disorder from a speckle field [46] and quasi-disorder from an incommensurate lattice [47]. However, the effects of disorder combined with strong interactions remains an outstanding challenge to theory.

1.4 Observations

The primary scientific result of this work, detailed in chapter 5, is the observation of a reversible decrease in condensate fraction for a BEC in the disordered lattice, indicating disorder-induced boson localization. Most quantitative predictions from theoretical work to date on the DBH model have addressed superfluid fraction, not condensate fraction. However, in many theoretical approaches, condensate fraction can be computed straightforwardly and work with cold atoms has begun to prompt such studies [21, 48].

Given the plethora of theories for the DBH and their divergent predictions, this work was essentially agnostic to theory. Our goal is to provide measurements for a fully characterized

and physically realizable disordered potential as a benchmark and aid for theorists. It is our hope that these results can be used to test current theories and guide the development of refined ones. Indeed, we have been in contact with a number of theorists engaged in these activities [49, 50]; this work is ongoing.

1.5 Future Work on the DBH Model

We recently completed a series of measurements studying transport in disordered lattices. This is done by observing the response of the atoms to an external force, playing the role of an electric field in an electronic system. In condensed matter systems, it is the transport properties of a system which are used to distinguish superfluid and insulating states.

Much work with granular superconductors, for example, has been done with thin films, to which the 2D instead of 3D DBH is more applicable. Thus, studying the DBH in lower dimensions is an important extension of this work. By greatly strengthening one or two of the three optical lattice beams, sets of 2D or 1D systems, respectively, can be produced.

Two approaches are under consideration for limiting the range of atomic densities present in these experiments, which would simplify the interpretation of results: imaging only the central portion of the atom cloud and using a blue-detuned optical lattice. By allowing the lattice depth and overall confining potential to be separately varied, this second approach also opens the possibility of measuring compressibility [51], one property which distinguishes the BG and MI phases [21].

1.6 Quantum Simulation

One would likely not, for example, refer to superfluid helium in porous media as a quantum simulator for the DBH model. However, because of the ability of cold atom systems to model a wide range of systems, with relatively minor experimental changes and with few non-idealities, the term quantum simulation is sometimes applied to the realization of a given Hamiltonian. Chapter 6 discusses a narrower construction of this term, showing how cold atoms in a lattice can be used to simulate any interacting spin Hamiltonian by a series of discrete steps.

The viability of a number of systems for quantum simulation is an active research topic, in part because a universal quantum computer can simulate any quantum system [52]. At present, however, cold atoms in an optical lattice are arguably the best candidate for quantum simulation of many-body quantum systems. Compared to other systems with high levels of quantum control, such as NMR [53] and ion traps [54], cold atoms offer much greater scalability. The level of control available with atoms in a lattice, while less than with trapped ions, is sufficient for quantum simulation. Since the optical potential is essentially perfect over the small region where the atoms are present, if each lattice site can

be populated with a single atom, we achieve a defect-free lattice. We are, of course, free to introduce defects such as vacant lattice sites, impurities, or disordered potentials. We also have a great level of control over interactions, through control of the polarization and intensity of the lattice light. Both long and short range interactions can be created.

Another key advantage offered by cold atom systems is the wide array of measurement tools available. In addition to probes which are available in one form or another for solids, such as correlation and transport measurements, cold atoms provide powerful new probes, such as the possibility of directly imaging spins. A number of groups are working toward site-resolved imaging, using either very high-resolution optics [55, 56] or a long period lattice [57]. On the horizon is the ability to control the quantum state of atoms on individual lattice sites.

As part of this dissertation, the major technical components needed to implement the quantum simulation scheme described in chapter 6 were constructed and a number of calculations were performed. The implementation of the scheme will fall to future students.

1.7 Outline

In the following chapters, as outlined below, we detail the path from an empty room to an apparatus for studying cold atoms in optical lattices. We will focus on the construction of the apparatus and our realization of the disordered Bose-Hubbard model. The apparatus has also been used to study transport in the Bose-Hubbard model, work that led to the observation of phase-slip induced dissipation, familiar in superfluid and superconducting systems but previously unobserved in cold atoms systems. More recent work has studied transport in the DBH model. In the same room, a second apparatus is presently under construction to study degenerate Fermi gases in optical lattices. This system borrows much of the design of the first, which may be considered high praise (or general laziness).

- Chapter 2 gives an overview of the portion of apparatus devoted to creating condensates, including the laser system, vacuum system, magnetic trapping system, optical dipole trap, and imaging system.
- Chapter 3 provides details about the hardware and software making up the control system for the apparatus and a custom, FPGA-based, high-performance digital output hardware system.
- Chapter 4 features theoretical discussion of the behavior of cold atoms in an optical lattice and experimental details about the optical lattice system in our apparatus.
- Chapter 5 presents the main result of this thesis, measurements on BEC in a disordered 3D optical lattice, created by superimposing an optical speckle field on the optical lattice.

- Chapter 6 discusses the possibility of simulating interacting spin Hamiltonians in an optical lattice, and theoretical and technical work toward this goal.

Chapter 2

BEC Apparatus

2.1 Introduction

The portion of the experimental system needed to produce a BEC of ^{87}Rb atom is based in part on the apparatus described in [58]. The basic principles of elements such as the vacuum system, magneto-optic trap, transport system, Ioffe-Pritchard magnetic trap, and imaging system have been described in detail in a number of texts [10, 59–65]. Our description in the following sections will constitute a brief overview, focusing on unique problems and solutions and technical information relevant to operation and maintenance that may also be helpful to other workers in the field. The hybrid optical/magnetic trap developed to replace the Ioffe-Pritchard trap is less standard and will be discussed in greater detail.

2.2 Atomic Properties

The electronic structure of ^{87}Rb relevant to this work, particularly the D2 transition at 780.24 nm, is shown in figure 2.1. For the weak magnetic field regime typical in cold atom experiments, the atomic state is labeled by the total atomic spin F , i.e., the sum of the electron spin, electron orbital angular momentum, and nuclear spin, and its projection along the direction of the field, m_F . The Zeeman splitting of the ground state hyperfine levels is given to lowest order by

$$\Delta E = m_F g_F \mu_B B, \quad (2.1)$$

where B is the magnitude of the magnetic field (which defines the quantization axis for F), $\mu_B = e\hbar/2m_e$ is the Bohr magneton, and g_F is the Lande g-factor, given in figure 2.1. Since only minima in B can be created in free space, only atoms in low-field seeking states ($m_F g_F > 0$ so that ΔE increases with increasing B) can be trapped magnetically.

Optical transitions (considering only electric dipole transitions) are subject to the selection rules $\Delta F = 0, \pm 1$ and $\Delta m_F = 1$ for σ^+ polarized light (right-hand circular relative to the quantization axis, defined by the magnetic field), $\Delta m_F = -1$ for σ^- polarized light (left-hand circular relative to the quantization axis), and $\Delta m_F = 0$ for π polarized light (linear and parallel to the quantization axis).

At low temperatures, atomic scattering is purely s -wave and can be parameterized by

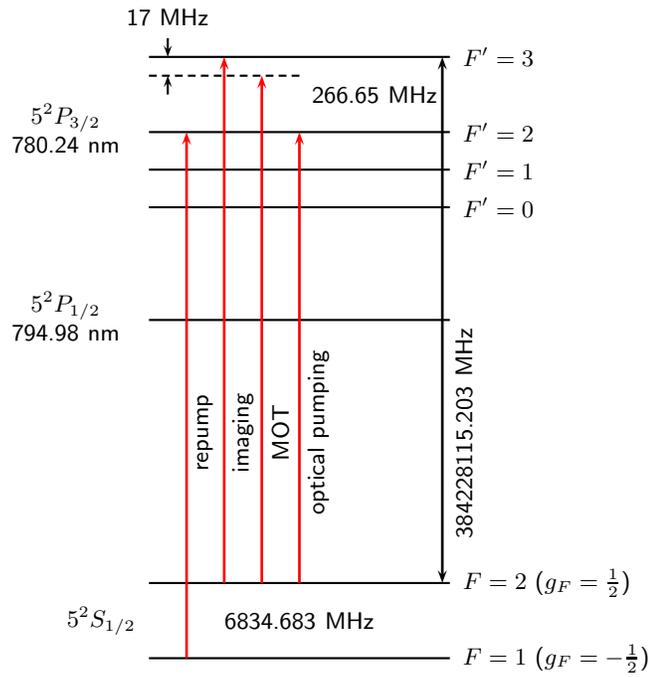


Figure 2.1 Electronic structure of ^{87}Rb . Laser frequencies used in the apparatus are shown for later reference. Note that F is typically used to label ground hyperfine states and F' to label excited hyperfine states. Only the low-field seeking ground states $|F = 1, m_F = -1\rangle$, $|F = 2, m_F = 2\rangle$, and $|F = 2, m_F = 1\rangle$ are magnetically trappable.

a single value, the s -wave scattering length, $a = 5.82$ nm for ^{87}Rb [66]. The effective inter-atomic potential is $U(\mathbf{r}, \mathbf{r}') = U_0\delta(\mathbf{r} - \mathbf{r}')$, where

$$U_0 = 4\pi\hbar^2 a/m. \quad (2.2)$$

The thermal de Broglie wavelength is

$$\lambda_{\text{dB}} = \frac{h}{\sqrt{2\pi m k_B T}} \quad (2.3)$$

When λ_{dB}^3 becomes on the order of the density n , the wavefunctions of atoms in the gas begin to overlap. This can be quantified by introducing the phase space density $\sigma = n\lambda_{\text{dB}}^3$, the average number of particles within a volume equal to the de Broglie wavelength cubed. For ^{87}Rb at a typical density near condensation of 10^{14} cm^{-3} , the phase space density becomes one at a temperature of 75 nK. For a uniform system, the BEC transition occurs at $\sigma = 2.6$; the transition temperature is denoted T_c . An excellent discussion of the BEC transition in harmonic traps can be found in [10].

2.3 Overview of BEC Apparatus

The apparatus operates on a cycle of one to two minutes. At the beginning of each cycle, several billion ^{87}Rb atoms are collected and cooled to less than 100 μK in a six-beam magneto-optic trap (MOT) [67] before being optically pumped to the $|F = 1, m_F = -1\rangle$ state and transferred to a spherical quadrupole magnetic trap, formed by two coils sharing an axis and carrying current in opposite directions (anti-Helmholtz configuration). The coils which create the quadrupole magnetic trap are moved on a linear translation stage approximately 0.75 meters to a lower pressure region of the vacuum system. Here, the atoms are transferred into a different trap (see sections 2.9 and 2.12) and forced evaporative cooling is used to produce a BEC. This is the starting point for the work described in chapter 4, where the BEC is loaded into an optical lattice.

An external cavity diode laser (ECDL) stabilized using a rubidium vapor cell provides an optical frequency reference. It is normally operated -17 MHz from the $F = 2 \rightarrow F' = 3$ transition, see figure 2.1. This master ECDL also provides light for imaging atoms; the light is switched on and off and frequency shifted onto resonance by a pair of acousto-optic modulators (AOMs). Another AOM shifts a portion of the light onto resonance with the $F = 2 \rightarrow F' = 2$ transition for optical pumping. The laser light used for the MOT trapping beams is supplied by three high-power laser diodes injection locked to the master ECDL. A second ECDL (the repump ECDL) referenced to the master ECDL provides light resonant with the $F = 1 \rightarrow F' = 2$ transition needed for the MOT. Figure 2.2 shows the laser system schematically.

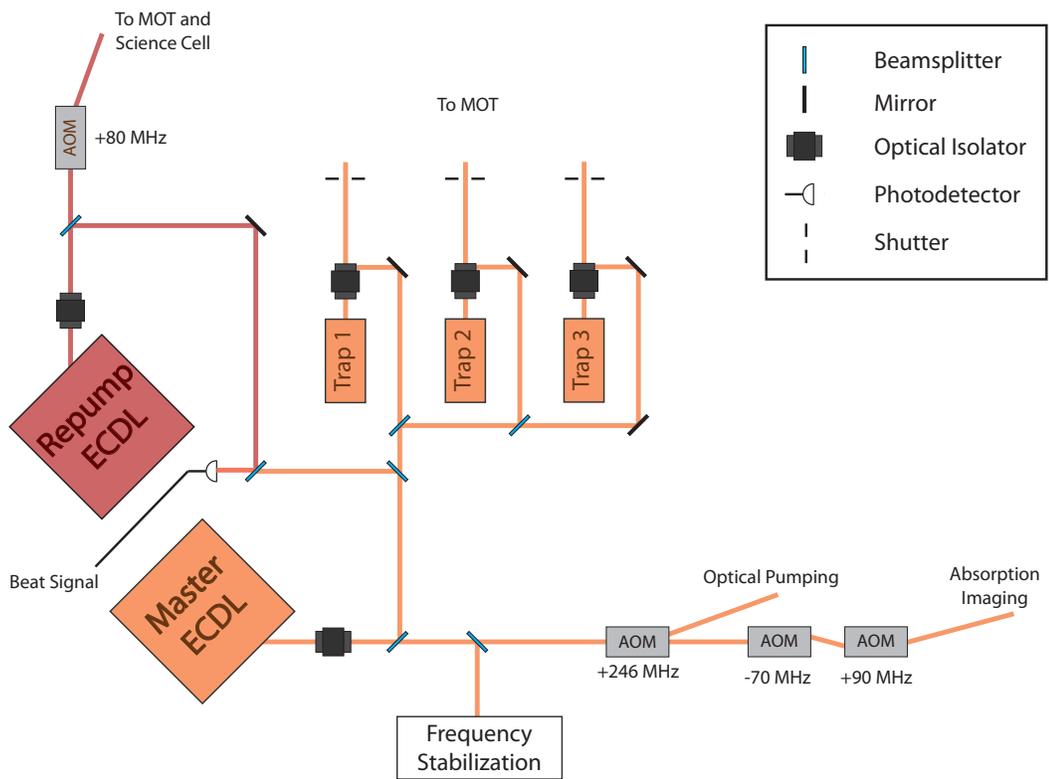


Figure 2.2 Schematic of diode laser system.

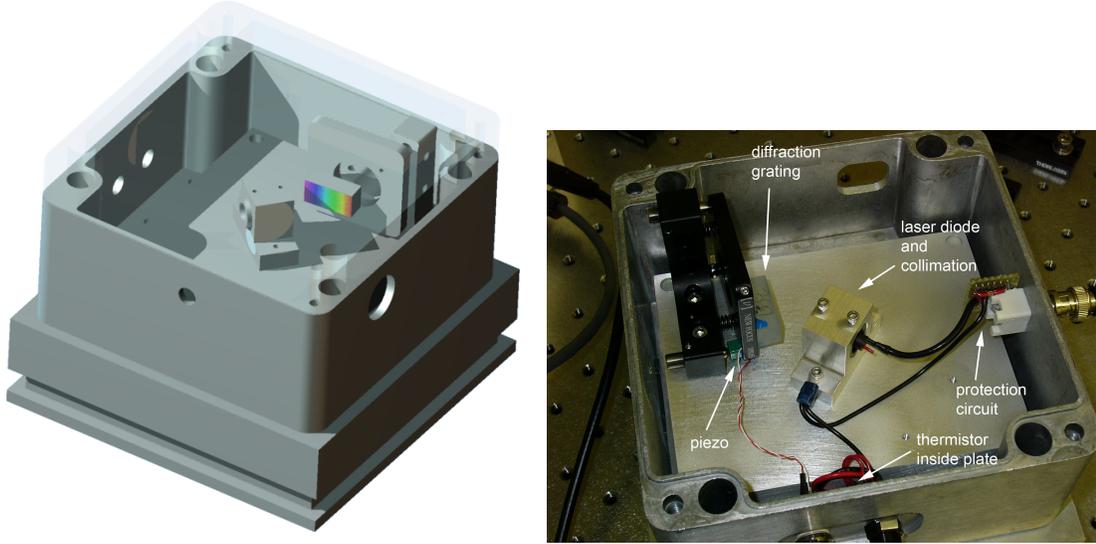


Figure 2.3 CAD model and photograph of external cavity diode laser design. A thermoelectric element between the enclosure and the plate to which the ECDL components are mounted is used to adjust the temperature of the plate.

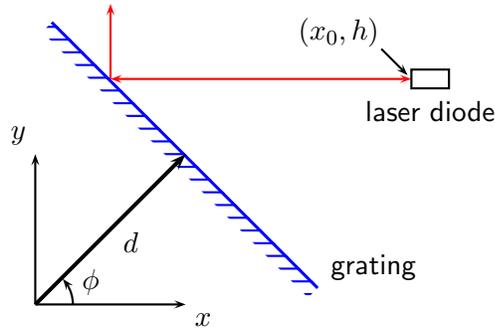


Figure 2.4 Schematic of improved ECDL design. The grating is parameterized by $y \sin \phi + x \cos \phi = d$ and the external cavity length is $l = x_0 - (d - h \sin \phi) / \cos \phi$. The wavelength selected by the grating angle is $\lambda = 2a \sin \phi$, where a is the grating period, while the wavelength selected by the cavity length is $\lambda = 2l/m$, where m is the number of waves in the cavity.

2.4 Laser System

2.4.1 External Cavity Diode Lasers

One of the key factors in the explosive growth of cold atom research following the realization of the first atomic BEC has been the ability to use low cost diode lasers for cooling and trapping a number of different atomic species [68]. The situation is especially fortuitous for rubidium. The wavelength of the rubidium D2 transition at 780 nm is very close to the nominal wavelength of diode lasers used in compact disk drives, 785 nm, allowing the use of off-the-shell parts. Furthermore, the push toward faster CD writers has resulted in the

development of 785 nm diode lasers capable of over 100 mW continuous output power.

The free-running linewidth (several MHz to several tens of MHz) of a 785 nm diode laser must be reduced to < 1 MHz and the center frequency stabilized to within a few hundred kHz for cooling and trapping rubidium. This is usually accomplished by placing the diode laser in an external cavity formed by a diffraction grating. Our design uses the Littrow configuration [69] (vs. Littman-Metcalf [70]), which allows for higher output power but has the disadvantage that the output direction changes slightly as the laser is tuned. The grating is installed so that the first diffracted order (15–25%) is sent back to the diode laser. By adjusting the angle of the grating, the wavelength of the light that will be sent back to the diode laser varies, tuning the cavity. The zeroth order (undiffracted) light is available for use. A small amount of light is lost to scattering and other diffracted orders. Because the optical gain of the diode laser is very high, a high Q cavity is not needed. Indeed, great care must be taken to prevent reflections from other optics from reaching the diode laser; such feedback to the diode can cause unstable behavior. Typically, the ECDL output passes first through a Faraday isolator to suppress unwanted optical feedback.

Competition between modes of the external cavity and the diode’s internal cavity can cause instability, manifested by sudden jumps in the laser frequency as it is tuned. This problem can be eliminated by applying an anti-reflection coating to the output facet of the diode, although at significant expense.

We use Arima ADL-78901TX diode lasers from Roithner Lasertechnik, operated at 50 mW; previously we used lower-powered Hitachi HL7851G diodes. When used with an external cavity, the maximum output power of the diode must be derated because of the increased intensity at the diode facet (thermal damage to the facet due to excess intensity is the primary cause of diode laser failure). A protection circuit consisting of several small signal diodes (1N914) protects the diode laser from overvoltage and reverse voltages.

Initial alignment of an ECDL is performed by first adjusting the vertical tilt of the grating with the laser near its lasing threshold. When the grating is aligned, the additional feedback will cause a large increase in the output power. The horizontal tilt (tuning) is then adjusted with the beam passing through a rubidium vapor cell. When fluorescence is seen in the vapor cell, the laser is within a few GHz of the desired frequency.

Figure 2.3 shows the ECDL design used in the apparatus. The piezoelectric actuator (piezo for short) tunes the ECDL by changing the angle of the grating and its distance from the diode laser, i.e., the cavity length. It is also possible to tune the laser with small changes in the current, as discussed below. The temperature of the laser and cavity also affects the frequency; it is monitored by a thermistor and controlled with a thermoelectric element (Melcor CP 1.4-71-10L-W6). Commercial parts are used for the laser current source (Thorlabs LDC202) and temperature controller (Thorlabs TED200).

A significant source of laser frequency noise was a mechanical resonance of the mirror mount (New Focus 9805-K) at 1150 Hz. It was found that this noise could be suppressed

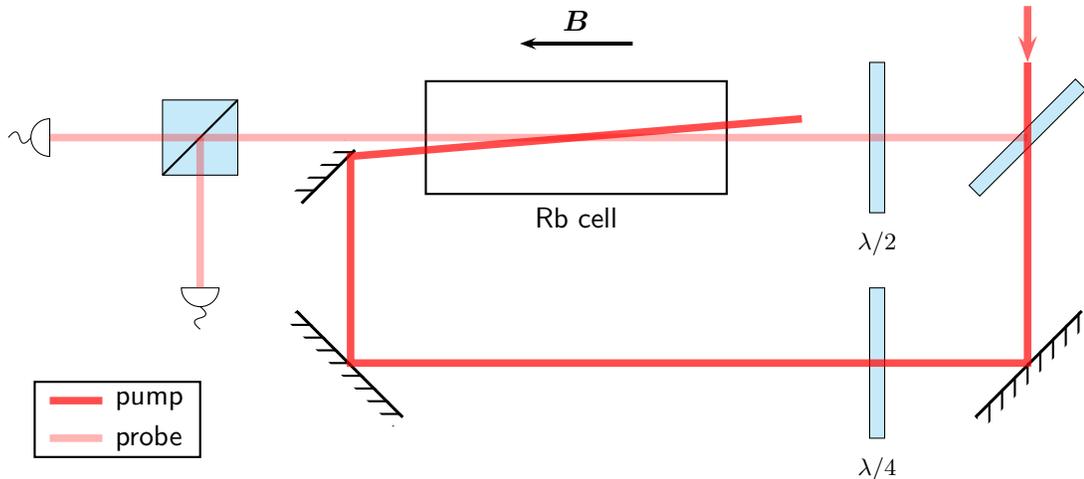


Figure 2.5 Polarization spectroscopy setup for laser stabilization. The pump beam is circularly polarized while the probe beam is linearly polarized.

by damping the mirror mount with foam pressed between the mount and the enclosure. A newer design using a different mirror mount (Newport U100-P) with less stiff springs displays a resonance at around 850 Hz.

The design in figure 2.3 contains a significant flaw: the tuning effect from the angle of the grating opposes the effect from the change in the distance of the grating from the diode. Recently, an improved version was built in which the pivot point of the grating is flipped. Here, the geometry, figure 2.4, is such that both effects should have the same sign and slope. This design shows a significantly larger mode-hop-free tuning range (about 8–10 GHz), especially when the diode laser current was adjusted in proportion to the piezo tuning voltage (feedforward). Using the same feedforward technique with the older design gives little to no improvement.

2.4.2 Laser Stabilization

Polarization Spectroscopy Lock

The master ECDL is stabilized by a polarization spectroscopy lock, figure 2.5, which uses pump-induced birefringence near a resonance to rotate the polarization of a probe beam [71, 72]. An AOM shifts the frequency of the light entering the locking setup to be resonant with the $F' = 2 \rightarrow F' = 3$ crossover peak when the laser is locked. The frequency of the AOM can be adjusted to tune the laser; the AOM is double-passed so that the beam's direction is independent of the AOM frequency. Direct digital synthesizers (section 3.2.3) are used as RF sources.

Figure 2.5 shows the optics used for the polarization spectroscopy lock. The σ^+ pump beam optically pumps atoms to Zeeman states with more positive m_F . The resulting

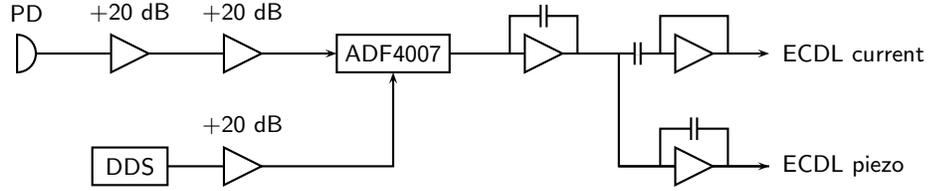


Figure 2.6 Phase-locked loop for stabilizing repump ECDL relative to master ECDL. The beat signal is detected by the photodiode (PD) at left, amplified, and compared to a reference frequency by an IC (Analog Devices ADF4007). The resulting error signal is filtered and then fed back to the repump ECDL piezo and current via a crossover network.

population imbalance results in birefringence, i.e., a difference in indices of refraction for σ^+ and σ^- light. One way to see this is that the population imbalance clearly results in different absorption coefficients for σ^+ and σ^- light, which implies different refractive indices via the Kramers-Kronig relation. The linearly polarized probe beam consists of equal amounts of σ^+ and σ^- light; the birefringence induced change of phase between these components represents a rotation of the linear polarization angle, changing the ratio of power incident on the photodetectors following the polarizing beam splitter. The signals from the two photodetectors are subtracted to produce an error signal which is amplified and fed back via a PI servo to the ECDL piezo for DC–100 Hz and to the diode laser current above 100 Hz.

Previously, saturated absorption spectroscopy (see [73] or [60]) was used for stabilization. This method is less sensitive to magnetic fields and less susceptible to drift, but is more difficult to implement because modulation and synchronous detection are required. Frequency modulation of the RF signal driving the AOM creates unwanted amplitude modulation which in turn creates offsets in the error signal.

Optical Phase-locked Loop

The repump ECDL is locked at an adjustable offset frequency from the master ECDL using an optical phase-locked loop (PLL), figure 2.6. About 1 mW of light with the same linear polarization from each ECDL is coincident on a high-speed unamplified GaAs photodetector (Electrooptic Technologies ET-4000). The photodetector is DC-coupled, which aids in initial alignment, a difficult task because high frequency operation necessitates a very small photodiode. The resulting weak electrical signal at the frequency difference between the lasers is amplified by two microwave amplifiers. A PLL IC (Analog Devices ADF4007) divides the microwave signal frequency by 64 and compares it to a reference frequency, which is adjusted to change the frequency offset between the lasers. The resulting error signal is fed back to the ECDL as described above.

The repump ECDL is usually kept on resonance throughout the experimental sequence, while the master ECDL must be tuned to several different frequencies. Because the optical

PLL can provide GHz of tuning (with better speed) while the saturated absorption lock a few tens of MHz at best, an improved design would interchange the stabilization setups.

2.4.3 Trapping Lasers

A single ECDL cannot provide enough light for a sizable MOT, so light from the master ECDL is used to injection lock [62] three high-power diode lasers (Sony GH0781JA2C). Each laser is tuned near the correct frequency by cooling it to 0–20°C (this number varies from part to part) and monitored with a Fabry-Perot spectrum analyzer as the current is adjusted slightly to achieve injection locking. Desiccant is placed inside the laser enclosure for operation at low temperature to prevent water condensation. Roughly 0.5 mW of light from the master ECDL is sent into the rejection port of the laser’s Faraday isolator which, when aligned, phase locks the laser to the master ECDL. The lasers are operated at 80–90 mW, well below their specification of 120 mW. Commercial OEM current and temperature controllers (Thorlabs ITC102) are used.

We have experienced failures of one or more of these diodes roughly every six months. Although we have been unable to determine the cause, we believe it is electrical in nature because often two or three diodes fail simultaneously, most often in winter. In response, we added an LC filter to the protection circuit since, unlike for the ECDLs, no feedback to the current is needed for the trapping lasers.

2.4.4 Beamlines

Once atoms are loaded into the magnetic trap and evaporative cooling begins, they must be carefully shielded from stray resonant light. Absorption of a single photon results in heating and a significant probability of ejection from the magnetic trap. Commercial shutters (Uniblitz LS3T2, LS6T2 with Uniblitz VMM-D3 driver) are used to block beams not in use. However, in response to observed heating and loss due to stray light, we completely enclosed the region of the optical table containing the laser system with black curtains and a cardboard sheet.

Acousto-optic modulators, which use a traveling RF acoustic wave in a crystal to diffract an incoming beam, are used throughout the apparatus to control light intensity (by changing the RF power), to steer beams (by changing the RF frequency), and to shift the frequency of beams (diffraction from the traveling wave results in a Doppler shift). Most AOMs in the apparatus require around 1 W of RF power to achieve maximum diffraction efficiency. We use cable TV distribution amplifiers (Freescale MHW6342T) in home-built enclosures as RF power amplifiers. Originally, the amplifiers were cooled by forced air, but subsequently passive cooling with a larger heat sink was found to be sufficient.

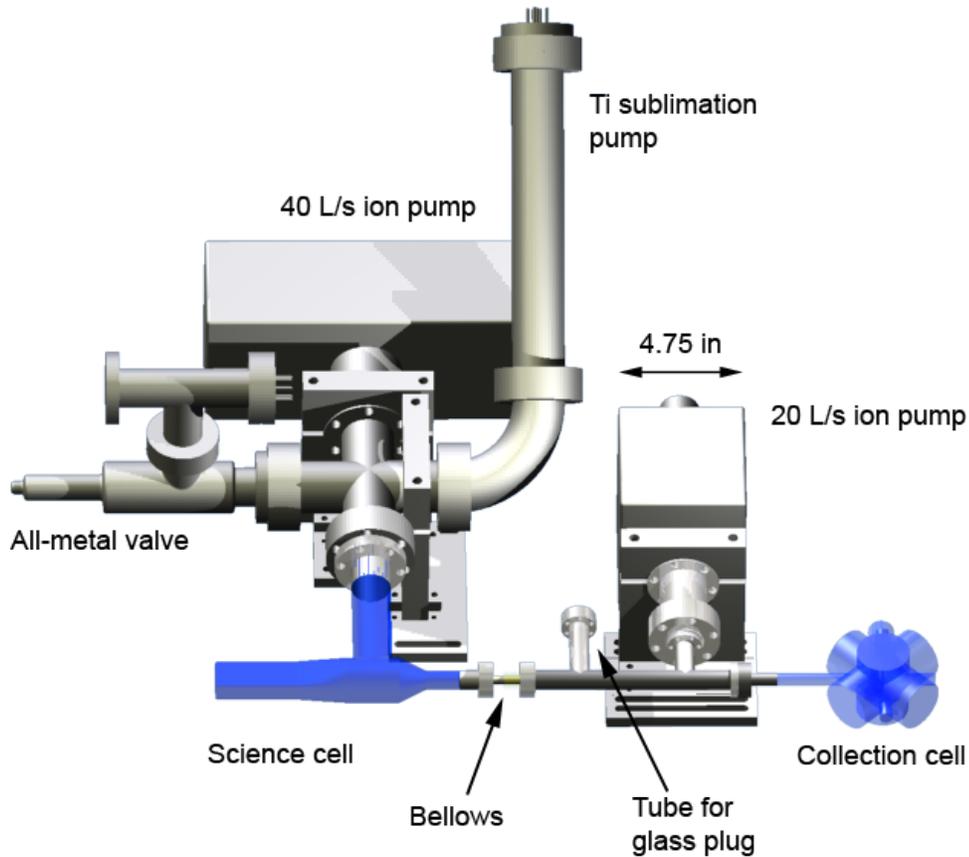


Figure 2.7 CAD design for ultra-high-vacuum system. The ion pumps are Varian StarCell pumps. The glass cells (blue) were custom made by Technical Glass.

2.5 Vacuum System

Figure 2.7 shows the core component of the apparatus, the vacuum system. The glass cells (blue in the figure) were custom made by Technical Glass Inc. (Boulder, CO). The collection cell contains four rubidium dispensers (only one is used at a time) which produce a room-temperature vapor from which atoms are captured in a magneto-optic trap (MOT), section 2.7. The vapor pressure of rubidium in the cell must be high enough (on the order of 10^{-9} torr) to collect several billion atoms in the MOT in a few tens of seconds, so it is important to limit the conductance between the collection cell and the 20 L/s ion pump. At this vapor pressure, the lifetime of magnetically trapped atoms in the collection cell is at most a few seconds due to collisions with room-temperature rubidium atoms. Thus, atoms captured in the MOT must be transferred to a lower pressure region, the science cell, shown in more detail in figure 2.8, for evaporative cooling to form BEC. The long, thin tube between the science cell and collection cell [58] and the wide tube from the science cell to the 40 L/s ion pump allow for a large pressure differential between the two cells.

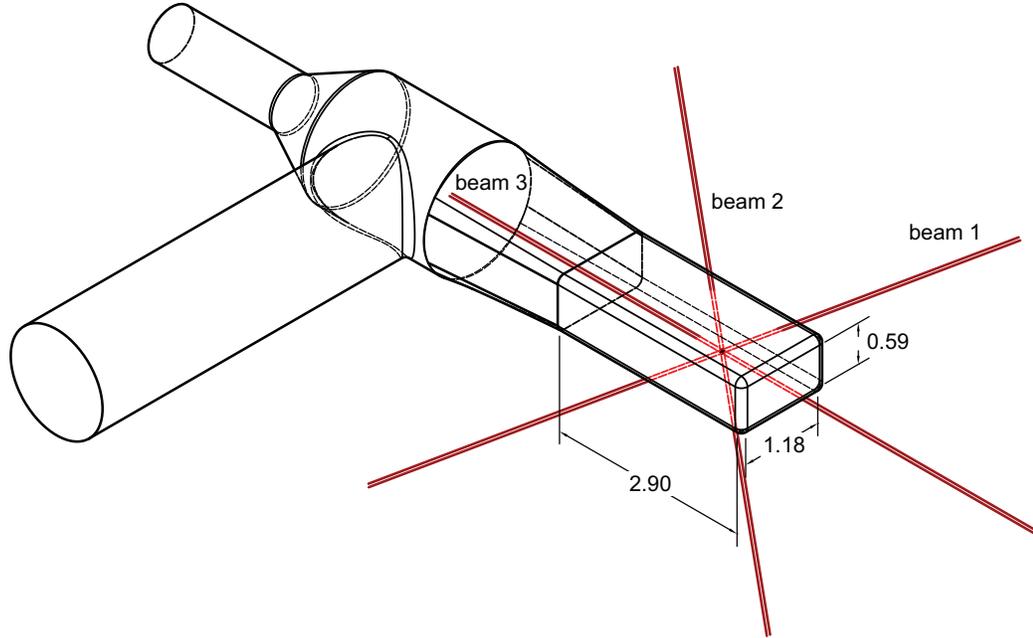


Figure 2.8 Isometric view from above of science cell; dimensions are in inches. The rectangular cross-section enables very high numerical aperture optical access from the top and bottom. This figure also shows the paths of the optical lattice beams (chapter 4).

Magnetically trapped atoms are moved to the science cell by physically moving the trapping coils (see below). When the apparatus is not in use, a glass encapsulated steel plug is moved by hand using a small magnet into the transfer tube to further reduce the conductance to the science cell.

The presence of two connections to the glass science cell required the use of a welded bellows between the science cell and the transfer tube to protect the science cell from mechanical stress. As discussed in section 2.8.4, the bellows interfere with the transfer of atoms to the science cell. In the vacuum system for the fermion apparatus, the tee-section of the science cell was replaced with a metal tee, eliminating the need for a bellows.

2.5.1 Assembly and Bake-out

Prior to assembly, metal components of the vacuum system were cleaned with deionized water and detergent in an ultrasonic cleaner and rinsed with water, acetone, and methanol, in that order. The pieces were then baked at 450°C in air for six hours to drive off gases and other volatiles and form an oxide layer to reduce outgassing.

After assembly, the vacuum system was baked at 300°C (limited by the rubidium dispensers) for five days while connected to a turbopump. The all-metal valve was then closed, the turbopump was disconnected, and each filament of the Ti sublimation pump was fired for 2 minutes (the Ti sublimation pump was also fired before the bake to degas it). The ion pumps were turned on and the system was baked for an additional two days at 200°C

(limited by the ion pump cables). After cooling the vacuum system to room temperature, the pressure was below the lower detection limit of the ion pump controller (3.6×10^{-10} torr).

2.6 Quadrupole Coils

2.6.1 Overview

The magnetic field required for the MOT and for transport of atoms from the collection cell to the science cell is provided by a pair of current-carrying coils connected in series in an anti-Helmholtz configuration to produce a quadrupole magnetic field.

Each of the coils consists of 48 (12×4) turns of 0.16 in. square, hollow copper tubing wrapped in Kapton plastic. They are mounted on anodized aluminum holders and cooled by passing water through the tubing itself. The cooling water system operates as a closed loop, with a heat exchanger to the campus chilled water system; the temperature is maintained at 62°C to prevent condensation on piping.

2.6.2 Translation Stage

The quadrupole coils are attached to a motorized linear translation system (Parker Hannifin; controller: GV6K-U3R, motor: CM231AR, table: 406700XRMS, 700 mm travel). The coils together with their carrier are referred to as the “cart”. The translation system uses a rotary resolver in a closed-loop system to provide $5 \mu\text{m}$ reproducibility in position. A small custom program running on the controller waits for a digital trigger and moves the cart between the MOT position and the science cell position (648 mm) in about one second. The cart is capable of moving faster, but one second is chosen to prevent excessive recoil of the optical table.

2.6.3 Current Servo

Figure 2.9 shows the quadrupole coil circuit. The current through the coils is stabilized and controlled using a simple PI type servo [74]. A digital signal selects the source of the current setpoint: one of several fixed voltage dividers or an analog input. Current is supplied from a 21 V, 240 A power supply (Agilent 6682A) and carried with 2/0 type stranded cable (0.5 in conductor diameter). Closed-loop Hall probe current sensors (F.W. Bell CLN-1000) are used to measure the current. The servo drives the gates of a bank of two power MOSFETs (Microsemi/APT APT10M07JVR), changing the drain-source current since the power supply is operated in constant voltage mode. The FETs are mounted to a water-cooled copper plate; the plate was home-built, however, water-cooled aluminum plates can be purchased inexpensively. Transient voltage suppressor Zener diodes ($2 \times$ Littelfuse

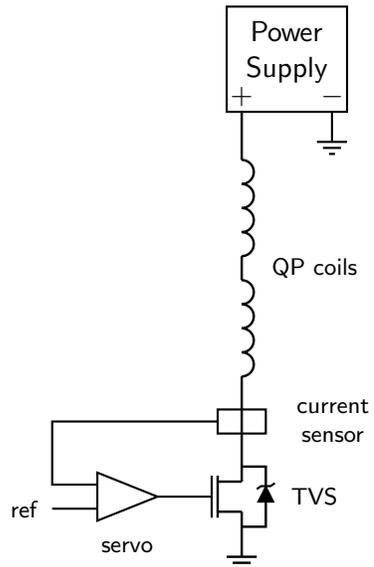


Figure 2.9 Schematic of quadrupole coil circuit.

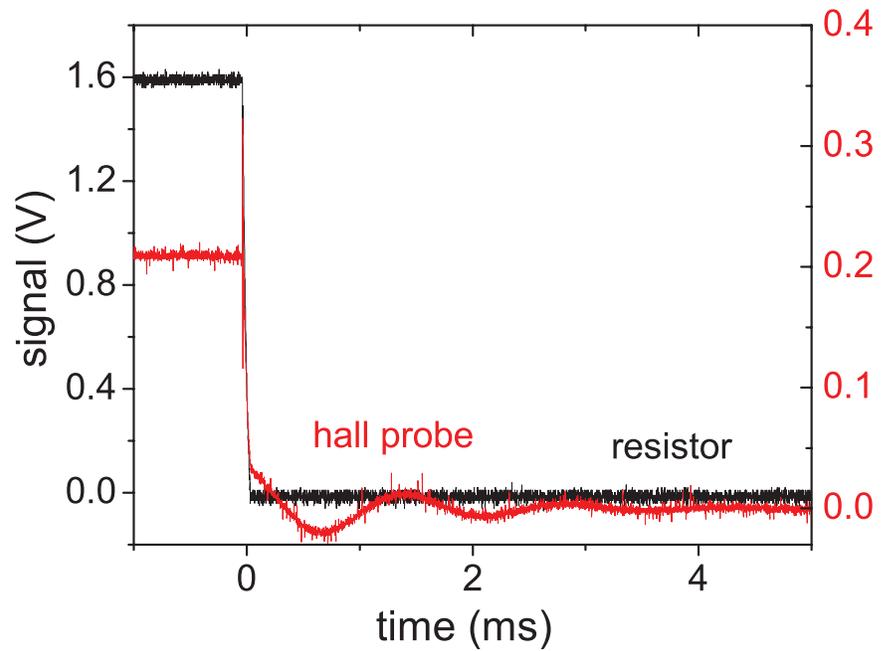


Figure 2.10 Ringing of closed-loop Hall probe current sensor in response to a sudden change in current. No ringing is detectable in the voltage drop across a sense resistor.

5KP58) are used to dissipate energy when the coils are turned off suddenly, although this is likely unnecessary.

The current sensors exhibit ringing after rapid changes in current; see figure 2.10. The sensors track the change itself faithfully and the frequency of the ringing is much slower than the bandwidth of the sensors. The manufacturer is aware of the problem, but was unable to provide any kind of explanation. Because the ringing was dependent on the position of the conductor within the aperture of the sensor, it was speculated that it could be related to some process within the soft magnetic core used for concentrating flux at the Hall sensor. Although not an issue when turning off the coils, the ringing is written onto the servoed current when the coils are switched on suddenly.

2.6.4 Rapid Turn-on Circuit

As originally built, the electronics for the quadrupole coils included a rapid turn-on circuit containing a capacitor ($60 \mu\text{F}$, 1200 V) charged to several hundred volts by a high-voltage DC/DC converter. The capacitor was discharged into the coils through a silicon-controlled rectifier activated by an optically isolated digital signal. A diode was used to prevent current from flowing backwards through the power supply during the discharge.

The purpose of the rapid turn-on circuit was to prevent the cloud from falling under gravity and expanding during the otherwise long turn-on of the coils. However, in practice we were not able to observe any difference in the number of atoms captured or their temperature with and without the rapid turn-on circuit, so the circuit is not used.

2.6.5 Interlock

The temperature of each of the coils is monitored with a thermistor. If either temperature exceeds a threshold, or if the cooling water flow to the apparatus drops below a threshold, a simple interlock circuit disables all high-power power supplies. The system is designed so that the power supplies are disabled if the interlock loses power or is disconnected from the power supplies or the water flow meter.

2.7 MOT

Rubidium atoms are captured, cooled, and trapped from the room temperature vapor in the collection cell using a standard magneto-optic trap [59]. Within a few tens of seconds, several 10^9 atoms are collected and cooled to $< 100 \mu\text{K}$. The optimal magnetic field gradient was determined to be 10 G/cm and the optimal detuning of the trapping beams -17 MHz .

About 10 mW of light from the repump ECDL is combined with one of the MOT beams using a PBS. Originally, we attempted to provide repump light in all three MOT beams by applying 6.8 GHz modulation, via a bias-tee, to the current for the trapping diode

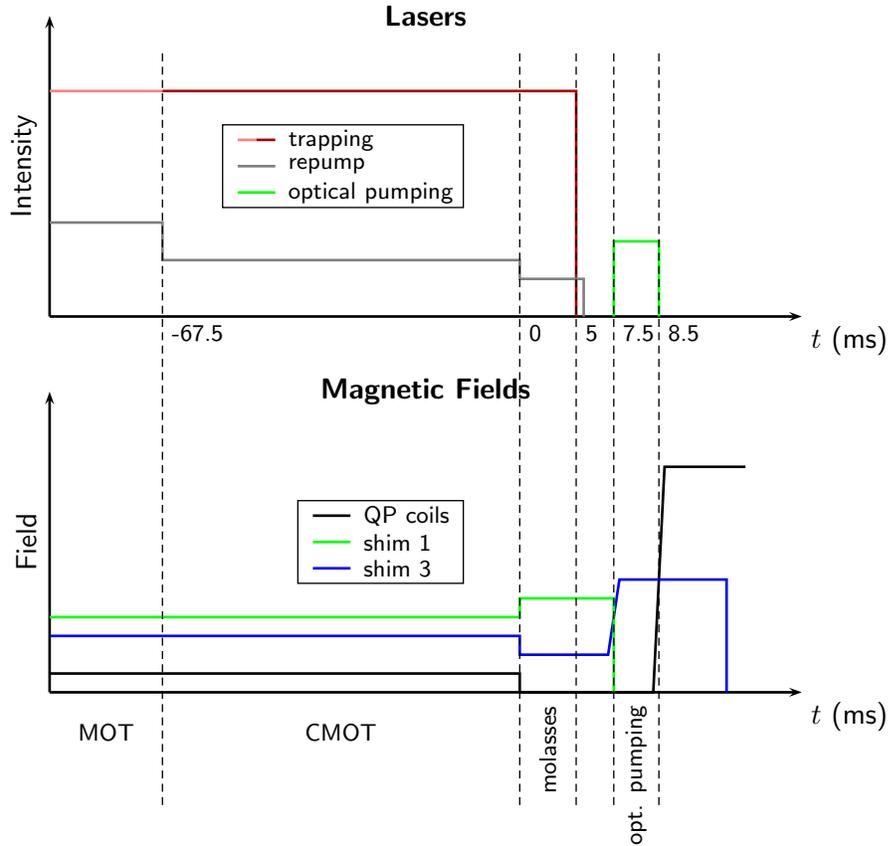


Figure 2.11 Sequence for transferring atoms into quadrupole trap from the MOT. The axes are not to scale. Note that the trapping lasers are detuned further at the start of the CMOT stage (indicated by the change in trace color).

lasers. Due to the dependence of the laser frequency on current, the modulation produces sidebands, one of which can be used for repumping [75]. However, it was found that the lasers would not remain injection locked when the modulation was turned off. This was most likely due to thermal effects, since several tens of mW at 6.8 GHz was needed to produce sufficient sideband power.

Shim coils are used to cancel unwanted stray magnetic fields from, e.g., the ion pump and the optical table. In theory, the shim coils should be set so that the center of the MOT matches the center of the quadrupole trap. In practice, the shim coils are adjusted to maximize the number of atoms after evaporation in the science cell.

2.8 Capture

Figure 2.11 shows the sequence for transferring atoms from the MOT to the quadrupole trap, discussed in more detail below. The laser frequencies used during this process are

shown in figure 2.1.

2.8.1 CMOT

The atom cloud should be made as localized and dense as possible to minimize the amount of energy imparted by loading into the quadrupole trap. The limiting factor to the density of the MOT is radiation pressure from rescattering of photons scattered by other atoms in the MOT. So, a compressed MOT (CMOT) stage is used in which the detuning of the trapping beams is increased from 17 MHz to 50 MHz (see figure 2.1) and the repump power is reduced to 75 μ W. The atoms then spend more time in the dark $F = 1$ ground state and radiation pressure is reduced. The CMOT stage decreases the number of atoms captured by around 10%, but this is more than compensated for by the lower temperature of the remaining atoms after loading into the quadrupole trap. Very recently, better performance was found when the repump power was not reduced during the CMOT stage, although this has not been thoroughly investigated.

2.8.2 Optical Molasses

To further cool the motion of the atoms, a step of polarization gradient cooling [76] is performed. The same power and detunings are used for the trapping and repump beams as during the CMOT stage, but the quadrupole field is turned off and different settings are used for the shim coils. For testing, the molasses stage can be left on for several seconds and the expansion of the atom cloud monitored by eye, using a CCD camera or infrared viewer. The shim coils settings for the molasses stage are adjusted to make this expansion as slow and isotropic as possible.

2.8.3 Optical Pumping

To transfer atoms to $|F = 1, m_F = -1\rangle$, the cloud is illuminated with σ^- polarized light resonant with the $F = 2 \rightarrow F' = 2$ transition. When well-optimized, the transfer efficiency (as measured by turning the MOT back on after 100 ms of hold time in the quadrupole trap) is 60%. To populate the $|F = 2, m_F = 2\rangle$ state, σ^+ light 30 MHz red-detuned from the $F = 2 \rightarrow F' = 2$ transition is used along with $F = 1 \rightarrow F' = 2$ repump light. A single shim coil is used to generate the quantization field (the quadrupole field remains off).

2.8.4 Transport to Science Cell

After optical pumping, the quadrupole coils are turned on to 192 A in 10 ms (10% to 90%), limited by the inductance of the coils. The cart is then moved to the science cell. By moving the cart various distances along the transfer tube then moving it back to the collection cell and turning the MOT back on, it was found that most atoms were lost in the vicinity of

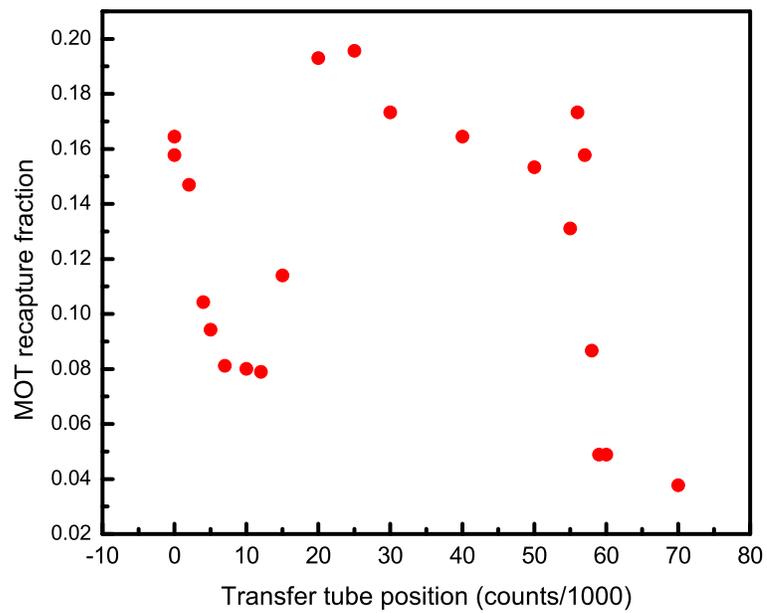


Figure 2.12 Fraction of atoms recaptured in MOT after transport to given position in transfer tube and return. The initial dip is due to stray light, which is blocked once the atoms are inside the transfer tube. The loss around 58000 counts, just at the entrance to the welded bellows connecting the transfer tube to the science cell, is believed to result from stray magnetic fields.

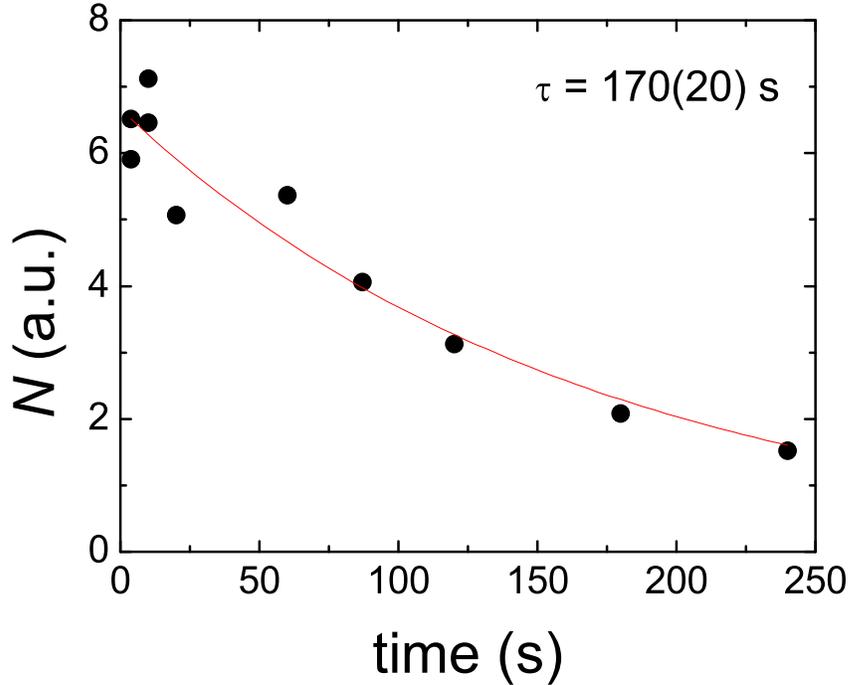


Figure 2.13 Lifetime of trapped atoms in the science cell. The red trace is an exponential fit, which yields a $1/e$ decay time τ of nearly three minutes. The temperature of the atoms is several hundred μK , high enough to make Majorana losses from the quadrupole trap negligible.

the welded bellows; see figure 2.12. We speculated that the loss may be due to magnetic fields from the bellows, as a significant field was measured near the exterior of the bellows. A coil was installed to try to cancel this field. It yielded only a small improvement but was left in place.

The ultimate measure of the quality of the vacuum in the science cell is the lifetime of trapped atoms there. Figure 2.13 shows the loss of atoms from the quadrupole trap in the science cell; the $1/e$ decay time is nearly three minutes, more than sufficient to allow evaporative cooling to BEC.

2.9 Ioffe-Pritchard Magnetic Trap

2.9.1 Overview

For the work described in chapter 5, BECs were created using RF evaporative cooling in a standard Ioffe-Pritchard (IP) magnetic trap [10, 77], figure 2.14. Taking the z axis through the center of the pinch coils (connected in a Helmholtz configuration) and parallel to the Ioffe bars (connected so that current in neighboring sets of bars flows in opposite directions),

the field magnitude near the center of the trap is

$$B = -b_{\text{bias}} + b_0 + \frac{1}{2} \left(\frac{c_1^2}{b_0 - b_{\text{bias}}} - b_2 \right) r^2 + b_2 z^2 + \dots, \quad (2.4)$$

where b_0 and b_2 are proportional to the current in the pinch coils and c_1 is proportional to the current in the Ioffe bars. A pair of larger coils, also connected in a Helmholtz configuration, produce a bias field b_{bias} to partially cancel the field b_0 from the pinch coils. The radial and axial trapping frequencies near the center of the trap are, respectively,

$$\omega_r = \sqrt{\frac{m_F g_F \mu_B}{m} \left(\frac{c_1^2}{b_0 - b_{\text{bias}}} - b_2 \right)}, \quad (2.5a)$$

$$\omega_z = \sqrt{\frac{2m_F g_F \mu_B}{m} b_2}. \quad (2.5b)$$

Notice that, holding all other trap currents constant, the radial frequency can be adjusted through a wide range by changing b_{bias} . In the radial direction, the trap is harmonic only in the region where $b_0 - b_{\text{bias}}$ is comparable to or larger than the field from the Ioffe bars; outside this region, the linear gradient of the Ioffe bars dominates (see figure 2.19). During evaporation, we typically set $b_0 - b_{\text{bias}}$ to around 2 gauss; for $|F = 1, m_F = -1\rangle$ this gave trapping frequencies of $220 \text{ Hz} \times 220 \text{ Hz} \times 14.5 \text{ Hz}$ and a final RF evaporation frequency around 1.4 MHz.

2.9.2 Construction

The Ioffe bars consisted of four sets of three 1/8 inch square pieces of hollow copper tubing (figure 2.15) electrically insulated with 0.0005 inch thick heat-shrink tubing and connected in a single series circuit with solid pieces of copper soldered between the copper tubes. To provide cooling water to a set of three tubes from a single 1/4 inch line without forming an electrical path between them, square holes were punched in Swagelok-type cap-end brass fluid connectors. A thin layer of epoxy (Tra-con BA-2151; matched thermal expansion coefficient) was applied to the inside of the holes and allowed to dry to ensure electrical insulation. The connectors were then glued on the end of the set of tubes with the same type of epoxy. To our surprise, none of these connections have failed to date (over three years), despite operation at 150 psi.

The pinch coils were wound from 1/8 inch circular, hollow copper tubing (4 turns) insulated with heat-shrink tubing. It was necessary to have the stock copper tubing annealed in a vacuum furnace to make it pliable enough to wind as tightly as required. The bias coils consist of 2×5 turns of the same hollow tubing used for the quadrupole coils.

As with the quadrupole coils, the temperature of the IP trap is monitored at several points with thermistors and interlocked to the power supplies.

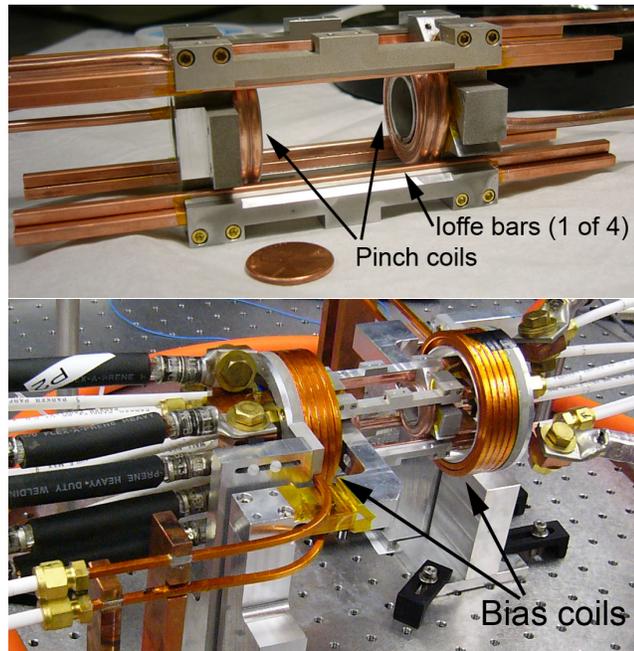


Figure 2.14 Core portion of Ioffe-Pritchard magnetic trap (top). Current flows in opposite directions in neighboring sets of bars. Assembled trap showing bias coils and connections for current and cooling water (bottom).

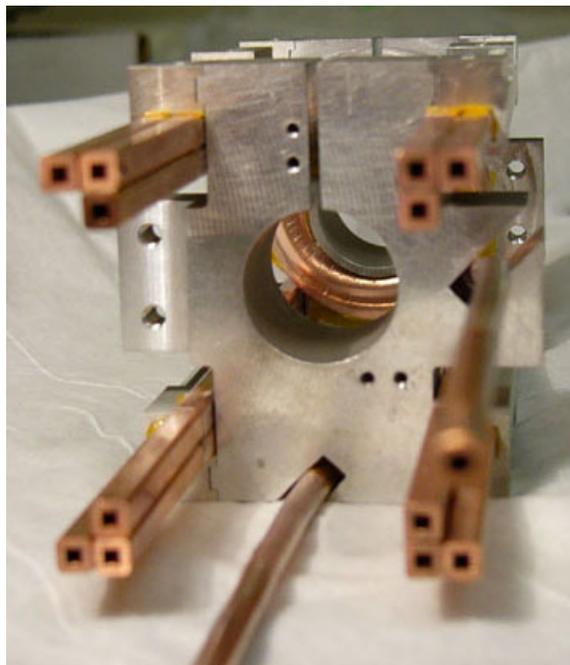


Figure 2.15 End-on view of Ioffe bars during construction.

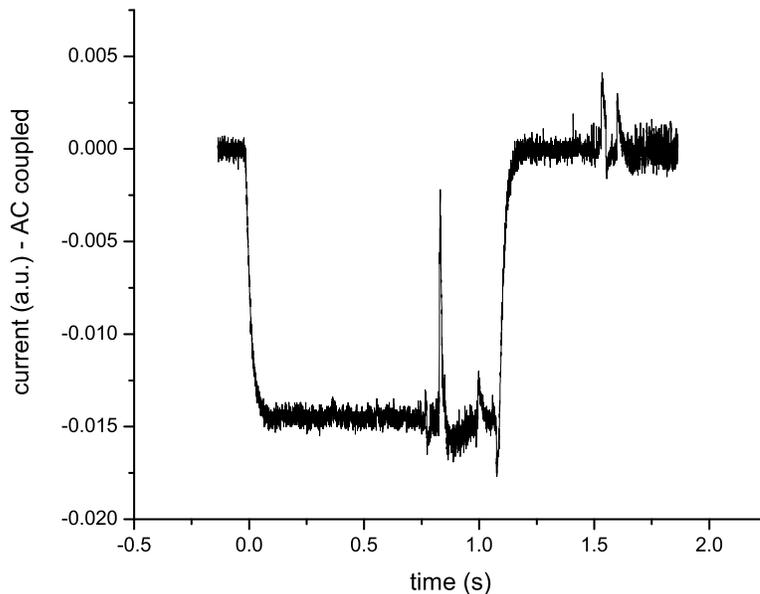


Figure 2.16 Sudden jumps in current through the Ioffe bars during and after a slow ramp down of the setpoint current. Note that the signal here is AC-coupled, so the ramp appears as a constant offset from zero between 0 s and 1.1 s.

Numerous difficulties were encountered during the construction of the Ioffe bars, particularly with the solder joints to the copper pieces connecting the individual tubes. We believe that thermal and mechanical (due to the strong magnetic force between different parts of the trap) cycling of the bars worked these connections loose. This happened several times during initial setup of the trap and required the redesign of one of the copper pieces. Slightly more than two years after the trap was installed, another solder joint failed, rendering the Ioffe bars unusable (open circuit).

We found that relaxing the IP trap confinement, necessary for optical lattice experiments, introduced significant motion to the BEC intermittently and randomly and that the position of the relaxed trap varied. Eventually, the problem was correlated with sudden (few ms) jumps in the voltage across the Ioffe bars at constant current—or, equivalently, sudden jumps in the bar current at constant voltage, see figure 2.16. It was suspected that this could be due to an intermittent partial short between two bars, which would cause the magnetic field to change as the amount of current diverted by the short changed. Changing the current through the Ioffe bars seemed to make the problem much worse, perhaps due to thermal relaxation or changes in the magnetic forces between parts of the trap, so the relaxation sequence was altered so that the current through Ioffe bars stayed constant and only b_{bias} was changed (the resulting final trap was slightly more confining than desired, however).

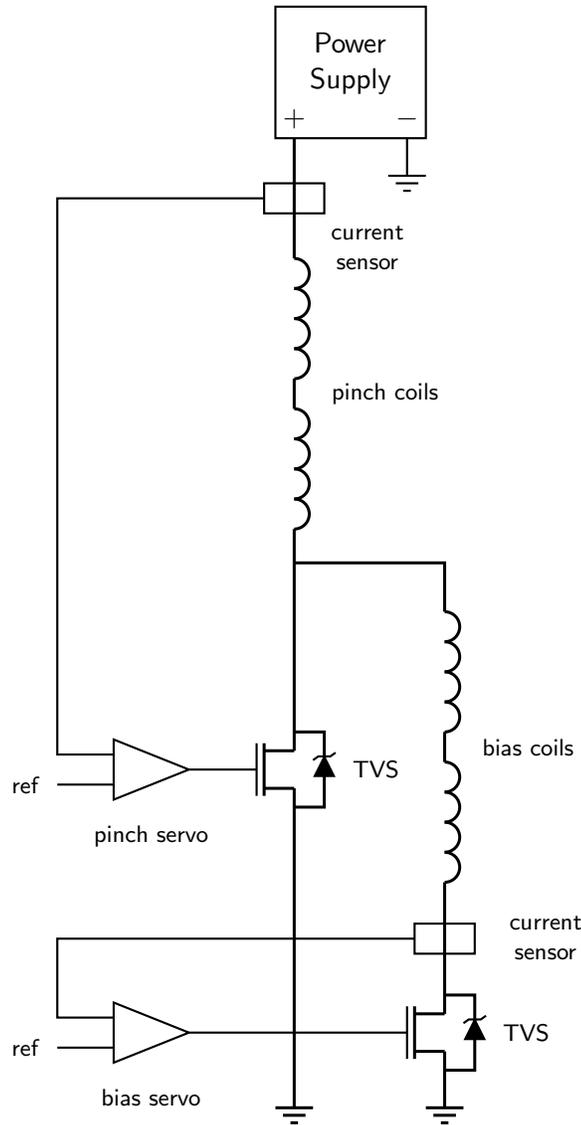


Figure 2.17 Schematic of IP trap circuit.

2.9.3 Electronics

The IP trap was designed to operate at high current with a large cooling water flow rate to minimize the amount of space taken up by conductors, allowing good optical access to the atoms. Current for the Ioffe bars was supplied from a 5 V, 875 A power supply (Agilent 6680A), carried with two 4/0 type cables (0.63 in conductor diameter), and regulated with a bank of six FETs. The circuit is otherwise identical to that used for the quadrupole coils, section 2.6.3. Current for the pinch and bias coils was supplied from an 8 V, 580 A power supply (Agilent 6681A), carried with a single 4/0 type cable, and regulated with two banks of four FETs; see figure 2.17. The FET bank for the pinch coil dissipates the most power

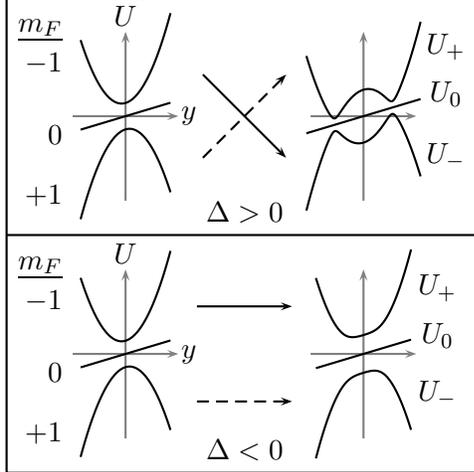


Figure 2.18 Adiabatic potentials for the $F = 1$ state in the presence of gravity. The potential U is shown along the \hat{y} direction for the $|+\rangle$, $|0\rangle$, and $|-\rangle$ states on the right and with no RF magnetic field present on the left. The arrows indicate how the m_F states transform when the RF field is turned on adiabatically. From [79].

per FET of any bank in the apparatus, though still well within the specified safe operating region, and failed several times. Most recently, a different part (IXYS IXFN340N07 power MOSFET) with lower on-resistance was installed and is still in operation at the time of writing.

2.10 Evaporative Cooling

The cloud of atoms loaded into the IP trap contains 10^9 or more atoms at several hundred μK . To cool this cloud to degeneracy, forced evaporated cooling (see, e.g., [78]) is performed by selectively removing the most energetic atoms from the cloud. A uniform RF magnetic field drives transitions from trapped to untrapped Zeeman levels (e.g., from $|F = 1, m_F = -1\rangle$ to $|F = 1, m_F = 0\rangle$ and $|F = 1, m_F = 1\rangle$ in ^{87}Rb). These transitions can only occur (with non-negligible probability) near a surface of constant magnetic field magnitude B defined by $\hbar\omega_{\text{RF}} = g_F m_F \mu_B B$. The RF frequency $\omega_{\text{RF}}/2\pi$ is chosen so that this surface (the RF “knife”) lies far enough from the center of the magnetic trap that only the most energetic atoms can reach it. As these atoms are lost from the trap, collisions between the remaining atoms reequilibrate them at a lower temperature. To continue the evaporation process, the RF frequency is decreased to bring the “knife” closer to the center of the cloud. Typically, the number of atoms must be reduced by a factor of 10^3 to 10^4 in order to reach degeneracy.

Forced evaporative cooling can also be visualized in terms of dressed potentials. In the presence of the RF magnetic field, the $|F, m_F\rangle$ states are no longer the energy eigenstates. Considering only the $F = 1$ manifold, we label the new eigenstates as $|+\rangle$, $|0\rangle$, and $|-\rangle$. The

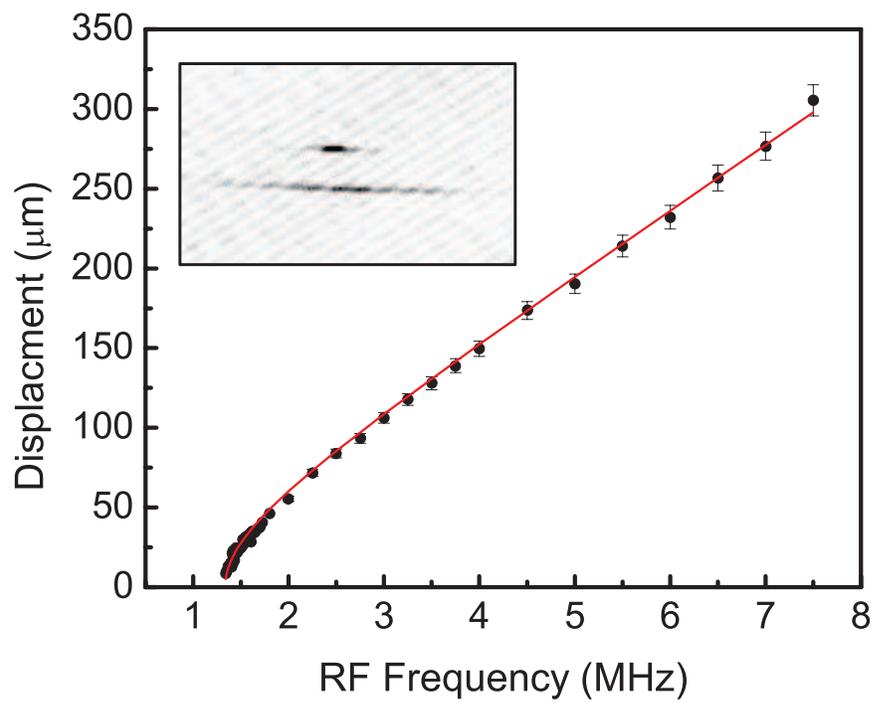


Figure 2.19 Displacement from the IP trap center of atoms trapped in U_+ potential vs. RF frequency. The red line shows the calculated displacement. The inset shows atoms loaded into both the U_- (upper cloud) and U_+ potentials at $\Delta = 2\pi \times 1$ MHz. The field of view for the insert is approximately $640 \mu\text{m}$ by $390 \mu\text{m}$. From [79].

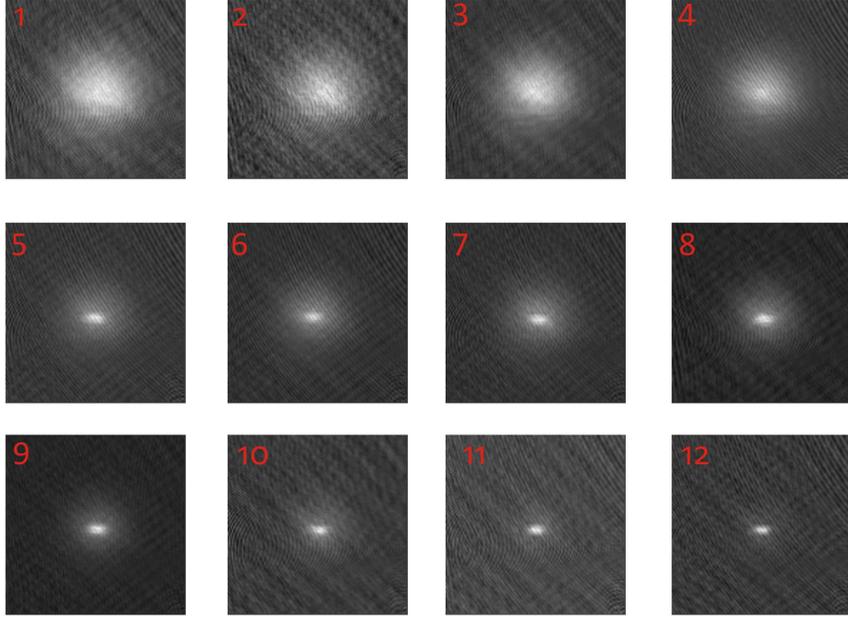


Figure 2.20 Formation of a BEC during evaporative cooling, viewed after several ms time-of-flight. The field of view is $800 \mu\text{m}$ square.

spatially dependent energies of these states are $U_{\pm}(\mathbf{r}) = \pm\sqrt{[g_F\mu_B B(\mathbf{r}) - \hbar\Delta]^2 + 2(\hbar\Omega)^2} + mgy$ and $U_0 = mgy$, where $\Delta = \omega_{\text{RF}} - g_F\mu_B B(0)$ and g is the local acceleration due to gravity. These dressed potentials are plotted in figure 2.18. Starting with $m_F = -1$ atoms, for $\Delta > 0$, we can see that the U_- potential is confining, but with a finite depth which decreases as the RF frequency, and hence Δ , is decreased. By starting with $\Delta < 0$ and sweeping to $\Delta > 0$, it is possible to trap atoms in the $\Delta > 0$ U_+ potential [80], as shown in figure 2.19. We explored this technique shortly after achieving BEC with our apparatus [79].

Efficient rethermalization, and thus a high elastic collision rate is essential to the process; since the collision rate increases with the strength of the trapping potential [74], a strong trap is desired. Competing undesirable heating and loss comes from collisions with background room-temperature atoms (minimized by having a good vacuum) and three-body collisions between ^{87}Rb atoms, in which two atoms form a molecule, with the third carrying away the extra energy. Three-body loss, the rate of which scales as the density squared, is only relevant at the high densities found during later stages of evaporation. At this point, the confinement can be relaxed to decrease the density without significantly slowing the evaporation process. Inelastic two-body collisions are impossible (to first approximation) in a cold, spin-polarized $|F = 1, m_F = -1\rangle$ ^{87}Rb gas.

Figure 2.20 shows images of one of the first BECs obtained in the lab. The two component distribution is indicative of a BEC, but the clearest sign is the asymmetry of the central cloud after expansion, reflecting the momentum distribution of a BEC in the ground

state of a nonisotropic potential.

2.11 Imaging

Before imaging, the IP trap is turned off suddenly and the atoms are typically allowed to expand long enough for the cloud to become several times larger than its size in trap. If not for interactions, the spatial distribution would then simply reflect the in-trap momentum distribution. We refer to the sudden turn off of all confining potentials as release, and the expansion process as time-of-flight. After time-of-flight, the spatial density distribution is measured using a standard absorption imaging technique, described below.

Limited by the constraint that the pinch coil turn off be slow enough to prevent an increase in the bias coil current (see figure 2.17), the IP trap turn off takes roughly 100 μ s. Before the IP trap magnetic field reaches zero, a uniform, few gauss bias field is applied using a pair of large coils to ensure that the atoms remain spin polarized during time-of-flight.

2.11.1 Procedure

The atoms are illuminated with a circularly polarized beam resonant with the $|F = 2, m_F = \pm 2\rangle$ to $|F' = 3, m'_F = \pm 3\rangle$ cycling transition and an image, S , of the beam is recorded, showing a dark region where light was scattered from the beam by the atoms. The bias field, parallel to the imaging beam, also provides a quantization axis relative to which the beam has σ^+ or σ^- polarization so that only the cycling transition is driven. Two additional images are taken: one with the beam on but without the atoms present, L , and one with the beam off to perform background subtraction, D . The optical density of the atom cloud is then calculated using

$$OD(x, y) = \ln \frac{L(x, y) - D(x, y)}{S(x, y) - D(x, y)}, \quad (2.6)$$

where we have taken the imaging axis along the z direction.

The optical density and spatial density distribution $n(x, y, z)$ are related by

$$OD(x, y) = \sigma \int n(x, y, z) dz, \quad (2.7)$$

where σ is the scattering cross-section for the imaging transition [59],

$$\sigma = \frac{3\lambda^2}{2\pi} \left(1 + 4 \frac{\delta^2}{\Gamma^2} + \frac{I}{I_{\text{sat}}} \right)^{-1}. \quad (2.8)$$

Here, Γ is the linewidth of the transition, I_{sat} is the saturation intensity of the transition, δ is the detuning from the transition, and I is the intensity. It is desirable to image with $\delta = 0$ to avoid lensing effects arising from the non-zero dispersion off resonance.

We typically use a Gaussian fit for thermal atoms (both spatial and momentum distributions) in a harmonic trap, although the correct Bose-Einstein distribution is sometimes used for very cold atoms. In the absence of interactions, the spatial distribution of a BEC in a harmonic trap would be that of the quantum ground state of the trap, a Gaussian. The repulsive interaction between atoms results in a more extended spatial distribution which, neglecting kinetic energy terms, takes the shape of an inverted parabola, called the Thomas-Fermi distribution [10] and characterized by widths w_x, w_y, w_z :

$$n_{\text{TF}}(x, y, z) = n_0 \left[1 - \left(\frac{x}{w_x} \right)^2 - \left(\frac{y}{w_y} \right)^2 - \left(\frac{z}{w_z} \right)^2 \right] \quad (2.9)$$

It can be shown [81] that, due to interactions, the shape of the spatial distribution remains the same during expansion, simply being enlarged by a scale factor. In agreement with this result, we find that the spatial distribution after expansion of the condensate is well-fit by a function of the form (2.9). Atoms clouds displaying both a thermal and condensed component are fit to the sum of a Thomas-Fermi function and a Gaussian (or Bose-Einstein function). Other specialized functions are used in various cases.

Just before imaging, light from the repump laser (resonant with the $F = 1$ to $F' = 2$ transition) is used to transfer atoms from $|F = 1, m_F = -1\rangle$ to the $F = 2$ manifold. For condensates too dense to image even after long time-of-flight, the duration of the repump pulse is shortened. To ensure the condensate is repumped evenly in this case, it is necessary to detune the repump light so that only a small fraction is absorbed passing through the cloud. This partial repumping technique has the disadvantage of being sensitive to changes in the repump beam power. A better approach is to use a microwave frequency sweep (section 6.7) to transfer atoms to the $F = 2$ manifold; the fraction of atoms transferred can be changed by varying the speed of the sweep or the microwave power.

2.11.2 Optics

Figure 2.21 shows a schematic of the high-magnification imaging system. A low magnification variant was used during initial setup of the apparatus. The camera used is a Princeton Instruments VersArray XP with a 1024×1024 pixel cooled CCD; the specified pixel size is $13 \mu\text{m}$ square. The imaging system was modeled using ZEMAX software, which predicted a magnification of 4.06, an RMS spot diameter of $3.7 \mu\text{m}$, and a depth of focus of $110 \mu\text{m}$. The actual magnification was measured to be 4.14 by allowing a cloud of atoms to fall freely under gravity and fitting the trajectory to $z = gt^2/2$; this corresponds to $3.14 \mu\text{m}$ per pixel. To ensure that stray magnetic fields do not affect the free fall, a quick RF sweep or pulse can be used to transfer some atoms to the magnetic field insensitive $m_F = 0$ state.

The imaging optics are aligned by measuring the size of a very small condensate vs. camera position for several different positions of the first lens. The position giving the

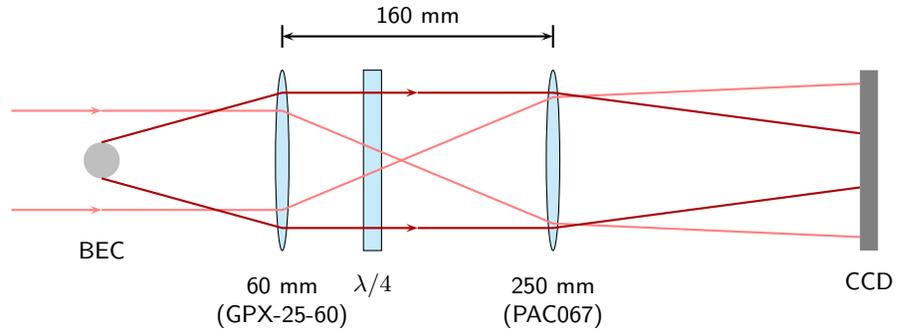


Figure 2.21 Optical system used for absorption imaging of atoms in science cell. The first lens is a Lightpath Gradium GPX-25-60, the second is a Newport PAC067 achromatic doublet. The predicted magnification is 4.06, with a spot diameter of $3.7 \mu\text{m}$. The quarter-wave plate (denoted $\lambda/4$) prevents etaloning between optics on either side of it.

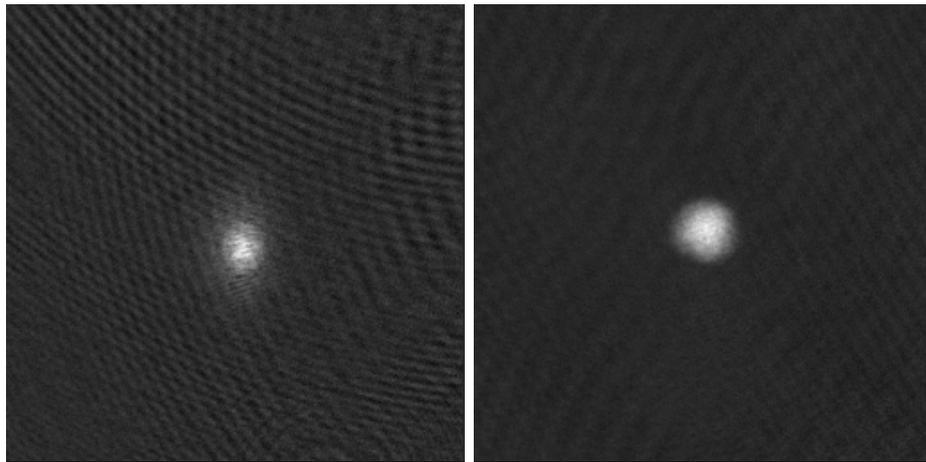


Figure 2.22 Sample images taken without (left) and with the quarter-wave plate installed in the imaging system.

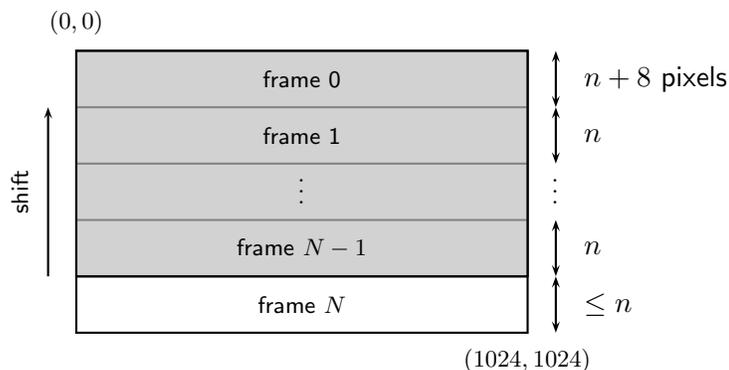


Figure 2.23 Layout of CCD in kinetics mode for a Princeton Instruments VersArray XP camera. The shaded region is shielded from light using, e.g., a razor blade, and is used to store images taken by exposing the unshaded region. The value n can be set using the camera’s software driver.

smallest size is chosen.

A quarter-wave plate (QWP) was added to the imaging system to prevent the formation of etalons between surfaces on either side of the QWP. Two passes through the QWP result in an orthogonal polarization (for the circularly polarized light used), preventing interference. The improvement in image quality was dramatic; see figure 2.22. If etaloning between the first lens and the glass cell was a concern, a very thin first-order waveplate could be placed between the two, instead of in the indicated position, without affecting the imaging quality.

2.11.3 Anti-gravity Coil

The field of view for the imaging system is 3.2 mm square, meaning that an atom cloud in free fall will leave the field of view in no more than 25 ms. It is often desirable to perform time-of-flight for longer times, for example, to see vortices and other topological features. A so-called anti-gravity coil is attached to the IP trap 1.5 inches below the center of the science cell to produce a vertical magnetic field gradient. When the coil is turned on at the beginning of time-of-flight the resulting force acts opposite to gravity, slowing the free fall of the atom cloud.

The coil consists of 80 turns of 24 gauge magnet wire and has an outer diameter of 2.5 inches. The current in the coil is controlled and stabilized with the same type of circuit as for the quadrupole coils (section 2.6.3). Neither the coil nor the FET (IRFB4410) is actively cooled, limiting the amount of time the coil can be on at high current to a few hundred ms; a fuse is used to break the circuit if the coil is left on for longer.

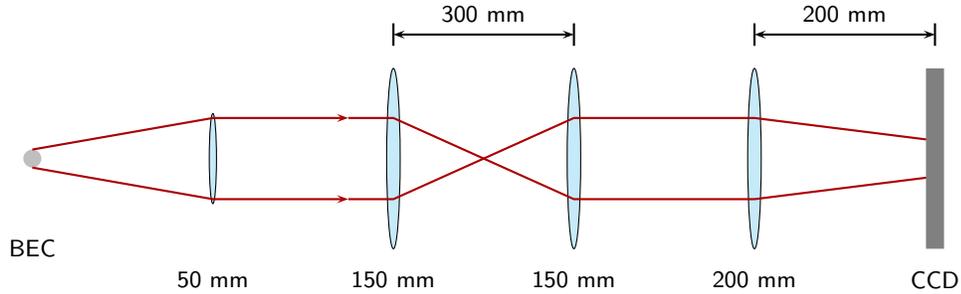


Figure 2.24 Auxiliary imaging system. The path contains several mirrors, not shown, to wind it toward the camera. For simplicity, the unscattered absorption imaging probe beam light is omitted from the figure.

2.11.4 Kinetics Mode

The kinetics mode of the CCD camera allows multiple images to be taken in rapid sequence. This is done by exposing part of the CCD, shifting the exposed pixels into a region masked from light, then repeating. The layout of the CCD for kinetics mode is shown in figure 2.23. Once the entire CCD has been filled, the camera shutter closes and the images are read out. With kinetics mode, the time delay between the shadow and light images for absorption imaging can be reduced from several hundred ms to a few ms, greatly reducing noise from changing etalon fringes.

We used kinetics mode to take reduced noise images for density-density correlation measurements 4.4.2, taking the shadow and light images in a single shot and separate background images for each in a second shot.

2.11.5 Auxiliary Imaging

To assist in alignment of the optical lattice, a second imaging system was set up in the vertical direction. The light from this imaging system was directed onto one half of the CCD using an edge mirror. The other half of CCD was used for the primary imaging system. Later, the auxiliary imaging direction was moved to the horizontal plane (still perpendicular to the primary imaging direction). A schematic of this imaging system is shown in figure 2.24.

2.12 Dipole Trap

After the loss of the IP trap, we explored several different trapping schemes. We will begin by discussing the current system, a hybrid trap combining a single beam optical dipole trap (see section 6.2) and a magnetic quadrupole field. Subsequently, we will refer to it simply as the dipole trap. The beam has a $120 \mu\text{m}$ waist and passes roughly $100 \mu\text{m}$ below the zero

of the quadrupole field. The quadrupole field is created by the pinch coils of the IP trap reconnected in an anti-Helmholtz configuration. The beam provides radial confinement, while the quadrupole field provides axial confinement and supports against gravity. The overall potential, taking gravity to be in the $-y$ direction and the beam along the z axis, is

$$V = -V_0 e^{-2(x^2+y^2)/w^2} + m_F g_F \mu_B |\mathbf{B}| + mgy, \quad (2.10)$$

where $-V_0$ is the maximum depth of the dipole potential and the second term arises from the magnetic field (see equation 2.1). The zero of the quadrupole magnetic field,

$$\mathbf{B} = B' \left(\frac{x}{2}, \frac{y - y_0}{2}, -z \right), \quad (2.11)$$

is offset a distance y_0 above the origin. We calculate the trapping frequencies for the potential numerically.

2.12.1 Fiber Laser

The light source for the dipole beam is an IPG Photonics YLR-10-1064-LP ytterbium fiber laser, which provides 10 W in a linearly polarized, single-transverse-mode beam at 1064 nm with 1 nm linewidth. Some groups [82] have reported severe anomalous heating in dipole traps using non-single-frequency fiber lasers. We observe a low heating rate of less than $0.02 \mu\text{K/s}$ for a low density thermal cloud at $0.5 \mu\text{K}$, comparable to that in the IP magnetic trap. Initially, a titanium:sapphire (Ti:Sa) laser (section 4.6.1) tuned to 820 nm was used for the dipole beam, but heating due to spontaneous scattering prevented cooling to BEC.

2.12.2 Optics

The optical system for the dipole trap is shown in figure 2.25. The optics to recycle the beam (with orthogonal polarization) can be used to create a crossed dipole trap, which allows for confinement in experiments where a uniform magnetic field is used. Otherwise, the beam path is blocked before these optics.

The undeflected beam from the AOM was originally incident on a beam dump placed just below the path of the deflected beam. The beam dump became quite hot, above 50°C , and it was suspected that thermal air currents might cause unwanted motion of the dipole beam. This was confirmed by watching the beam on a camera. To correct the problem, the undeflected beam was reflected to a beam dump away from the rest of the dipole beam path.

The dipole beam intensity is stabilized and controlled via the AOM using the same scheme as for the optical lattice beams, described in section 4.7. Because of the far greater power in the dipole beam, an anti-reflection coated window is used to sample the beam instead of an uncoated beam sampler.

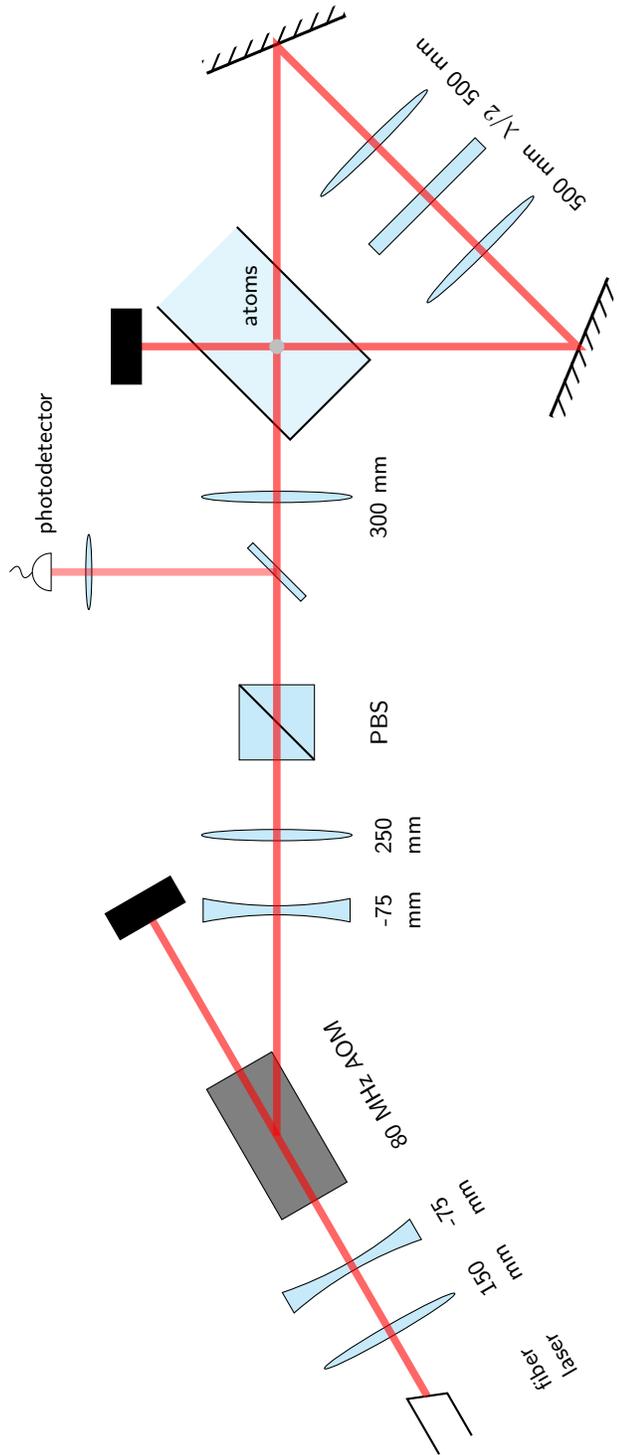


Figure 2.25 Schematic of optical dipole trap system. The beam is blocked before the first 500 mm lens unless a crossed dipole trap is desired, in which case the half-wave plate ($\lambda/2$) is used to make the cross beam polarization orthogonal to the first pass beam polarization.

2.12.3 Alignment

Coarse alignment of the dipole beam is performed by overlapping it (by eye) with a resonant beam aligned using fluorescence (see section 4.9). This is sufficient to load some atoms into the trap. Next, the focus is positioned by minimizing the width of the atom cloud trapped by the dipole beam. This can be checked by measuring the transverse harmonic oscillator frequency at several different positions. Finally, the vertical position is adjusted to maximize the number of atoms loaded into the dipole trap and the horizontal position adjusted to match that of the quadrupole trap prior to loading, since the fraction of atoms loaded into the dipole trap is much less dependent on the horizontal position than on the vertical position.

2.12.4 Transfer and Evaporation

Figure 2.26 illustrates loading of the dipole trap and evaporative cooling, along with the trapping potential for various stages. Atoms are loaded into the dipole trap after a stage of evaporative cooling in the magnetic quadrupole trap. The dipole beam is then turned up in 500 ms while the quadrupole trap is turned down until it is just slightly weaker than gravity. The magnetic axial confinement helps the cloud stay spin polarized, although an atom scattering light could decay to the $|F = 2, m_F = 1\rangle$ state, which is also trapped.

Evaporation in optical dipole traps is often performed by simply reducing the optical power. While this has the desired effect of lowering the trap depth to allow for evaporation, it also has the undesired effect of weakening the trapping potential: the trap frequency scales as the square root of the intensity for constant beam profile. If an additional force is applied to the atoms, “tilting” the overall potential (see the middle trace in figure 2.26), the trap depth can be lowered with minimal decrease in the trap frequencies [83]. We found that this tilt-evaporation method performed far better than lowering the dipole beam power. Figure 2.27 shows the evaporation trajectory in the magnetic quadrupole trap and dipole trap.

After loading atoms into the dipole trap, we perform tilt-evaporation by strengthening the quadrupole potential. The anti-gravity coil (section 2.11.3) is used to generate an additional field to keep the center of the quadrupole potential from changing as it is strengthened (due to non-idealities). Because the quadrupole field itself is confining, we use a 200 kHz RF field to remove the atoms which have been evaporated from the dipole trap into the quadrupole trap.

After evaporation to roughly T_c , we reduce the dipole beam power and the quadrupole field strength to make the trap less confining. We then perform a final stage of evaporation by gradually decreasing the dipole beam power in a slightly tilted trap. We find that performing the final stage of evaporation in a weaker trap allows us to achieve higher condensate fraction and remove any motion introduced during relaxation. Following this

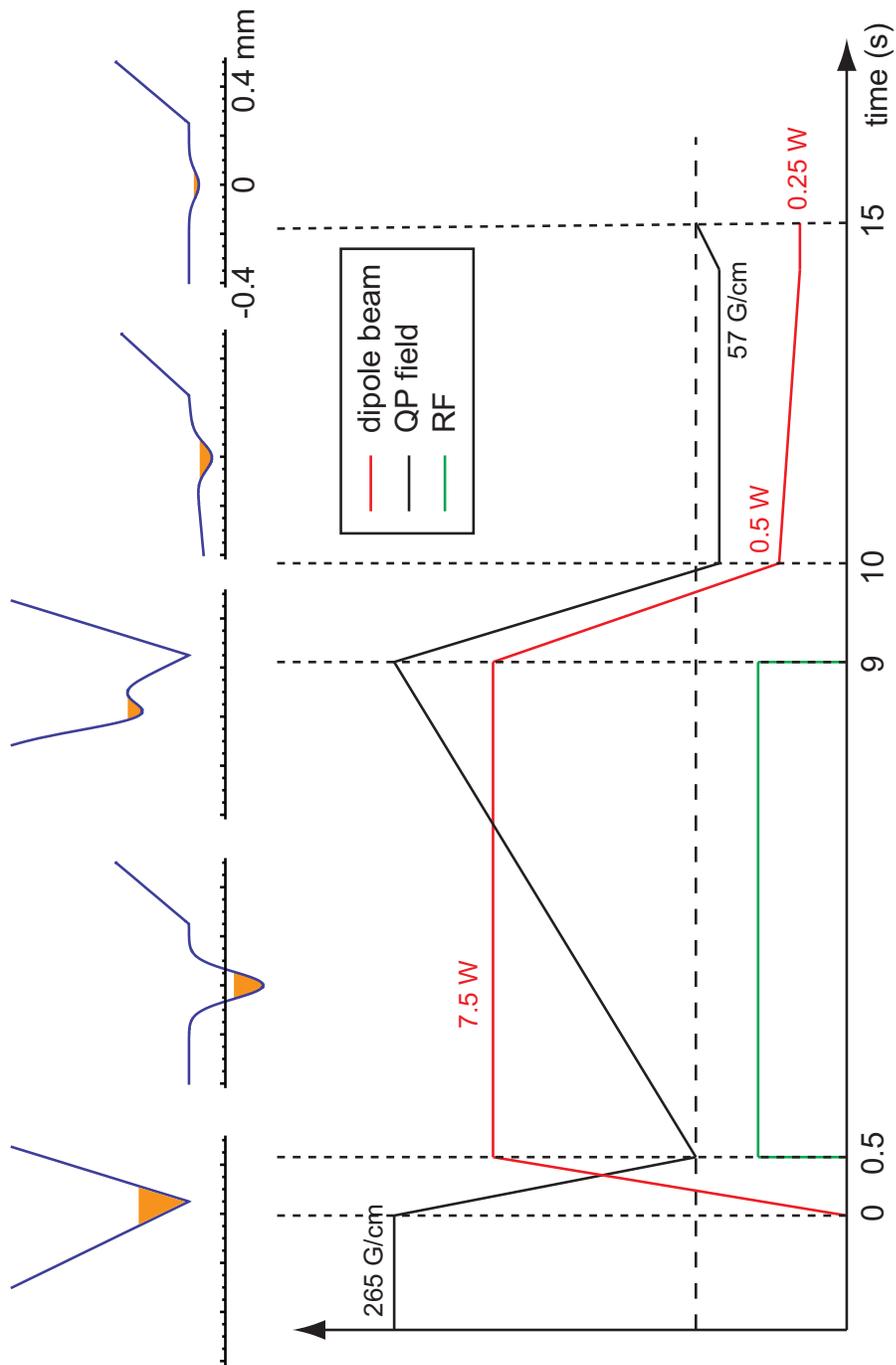


Figure 2.26 Evaporation sequence for producing nearly pure BEC in the hybrid optical/magnetic trap. The dashed horizontal line indicates the level at which the quadrupole field cancels gravity (61 G/cm). The five small plots show the potential experienced by the atoms along the vertical direction (including gravity, which points in the negative direction) at the five points marked by dashed vertical lines. Atoms in the trap are represented in orange.

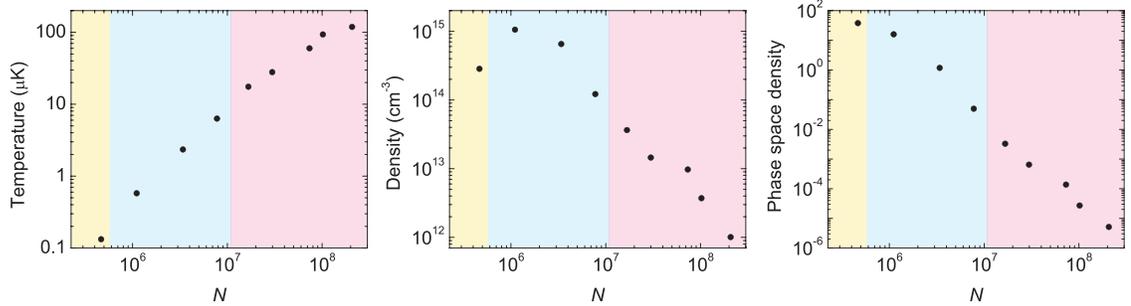


Figure 2.27 Temperature, in-trap density, and phase space density vs. total atom number during evaporation sequence. The colored regions correspond to evaporation in, from right to left, the quadrupole trap, the strong dipole trap, and the weak dipole trap. The decrease in density for the final stage is due to the relaxation of the dipole trap.

final stage of evaporation, the quadrupole field is strengthened so that it cancels gravity as closely as possible.

2.13 Plugged Quadrupole Trap

The first BECs were created in quadrupole magnetic traps using two different approaches to avoid atom loss due to Majorana spin flips in the weak field region: a time-orbiting potential (TOP) trap [84] and an optically-plugged trap (OPT) [2]. Prior to implementing the dipole trap, we experimented with the OPT method. In this method, a blue-detuned beam is focused at the center of the quadrupole trap to repel atoms from the weak field region. The standard RF evaporation technique (section 2.10) is used. Only a few groups are known to be using an OPT at present [85, 86].

We used 300 mW of 770 nm light from a Ti:Sa laser focused to a 40 μm waist and directed horizontally through the center of the quadrupole trap. Coarse alignment of the plug beam was done by tuning the laser to resonance and using fluorescence (section 4.9). The waist was positioned by tuning the laser to 783 nm to form a strong dipole trap and minimizing the width of an atom cloud loaded into the dipole trap. Next, the laser was tuned to 778 nm and the beam position adjusted to punch a hole in the middle of a cold thermal cloud. Finally, the laser was tuned to 770 nm and the beam position adjusted to maximize the number of atoms obtained after RF evaporation.

After alignment, RF evaporation to 0.4 MHz with a 200 G/cm quadrupole gradient reliably produced BEC. Since the plug beam is in the horizontal plane only one condensate forms, below the beam. However, the spatial distribution of the BEC after time-of-flight was irregular and fluctuated somewhat from shot to shot.

Because atoms are repelled from the plug beam, the scattering rate is much lower than for beam detuned an equal amount to the red; figure 2.28 shows the increase in lifetime due to the plug beam and that increasing the plug beam power does not decrease the lifetime,

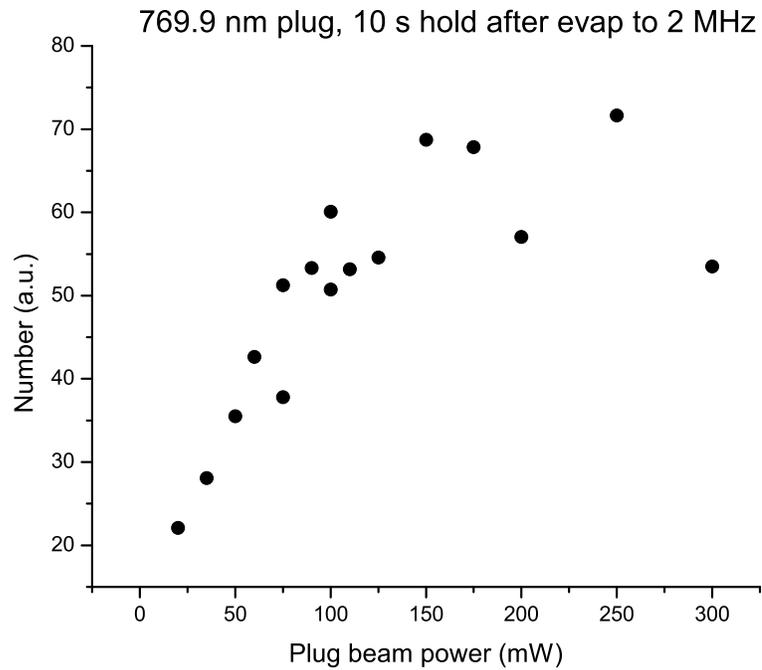
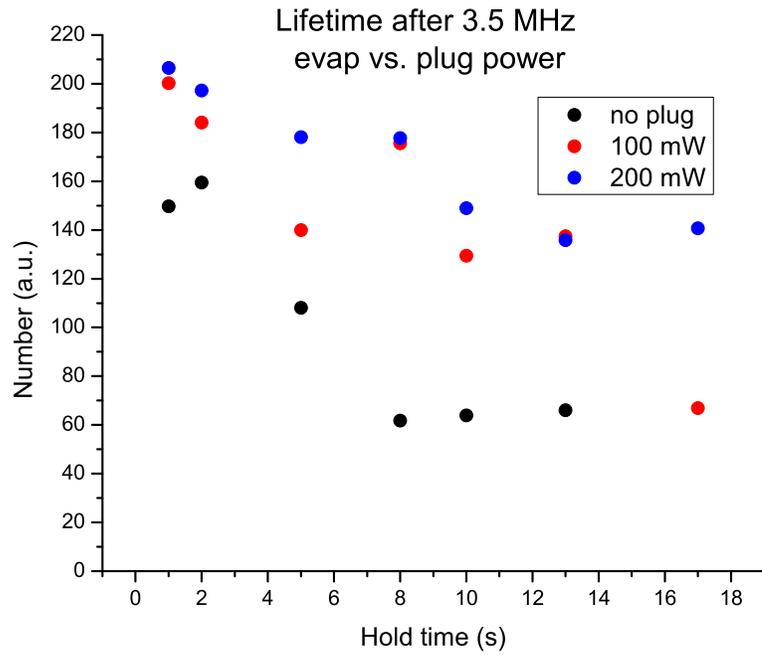


Figure 2.28 Lifetime of cold atoms (after RF evaporation) in the plugged quadrupole trap vs. plug beam power. The beam is detuned by only -10 nm from the D2 resonance.

as would be expected if scattering was significant.

One possible improvement to the OPT technique is the use of a diffractive optical element to create a hollow plug beam, reducing the optical power needed to obtain a potential barrier with a given height and width (vs. a Gaussian beam).

2.13.1 Broad-area Laser Diode

After our initial success with the Ti:Sa laser, we attempted to use a 760 nm, 2.5 W broad area laser diode (LDX Optronics LDX-3215-760) to generate the plug beam. Fiber coupling light from the diode was much more difficult than originally thought, and we unable to approach the 600 mW needed to match the Ti:Sa plug beam. To obtain a satisfactory beam, we used a pinhole as a spatial filter and imaged the pinhole onto the desired position of the plug. However, we found that this plug beam caused rapid atom loss when turned on. We decided to move forward with an optical dipole trap instead of investigating this problem carefully. Using an optical spectrum analyzer (scanning grating type), we were unable to detect any near resonant light in the beam, and we found that the atom loss was not strongly dependent on beam alignment, ruling out loss due to trapping of atoms near the zero of the quadrupole field.

Chapter 3

Control System

3.1 Control Software

3.1.1 Overview

The most popular approach to control software for cold atom experiments is a visual timeline system (see, e.g., [87]) which displays the value of each digital or analog output vs. time. The time axis is usually compressed by showing only time steps where a change occurs. We chose an alternate approach wherein the sequence of commands for the experiment is specified in a simple programming language which is translated to C code by a short perl [88] script, `seqmeta`. The C code is compiled to an executable, which in turn produces raw analog and digital output data on `stdout` to be used by the hardware interface program, `out_pc`. The entire chain of events takes less than one second.

Initially, a custom parser was used for a subset of the sequence language. When considering how to add more features, it was realized that writing a simple translator to C and using the capabilities of a mature C compiler would be much easier and more powerful than extending the custom parser.

3.1.2 Sequence Syntax

A sequence consists of one or more processes which ultimately specify analog and digital signal output commands and the time at which they are to occur relative to the beginning of the sequence. The name and any arguments of a process are enclosed in brackets “[]” on a new line. All lines of text until the next process are part of the process.

The basic building blocks of a process are expressions which evaluate to a time. Typical expressions are numeric literals (e.g. `10.000`), simple expressions (e.g., `ii*750*ms`), and calls to other processes (e.g., `MyProcess()`). As the sequence file is processed, the current sequence time is tracked by an accumulator.

An example sequence and the generated C code are shown in figure 3.1, along with a timeline of events in the sequence. The number above the timeline is the value of the accumulator when the processes below the timeline are called. Figure 3.2 shows several examples of more advanced techniques. Notice that lines of plain C code can be included

```
[HybridDipoleTightRelease]
0.00 AG_Coil_Off(), DipoleBeamOff()
0.00 PinchQP_Off()
0.20 Imaging_Coil_On()
\ 50 Imaging_Coil_Off()
0.00
```

```
[AG_Coil_Off]
-0.1 DigitalOut(56, 56, 0)
0 AnalogOut(AO_AGSUPPLY, 0)
0 AnalogOut(AO_AGSERVO, 0)
```

Generated C code:

```
time_t HybridDipoleTightRelease(void) {
    PREFIX;
    /* 0.00 AG_Coil_Off(), DipoleBeamOff() */
    RESET; CHAIN_PROCESS(0.00); DO_PROCESS(AG_Coil_Off()); CHAIN_PROCESS(DipoleBeamOff());
    /* 0.00 PinchQP_Off() */
    RESET; CHAIN_PROCESS(0.00); CHAIN_PROCESS(PinchQP_Off());
    /* 0.20 Imaging_Coil_On() */
    RESET; CHAIN_PROCESS(0.20); CHAIN_PROCESS(Imaging_Coil_On());
    /* \ 50 Imaging_Coil_Off() */
    CHAIN_PROCESS(50); CHAIN_PROCESS(Imaging_Coil_Off());
    /* 0.00 */
    RESET; CHAIN_PROCESS(0.00);
    SUFFIX;
}

time_t AG_Coil_Off(void) {
    PREFIX;
    /* -0.1 DigitalOut(56, 56, 0) */
    RESET; CHAIN_PROCESS(-0.1); CHAIN_PROCESS(DigitalOut(56, 56, 0));
    /* 0 AnalogOut(AO_AGSUPPLY,0) */
    RESET; CHAIN_PROCESS(0); CHAIN_PROCESS(AnalogOut(AO_AGSUPPLY,0));
    /* 0 AnalogOut(AO_AGSERVO,0) */
    RESET; CHAIN_PROCESS(0); CHAIN_PROCESS(AnalogOut(AO_AGSERVO,0));
    SUFFIX;
}
```

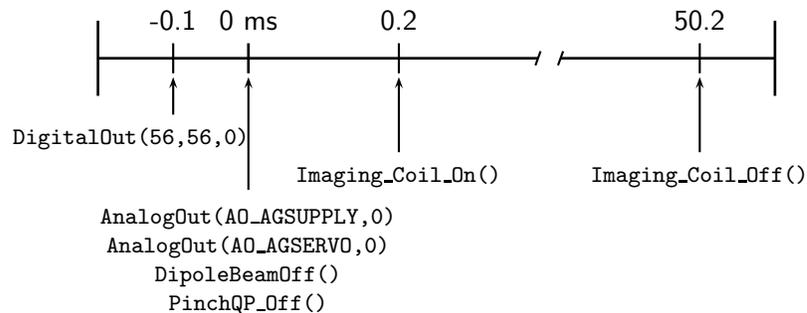


Figure 3.1 Sample of sequencer code, generated C code, and corresponding timeline.

```

[TestPinchEvap1]
; use a big ROI to fit large thermal cloud
MATLAB("runfit('gauss', 'fix', {'offset'}, 'AutoROI', [200, 200])");
MATLAB("showres()");
;; sequence
0 PinchQP_TightEvap()
\ 100 HybridDipoleTightRelease()
\ 10 ImageAndRepeatExt(200*us, 0, IMG_FullFrame)

[PinchQP_TightEvap]
f64 fstart1 = 80.000*MHz;
f64 fstop1 = 40.000*MHz;
;; sequence
0 SciCellLoad()
\ 200 Linear_Evap(fstart1, fstop1, 2*MHz/sec)
\ 100 PinchQP_Transfer()
\ 200 Move_Cart_PRUN(), Pinch_Bias_Ramp(1000*ms, 585, 585, 52.5, 0) Bias_SetPt(0)
\ 200 PinchQP_Evap3()

[KineticsImage(time_t repump_time, f64 frames[])]
int ii = 0;

dprintf(PRINT_VARS, "repump_time_%d = %f\n", kin_ctr, repump_time);
dprintf(PRINT_VARS, "kinetics_offset_%d = %f\n", kin_ctr++, timeAbs - kin_base_time);

for(; ii < 3; ii++) {
    if(frames[ii] == 1) {
        ii*750*ms Kinetics_Pulse_F1(repump_time)
    }
    else if(frames[ii] == 0) {
        ii*750*ms Kinetics_Bkgnd(repump_time)
    }
}

```

Figure 3.2 Another sample of sequencer code, showing the ability to interleave C code. In `KineticsImage`, only the two lines beginning with `ii*750*ms` are sequence code. `TestPinchEvap1` is a top level process; notice the two commands to the fitter. Lines beginning with “;” are comments.

arbitrarily for initializing variables, creating loops, and so on. The `MATLAB()` commands control processing the of image data by the fitter (section 3.5).

The concept behind the sequence syntax is that expressions on the same line of text separated by spaces run in series from left to right, while each line of text runs in parallel. Specifically, the return value of each expression is added to the sequence time accumulator when the expression is evaluated. The accumulator is reset to its value at the beginning of the current process when a new line begins. If an expression is followed by a comma, its return value is not added to the accumulator. A backslash (`\`) is used as a line continuation character (as in many computer languages); it prevents the accumulator from being reset for the new line. When placed at the beginning of a line, a backslash serves to continue in series from the previously run expression, regardless of location within the text, so it works within loops and control structures. The (time) value returned by a process is the value of the accumulator when evaluation of the process completes (any trailing new lines are ignored).

The use of spaces as expression separators introduces indeterminate cases, e.g., `a() - 10 - b()`. Because of the rarity of these cases and the keystrokes saved by using spaces as separators, the general rule is to use parenthesis to clarify any such cases.

Inline (i.e., anonymous) processes can be defined by enclosing code in double brackets: “[`[]`”]. These behave exactly as if the enclosed code had been placed in a separate process. We rarely use this feature, typically only in lower level processes.

3.1.3 Components

The program to control execution of sequences, `ctrl`, integrates a text editor so that the sequence to be executed can be selected simply by clicking next to its definition in the sequence file. When a sequence is selected in `ctrl`, the standard `make` utility is used to process only files that have been modified and their dependents, generating the executable to be sent to `out_pc`. `make` also invokes revision control software (Mercurial [89]) to record changes made to the sequence files. This provides a complete record of every sequence executed in the experiment. Two special processes, `RELOAD()` and `RESTART()`, can be used to run `make` and start the next cycle of the experiment automatically. Figure 3.8, below, shows the software connections schematically.

`out_pc` sorts the output commands from the sequence executable by absolute time and waits for a start signal from `ctrl`. Each output hardware component (card) has a separate execution thread responsible for monitoring the state of the card’s memory buffer and refilling it according to the output commands. `out_pc` and `ctrl` communicate using TCP/IP and need not run on the same computer. Because of the acoustic noise generated by the apparatus and the lack of sunlight, it is often desirable to operate the experiment from a different room.

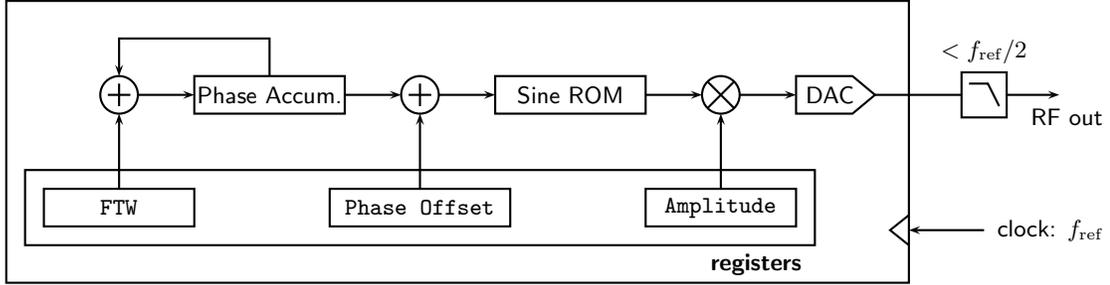


Figure 3.3 Schematic of a direct digital synthesizer. The FTW register determines the output frequency according to $f = f_{\text{ref}} \times \text{FTW}/2^N$, where N is the width of the FTW register. To keep the size of the sine ROM to manageable level, typically only the most significant 15–19 bits of the phase value are used.

3.2 Control Hardware

3.2.1 Analog and Digital Output

The analog output card (National Instruments PCI-6733) has eight 16-bit channels operating at up to 750 kHz (1 MHz if not all channels are used). We use the buffered output mode where the card continuously reads data from a memory buffer which is periodically refilled by `out_pc`.

The 64-bit digital I/O card (UEI PDL-DIO-64TS) includes a time-sequence mode (see section 3.6), however, this mode proved to be unusable due to non-deterministic behavior. Random jitter (in multiples of the card’s processor clock period) in the timing of digital edges was observed. We fell back to using the buffered output mode (identical to the analog output card’s buffered mode), which is limited to 400 kHz operation.

3.2.2 Other Analog Sources

Dedicating one of the relatively few channels of the analog output card to a signal which only needs to change between a handful of static values is certainly not desirable. To drive these signals, a circuit utilizing a crosspoint switch (Intersil CD22M3494) was constructed to make connections between an array of potentiometers and output channels. Using a crosspoint switch, which can form a connection between any input wire and any output wire, instead of an array of multiplexers is more compact and versatile.

A potentially better solution, being used in the fermion apparatus, is a many-channel digital-to-analog converter, such as the Analog Devices AD5372 which provides 32 independent 16-bit DACs on a single IC with a simple serial data interface.

3.2.3 Direct Digital Synthesizers

Over ten (and growing) different RF signals are needed throughout the apparatus, primarily for driving acousto-optic modulators. Even low-end commercial RF sources are surprisingly expensive and often do not allow for sufficient control, leading many groups to use home-built sources. Although we followed this route initially, we now use direct digital synthesizer (DDS) chips almost exclusively. A DDS generates a digital approximation to a sine wave, which is sent to a multi-bit digital-to-analog converter, see figure 3.3, and then filtered off-chip to remove frequency components above the Nyquist limit (half the conversion rate). A serial data connection provides real-time control over the frequency, phase, and amplitude of the RF signal with digital precision and reproducibility.

The Analog Devices AD9959 IC, which combines 4 independent RF sources of up to 200 MHz on a single chip, is of particular note. An evaluation board including the chip, filters, and data connectors is available for \$250. Because all four sources share a common clock, they can be used in situations where a PLL might otherwise be required, such as synchronous detection.

3.2.4 Frequency Reference

Frequency precision and stability are not a concern for most of the RF signals in the apparatus. One exception is the microwave signal for driving hyperfine transitions, section 6.7, which was the motivation for the purchase of a commercial frequency reference. The reference (SRS FS700) uses timing information from the LORAN-C radio navigation system to produce a precise 10 MHz signal. A 1 GHz voltage-controlled SAW oscillator (Crystek CVCSO-914-1000) is phase locked to the 10 MHz signal using an Analog Devices AD9510 IC, which also distributes multiple copies of the 1 GHz signal or subharmonics. RF transformers are used to convert the four differential signal outputs of the IC to isolated single-ended signals for distribution over coax. Because of the large number of outputs available, some are used to provide the clocks for DDS ICs.

3.3 Isolation

The digital and analog control systems are unique in that they connect to most of the major electronic components of the apparatus from a central point. Since all these components already share a common ground connection through their power connections, introducing another shared ground from the control system creates ground loops [90]. Induced electromotive force from 60 Hz AC power and lighting, return currents, and other sources causes different parts of the loop to be at different potentials. Since all single-ended signals are referenced to the local “ground”, the ground loop noise effectively appears on all signals referenced to the noisy ground, not, e.g., just the control signals.

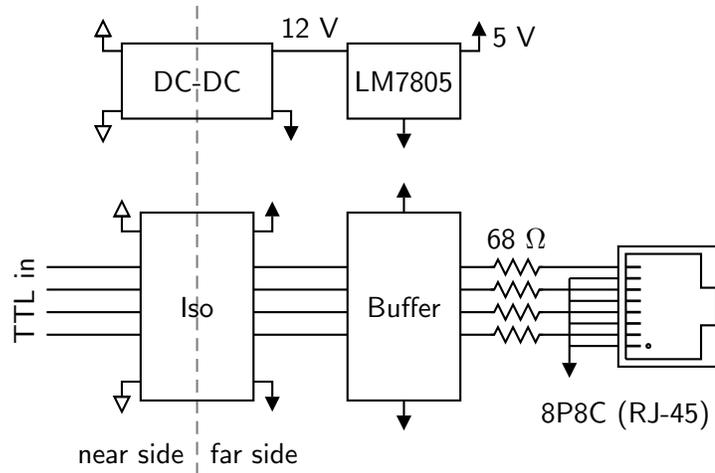


Figure 3.4 Isolation circuit for digital signals. The DC-DC converter is a Texas Instruments DCP010512, the isolator is an Analog Devices ADuM1400, and the buffer is a generic 74ACT573.

Even if ground isolation is not desired or necessary, some type of buffer is advisable to protect the analog and digital output cards from damage—replacing a component on a home-built circuit is much faster and cheaper than sending a card back for repair or replacement (as we discovered first-hand).

In light of the fact that many ground loops are unavoidably formed by various other connections between electronic components, the necessity of ground isolation for control signals is questionable. In the fermion apparatus, it is hoped that placing the control system in a central location that minimizes the size of ground loops formed will eliminate the need for isolation (a buffer is still used); another option for digital signals is simply AC-coupling the far-side and near-side grounds (10–100 nF works well).

3.3.1 Digital Isolation

Isolation and buffering for digital signals, in groups of four, is provided by the circuit in figure 3.4. The output from an isolated DC-DC converter in a standard DIP package is filtered and regulated with a linear regulator and used to power the output side of a digital isolation IC and a buffer used to provide more drive strength. Single ICs that combine the DC-DC converter and digital isolation are now available (Analog Devices ADuM540x series).

Many electronic components of the apparatus require several digital signals. To reduce the number of cables needed and keep costs low, standard ethernet cabling (8-conductor shielded CAT5e cable) and connectors (8P8C / RJ-45) are used to carry up to four digital signals on a single cable. The cable shield is connected only at the source end. Unlike many other types of cable, assembling these cables by hand is easy and inexpensive. One disadvantage of this type of cable is crosstalk between signals, although the four twisted pairs in

cable use different twist periods to minimize coupling. In latest generation of ethernet cable (CAT7), this problem is greatly reduced because each twisted pair is individually shielded.

Source termination, i.e., a resistor in series with the signal source, is used instead of load termination to minimize power dissipation and allow for direct connection to high impedance logic inputs. The value of the resistor is chosen to make the total source impedance, including the output impedance of the driver, equal to the characteristic impedance of the transmission line, 100 ohms for CAT5e. For the drivers used, the source impedances for low and high logic level outputs were different; the termination was chosen to give the best match for low output, to which the logic level threshold is closer.

3.3.2 Analog Isolation

The analog isolation scheme also called for an isolated power supply for each channel to power a differential amplifier buffer. Because the DC-DC converter used for the digital isolation circuit was far too noisy for the analog isolation circuit, we endeavored to build a custom DC-DC converter system using a sinusoidal signal to greatly improve noise performance at the cost of efficiency. Constructing the system took longer than expected due to a number of ancillary problems. Although the circuit showed great promise after these were overcome (the remaining problem being due to the noise spikes during reverse recovery of the rectifiers; a slower recovery rectifier such as a simple p-n junction diode likely would have corrected this), a different solution was implemented in the interest of expediency: eight separate linear AC-DC power supplies referenced to the far-side grounds.

3.4 Experiment Monitor

The sheer number of components in the apparatus, especially free-space optics components which inevitably drift and become misaligned, make technical problems a frequent occurrence. In response, we purchased a data acquisition card (National Instruments PCI-6255) to monitor various signals throughout the apparatus. However, we were unable to find suitable software for displaying these signals in real-time.

We wrote a simple multi-channel, multi-display oscilloscope in MATLAB, figure 3.5. Each display can be configured to show a continuously scrolling view of any number of channels, like a chart recorder, or a snapshot of any number of channels when triggered, like a standard digital oscilloscope. For triggering, a user-defined arbitrary function of any number of channels, not restricted to channels shown on the display being triggered, is used. Arbitrary measurement functions may be defined to display mean values, variances, pulse widths, rise times, etc. on each display. In addition to alerting us of problems in the apparatus, this software oscilloscope has proved invaluable for providing confirmation that commands are being followed as expected by the operator.

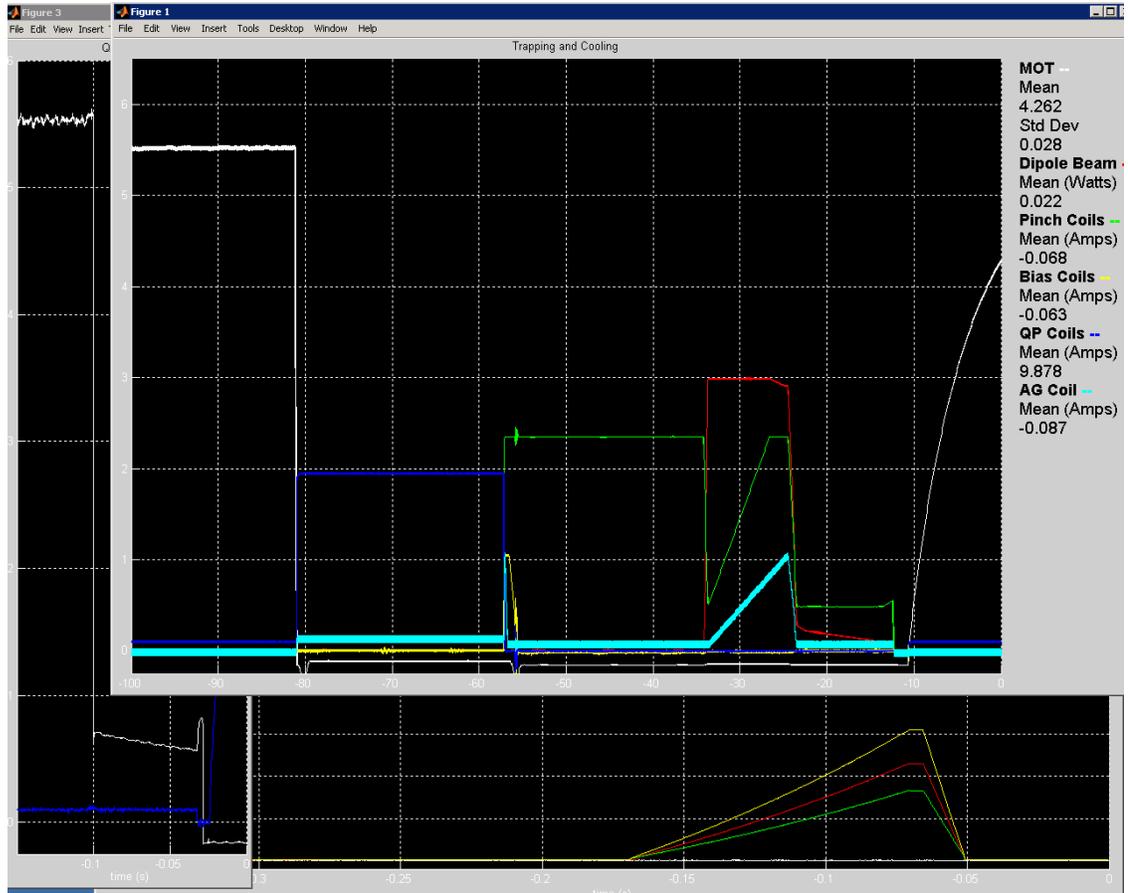


Figure 3.5 Screenshot of experiment monitor software. In the layout shown here, the main window displays the previous 100 seconds of operation, including one cycle of an experiment loading a BEC into a 3D optical lattice. The lower window shows the lattice beam intensity for +50 ms, -300 ms (partially obscured) around the release of the atoms. The left window shows the MOT to quadrupole trap transfer (section 2.8).

3.5 Image Processing

Custom software written in MATLAB, `fitter`, is used to fit images to density distribution functions, section 2.11.1. The software attempts to automatically find the atom cloud within a specified region and create initial guesses for the least-squares fit routine. For each fit function, various measurements such as atom number and condensate fraction can be calculated once the fit is completed. Images can be saved along with text files containing the fit results and other information.

Commands to the fitter can be specified in the sequence, as seen in figure 3.2. `ctrl` also sends the sequence executable to `fitter`, which runs it with a special command line argument to obtain these commands as well as variables to be saved with the image.

3.6 FPGA Sequencer

3.6.1 Overview

The nature of cold atom experiments is such that a control signal may be active (i.e., changing) for only a small fraction of the total run time. Combined with the need for precise and accurate timing on sub- μ s timescales for some signals, this means that vastly more data must be generated and sent to a typical buffered digital output card than is necessary. A simple compression scheme that sends to hardware only a list of times at which each bit changes can yield a massive improvement. Although hardware with such a time-sequence mode is available, most implementations are buggy or overpriced and provide a relatively small number of channels compared to what is required for a cold atom experiment, especially one which relies heavily on digital systems such as DDS ICs. A modern field programmable gate array (FPGA) can provide many more channels, faster performance, and, coupled with inexpensive RAM, deeper memory at a much lower price. An FPGA is an IC consisting of look-up tables (to provide combinatorial logic functions), registers, and special functional blocks such as clock generators which can be interconnected almost arbitrarily using a switching fabric (which accounts for a large majority of the area of the IC). Design components for configuring an FPGA are typically called intellectual property cores, or simply cores. Below, we will discuss the development of a time-sequence digital output system with 256 channels and a 100 MHz clock using an evaluation board (Avnet ADS-XLX-V4LX-EVL60-G; Xilinx Virex-4 FPGA) available for \$700. Similar performance can be obtained from evaluation boards costing half as much.

The design consists of custom software, `dio_app`, running on a CPU core (Xilinx Microblaze) and peripheral components provided by Xilinx combined with custom cores. An on-chip bus (PLB bus) connects all components, while specialized point-to-point connections are used for high-speed data transfers. Communication to the outside world is provided by a 10/100 Mbps ethernet interface on the evaluation board. The board has a 32 MB RAM

chip connected to the FPGA. The lower 1 MB of the RAM is used to store `dio_app`; the rest holds the time sequence data.

The entire design operates from a single 100 MHz clock. The time-sequence core can be clocked independently and at a higher rate, but the useful signaling rate is limited to 100 MHz or even less in the targeted application by the long cables carrying the signals.

3.6.2 Hardware and Firmware

Figure 3.6 shows the structure of the FPGA sequencer. The custom time-sequence core, `plb_fastdio`, receives commands from and reports status to `dio_app` via a PLB bus connection. The actual sequence data is streamed to the core directly from the evaluation board's RAM through another custom core, `npi_wrapper`, which interfaces with the Xilinx memory controller core.

The time-sequence core contains 16 identical 16-bit output elements; both the data width and the number of elements are easily modified. Each element includes a buffer, refilled by the top-level part of the core, from which entries (see figure 3.7) are drawn. The data bits of the entry form the output of the element, while the delay bits of the entry are loaded into a down counter decremented every clock cycle. When the counter reaches zero, the next entry from the buffer is loaded. If all delay bits of an entry are set, all 32 bits of the next entry are treated as a delay, since a 16-bit delay is limited to less than 1 ms with a 100 MHz clock. If the buffer is empty when the down counter reaches zero (underflow), output from all elements halts until none are in an underflow condition. The simultaneous underflow of all elements is used to signal completion of the sequence.

The format for a block of time-sequence data is shown in figure 3.7. The top-level of the core uses the first four bytes of the block to select the elements to which the four entries should be distributed. If an element is not ready to accept the entry, distribution halts until it is ready. Entries for all elements are combined and ordered chronologically in RAM. For some sequences, slightly better performance could be obtained with a modified ordering, but at the cost of significantly increased software complexity.

The maximum output data rate for `plb_fastdio`, 25 Gb/s, far exceeds the rate data can be transferred from RAM, so only short bursts are possible at very high data rates without underflow occurring. The maximum sustained output rate, as limited by the RAM, is over 2 Gb/s, much faster than most commercially available digital output cards.

3.6.3 Software

The power-on FPGA configuration bitstream includes network bootloader software which requests an IP address from a DHCP server and then attempts to download a software file, specified in the DHCP reply, to RAM using the TFTP protocol. This allows the software to be modified without changing the FPGA bitstream. The bootloader is based on an open-

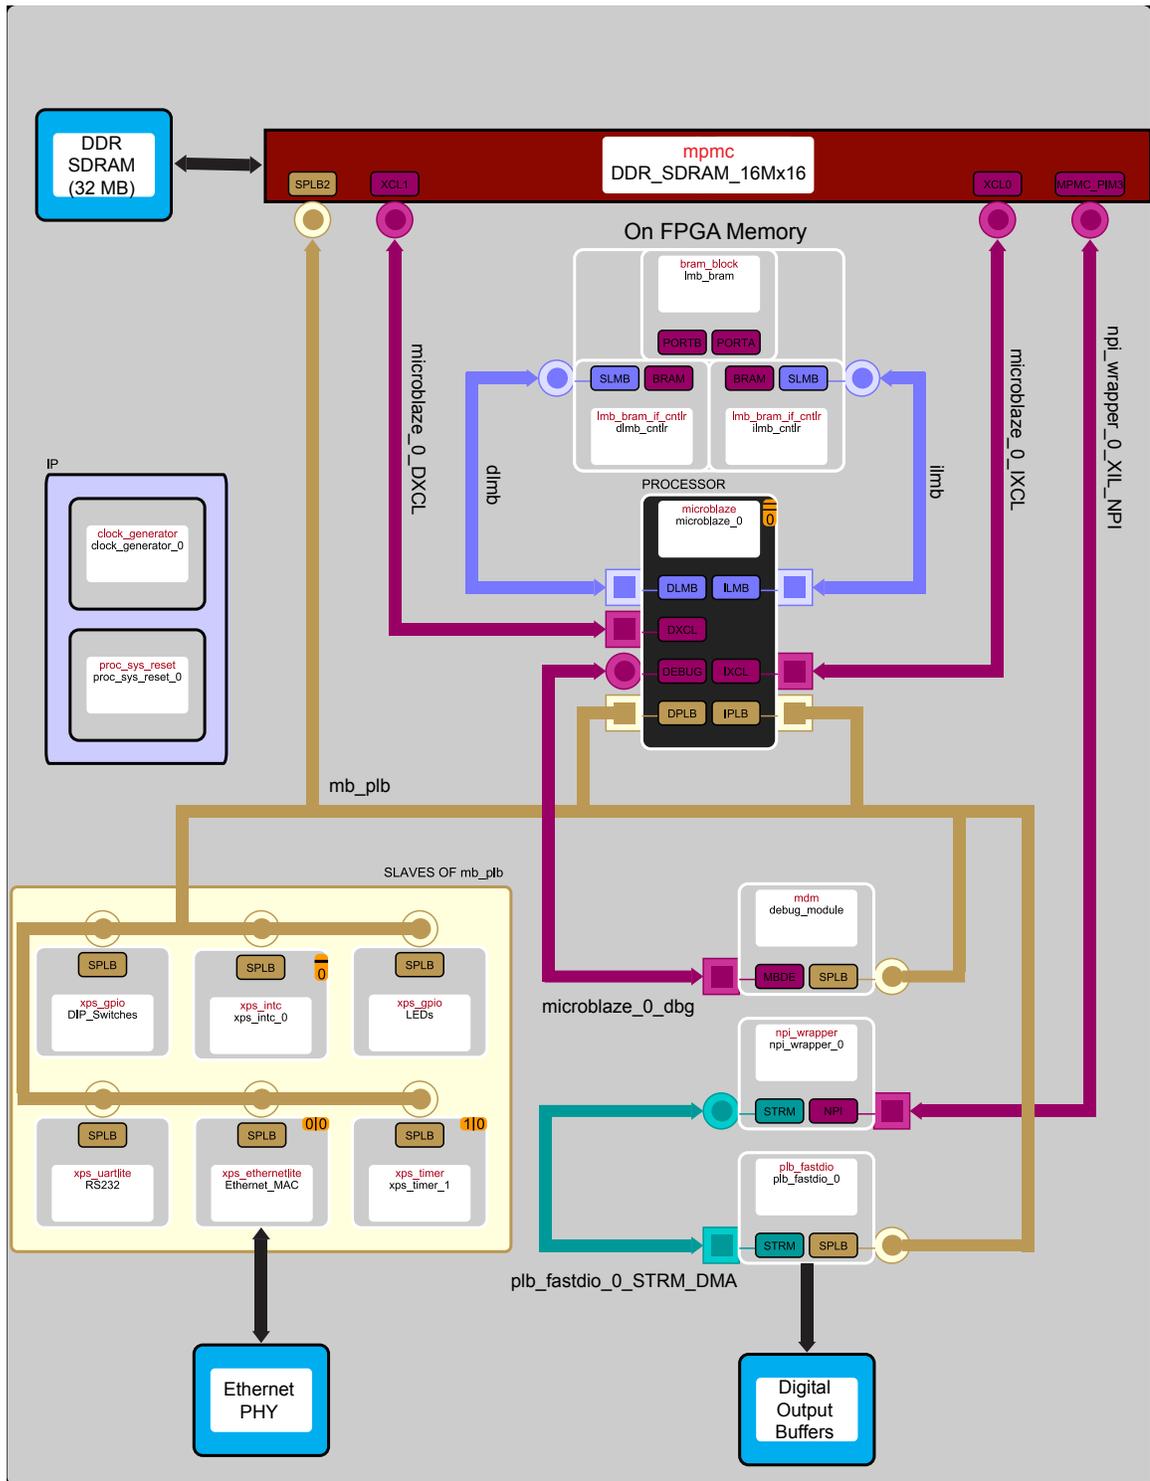


Figure 3.6 Structure of FPGA-based digital sequencer.

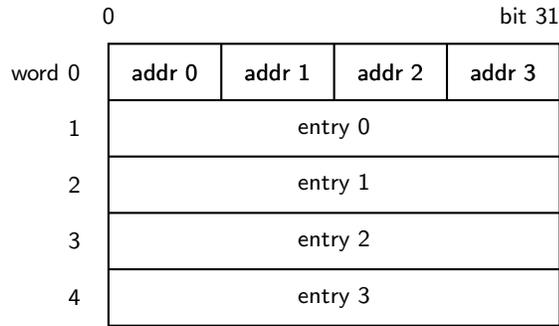


Figure 3.7 Data format generated by `out_fpga` and used by `plb_fastdio` core.

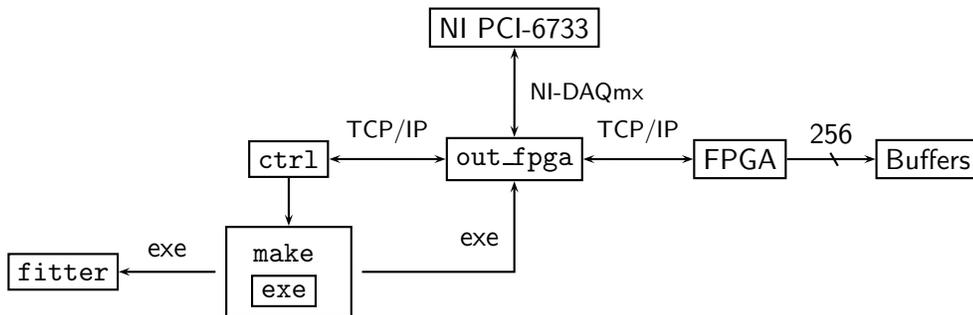


Figure 3.8 Schematic of control system software and FPGA sequencer. NI-DAQmx is the software driver for the PC-based analog output hardware (NI PCI-6733). For the description in section 3.1.3, `out_fpga` is replaced by `out_pc`, which communicates directly with driver for the PC-based digital output hardware (UEI PDL-DIO-64TS).

source project `tftpnaive` [91] but was extensively modified to work with the Microblaze processor.

The time-sequencer software, `dio_app`, communicates with `out_fpga`, a modified version of `out_pc`, (section 3.1.3) running on a computer via TCP/IP over ethernet (using an open-source network stack, `lwip`, modified by Xilinx). `out_fpga` converts digital output data into the proper format (figure 3.7) and transmits it to `dio_app`, which copies the data to RAM. `dio_app` then waits for a start command from `out_fpga` which it passes along to the time-sequence core. `out_fpga` polls `dio_app` for status information at regular intervals until the sequence completes or an error occurs. This software setup allows the FPGA sequencer to function as a drop-in replacement for the digital output card in the apparatus.

Chapter 4

Optical Lattice

Even before the availability of BEC, the behavior of cold atoms in periodic optical potentials (optical lattices) was a topic of significant interest, particularly the diffraction of atoms by 1D lattices [92, 93] and atom interferometry [94]. Methods of cooling using optical lattices were also explored [95–97]. Starting with a BEC, it is possible to load atoms directly into the ground state of an optical lattice. This approach enabled the study of, for example, Josephson dynamics and superfluid dissipation, familiar in condensed matter systems and, because of the essentially defect-free lattice potential, effects such as Bloch oscillations which are difficult to observe in condensed matter systems [98–100].

The realization of the superfluid/Mott insulator phase transition [9] accelerated interest in cold atoms in optical lattices as a model for strongly correlated condensed matter systems, with the hope of not just to observing well-understood effects, but gaining new insight into open questions. The addition of disorder to optical lattice potentials, discussed in the next chapter, opens new avenues in this direction. The creation of many-particle entanglement (see chapter 6) stirred interest in applications of optical lattices to quantum information [101].

In this chapter, we will begin with an overview of the relevant theory. We will then describe the design and implementation of a 3D optical lattice.

4.1 Optical Lattices

An atom in an electric field, assumed to be of single frequency ω , gains an induced dipole moment related to the field by the polarizability $\alpha(\omega)$:

$$\mathbf{d}(t) = \alpha(\omega)\mathbf{E}(t). \quad (4.1)$$

Note that $\alpha(\omega)$ will in general depend on the polarization of the electric field. See section 6.2 for further discussion. The interaction between the induced dipole and the field results in an energy shift

$$U = -\frac{1}{2}\langle \mathbf{d} \cdot \mathbf{E} \rangle = -\frac{1}{2}\text{Re} \alpha(\omega)\langle E^2 \rangle, \quad (4.2)$$

where the average is a cycle (time) average. We can see that the energy shift is proportional to the optical intensity.

In the simple model of a two-level atom with resonant frequency ω_0 we can write $\omega = \omega_0 + \delta$, where δ is the detuning. For red detuning, $\delta < 0$, the induced dipole is in phase with the electric field (in analogy with a forced classical harmonic oscillator) so the energy shift is negative; the atom is attracted to regions of higher optical intensity. For blue detuning, $\delta > 0$, the induced dipole is out of phase with the field and the energy shift is positive.

Now consider the combined electric field from a pair of identical, linearly polarized, counterpropagating optical beams:

$$E_z = E_0 \cos(kx - \omega t) + E_0 \cos(-kx - \omega t) = 2E_0 \cos kx \cos \omega t, \quad (4.3)$$

where $k = 2\pi/\lambda$ is the wavevector. The time average of the square of the electric field is given by

$$\langle E_x^2 \rangle = 2E_0^2 \cos^2 kx. \quad (4.4)$$

This periodic intensity gives rise to a periodic energy shift—an optical lattice. Additional pairs of beams with orthogonal polarizations can be added along the y and z directions to create a simple cubic 3D optical lattice. Multiple beams can be interfered in other ways to create, for example, triangular lattices or square lattices with twice the period [102].

We will typically write the potential arising from a 1D lattice as $U(x) = -V_0 \cos^2 kx$, where $V_0 > 0$ for a red-detuned lattice. For a 3D lattice,

$$U(\mathbf{r}) = -V_0(\cos^2 kx + \cos^2 ky + \cos^2 kz), \quad (4.5)$$

assuming that each beam has approximately the same wavelength. Note that the potential minima have a depth of $-3V_0$ in the 3D case. A natural energy scale for V_0 is

$$E_R = \frac{\hbar^2 k^2}{2m}. \quad (4.6)$$

This is the recoil energy—the kinetic energy change for an atom absorbing or emitting a photon with wavevector k —but in this context it is not related to light scattering.

4.2 Band Structure

4.2.1 Bloch States

The usual single-particle band structure description for electrons in a crystalline solid applies equally well to non-interacting atoms in the optical lattice. Invoking Bloch's theorem [103], the energy eigenstates can be written as the product of a plane wave $e^{i\mathbf{q}\cdot\mathbf{r}}$ and a function u

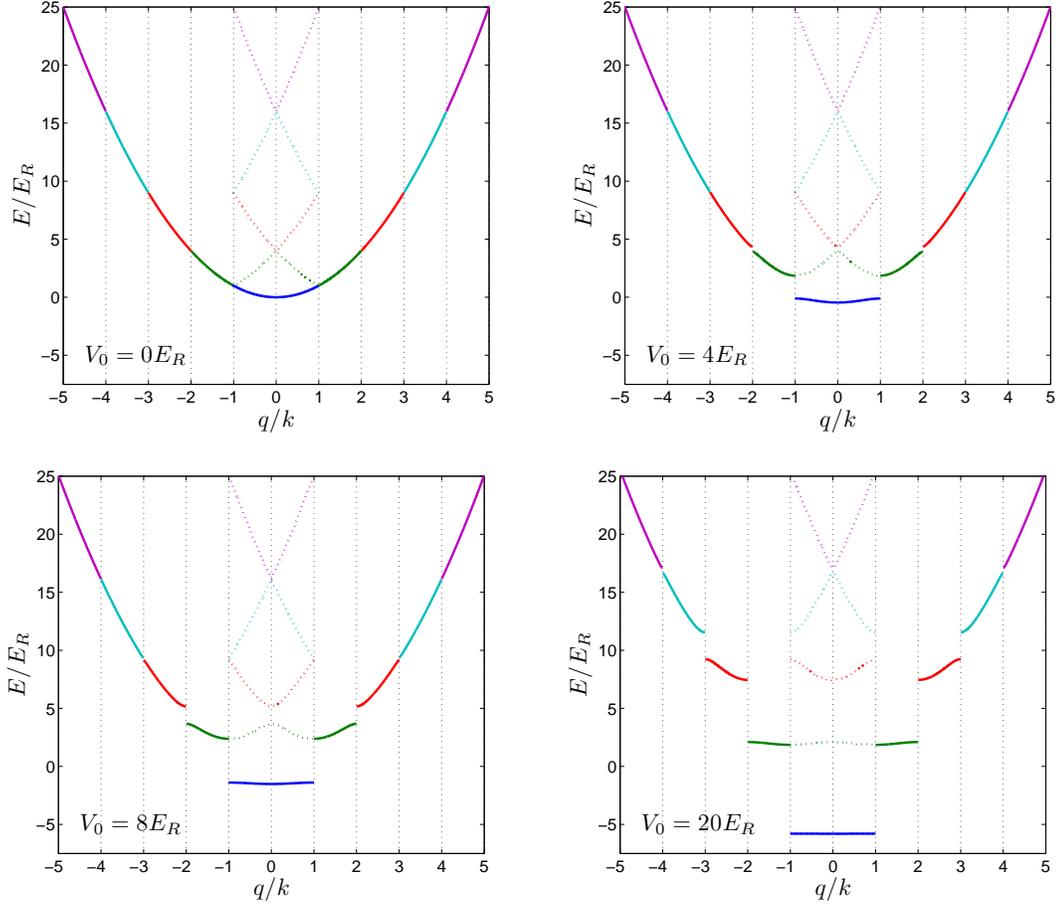


Figure 4.1 Bloch bands for a red-detuned 1D optical lattice potential $-V_0 \cos^2 kx$, with $E_r = \hbar^2 k^2 / 2m$. The bands are shown in the extended zone (solid lines) and reduced zone (broken lines) schemes. For the extended zone scheme, the lowest band occupies the first Brillouin zone, the first excited band occupies the second Brillouin zone, and so on.

with the periodicity of the lattice as

$$\psi(\mathbf{r}) = e^{i\mathbf{q}\cdot\mathbf{r}} u_{n,\mathbf{q}}(\mathbf{r}), \quad (4.7)$$

where \mathbf{q} is called the quasi-momentum or crystal momentum. Expanding both u and the lattice potential U into plane waves with the periodicity of the lattice yields

$$u_{n,\mathbf{q}}(\mathbf{r}) = \sum_{\mathbf{G}} c_{\mathbf{G}}(\mathbf{q}) e^{i\mathbf{G}\cdot\mathbf{r}}, \quad U(\mathbf{r}) = \sum_{\mathbf{G}} U_{\mathbf{G}} e^{i\mathbf{G}\cdot\mathbf{r}},$$

where \mathbf{G} are reciprocal lattice vectors. Substituting into the Schrodinger equation yields an infinite set of equations [103]:

$$\left(\frac{\hbar^2 |\mathbf{q} + \mathbf{G}|^2}{2m} - E(\mathbf{q}) \right) c_{\mathbf{G}}(\mathbf{q}) + \sum_{\mathbf{G}'} U_{\mathbf{G}'} c_{\mathbf{G}-\mathbf{G}'}(\mathbf{q}) = 0. \quad (4.8)$$

For a 1D optical lattice with $U(x) = -V_0 \cos^2 kx$, only the $G = \pm 2k$ components of the potential are non-zero; we can set the $G = 0$ component to zero since this amounts to a constant energy offset. So equation (4.8) becomes

$$\left(\frac{\hbar^2 (q + G)^2}{2m} - E(q) \right) c_G(q) + U_{2k} c_{G-2k}(q) + U_{-2k} c_{G+2k}(q) = 0, \quad (4.9)$$

with $U_{\pm 2k} = -V_0/4$. Here, $G = 2nk$, n integer, can be truncated (we use $-5 \leq n \leq 5$ in our calculations) to obtain a system of linear equations to be solved for $E(q)$. Figure 4.1 shows the result for several different lattice depths. The extended zone scheme is shown to indicate the mapping between Bloch states and plane waves when the optical lattice is turned on or off slowly (section 4.11).

The effective mass is a useful parameter for describing the effect of a shallow lattice on transport [104] and can be calculated in the usual manner as

$$m^* = \hbar^2 \left(\frac{\partial^2}{\partial q^2} E(q) \right)^{-1}. \quad (4.10)$$

From figure 4.1, we can see that as the lattice depth increases, the bands become flatter, making the effective mass larger.

4.2.2 Deep Lattice Case

We will now discuss the tight-binding approximation [104], applicable to the deep lattice case where it is more convenient to consider wavefunctions localized on lattice sites rather than Bloch waves spread over the entire lattice. Bloch wavefunctions can be written in a

basis of the maximally localized functions as

$$\psi_{n,q}(x) = N^{-1/2} \sum_{\text{sites } i} e^{iqx_i} w_n(x - x_i), \quad (4.11)$$

where N is a normalization constant and

$$w_n(x - x_i) = N^{-1/2} \sum_q e^{-iqx_i} \psi_{n,q}(x) = \sum_q e^{iq(x-x_i)} u_{n,q}(x) \quad (4.12)$$

is called a Wannier function and is localized around x_i . The Wannier states become energy eigenstates in the limit of infinite lattice depth. Subsequently, we will drop the band index n .

For a sufficiently deep lattice, $w(x - x_i)$ falls off quickly away from x_i and we can calculate the energy of the Bloch states in the tight-binding approximation by retaining only on-site and nearest neighbor terms:

$$\begin{aligned} E(q) &= \langle \psi_q | \hat{H} | \psi_q \rangle \\ &= N^{-1} \sum_{i,j} e^{iq(x_j - x_i)} \int_{-\infty}^{\infty} w^*(x - x_i) \hat{H} w(x - x_j) dx \\ &\approx N^{-1} \sum_i \left[\int_{-\infty}^{\infty} w^*(x) \hat{H} w(x) dx + (e^{iqa} + e^{-iqa}) \int_{-\infty}^{\infty} w^*(x - a) \hat{H} w(x) dx \right] \\ &= \varepsilon - 2J \cos qa, \end{aligned} \quad (4.13)$$

where $a = \pi/k$ is the lattice period and we have defined

$$\varepsilon = \int_{-\infty}^{\infty} w^*(x) \hat{H} w(x) dx, \quad (4.14a)$$

$$J = \int_{-\infty}^{\infty} w^*(x - a) \hat{H} w(x) dx. \quad (4.14b)$$

The tight-binding approximation is valid when J is much larger than the next-nearest neighbor term J' ; for $V_0/E_R = 6$, $J'/J = 0.035$. From (4.13), we can see that J is related to the width of the band by $J = (E_{\max} - E_{\min})/4 = [E(k) - E(0)]/4$ in the tight-binding approximation.

As the lattice is strengthened, the Wannier functions approach harmonic oscillator ground state functions (i.e., Gaussians), and the energy bands (figure 4.1) flatten and approach the harmonic oscillator energy levels for a single lattice site. Linearizing around the minima of the lattice potential, we find for the on-site trapping frequency

$$\omega_{\text{site}} = \sqrt{\frac{2k^2 V_0}{m}} = 2\sqrt{\frac{V_0}{E_R} \frac{E_R}{\hbar}}. \quad (4.15)$$

4.2.3 Diffraction from Lattice

The method used to calibrate the depth of the lattice potential (section 4.9.4) involves sudden loading into and release from the lattice [105]. Immediately after a 1D lattice is switched on, the state of the atoms can be written in terms of Bloch waves as

$$|\psi\rangle = |\phi_q\rangle = \sum_n |n, q\rangle \langle n, q | \phi_q\rangle,$$

where we have assumed that the atoms are initially in a plane wave state with wavevector q . After evolution for time τ , the state is

$$|\psi(\tau)\rangle = \sum_n |n, q\rangle \langle n, q | \phi_q\rangle e^{-iE_n(q)\tau/\hbar},$$

where $E_n(q)$ is the energy of the n th band. If the lattice is suddenly turned off at this point, the projection back to plane wave states is

$$|\psi(\tau)\rangle = \sum_m \sum_n |\phi_{q+2m\hbar k}\rangle \langle \phi_{q+2m\hbar k} | n, q\rangle \langle n, q | \phi_q\rangle e^{-iE_n(q)\tau/\hbar}. \quad (4.16)$$

For a sufficiently weak lattice, only the $n = 0$ and $n = 2$ bands are significantly populated; from the above equation, we can see that when τ is a multiple of $(E_2(q) - E_0(q))/\hbar$, the original wavefunction is recovered.

If the lattice is turned on adiabatically but turned off suddenly, the projection to plane wave states is simply

$$|\psi\rangle = \sum_m |\phi_{q+2m\hbar k}\rangle \langle \phi_{q+2m\hbar k} | n, q\rangle. \quad (4.17)$$

The $\langle \phi_{q+2m\hbar k} | n, q\rangle$ terms (for $m \neq 0$) result in the appearance of a diffraction pattern in this case (independent of the hold time, of course), as can be seen in figure 4.11.

Diffraction of atoms from an optical lattice can also be treated as diffraction from a grating, without reference to band structure or Bloch states [93]. If the grating (i.e., lattice) period is $d = \lambda/2 = \pi/k$, the possible momentum changes are $n \times 2\hbar k$, where n is an integer. Diffraction produced when the lattice is pulsed on for a sufficiently short time that the maximum phase shift is much less than 2π is said to be in the Raman-Nath regime. The fraction of atoms diffracted into various orders can be used to calculate the optical lattice depth [92], which provides an alternative calibration method. When applied to matter waves diffracting from a standing wave of light, diffraction in this regime is often called Kapitza-Dirac scattering. The original paper from Kapitza and Dirac [106] (in [107]) specifically discusses the diffraction of electrons from a standing light wave.

4.3 Bose-Hubbard Hamiltonian

4.3.1 Derivation

Jaksch *et. al.* [33] showed that cold bosonic atoms in an optical lattice are an almost ideal realization of the Bose-Hubbard model over much of the phase diagram. This is a non-trivial observation because it depends on the weak, short-range contact nature of the bare atom-atom interaction and the fact that, using a BEC, atoms can be loaded into the lowest band of lattice. Thus, to an excellent approximation, off-site interactions, beyond-nearest-neighbor tunneling, and higher bands can be ignored. We will see the role of each of these approximations in the following derivation.

The general Hamiltonian for a system of N bosons with contact (“hard-core”) interactions can be written

$$\hat{H} = \sum_{i=1}^N \left[\frac{\hat{\mathbf{p}}_i^2}{2m} + V(\mathbf{r}_i) \right] + \frac{U_0}{2} \sum_{i,j} \delta(\mathbf{r}_i - \mathbf{r}_j), \quad (4.18)$$

where $\hat{\mathbf{p}}^2 = -\hbar^2 \nabla^2$ and the effective interaction energy U_0 is given in equation (2.2). We will take the potential to consist of a rapidly varying lattice part V_0 and a slowly varying part V_T , which could, for example, represent a harmonic trapping potential, so $V(\mathbf{r}) = V_0(\mathbf{r}) + V_T(\mathbf{r})$. The Hamiltonian can be written in terms of a general set of field operators as

$$\hat{H} = \int \hat{\psi}^\dagger(\mathbf{r}) \left[\frac{\hat{\mathbf{p}}^2}{2m} + V(\mathbf{r}) \right] \hat{\psi}(\mathbf{r}) d^3r + \frac{U_0}{2} \int \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}(\mathbf{r}) \hat{\psi}(\mathbf{r}) d^3r, \quad (4.19)$$

where we have integrated over the delta function. We now choose the Wannier functions of the lowest band as the basis for the field operators, i.e., $\hat{\psi}(\mathbf{r}) = \sum_i \hat{b}_i w(\mathbf{r} - \mathbf{r}_i)$, where \hat{b}_i destroys a particle in the state with wavefunction $w(\mathbf{r} - \mathbf{r}_i)$. We can insert this into the above expression and make the tight-binding approximation by retaining only on-site and lowest order nearest neighbor terms, i.e., w_i^n and $w_i w_j$, i, j nearest neighbors. This gives

$$\begin{aligned} \hat{H} = \sum_{\langle i,j \rangle, i=j} \hat{b}_i^\dagger \hat{b}_j \int w^*(\mathbf{r} - \mathbf{r}_i) \left[\frac{\hat{\mathbf{p}}^2}{2m} + V(\mathbf{r}) \right] w(\mathbf{r} - \mathbf{r}_j) d^3r \\ + \frac{U_0}{2} \int \sum_i \hat{b}_i^\dagger \hat{b}_i^\dagger \hat{b}_i \hat{b}_i |w^*(\mathbf{r} - \mathbf{r}_i)|^4 d^3r. \end{aligned} \quad (4.20)$$

We now use $V(\mathbf{r}) = V_0(\mathbf{r}) + V_T(\mathbf{r})$. For $i = j$, the first term above becomes

$$\sum_i \hat{b}_i^\dagger \hat{b}_i \int w^*(\mathbf{r} - \mathbf{r}_i) V_T(\mathbf{r}) w(\mathbf{r} - \mathbf{r}_i) d^3r = \sum_i \hat{b}_i^\dagger \hat{b}_i V_T(\mathbf{r}), \quad (4.21)$$

where we have discarded the terms which do not vary with i , since they represent only

a constant energy offset. To arrive at the final expression, we used the assumption that V_T varies slowly (compared to the extent of the Wannier functions) to remove it from the integral and finally that the Wannier functions are normalized. For nearest-neighbor sites, the integral becomes instead

$$\int w^*(\mathbf{r} - \mathbf{r}_i) \left[\frac{\hat{\mathbf{p}}^2}{2m} + V_0(\mathbf{r}) \right] w(\mathbf{r} - \mathbf{r}_j) d^3r = \int w^*(\mathbf{r}) \left[\frac{\hat{\mathbf{p}}^2}{2m} + V_0(\mathbf{r}) \right] w(\mathbf{r} - a\hat{\mathbf{x}}) d^3r.$$

Here, we have again used the fact that V_T varies slowly to remove it from the integral, which then becomes zero because of the orthogonality of the Wannier functions. To arrive at the final result, we assume a cubic lattice with lattice period a along each direction; $\hat{\mathbf{x}}$ is the unit vector along the x direction.

With the following definitions,

$$J \equiv - \int w^*(\mathbf{r}) \left[\frac{\hat{\mathbf{p}}^2}{2m} + V_0(\mathbf{r}) \right] w(\mathbf{r} - a\hat{\mathbf{x}}) d^3r, \quad (4.22a)$$

$$U \equiv U_0 \int |w^*(\mathbf{r})|^4 d^3r, \quad (4.22b)$$

$$\varepsilon_i \equiv V_T(\mathbf{r}_i), \quad (4.22c)$$

and introducing the chemical potential μ to fix the total number atoms, we arrive at the traditional form of the Bose-Hubbard Hamiltonian,

$$\begin{aligned} \hat{H} &= -J \sum_{\langle i,j \rangle} \hat{b}_i^\dagger \hat{b}_j + \sum_i (\varepsilon_i - \mu) \hat{b}_i^\dagger \hat{b}_i + \frac{U}{2} \sum_i \hat{b}_i^\dagger \hat{b}_i^\dagger \hat{b}_i \hat{b}_i \\ &= -J \sum_{\langle i,j \rangle} \hat{b}_i^\dagger \hat{b}_j + \sum_i (\varepsilon_i - \mu) \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i (\hat{n}_i - 1), \end{aligned} \quad (4.23)$$

where $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$ counts the number of particles on site i . Note that the definitions for J and ε in equation (4.22) differ from those in (4.14) only in the additional approximations made here.

Let us consider the tunneling matrix element J , without the nearest-neighbor assumption, in a cubic lattice so that $V_0(\mathbf{r}) = V_0(x) + V_0(y) + V_0(z)$ and $w(\mathbf{r}) = w(x)w(y)w(z)$. We find, writing only the term for x explicitly,

$$\begin{aligned} -J_{ij} &= \int w^*(x - x_i) \left[\frac{\hat{p}_x^2}{2m} + V_0(x) \right] w(x - x_j) dx \\ &\quad \times \int w^*(y - y_i) w(y - y_j) dy \int w^*(z - z_i) w(z - z_j) dz + \dots \end{aligned}$$

Due to the orthogonality of the Wannier functions, the second factor is zero unless $y_i = y_j$; the third is zero unless $z_i = z_j$. We see that J_{ij} is zero if more than one component of

\mathbf{r}_i and \mathbf{r}_j differ, so that tunneling along diagonal directions is forbidden. Thus, there is no next-nearest neighbor tunneling in a 3D cubic lattice. For nearest neighbors in a cubic lattice, the matrix element reduces to

$$J = - \int w^*(x) \left[\frac{\hat{p}_x^2}{2m} + V_0(x) \right] w(x-a) dx. \quad (4.24)$$

The interaction energy also reduces to a one-dimensional integral in this case,

$$U = U_0 \left[\int_{-\infty}^{\infty} |w(x)|^4 dx \right]^3. \quad (4.25)$$

In the limit of a deep lattice, we can calculate U using Gaussian ground states, which yields

$$U = U_0 \left(\frac{m\omega_{\text{site}}}{\hbar} \right)^{3/2} = U_0 \left(\frac{k^2}{2\pi} \sqrt{\frac{V_0}{E_R}} \right)^{3/2},$$

where ω_{site} is given by (4.15). In the same limit, J can be calculated from the band width determined from the Mathieu equation as [108]

$$J = \frac{4E_R}{\sqrt{\pi}} \left(\frac{V_0}{E_R} \right)^{3/4} e^{-2\sqrt{V_0/E_R}}.$$

From these approximate equations, it is clear that J and U depend on both the lattice depth and period.

4.3.2 Double Well Model

Before discussing the mean-field solution to the Bose-Hubbard Hamiltonian, it is instructive to consider a system of just two lattice sites, i.e., a double well system [10]. We assume that particles in well i can only occupy the state $|\psi_i\rangle$, with corresponding annihilation operator \hat{a}_i . The second-quantized Hamiltonian for this system is

$$\hat{H} = -J(\hat{a}_2^\dagger \hat{a}_1 + \hat{a}_1^\dagger \hat{a}_2) + \frac{U}{2} \sum_{i=1,2} \hat{a}_i^\dagger \hat{a}_i^\dagger \hat{a}_i \hat{a}_i = -J(\hat{a}_2^\dagger \hat{a}_1 + \hat{a}_1^\dagger \hat{a}_2) + \frac{U}{2} \sum_{i=1,2} \hat{n}_i(\hat{n}_i - 1), \quad (4.26)$$

where $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ counts the number of particles in a well and U and J are defined as in (4.22). In the non-interacting ($U = 0$) case, the Hamiltonian can be written as

$$\hat{H} = -J\hat{a}_+^\dagger \hat{a}_+ + J\hat{a}_-^\dagger \hat{a}_-, \quad (4.27)$$

where $\hat{a}_+ = (\hat{a}_1 + \hat{a}_2)/\sqrt{2}$ and $\hat{a}_- = (\hat{a}_1 - \hat{a}_2)/\sqrt{2}$ destroy particles in the familiar symmetric and anti-symmetric states, respectively. We can see that these states, for which the number of particles in a given well is indeterminate, are the eigenstates of the non-interacting

Hamiltonian. On the other hand, the eigenstates in the opposite limit $J \rightarrow 0$ are clearly number states on each well.

In the Heisenberg picture, we have from (4.26)

$$i\hbar \frac{d\hat{a}_1}{dt} = [\hat{a}_1, \hat{H}] = U\hat{n}_1\hat{a}_1 - J\hat{a}_2, \quad (4.28a)$$

$$i\hbar \frac{d\hat{a}_2}{dt} = [\hat{a}_2, \hat{H}] = U\hat{n}_2\hat{a}_2 - J\hat{a}_1. \quad (4.28b)$$

In the limit of a large number of particles, we can replace the operators with their expectation values, so that, e.g., $n_i = |a_i|^2$. Introducing phases ϕ_i by writing $a_i = \sqrt{n_i}e^{i\phi_i}$ and manipulating the above equations, we obtain

$$\frac{dn_1}{dt} = -\frac{dn_2}{dt} = \frac{2J}{\hbar} \sqrt{n_1 n_2} \sin(\phi_1 - \phi_2), \quad (4.29)$$

a result first found in the context of Josephson junctions. By considering the Josephson oscillations described by the linearized version of this equation, it can be shown that the fluctuations in phase and number are given by

$$\langle (\phi_1 - \phi_2)^2 \rangle = \sqrt{\frac{U + 2J/N}{2JN}}, \quad (4.30)$$

$$\langle (n_1 - n_2)^2 \rangle = \sqrt{\frac{2JN}{U + 2J/N}}, \quad (4.31)$$

where $N = n_1 + n_2$ and the phase fluctuations are assumed to be small. We see that increasing J relative to U suppresses fluctuations in the relative phase between the wells, but increases fluctuations in the number of atoms in each well. This captures the essence of the superfluid/Mott insulator phase transition.

4.3.3 Bose-Hubbard Mean-Field Theory

A variety of methods have been employed to calculate the phase diagram for the Bose-Hubbard Hamiltonian (4.23), both at $T = 0$ and finite temperature. These include quantum Monte Carlo [29, 32] and perturbation theory [109] methods. Here, we will consider what is arguably the simplest approach, site-decoupled mean-field theory [10, 20, 110] at zero temperature. In three dimensions, mean-field theory agrees well with results obtained by other methods, and it becomes exact in the limit of infinite dimensions. Experimental results from our group and others [111, 112] also agree well with mean-field theory results.

We begin by expanding the hopping term in (4.23) as follows:

$$\hat{b}_i^\dagger \hat{b}_j = \langle \hat{b}_i^\dagger \rangle \hat{b}_j + \hat{b}_i^\dagger \langle \hat{b}_j \rangle - \langle \hat{b}_i^\dagger \rangle \langle \hat{b}_j \rangle + (\hat{b}_i^\dagger - \langle \hat{b}_i^\dagger \rangle)(\hat{b}_j - \langle \hat{b}_j \rangle). \quad (4.32)$$

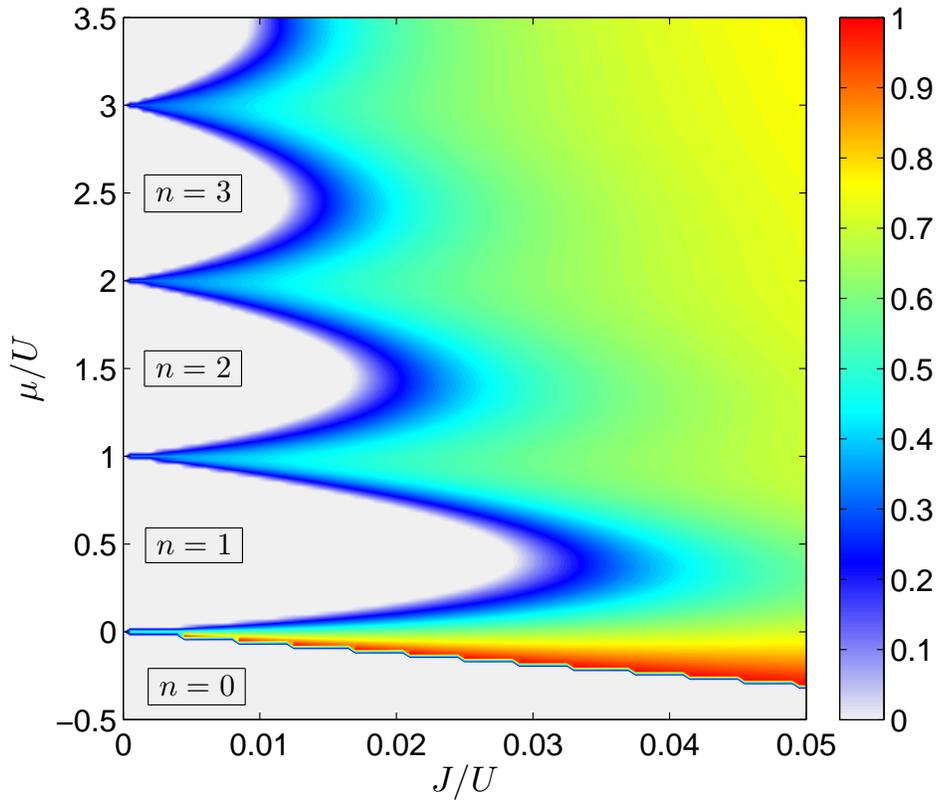


Figure 4.2 Mean-field theory calculation of ground state condensate fraction for the Bose-Hubbard model in 3D. The $n = 0$ through $n = 3$ Mott insulator lobes are labeled.

The mean-field approximation consists of dropping the second order term $(\hat{b}_i^\dagger - \langle \hat{b}_i^\dagger \rangle)(\hat{b}_j - \langle \hat{b}_j \rangle)$. This allows us to separate the Hamiltonian into a sum of single site Hamiltonians, $\hat{H}_{\text{MF}} = \sum_i \hat{H}_i$, with

$$\hat{H}_i = -zJ(\psi^* \hat{b}_i + \psi \hat{b}_i^\dagger - |\psi|^2) + (\varepsilon_i - \mu) \hat{n}_i + \frac{U}{2} \hat{n}_i(\hat{n}_i - 1), \quad (4.33)$$

where z is the number of nearest-neighbor sites ($z = 6$ for a 3D cubic lattice) and the field $\psi \equiv \langle \hat{b} \rangle = \langle \hat{b}^\dagger \rangle^*$ is a free parameter. Then, \hat{H}_i is diagonalized in the on-site number basis truncated at a sufficiently large n , $|0\rangle, |1\rangle, \dots, |n\rangle$, and the value of ψ which minimizes the ground state energy is determined. If this ψ is non-zero, we label the state superfluid, otherwise, (Mott) insulator. We identify $|\psi|^2 / \langle \hat{n}_i \rangle$ as the condensate fraction, where $\langle \hat{n} \rangle$ is the expectation value of the number operator in the ground state. This choice follows from the Penrose/Onsager definition [113, 114] of condensate fraction as the largest eigenvalue of the single particle density matrix $\rho_{ij} = \langle \hat{b}_i^\dagger \hat{b}_j \rangle / N$, where N is the total number of particles. The matrix has dimensions $N_s \times N_s$, where N_s is the number of lattice sites. In the uniform mean-field calculation, every diagonal element of the matrix is $\langle \hat{n} \rangle$ and every off-diagonal element is $|\psi|^2$. For a matrix of this form, only one eigenvalue is of order one, $[(N_s - 1)|\psi|^2 + \langle \hat{n} \rangle] / N \approx |\psi|^2 / \langle \hat{n} \rangle$, since $N / N_s = \langle \hat{n} \rangle$.

The reduction of condensate fraction below one in the superfluid state is called quantum depletion and can be attributed to the ground state of the interacting Hamiltonian containing overlap with excited states of the non-interacting ($U = 0$) Hamiltonian. Figure 4.2 shows the mean-field result for condensate fraction in a uniform system ($\varepsilon_i = 0$ for all i).

Alternately, the mean-field result can be obtained variationally using the Gutzwiller state,

$$|\Psi\rangle = \prod_{\text{sites } i} \sum_n \frac{c_n}{\sqrt{n!}} (\hat{b}_i^\dagger)^n |0\rangle, \quad (4.34)$$

where n must be truncated for numerical calculations, as above. We can compute the expectation value of the Bose-Hubbard Hamiltonian (4.23) for this state, $\langle \Psi | \hat{H} | \Psi \rangle$, and minimize it with respect to the values c_n to find the mean-field ground state (with the constraint $\langle \Psi | \Psi \rangle = 1$).

The Gutzwiller state, consisting of independent single-site quantum states, should be contrasted to the maximally delocalized superfluid state, a product over Bloch states rather than single-site states:

$$|\Psi_{\text{SF}}\rangle \propto \left(\sum_{\text{sites } i} \hat{b}_i^\dagger \right)^N |0\rangle, \quad (4.35)$$

where N is the total number of particles. Note that the Gutzwiller state, which is used in mean-field theory to describe the superfluid state as well as the insulating state, omits all the multi-site terms found in $|\Psi_{\text{SF}}\rangle$.

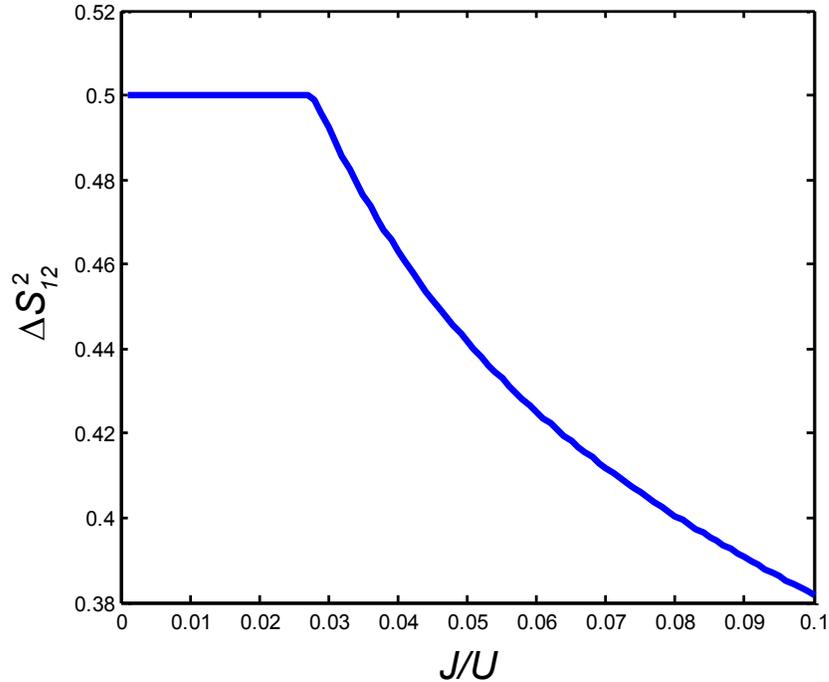


Figure 4.3 Variance of \hat{S}_{12} between two lattice sites calculated using mean-field ground state for $\mu/U = 0.5$ (in 3D).

The superfluid/Mott insulator phase boundary can be found analytically by treating the mean-field terms as a perturbation to the $J = 0$ Hamiltonian [10]. The boundary of the n filling Mott insulator lobe is given by

$$zJ = \frac{(nU - \mu)(\mu - (n - 1)U)}{\mu + U}. \quad (4.36)$$

The tip of lobe is the point specified by $U/J = z(2n + 1 + 2\sqrt{n(n + 1)})$.

4.3.4 Phase Fluctuations

Motivated by the double well model, the superfluid/Mott insulator phase transition is often described qualitatively in terms of increasing fluctuations of the relative phase of atoms on different lattice sites. Although no phase operator can be defined in quantum mechanics, we can introduce $\hat{S}_{12} = \hat{S}_1\hat{C}_2 - \hat{S}_2\hat{C}_1$, which corresponds to the sine of the phase difference between two sites in the classical limit [115]. Here $\hat{C} = (\hat{E}_- + \hat{E}_+)/2$ and $\hat{S} = (\hat{E}_- - \hat{E}_+)/2i$ are Hermitian operators. The operators \hat{E}_\pm , corresponding to $\exp(\mp i\phi)$ in the classical limit, are defined by $\hat{E}_+|n\rangle = |n + 1\rangle$, all n , and $\hat{E}_-|n\rangle = |n - 1\rangle$, $n \geq 1$, $\hat{E}_-|0\rangle = 0$. As shown in figure 4.3, the variance of \hat{S}_{12} , $\Delta\hat{S}_{12}^2 = \langle\hat{S}_{12}^2\rangle$, for the ground state of the mean-field calculation increases as the superfluid/insulator phase transition is approached and saturates at its maximal value of 0.5 in the Mott insulator state, where the ground state of

the mean-field calculation is a number state.

4.3.5 External Confinement

For the optical lattice in our apparatus, and most others at the time of writing, red-detuned, single Gaussian mode beams are used to create the lattice. To lowest order, the harmonic oscillator frequency for the transverse potential $V(r) = V_0 \exp(-2r^2/w^2)$ arising from a Gaussian beam with waist w is

$$\omega_{\text{ext}} = \frac{2}{w} \sqrt{\frac{V_0}{m}}. \quad (4.37)$$

This, taken together with a harmonic trapping potential along the same direction ω_0 (see, e.g., section 2.9), results in an overall harmonic confinement V_T characterized by $\omega = \sqrt{\omega_0^2 + \omega_{\text{ext}}^2}$. A more exact (and smaller) result can be obtained by evaluating the integral in (4.21). To avoid the complication of dealing with changes in effective mass (section 4.2.1), it may sometimes be desirable to characterize the trapping potential by its spring constant $m\omega^2$.

The potential V_T results in a site-dependent energy ε_i . In the local density approximation, we calculate the superfluid fraction and condensate fraction by taking a simple weighted average of the uniform mean-field results over the range of local chemical potentials $\mu - \varepsilon_i$ present, from the value at the edge of the cloud (the boundary of the $n = 0$ lobe in figure 4.2) to the global chemical potential μ at the center. In other words, we assume that the slowly varying confining potential does not locally affect the mean-field result. The procedure can be shown schematically by drawing a vertical line on the figure (see figure 5.16). Spatially, this represents a shell structure of alternating regions of Mott insulator and superfluid for sufficiently high lattice depth. As the lattice depth is increased, the Mott insulator regions grow while the superfluid regions shrink. We have observed good agreement between the condensate fraction predicted by this calculation and experimental results.

In analogy with the electrochemical potential often used in the context of semiconductors [104], we can introduce an effective chemical potential $\tilde{\mu}_i = \mu - V_T(\mathbf{r}_i)$ to remove the dependence on V_T of ε_i . This is useful when another site-dependent contribution to ε_i is present, for example, a disordered potential.

4.4 Expansion from the Lattice

By considering the free expansion of the wavefunctions on each lattice site, we can produce a more quantitative picture of what is observed when the lattice is turned off suddenly after adiabatic loading into the ground state. This approach also allows us to calculate properties of higher-order density correlations, providing a useful measurement technique for the Mott insulator regime (see sections 5.6 and 6.6.5).

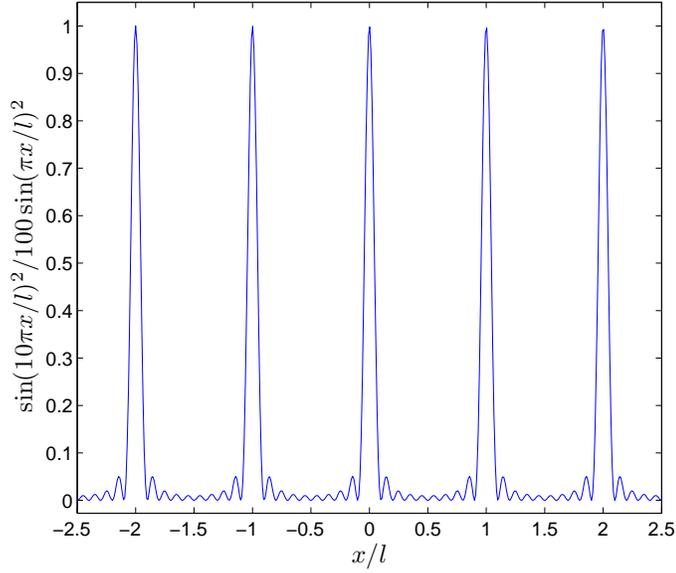


Figure 4.4 Single particle density (without envelope and normalized) after expansion from 1D lattice with 10 sites for superfluid regime.

In the tight binding regime, the state of an atom on a lattice site i can be approximated by a harmonic oscillator state, which expands as a Gaussian wave packet [116] if the lattice is turned off suddenly:

$$\psi_i(\mathbf{r}, t) = w(\mathbf{r} - \mathbf{r}_i, t) e^{im|\mathbf{r} - \mathbf{r}_i|^2 / 2\hbar t}, \quad (4.38)$$

where $w(\mathbf{r} - \mathbf{r}_i, t)$ is the envelope. To simplify the phase factor, we have assumed t is large enough that the width of $\psi_i(\mathbf{r}, t)$ is much greater than the initial width. We further assume that the width of $\psi_i(\mathbf{r}, t)$ is much larger than $|\mathbf{r}_i|$, all i , so that we can take $w(\mathbf{r} - \mathbf{r}_i, t) = w(\mathbf{r}, t)$. In other words, the atom cloud has expanded to be much larger than its size in the lattice. This assumption is well satisfied by the long time-of-flight used for these measurements. For a shallower lattice, the appropriate Wannier function can be used as the envelope $w(\mathbf{r} - \mathbf{r}_i, t)$.

4.4.1 Single Particle Density

Including contributions for each lattice site, the field operator in second quantized notation is

$$\hat{\psi}(\mathbf{r}, t) = \sum_i \psi_i(\mathbf{r}, t) \hat{b}_i. \quad (4.39)$$

Using (4.38) and associated assumptions, the density operator $\hat{n}(\mathbf{r}, t) = \hat{\psi}^\dagger(\mathbf{r}, t) \hat{\psi}(\mathbf{r}, t)$ can be written

$$\hat{n}(\mathbf{r}, t) = |w(\mathbf{r}, t)|^2 \sum_{i,j} e^{im(\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{r} / \hbar t} \hat{b}_i^\dagger \hat{b}_j, \quad (4.40)$$

where we have discarded a term proportional to $|\mathbf{r}_j|^2 - |\mathbf{r}_i|^2$ in the exponential. To calculate $\langle \hat{n}(\mathbf{r}, t) \rangle$ in the superfluid regime, we use the mean-field result that the on-site state becomes a coherent state for $U/J \rightarrow 0$. Thus, $\langle \hat{b}_i \rangle = \sqrt{\bar{n}_i}$ and we have

$$\langle \hat{n}(\mathbf{r}, t) \rangle = |w(\mathbf{r}, t)|^2 \sum_{i,j} e^{im(\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{r} / \hbar t} \sqrt{\bar{n}_i} \sqrt{\bar{n}_j}. \quad (4.41)$$

Considering just a 1D lattice with period a , we can write $\mathbf{r}_i = i \times a$ so

$$\langle \hat{n}(x, t) \rangle = |w(x, t)|^2 \sum_{i,j} e^{2\pi i(i-j)x/l} \sqrt{\bar{n}_i} \sqrt{\bar{n}_j}, \quad (4.42)$$

where $l = 2\hbar kt/m$. Notice that l is the distance an atom with momentum $2\hbar k$ and mass m would travel in time t . Assuming N lattice sites with an average of one atom on each site (so, $\sqrt{\bar{n}_i} = 1$, all i), the geometric sums in (4.42) can be evaluated to give

$$\langle \hat{n}(x, t) \rangle = |w(x, t)|^2 \frac{\sin^2(\pi N x/l)}{\sin^2(\pi x/l)}. \quad (4.43)$$

This function is plotted without the envelope in figure 4.4. The peaks at integer x/l are exactly as we would expect from section 4.2.3.

For the Mott insulator, $\langle \hat{b}_i^\dagger \hat{b}_j \rangle = \delta_{ij} n_i$ so

$$\langle \hat{n}(\mathbf{r}, t) \rangle = |w(\mathbf{r}, t)|^2 \sum_{i,j} e^{im(\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{r} / \hbar t} \delta_{ij} \hat{n}_i = N |w(\mathbf{r}, t)|^2. \quad (4.44)$$

We see that only the unmodulated envelope remains.

4.4.2 Density-Density Correlations

Hidden in the seemingly featureless cloud produced by atoms expanding from the lattice in the Mott insulator regime are Hanbury Brown-Twiss type correlations [117–119]. These are sometimes called atom shot noise correlations since they appear within the atom shot noise which is present even without the lattice (or a BEC). Specifically, if the average number of atoms within the volume imaged onto a pixel of the camera is n , we expect shot noise fluctuations on the order of \sqrt{n} if interactions do not affect the expansion.

Considering two atoms on different lattice sites, a and b , and two pixels, 1 and 2, there are two possible paths: $a \rightarrow 1, b \rightarrow 2$ and $a \rightarrow 2, b \rightarrow 1$. Since the atoms are indistinguishable, the complex probability amplitudes of the two paths must be added and destructive or constructive interference will occur depending on the distance between the two pixels. As we show below, this higher-order interference pattern can be extracted from

the images using the following observable:

$$C(\mathbf{d}) = \frac{\int \langle \hat{n}(\mathbf{r}) \hat{n}(\mathbf{r} - \mathbf{d}) \rangle d^3r}{\int \langle \hat{n}(\mathbf{r}) \rangle \langle \hat{n}(\mathbf{r} - \mathbf{d}) \rangle d^3r} - 1, \quad (4.45)$$

where \mathbf{d} is the distance between pixels, \hat{n} is the density operator after time-of-flight, and the expectation value is for the state in the lattice. This expression is simply the average of the autocorrelation of each density distribution, i.e., image, normalized by the autocorrelation of the average density distribution, i.e., the average of all images.

In the previous section, we considered the single-particle density operator. Now we consider the two-particle density-density correlation operator:

$$\begin{aligned} \hat{n}(\mathbf{r}_1, t) \hat{n}(\mathbf{r}_2, t) &= |w(\mathbf{r}_1, t)|^2 |w(\mathbf{r}_2, t)|^2 \\ &\times \sum_{i,j,k,l} e^{-im|\mathbf{r}_1 - \mathbf{r}_i|^2/2\hbar t} e^{im|\mathbf{r}_1 - \mathbf{r}_j|^2/2\hbar t} e^{-im|\mathbf{r}_2 - \mathbf{r}_k|^2/2\hbar t} e^{im|\mathbf{r}_2 - \mathbf{r}_l|^2/2\hbar t} \hat{b}_i^\dagger \hat{b}_j \hat{b}_k^\dagger \hat{b}_l. \end{aligned} \quad (4.46)$$

For a Mott insulator, we have $\hat{b}_i^\dagger \hat{b}_j = \delta_{ij} \hat{n}_i$ so

$$\hat{b}_i^\dagger \hat{b}_j \hat{b}_k^\dagger \hat{b}_l = \hat{b}_i^\dagger (\hat{b}_k^\dagger \hat{b}_j + \delta_{jk}) \hat{b}_l = \delta_{ij} \delta_{kl} \hat{n}_i \hat{n}_k + \delta_{il} \delta_{jk} \hat{n}_i \hat{n}_k + \dots, \quad (4.47)$$

where we have omitted the normal ordering and $\delta_{ij} \delta_{jk} \delta_{kl}$ terms in the last expression. These terms do not contain information about spatial ordering.

Substituting (4.47) into (4.46), the $\delta_{ij} \delta_{kl}$ term gives $\sum_{i,k} \hat{n}_i \hat{n}_k = N^2$ while the $\delta_{il} \delta_{jk}$ term yields

$$\begin{aligned} \sum_{i,k} e^{-im|\mathbf{r}_1 - \mathbf{r}_i|^2/2\hbar t} e^{im|\mathbf{r}_1 - \mathbf{r}_k|^2/2\hbar t} e^{-im|\mathbf{r}_2 - \mathbf{r}_k|^2/2\hbar t} e^{im|\mathbf{r}_2 - \mathbf{r}_i|^2/2\hbar t} \hat{n}_i \hat{n}_k \\ = \sum_{i,k} e^{im(\mathbf{r}_i - \mathbf{r}_k) \cdot (\mathbf{r}_1 - \mathbf{r}_2) / \hbar t} \hat{n}_i \hat{n}_k \end{aligned}$$

Inserting these results into (4.45) and defining $\mathbf{r} = \mathbf{r}_1$ and $\mathbf{d} = \mathbf{r}_1 - \mathbf{r}_2$ yields

$$C(\mathbf{d}) = \frac{1}{N^2} \sum_{i,k} e^{im(\mathbf{r}_i - \mathbf{r}_k) \cdot \mathbf{d} / \hbar t} n_i n_k. \quad (4.48)$$

This expression is identical in form to (4.41). If we again consider a 1D system of N lattice sites with one atom on each site, so $n_i = 1$ for all i , we obtain

$$C(d) = \frac{1}{N^2} \frac{\sin^2(\pi N d / l)}{\sin^2(\pi d / l)}. \quad (4.49)$$

Thus, the density-density correlation $C(\mathbf{d})$ for the Mott insulator will look just like the single particle density distribution for the superfluid but without the envelope, normalized

out in our definition of $C(\mathbf{d})$. A measured density-density correlation image can be found in section 5.6, where the effect of disorder on $C(\mathbf{d})$ is discussed.

4.5 Technical Overview

The optical lattice light, up to several hundred mW for each beam, is generated by a titanium:sapphire (Ti:Sa) laser. We usually work at a lattice wavelength of 812 nm; this wavelength and the corresponding value of E_R should be assumed unless otherwise stated.

The trajectories of the lattice beams, shown in figure 2.8, were chosen with the work discussed in chapter 6 in mind. Furthermore, because the glass cell is not anti-reflection coated, it is important that none of the beams be near normal incidence in order to prevent light reflected from the glass from overlapping with the atoms. The beams are retro-reflected by mirrors mounted to the optical table to create standing waves.

To prevent unwanted interference between different lattice beams, the polarizations of the beams are chosen to be mutually orthogonal (there are two ways to do this, related by rotating each polarization 90 degrees) and the frequencies of the beams are offset by several tens of MHz so that any remaining interference averages out on all timescales relevant to the atoms.

4.6 Light Source

Figure 4.5 shows the system used to provide light for the optical lattice. The light for each beam is carried in an optical fiber to the science cell region.

4.6.1 Ti:Sapphire Laser

The Ti:Sa laser (Tekhnoscan TIS-SF-077) is pumped with 10 W of 532 nm light from a frequency-doubled Nd:YVO₄ laser (Coherent Verdi V-18). A home-built interlock circuit prevents the Verdi from operating if the Ti:Sa cooling water flow drops below a certain threshold.

Before purchasing the Verdi, we attempted to use a less expensive pump laser (ELS Monodisk). However, this device suffered from severe intensity fluctuations, often shutting off completely for hundreds of microseconds. After two replacements with similar problems and a visit from the president of the company, we gave up on the Monodisk. Unfortunately, the Verdi was not without its own (expensive) problems: the laser head had to be replaced after just over a year.

The Ti:Sa is tuned using, in order of increasing resolution, a birefringent filter, a thin etalon, a thick etalon, and cavity mirror piezos (which change the cavity length). The laser can be locked to a provided external reference cavity to reduce its linewidth to below 100 kHz. This does not have any measurable effect on the performance of the lattice.

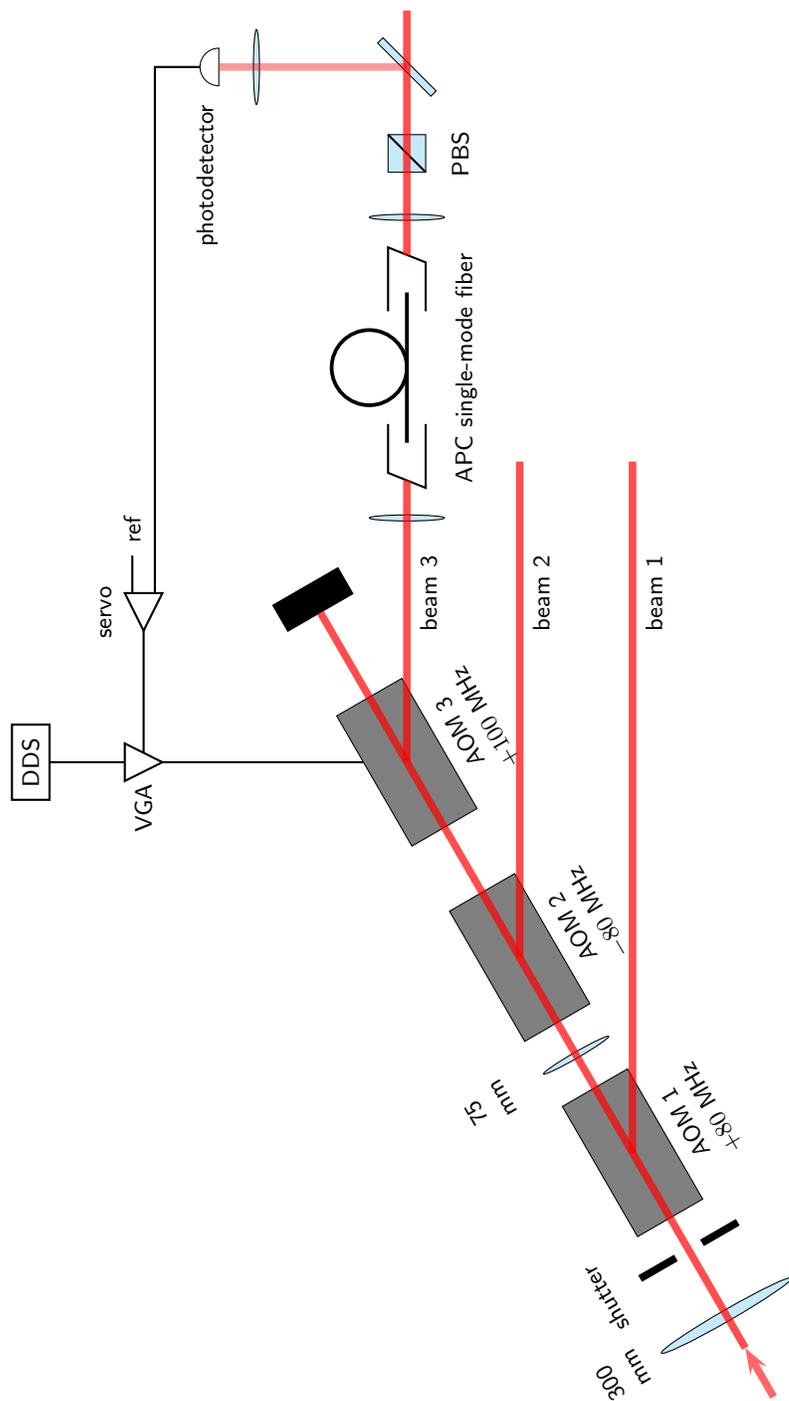


Figure 4.5 Distribution of light for the optical lattice. Light from the Ti:Sa laser is picked off and frequency shifted by three AOMs. The intensity exiting the fiber is monitored by a photodetector and stabilized by feeding back to the RF power driving the AOM, as shown for beam 3. Identical systems (not shown) are used for beams 1 and 2. The DDS can be used to control the RF amplitude during open loop operation, e.g., for fast pulses.

4.6.2 Hard Drive Shutter

A mechanical shutter is used to block the light from the Ti:Sa laser when the optical lattice is off. The shutter is constructed from a computer hard disk drive, based on the design in [120, 121]. The rise and fall time of the shutter is about 1 ms, comparable to the commercial shutters used elsewhere in the apparatus. The shutter flag is a polished piece of sheet metal that reflects the beam into the hard drive case, which is left intact except for entry and exit holes, unlike in [120]. This minimizes the amount of light leaking from the shutter in either direction and prevents the flag from heating significantly, even with over a watt of light incident.

4.6.3 AOMs and Fiber Coupling

The shutter is followed by three AOMs in series which deflect light into three optical fibers, one for each lattice direction. The AOMs (Neos 23080; nominally 80 MHz) are required for two reasons. First, they provide slight frequency offsets for each of the beams. This helps prevent interference between different lattice beams. The frequency shifts for lattice beams 1, 2, and 3 are 80 MHz, -80 MHz, and 100 MHz, respectively. These specific frequencies were chosen to minimize DDS spurs (since the DDS clock is an even 500 MHz). Secondly, the AOMs are used to servo the intensity of the lattice, as discussed below.

The beam deflected from each AOM is coupled into a single-mode polarization-maintaining fiber using a New Focus 9091 mount. This particular mount has several major drawbacks. It is impossible to unscrew or rotate the fiber after installation (the fiber is attached to small, recessed part held by set screw), necessitating a half-wave plate before the fiber. The coarse adjustment screws have severe backlash while the fine adjustment screws do have enough travel to be useful.

We have been able to achieve approximately 65% coupling efficiency with minimal effort (for example, the input beam size has not been optimized). Occasional realignment (once every few months) is necessary to maintaining sufficient coupling. The fiber has FC/APC (angled-polished) connectors at both ends. This prevents reflections back into the Ti:Sa at the input end and the formation of a cavity between fiber face and the lattice retro mirror at the output end. However, one can use a non-angle-polished fiber and make use of the resulting cavity to measure the lattice beam path stability. Note that while this will measure the stability of the entire path from fiber to retro mirror, only the stability of the path from the atoms to the retro mirror matters. Using this technique, we measured position noise of roughly $5 \times 10^{-12} \text{m}/\sqrt{\text{Hz}}$ in the 10–30 kHz band, corresponding to the on-site harmonic oscillator frequencies for lattices between $2 E_R$ and $17 E_R$. No major peaks were observed in the noise spectrum.

4.7 Intensity Stabilization

The intensity of each lattice beam is stabilized by adjusting the AOM RF power while monitoring the power exiting the optical fiber (figure 4.5). A simple PI controller (~ 10 kHz bandwidth) is used.

4.7.1 Measurement

A few percent of the power in each beam is directed onto a photodiode using a beam sampler (Thorlabs BSP05-B1), simply a glass window with anti-reflection coating on only one side. Each photodiode is shrouded using a 0.5-inch lens tube to block ambient light and is connected to a small transimpedance amplifier circuit (although monolithic photodiode/amplifier packages are available, e.g., Advanced Photonix SD100-41-21-231). For future use, we have designed a compact photodiode/amplifier board which fits conveniently within a 0.5-inch lens tube.

It is, of course, essential that the fraction of light sampled from the beam remain constant. This fraction is surprisingly sensitive to small changes in the polarization of the beam since the beams are, in general, not purely s or p polarized relative to the beam sampler. The polarization of the light exiting the optical fiber fluctuated enough to cause significant (several percent) noise in the stabilized beam intensity. The problem was corrected by placing a PBS after the fiber.

4.7.2 Variable Gain Amplifier

The stabilization circuit controls the RF power driving the lattice beam AOM by way of an analog variable gain amplifier (VGA; Analog Devices ADL5330) which follows the DDS. This part was selected instead of a variable attenuator since it can provide positive gain—the maximum DDS output power is not sufficient to drive the AOM RF power amplifier directly. A fixed attenuator was placed after the VGA to limit the final output power to a safe level for the AOM. A better solution would be to place a diode clamp on the VGA control input and use a smaller attenuator to keep the VGA away from its saturation region. The gain of the VGA in dB is proportional to the input control voltage. This makes the controller nonlinear in the sense that the lattice beam power is not proportional to the VGA control voltage. This nonlinearity is one limitation on the speed of the servo. The other limitation is the response of the AOM. By placing a log amp (e.g., Analog Devices AD8305) in front the VGA control input, a linear servo could be obtained.

4.8 Lattice Beam Path

Figure 4.6 shows the beam path for one lattice beam, which will be explained in more detail in the following sections. This figure does not show the mirrors used to steer the

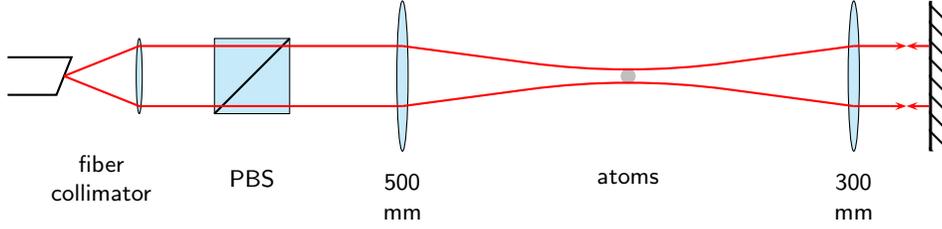


Figure 4.6 Optical system for one lattice beam.

beam. Due to the importance of maintaining the lattice beam polarization, gold mirrors are used instead of dielectric mirrors since, in general, the lattice beams are not purely s or p polarized relative to the steering mirrors.

The acute angles of the lattice beams on the science cell and the restriction on their relative polarization mean that a significant fraction of the light is lost due to reflections from the glass of the science cell. One solution to this problem would be to apply an anti-reflection coating to the glass cell during manufacture. However, finding a coating that works across a wide range of wavelengths, 532 nm, 780 nm, 812 nm, and 1064 nm in our system, and incidence angles, 30° to 90° , is difficult. Furthermore, although the loss increases the optical power required, it does not significantly affect the contrast of the optical lattice. This perhaps surprising result can be shown by a simple calculation of the lattice contrast. With $E_1 = \sqrt{(2I_0/c\varepsilon_0)}e^{-ikz}$ and $E_2 = \sqrt{(2\alpha I_0/c\varepsilon_0)}e^{-ikz}$, where $\alpha < 1$ is the ratio of intensities in the forward and retro beams, the standing wave intensity is

$$I = |E_1 + E_2|^2 = \frac{2I_0}{c\varepsilon_0} (1 + \alpha - 2\sqrt{\alpha} + 4\sqrt{\alpha} \cos^2 kz). \quad (4.50)$$

The contrast is

$$\frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}} = \frac{2\sqrt{\alpha}}{1 + \alpha}. \quad (4.51)$$

Even for $\alpha = 0.5$, the contrast is still very high, over 94%.

4.8.1 Forward Beam Optics

The forward beam optics are mounted on a 0.5-inch thick aluminum platform, approximately 12 inches above the table. Holes were cut for the beams to pass through to the science cell. The beam is coupled from the fiber using a fiber collimator (Oz Optics HPUC-2A3A series; Newport F-CS-F2-780 is now our preferred part). A PBS cleans up the polarization and a 500 mm lens focuses to a $120 \mu\text{m}$ waist at the atoms. These three components are mounted in a rigid cage system. Two mirrors guide the beam to the vacuum system. For two of the beams, the second mirror is mounted on the bottom of the platform.

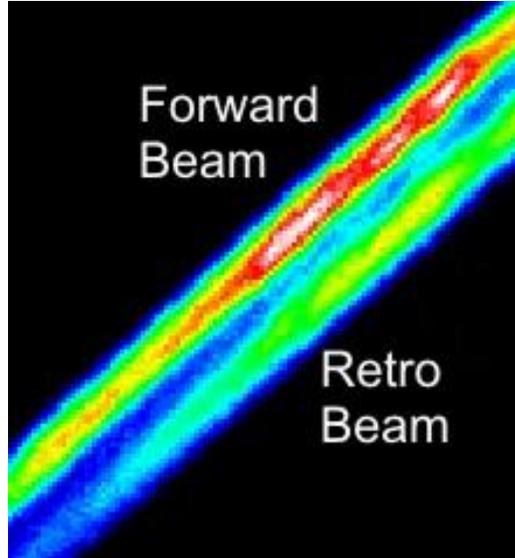


Figure 4.7 Fluorescence from one lattice beam on resonance passing through hot atom cloud. The retro beam is not yet aligned.

4.8.2 Retro Beam Optics

After the beam exits the vacuum system, it is collimated by a lens and reflected back along the same path. The position of the retro beam waist (along the beam direction) is extremely sensitive to the collimating lens position.

The retro mirror mounts (Thorlabs VM1) use 100 thread-per-inch (TPI) adjustment screws combined with 30 degree inclined planes to give an effective TPI of 200. This enhanced resolution has proven to be important. The adjustment screws are graduated (i.e., labeled with a numbered scale), which is essential for accurate alignment of the retro beam.

4.9 Alignment and Calibration

Alignment of lattice beams to the BEC position involves several steps, outlined below. The procedure requires imaging along two different directions. After alignment, the intensity of the lattice beams must be calibrated to determine the relation between optical power and lattice depth.

4.9.1 Coarse Alignment of Beams

Initial alignment of the lattice beams is performed by tuning the Ti:Sa laser to the D2 transition (a wavemeter provides sufficient precision) and looking for fluorescence along the path of the beam through a large, hot cloud. The beam is turned on for the entire 10 ms exposure of the CCD, beginning as soon as the atoms are released from the trap. First, the

forward beam alignment is adjusted to overlap the beam with the previously determined BEC position for both imaging directions, with the retro beam blocked. Next, the forward beam is scanned in all directions around the aligned position to check for clipping, since the beams normally pass very close to the IP trap structure. If clipping is found, the incident angle of the beam is adjusted slightly and the alignment procedure repeated. Finally, the retro beam is unblocked and aligned in the same manner, figure 4.7. Note this technique does not work if the lattice beam is in the plane defined by the two imaging axes. If this is unavoidable, the atom cloud can be made smaller by evaporation and the beam adjusted so that fluorescence is still visible. This iterative procedure takes much longer. Because of the long Rayleigh length of the lattice beams, no in-situ adjustment of the focusing lens position is necessary.

4.9.2 Fine Alignment of Forward Beam

If for no other reason, the fluorescence alignment step alone is insufficient because it can only be performed on resonance; the path of beam will be slightly different at the wavelength used for the lattice. The alignment of the forward beam, again with the retro beam blocked, is completed by turning the beam on slowly (in about 100 ms) to load atoms into the dipole trap created by the beam. For a red-detuned lattice, this is done first with the Ti:Sa tuned to 796 nm, just red of the D1 transition, then checked at the final lattice wavelength. The location of the BEC in this dipole trap should match the location of the BEC without the lattice beam on. An alternate method is to pulse the beam on for a short time just before release and then look at the BEC position after time-of-flight. When properly aligned, the beam will not impart any center-of-mass motion to the atoms, so the position after time-of-flight will be unchanged.

4.9.3 Fine Alignment of Retro Beam

The final step is fine alignment of the retro beam. The lattice beam is pulsed on at maximum power for roughly 10 μ s to diffract atoms from a BEC just before release. If no diffraction is observed, the pulse time can be increased, or the laser can be tuned closer to resonance to start the alignment. Once diffraction can be observed, the retro beam alignment is adjusted to maximize the percentage of atoms diffracted while decreasing the beam power as needed to ensure that the pulse length is always less than half the oscillation period (section 4.2.3). These pulses are too fast to be stabilized with the intensity servo, so enough points must be taken to suppress noise from intensity fluctuations.

4.9.4 Calibration

Calibrating a lattice beam involves measuring the lattice depth seen by the atoms (in E_R) versus the optical power in the lattice beam. A number of different methods can be used

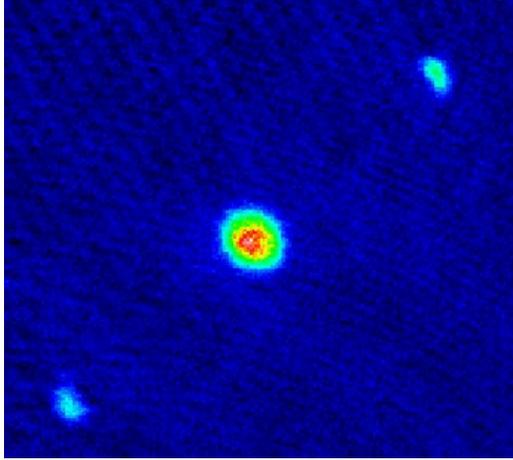


Figure 4.8 Diffraction atoms using a single lattice beam.

[122]. We have found that the diffraction resulting from pulsing the lattice on for a few tens of μs provides the most reliable method. As explained in section 4.2.3, the period of oscillations in the fraction of atoms diffracted vs. pulse time can be used to determine the lattice depth. Because the shot-to-shot intensity stability of the lattice beams is quite poor without the servo, another approach is to fix the pulse time near where a minimum in the fraction of diffracted atoms occurs and plot the diffracted fraction vs. the pulse intensity (as measured by the photodiode) for some number of shots, then fit this to find the intensity which gives the minimum.

4.10 Loading into Optical Lattice

Adiabatic loading of atoms into the optical lattice is accomplished by increasing the lattice intensity from zero to the final value in 100 ms using an exponential or sigmoidal ramp. The confinement from the IP trap or dipole trap is not changed during this time. It is essential that the loading be adiabatic with respect to both the band gap (sub-ms timescale) and the additional confinement arising from the profile of the lattice beams [123], for which the characteristic timescale is given by the harmonic oscillator frequency, on the order of 100 Hz.

For most optical lattice experiments, it is desirable first to keep the density below three atoms per lattice site to suppress loss from three-body recombination and, second, to load the lattice from a spherical trap to match the spherical confinement arising from the Gaussian profiles of the three lattice beams. Therefore, the strong and usually highly anisotropic confinement used during evaporative cooling must be relaxed before turning on the optical lattice. We perform evaporative cooling to near the BEC transition temperature in a strongly confining trap, then relax (typically in one second) to the trap from which the optical lattice will be loaded ($14.5 \text{ Hz} \times 22 \text{ Hz} \times 22 \text{ Hz}$ for the IP trap, slightly stronger

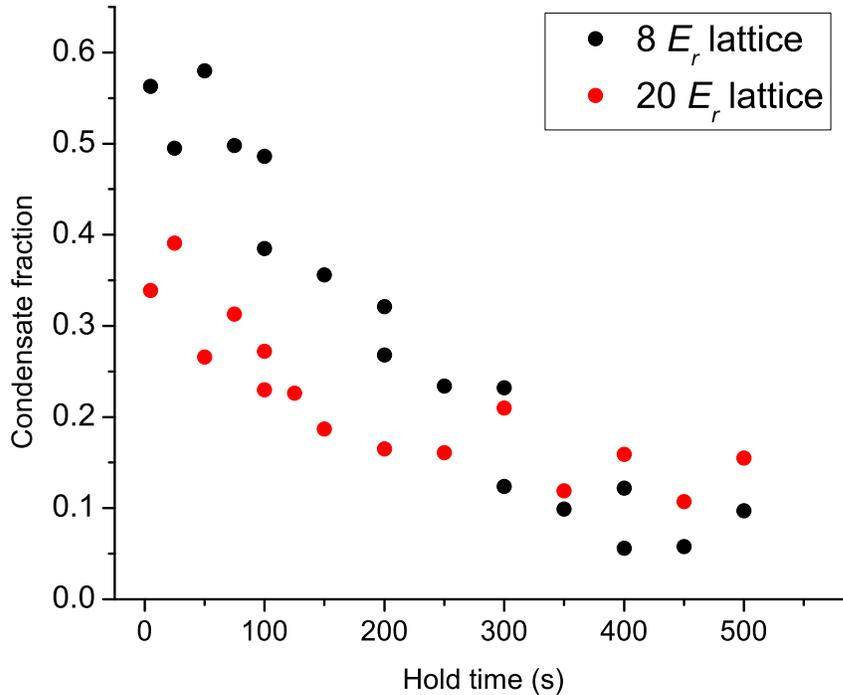


Figure 4.9 Lifetime of BEC in the optical lattice, for two different lattice depths. The hold time does not include the 100 ms turn on and 15 ms turn off of the lattice. Fits to decaying exponentials give $1/e$ times of 330 ms for 8 E_R and 450 ms for 20 E_R .

for the dipole trap) and finish evaporative cooling over several seconds to yield a BEC with condensate fraction over 90% (see section 2.12.4). This also helps remove any motion introduced during the relaxation to the weak trap. Any center-of-mass motion of the BEC remaining when the lattice is tuned on will be converted to heat, particularly when going through the Mott insulator phase transition. The random motion and center position shifts introduced by relaxation of the IP trap (see section 2.9) initially prevented us from loading atoms into the optical lattice without massive heating.

4.10.1 Heating and Loss

Heating and atom loss in the optical lattice can be attributed to scattering of lattice light, parametric heating due to fluctuations in the lattice depth and position, and enhanced three-body recombination when the density exceeds two atoms per lattice site. Figure 4.9 shows the decay of condensate fraction for a nearly pure BEC of about 10^5 atoms loaded into an 8 E_R 3D optical lattice and a 20 E_R (well into the Mott insulator regime) lattice. Because the observed heating rate was low enough for our work and comparable to that observed by other groups [124], we did not undertake any study of the individual contributions to the heating rate.

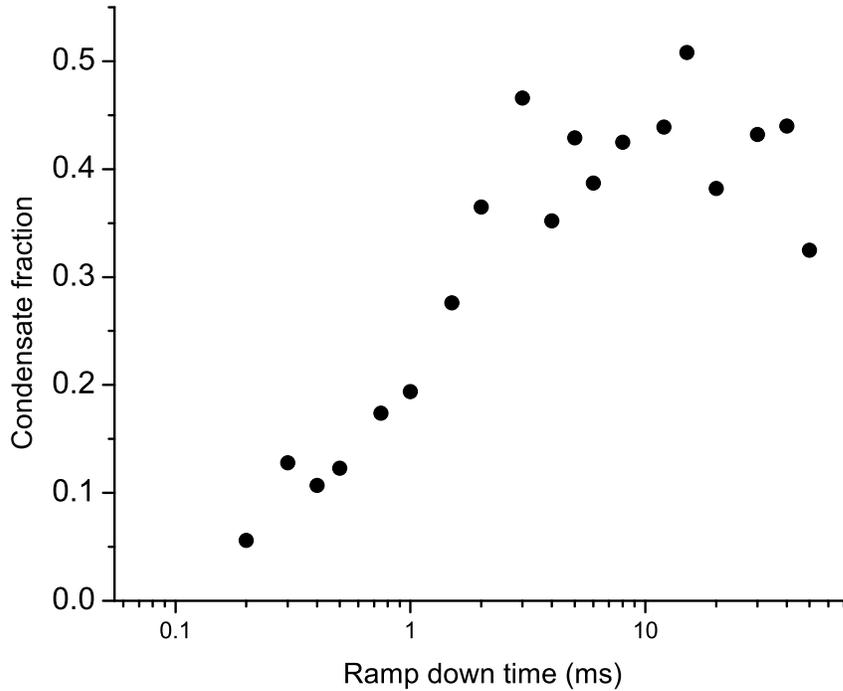


Figure 4.10 Condensate fraction recovered vs. (linear) ramp down time of $20 E_R$ lattice potential.

4.11 Release from Optical Lattice

Sudden turn off of the lattice beams, compared to the bandgap, results in a diffraction pattern for a shallow lattice, as discussed in section 4.2.3. For a deep lattice, in the Mott insulator regime, the lack of phase coherence between sites prevents the formation of such a pattern. If instead the lattice is ramped off slowly (≥ 5 ms), phase coherence is restored and the system returns to a superfluid state, although with a higher fraction of thermal atoms due to heating in the lattice. Alternately, the lattice can be ramped off in an intermediate timescale, typically a 0.2–0.8 ms line, a process called bandmapping. This time is slow compared to the bandgap, but too fast to restore significant phase coherence to a Mott insulator state. Figure 4.10 shows the extent of recovery of coherence for different lattice ramp down times for a $20 E_R$ lattice. Our group recently completed an in-depth experimental and theoretical study of the bandmapping process [125].

Because we do not image along a lattice beam direction, interpretation of images can be more complex. In the imaging system coordinates, where the imaging axis is along z and x is along gravity, the lattice beam directions (see figure 2.8) are given by the unit vectors

$$\hat{\mathbf{k}}_1 = \left(\frac{-1}{2}, \frac{1}{2}, \frac{-1}{\sqrt{2}} \right), \quad \hat{\mathbf{k}}_2 = \left(\frac{-1}{2}, \frac{1}{2}, \frac{1}{\sqrt{2}} \right), \quad \hat{\mathbf{k}}_3 = \left(\frac{-1}{\sqrt{2}}, \frac{-1}{\sqrt{2}}, 0 \right). \quad (4.52)$$

Imaging along a lattice direction can be accomplished using dichroic optics to separate the lattice and imaging light.

4.12 Superfluid/Mott insulator Phase Transition

The ability to observe the superfluid/Mott insulator phase transition is a sensitive measure of the quality of the optical lattice alignment. Figure 4.11 shows a series of absorption images of atoms released suddenly from lattices of increasing depth at 811 nm. For the last image, the lattice was ramped up to $22 E_R$ in 100 ms, then down to $4 E_R$ in 10 ms to restore phase coherence between atoms on different lattice sites.

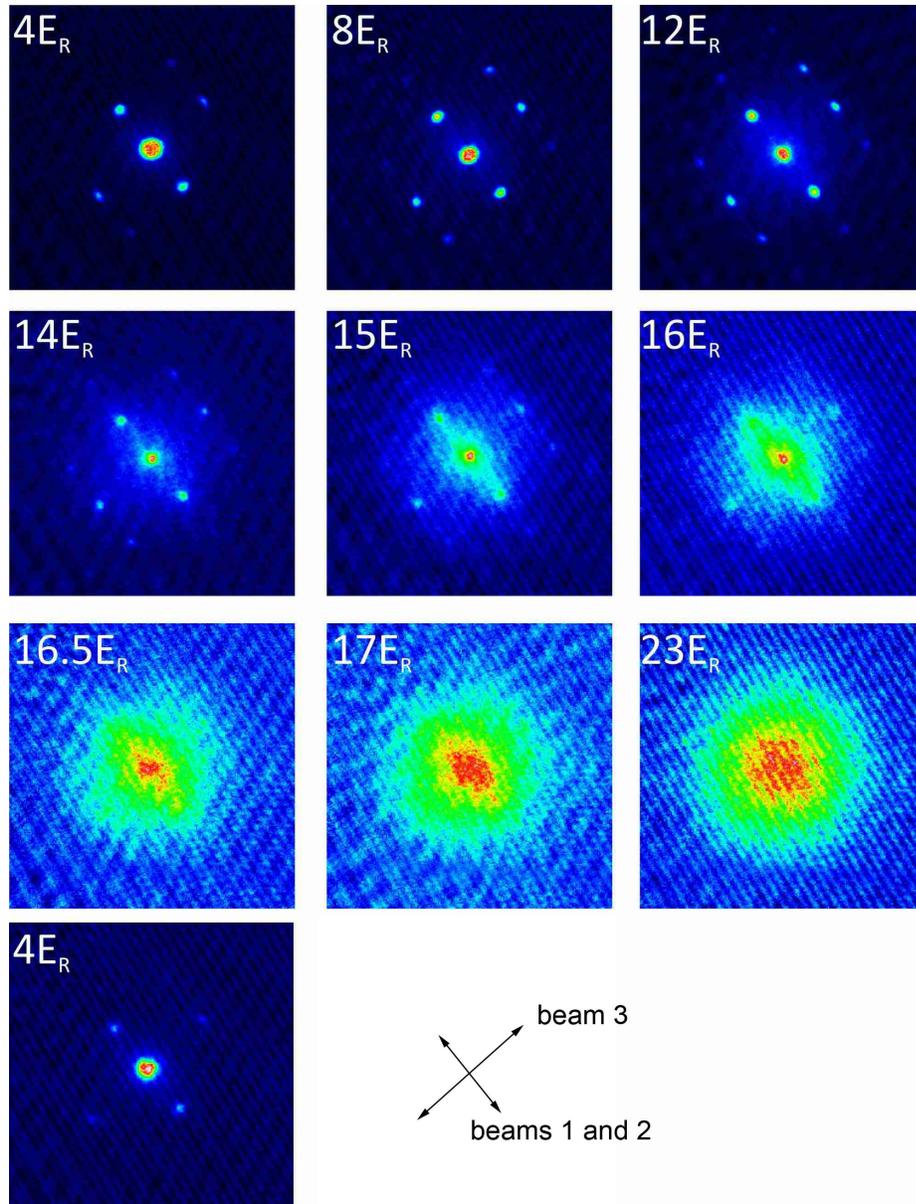


Figure 4.11 Observation of superfluid to Mott insulator phase transition. The last image shows the recovery of coherence after the lattice is ramped down, which distinguishes the phase transition from irreversible heating.

Chapter 5

Disordered Lattice

The work described in this chapter was published as: M. White, M. Pasienski, D. McKay, S. Q. Zhou, D. Ceperley and B. DeMarco. Strongly interacting bosons in a disordered optical lattice. *Phys. Rev. Lett.* **102**, 055301 (2009).

5.1 Introduction

While cold atom gases have been used previously to study disorder, the regime of strong interactions and random, fine-grain disorder, needed to explore the disordered Bose-Hubbard (DBH) model, has not been previously accessed. We have implemented the 3D DBH model in a cold atom system by combining a 3D optical lattice and an optical speckle field possessing a length scale on the order of the lattice period. Incommensurate optical lattices have been used to study quasi-disorder in 1D [34, 35], realizing the Aubry-André Hamiltonian, not the DBH Hamiltonian. Although optical speckle fields have been used to create true disorder [38–40, 46], only one study [41], in 1D, has included an optical lattice. In [41], the lattice depth was not sufficient to reach the strongly interacting regime and the characteristic speckle length scale was nearly an order of magnitude larger than lattice period, whereas most theories for the DBH model assume weak correlations in the single-particle properties at neighboring lattice sites. In [46], fine-grain speckle, comparable to that utilized in this work, was used without a lattice to study Anderson localization in the weakly-interacting regime. The present work represents the first observation of disorder-induced boson localization in the strongly-interacting regime.

5.2 Overview of Experiment

We load BECs containing $(3.2 \pm 0.6) \times 10^5$ atoms into a 3D optical lattice with 406 nm between sites as described in section 4.10. In this chapter, we will quantify the depth of the lattice potential with a dimensionless parameter s , so that the lattice depth is sE_R (E_R is the recoil energy at 812 nm). The experiment discussed here was performed using the IP trap and optical lattice, but work with disorder has continued using the dipole trap (after the failure of the IP trap). Nearly pure BECs are created by performing a long final evaporation stage in a weakened trap. At high condensate fractions, absorption imaging

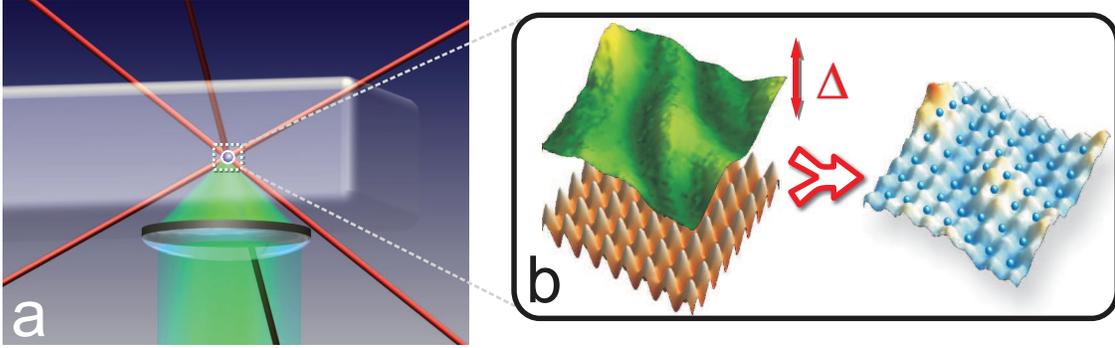


Figure 5.1 Schematic view of (a) science cell showing speckle beam and optics and optical lattice beams (red) and (b) creation of disordered lattice by combining speckle field and clean lattice.

lacks sufficient dynamic range to measure both the condensed and thermal component with a single image. The average condensate fraction can be determined by alternately taking partially repumped images (section 2.11) to measure the number of condensed atoms and fully repumped images to measure the number of thermal atoms. We choose evaporation parameters so that no thermal atoms were visible, even in the fully repumped images. This corresponds to condensate fractions of 95% or more; this lower bound is set by the signal-to-noise ratio of the images. Due to drift, the condensate fraction may occasionally be slightly lower; a lower bound of 90% is conservative.

Fine-grained disorder is superimposed on the optical lattice by passing $\lambda = 532$ nm light (from the same laser used to pump the optical lattice Ti:Sa) through a holographic diffuser (Luminit LLC) to generate an optical speckle field. The geometry is shown in figure 5.1. The diffuser, a thin plastic film, randomly scatters the light through a 0.5° range of angles. Just before the diffuser the light passes through a 0.9 numerical aperture lens (Lightpath Industries GPX-15-15), leading to a random distribution of intensities around the focal plane covering a Gaussian envelope with a $160 \mu\text{m } 1/e^2$ radius. The lens is a Gradium brand singlet, which provides performance comparable to a multi-element lens. The lens and diffuser assembly are approximately 13 mm from the atoms. Atoms experience a potential energy shift proportional to the intensity of this light, which, in combination with the optical lattice, produces a disordered lattice potential. The disorder strength is quantified by the average potential energy shift due to the speckle field, Δ , measured in E_R (at the lattice wavelength, not 532 nm).

5.3 Optics

Figure 5.2 shows the optics making up the speckle beamline. With 2–3 W incident, the optical intensity in the AOM used initially (Neos 23080-1) was well within the manufacturer’s specifications. However, we found that the AOM was rapidly damaged (indicated by a marked drop in diffraction efficiency) by the light. After contacting the manufacturer, we

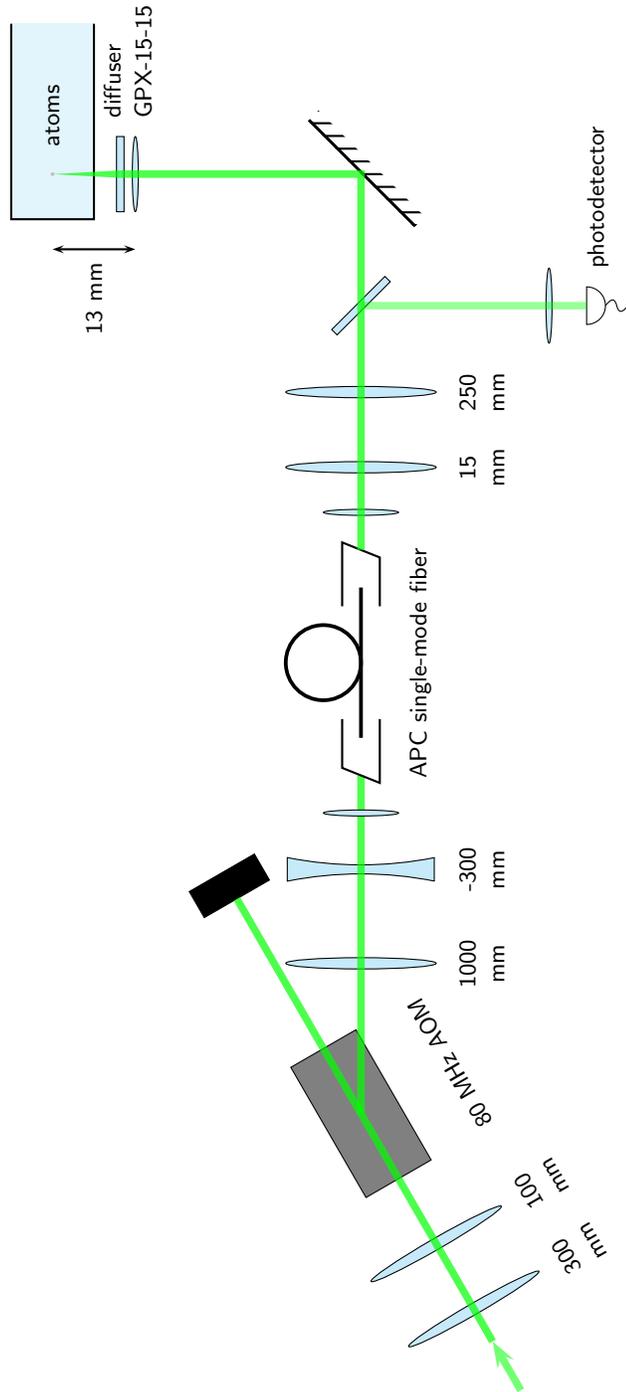


Figure 5.2 Speckle beam optics. The photodetector is used to monitor and stabilize the beam power (via feedback to the AOM RF power). Distances are not to scale.

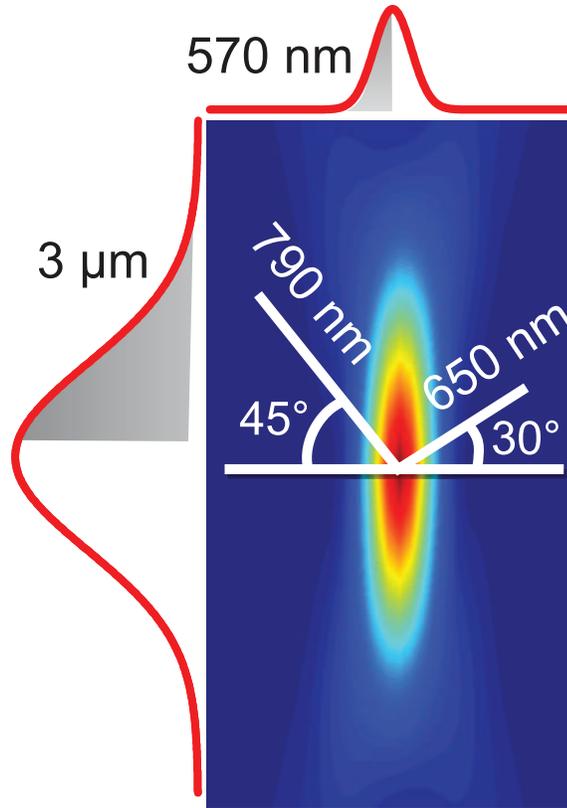


Figure 5.3 Autocorrelation function (false color) of the optical speckle field determined from ex situ measurement. The directions of the lattice beams and corresponding autocorrelation widths are overlaid.

learned that the specifications were incorrect—the TeO_2 crystal used in the AOM is very susceptible to damage at 532 nm. The AOM was replaced by another AOM (Neos 35086-3) using a silica crystal and requiring 6 W of RF power (vs. 1 W).

The same method is used to stabilize and control the speckle beam power as for the lattice and dipole beams, section 4.7. Because the polarization of the speckle beam does not need to be extremely pure, no PBS is used after the optical fiber. This means that the beam cannot be sampled in the same way as the optical lattice beams to monitor power. Instead, the beam is sampled spatially—a slice of light from the edge of beam is directed through a lens onto the photodetector. Since the beam is sampled a short distance from the output of a single-mode fiber, minimizing fluctuations in beam shape and position, the fraction of the beam power sampled is constant.

5.4 Speckle Characterization

Let us consider the problem of a diffuser placed immediately after a lens of focal length f . Assuming the spread in angles from the diffuser is sufficiently large, ray optics can be used

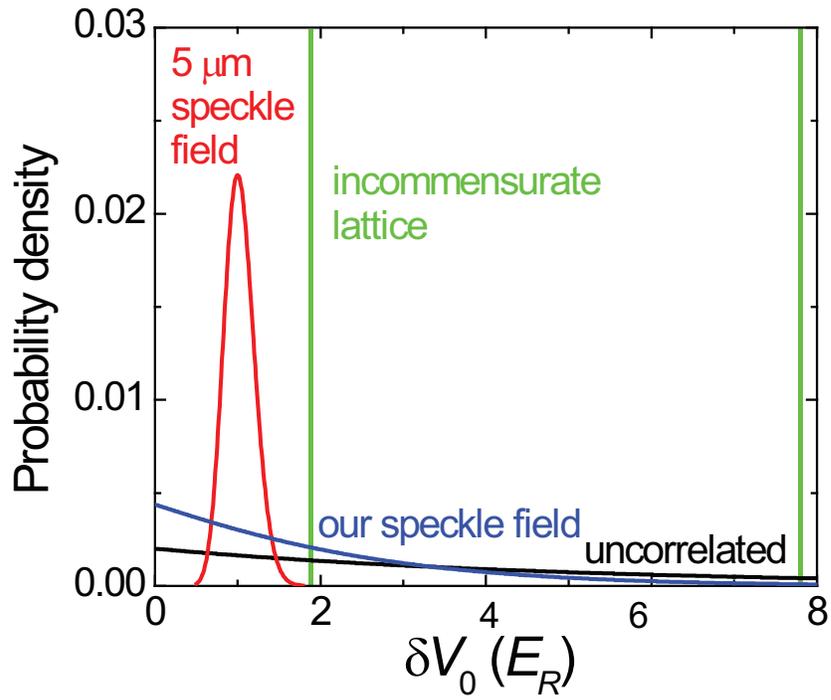


Figure 5.4 Joint probability of obtaining site energy offset δV_0 for a site adjacent to a site with offset $\delta V_0 = 1 E_R$, shown for several different disorder potentials which have been realized experimentally. The mean site energy shift over the entire lattice is $\Delta = 5E_R$ for each disorder potential.

to calculate the waist size at the focal point (a distance f from the lens): $w_0 = \theta f$, where $\pm\theta$ is the full angular range of the diffuser. This expression is valid provided w_0 is much larger than the waist that would be obtained without the diffuser. For our system, $\theta = 0.5^\circ$ and $f = 15$ mm, so $w_0 = 130$ μm . The length scale of the speckle field can be characterized by the width of the intensity autocorrelation function,

$$AC(\delta\mathbf{r}) = \langle I(\mathbf{r})I(\mathbf{r} + \delta\mathbf{r}) \rangle, \quad (5.1)$$

where the average is over \mathbf{r} . The autocorrelation function of the speckle field is well-approximated by a Gaussian. The $1/e^2$ radius of the autocorrelation function in the transverse direction a distance z from a circular aperture with diameter D is $1.2\lambda(z/D)$ [126]. For the longitudinal direction, the $1/e^2$ radius is $5.7\lambda(z/D)^2$. In our system, $z/D = 0.84$, which gives 530 nm and 2.1 μm for the transverse and longitudinal autocorrelation radii, respectively.

We determined the intensity autocorrelation function, figure 5.3, from a measurement of the speckle intensity in a $10 \times 100 \times 80$ μm volume using a scanning, high-resolution optical microscope. The characterization was done ex situ using the same optical setup as in the experiment. The microscope consisted of a microscope objective and CCD camera mounted to a precision 1D motorized translation stage. A fit of the measured autocorrelation distribution to a Gaussian gives $1/e^2$ radii of 570 nm and 3 μm along the transverse and speckle propagation directions, in good agreement with the theoretical values. Along the lattice directions, the autocorrelation radii are 650 nm (beams 1 and 2) and 790 nm (beam 3), less than two lattice periods (figure 5.3).

Figure 5.4 shows the probability of obtaining a given difference δV_0 in site energies between adjacent sites for several disorder potentials which have implemented experimentally [34, 127]. In contrast to previous experiments, the fine-grain disorder used in this work allows access to a qualitatively different regime, approaching the case of uncorrelated site-to-site disorder.

5.5 Alignment

The lens and diffuser were mounted in a tube threaded through a custom 2D flexure type translation stage (with 80 TPI adjusters) providing positioning in the horizontal plane, figure 5.5. The threaded tube (32 TPI) provided vertical position adjustment.

As with other beams, alignment of the speckle beam is accomplished in several steps. Coarse alignment uses fluorescence from a copropagating resonant beam, as with the dipole beam (section 2.12.3). Next, the beam is pulsed on for 1 to 3 ms immediately after release of a BEC. From an image taken after 20 ms time-of-flight, the width of the condensate in the direction normal to the speckle beam axis is measured. As the beam center is moved

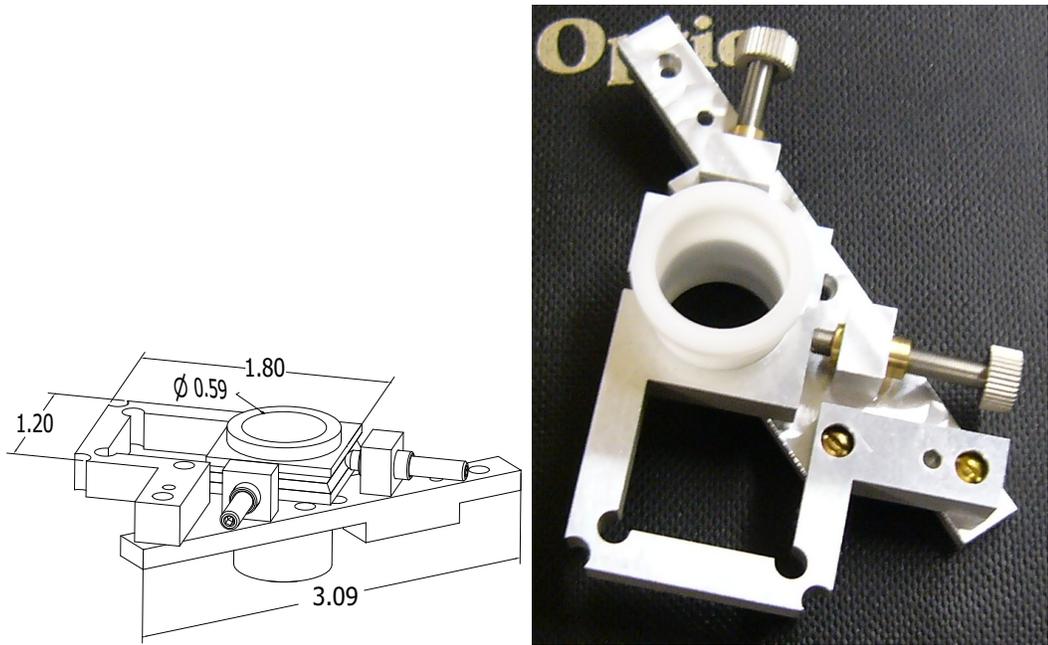


Figure 5.5 CAD design of flexure mount for speckle lens and diffuser and image of finished mount. Representative dimensions (in inches) given to suggest scale.

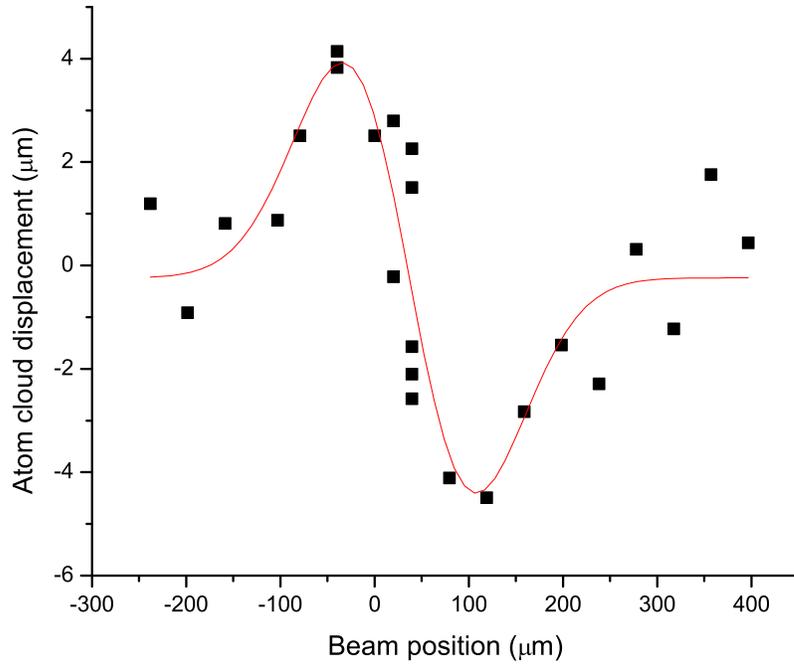


Figure 5.6 Displacement of atom cloud due to speckle beam. The red trace is a fit to the derivative of a Gaussian and gives a waist of $142 \pm 17.5 \mu\text{m}$.

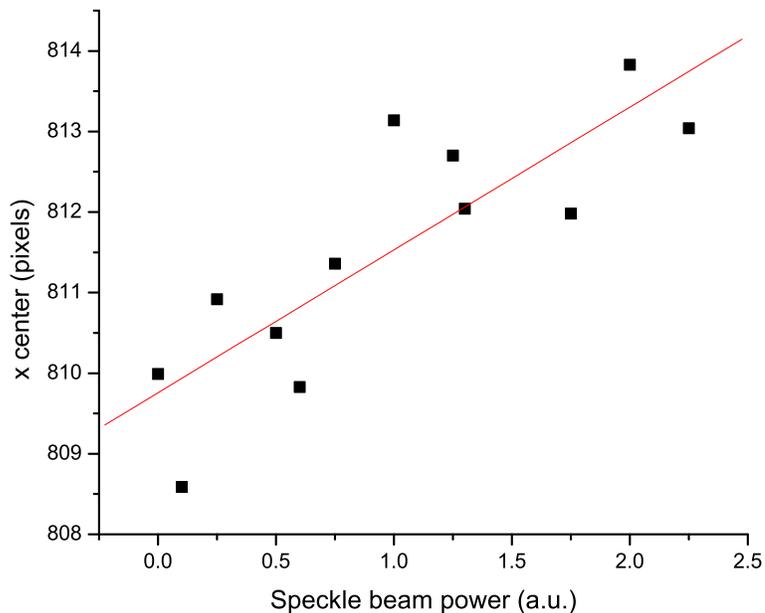


Figure 5.7 Displacement of a 500 nK thermal cloud vs. speckle intensity.

closer to the BEC, this width increases dramatically. Although displacement of the cloud is also observed during this step, it is always in the same direction and cannot be used as a signal for alignment. The final stage of alignment uses a thermal cloud at 500 nK; the speckle beam is turned on slowly (in 200 ms) and the position of the cloud is measured 1 ms after release with both the main and auxiliary imaging systems; see figure 5.6. This method allows for measurement of the beam waist and is repeated at several vertical positions of the lens to overlap the beam’s focus with the atoms. A thermal cloud is used instead of a BEC because high intensity speckle was observed to pin the BEC, preventing it from moving. We confirmed that 500 nK is sufficient to overcome the pinning by measuring a linear increase in the displacement with increasing speckle beam intensity, figure 5.7. The amount of deflection also provides an in situ measurement of the disorder strength Δ (estimated to be good within 15%).

An alternative alignment method that can be used is to set up an imaging system with the BEC at its focus. Then the beam to be aligned is viewed on the imaging system and adjusted to be concentric with the BEC position. The focus is moved along the beam axis until the beam appears smallest on the camera. Because this method does not require cold atoms, it can be done very quickly once the imaging system is in place.

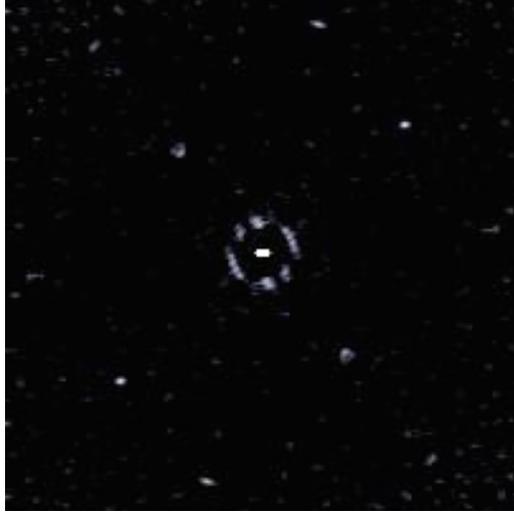


Figure 5.8 Density-density correlation image for the disordered lattice with $s = 25$ and $\Delta = 4E_R$. Although the image is noisy, the correlation peaks are clearly visible; compare their locations to figure 4.11. The correlation image for $\Delta = 0$ is essential identical.

5.6 Density-Density Correlations with Disorder

The density-density correlation measurement discussed in section 4.4.2 can potentially provide a method for in-situ characterization of the disordered lattice. We consider the case of a disordered 1D lattice with $\mathbf{r}_i = (i + \delta_i) \times a$, where $\delta_i \times a$ is the site position offset due to the disordering potential (recall that a is the lattice period). Equation (4.48) becomes

$$\langle C(d) \rangle = \frac{1}{N^2} \sum_{i,k} e^{2\pi i(i-k)d/l} \langle e^{2\pi i(\delta_i - \delta_k)d/l} \rangle n_i n_k, \quad (5.2)$$

where the average is over many realizations of the disorder. Here, we are only considering the effect of the disorder on the site positions, not, e.g., on the site wavefunctions.

A numerical calculation with our disordered potential yields an approximately Gaussian distribution for δ with $\sigma = 0.01$. Using $p(\delta) = (1/\sqrt{2\pi}\sigma) \exp(-\delta^2/2\sigma^2)$, the expectation value in the above expression becomes

$$\langle e^{2\pi i(\delta_i - \delta_k)d/l} \rangle = e^{-(2\pi\sigma d/l)^2}, \quad (5.3)$$

which can be brought outside the sum. We can see that the disordering potential simply decreases the height of the peaks in $C(d)$, without, e.g., broadening them.

Figure 5.8 shows a density-density correlation image generated from about 100 images for a 25 E_R lattice with $\Delta = 4E_R$. The kinetics mode of the camera (section 2.11.4) was used to minimize fringes in the images, but it was still necessary to notch out several sets of fringes in post-processing to reveal the correlation peaks. Given the low signal-to-noise ratio in our data, the small predicted decrease in peak height would not be detectable. It

is of note that no additional peaks appear in the density-density correlation image, unlike in the case of quasi-disorder produced with an incommensurate lattice.

By working at a smaller s and larger Δ and improving image signal-to-noise, it may be possible to detect the effect of the disordered potential on density-density correlation measurements. The failure of the IP trap (section 2.9) prevented us from pursuing further work at the time. Nevertheless, these measurements are an important proof-of-principle, since the technique is expected to be useful for readout of stroboscopic quantum simulation, chapter 6.

5.7 Disordered Bose-Hubbard Hamiltonian

The presence of disorder modifies the Bose-Hubbard Hamiltonian (section 4.3) by making the parameters ε , J , and U site dependent:

$$\hat{H} = - \sum_{\langle ij \rangle} J_{ij} \hat{b}_i^\dagger \hat{b}_j + \sum_i (\varepsilon_i - \tilde{\mu}_i) \hat{n}_i + \sum_i U_i \hat{n}_i (\hat{n}_i - 1)/2. \quad (5.4)$$

The site energies are disordered by potential energy shifts, the tunneling energy by changes in the potential barrier between sites, and the interaction energy by modification of the curvature of the potential near the center of each lattice site. We have used the effective chemical potential μ_i discussed in section 4.3.5 so that ε_i varies only due to disorder.

In typical theoretical treatments of the DBH model, the disorder is limited to either the site or tunneling energies, which are generally assumed to be uniformly or Gaussian distributed. Because the statistical properties of the disordering potential are known in this case (figure 5.3), we were able to enlist the help of computational physics colleagues in David Ceperley’s group to calculate the distributions for ε , J , and U [128]. Local basis functions (analogous to the Wannier functions for the clean lattice) were computed for a 3D disordered lattice potential with $7 \times 7 \times 7$ lattice sites that reproduces the geometry shown in figure 5.1; the parabolic confining potential was not included in the calculation. From the basis functions, ε , J , and U are calculated using their definitions, equation (4.22).

Computed probability distribution functions for the DBH parameters are shown in figure 5.9. While the disorder results in distributions for J and U that are broadened to higher and lower values, the distribution for ε is one-sided, since the blue-detuned speckle potential can only increase the potential energy of a lattice site. These distributions show that the speckle induces the strongest disorder in J and ε compared with U —the relative scales of the disorder in each parameter are determined by the width of the distributions compared with their mean. The disorder does not shift the most probable values of J and U , even though the mean values of these parameters change with increasing Δ . We have verified that the width σ_ε of the site energy distribution is proportional to Δ , as expected from the distribution of speckle intensities $P(I) \propto e^{-I/\bar{I}}$ (\bar{I} is the mean intensity and $\Delta \propto \bar{I}$).

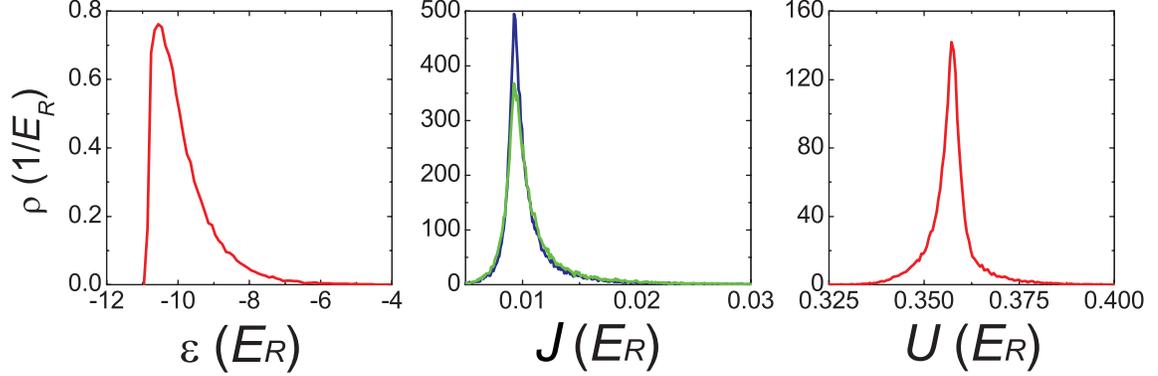


Figure 5.9 Probability density ρ of DBH parameters ε , J , and U for $s = 14$ and $\Delta = 1E_R$. In (b), the tunneling energy distributions are computed separately along the x (blue) and z (green) directions. The $\Delta = 0$ values are $\varepsilon = -10.85E_R$, $J = 0.0095E_R$, and $U = 0.360E_R$.

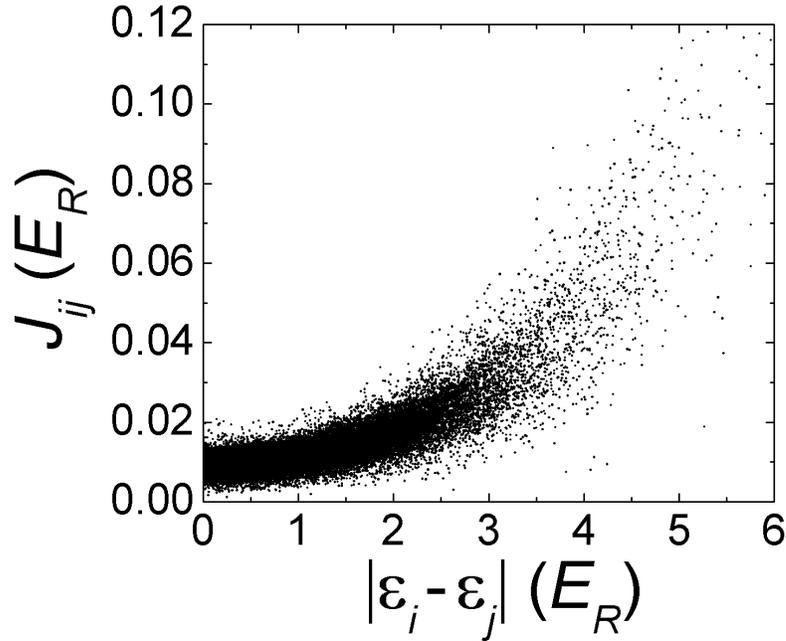


Figure 5.10 Correlation between nearest-neighbor energy offset $|\varepsilon_i - \varepsilon_j|$ and tunneling matrix element J_{ij} for the realization of the disordered Bose-Hubbard model in our system.

The similarity of the J distributions shown separately along the x and z directions in figure 5.9 confirms that the asymmetry of the speckle field with respect to the lattice directions does not play an important role in the distribution of tunneling energies. While the next-neighbor (i.e., diagonal) tunneling energies are strictly zero in the pure lattice, we find that disorder leads to finite next-neighbor tunneling energies, but smaller by at least a factor of 10^{-3} compared with the nearest-neighbor tunneling energies.

Another aspect of the statistical properties of the DBH model is the correlations between

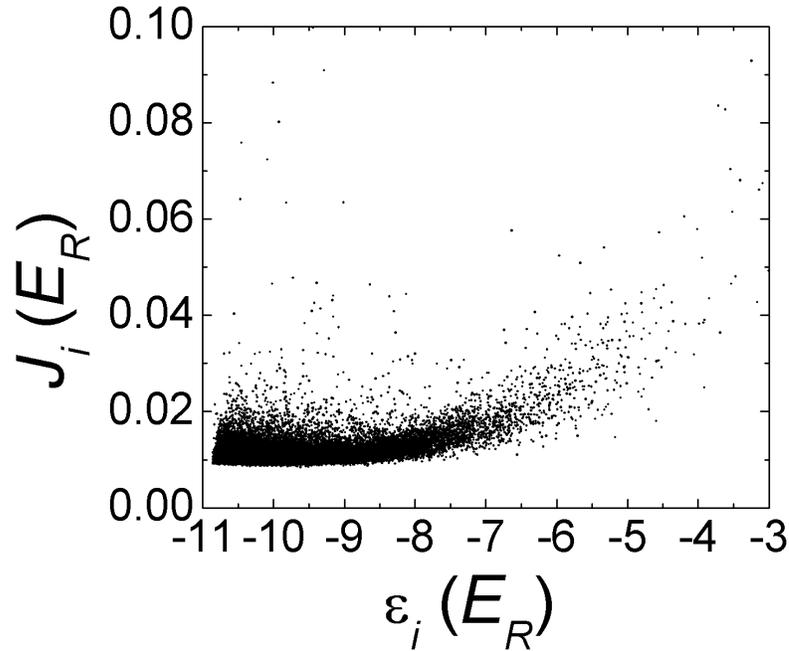


Figure 5.11 Correlation between site energy ϵ_i and tunneling matrix elements averaged over all nearest neighbors, J_i .

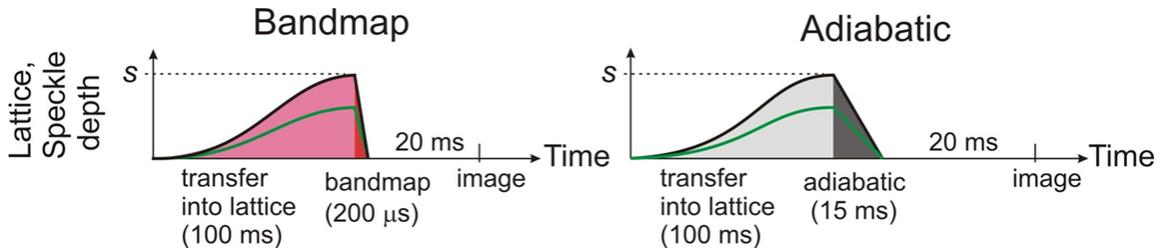


Figure 5.12 Lattice and speckle intensity vs. time for disorder experiments.

the DBH model parameters. The correlation between J and ϵ is shown in scatter plots in figures 5.10 and 5.11. While J for a site averaged over all nearest neighbors (J_i in figure 5.11) is not strongly correlated to ϵ , J_{ij} and the difference between site energies $|\epsilon_i - \epsilon_j|$ for nearest neighbors are directly correlated. This behavior is anticipated for fine-grained disorder, which induces local changes to the site energies coinciding with modifications to the potential barrier between neighboring sites.

5.8 Results

As shown in figure 5.12, the optical lattice and speckle are turned on together in 100 ms using a sigmoidal ramp, held constant for 5 ms, then turned down in 200 μ s (bandmap) or 15 ms (adiabatic), see section 4.11; the ratio s/Δ is kept fixed at all times. The atoms

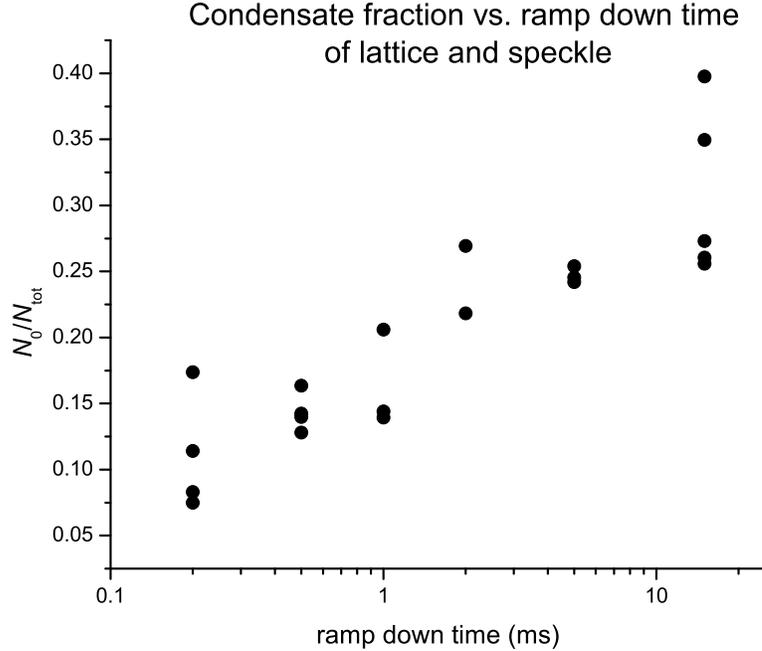


Figure 5.13 Recovery of condensate fraction vs. ramp down time of disordered lattice with $s = 12$ and $\Delta = 6E_R$.

are then immediately released and an absorption image is taken 20 ms later to measure the fraction of condensed atoms. Two-component clouds are observed for the entire range of parameters explored and the images are fit to the sum of a Thomas-Fermi (TF) profile and a broader Gaussian, as described in section 2.11.1. The number of atoms in the TF component is denoted by N_0 ; the number of atoms in the Gaussian component by N_{nc} . The total number of atoms is $N = N_0 + N_{\text{nc}}$. The Gaussian component contains thermal atoms and, in the case of bandmap release, atoms in localized states (corresponding to large momentum spreads). The possible nature of these states and states without coherence across the entire cloud that may contribute to the TF component are discussed below.

Finite signal-to-noise ratio in the imaging system causes a systematic error for measurements of low N_0/N at $s = 12$ and 14. This effect was modeled by generating clean images of two-component clouds, adding noise matching that present in our system, and then using the same fitting procedure as for real data; we found that the systematic error in N_0/N was less than 0.03.

Figure 5.13 shows the dependence of recovered condensate fraction on the ramp down time for the case of a deep lattice and strong disorder. Release times longer than 15 ms do not increase the recovered condensate fraction, in agreement with the clean lattice case (section 4.11). Hold times in the disordered lattice greater than 5 ms have no effect besides

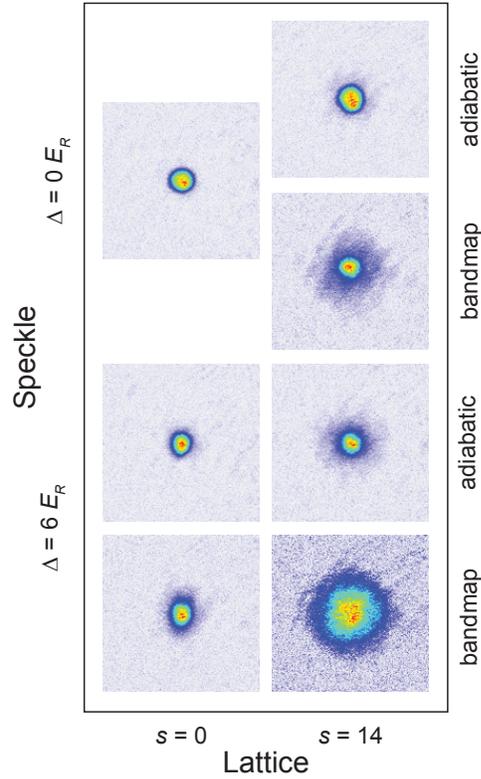
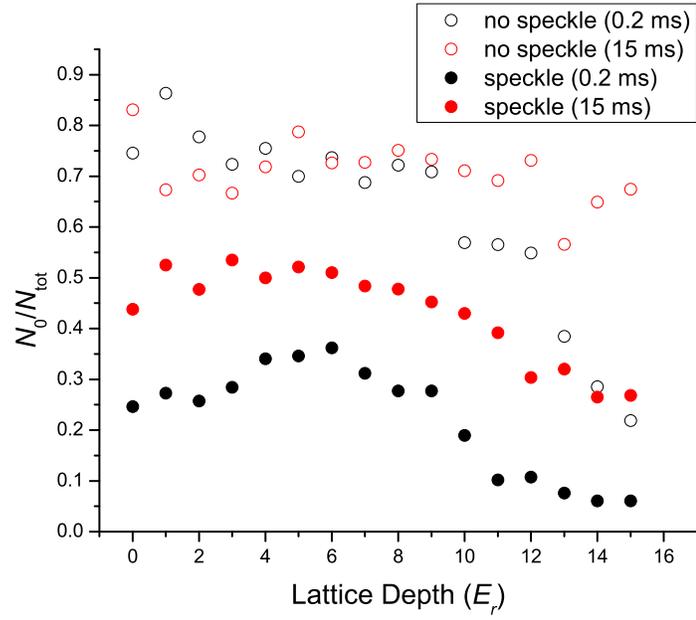


Figure 5.14 Effect of strong, fine-grain disorder, $\Delta = 6E_R$, on condensate fraction across a range of lattice depths, comparing adiabatic and bandmap turn off of disordered and clean lattices. For condensate fractions above 0.75, the point should be considered a lower limit. The bottom figure shows representative images.

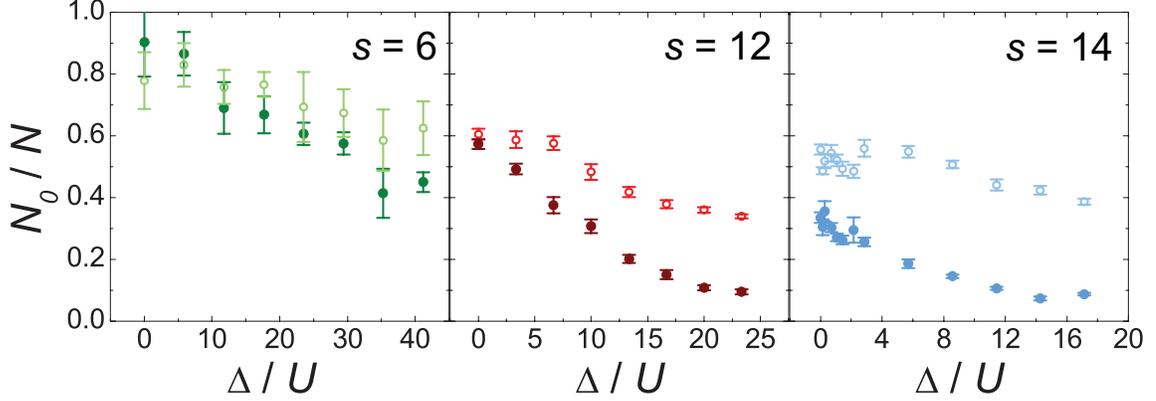


Figure 5.15 Change in condensate fraction for lattice depths of 6, 12, and 14 E_R vs. speckle intensity. The disorder strength normalized to U corresponds to $\Delta = 0-7 E_R$ for $s = 6$ and 12 and $\Delta = 0-6 E_R$ for $s = 14$. The hollow (solid) points show the condensate fraction measured after adiabatic (bandmap) release from the lattice. Each data point is the average of seven or more individual measurements, the error bars showing the statistical uncertainty.

increasing the amount of heating.

Figure 5.14 shows the effect of strong disorder ($\Delta = 6E_R$) on condensate fraction for bandmap and adiabatic release across a range of lattice depths. Two features stand out. First, disorder causes a partially reversible change in condensate fraction, the fractional change becoming larger at larger lattice depths but present even in the absence of a lattice. Second, disorder, or rather the speckle beam, causes irreversible heating. The cause of the heating was not determined. It may result from vibration of the speckle field relative to the lattice potential or long adiabatic timescales related to the disappearance of the band gap for strong disorder. A calculation of the density of states using the basis set constructed on the disordered lattice indicates that the gap between the ground and first excited band disappears for approximately $\Delta = 0.8, 1,$ and $1.75 E_R$ (corresponding to $\Delta/U = 4.7, 3.3,$ and 5) for $s = 6, 12,$ and 14 .

The dependence of condensate fraction on disorder strength at $s = 6, 12,$ and 14 is shown in figure 5.15, again for both adiabatic and bandmap release. The larger error bars for $s = 6$ arise from shot-to-shot number fluctuations because two images are necessary to measure condensate fractions above approximately 0.75. The minimal heating at low disorder strengths is consistent with heating related to the disappearance of the bandgap for sufficiently high disorder. Figure 5.16 shows the regions of the (clean) Bose-Hubbard phase diagram sampled in this data, where J is calculated in the tight-binding approximation; The border of the $n = 0$ region marks the edge of the cloud. Note that for $s = 14$, some atoms are present in the Mott insulator phase in the clean lattice.

To isolate the effect of the disordered lattice on condensate fraction from changes caused

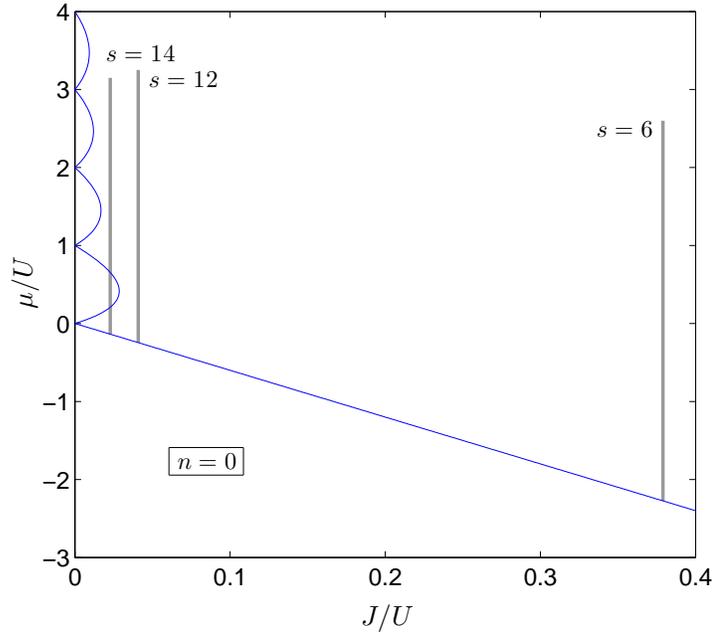


Figure 5.16 Regions of the Bose-Hubbard phase diagram explored by the data in figure 5.15.

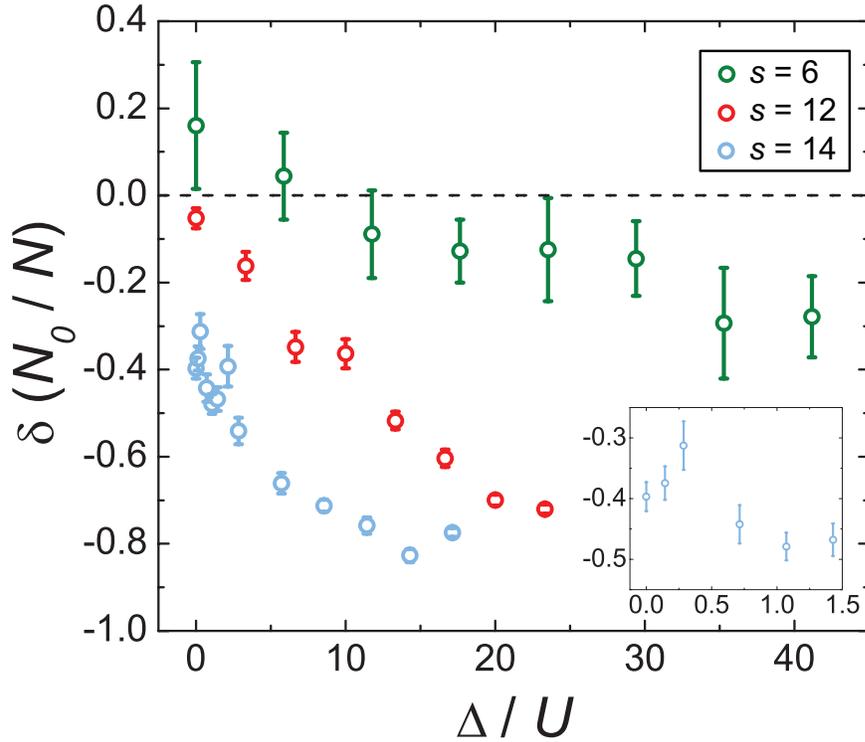


Figure 5.17 Reversible change in condensate fraction induced by disorder. Each data point is derived from a pair of adiabatic and bandmap release points from figure 5.15; the error bars are computed from the statistical uncertainties in those points. The inset is an expanded view of the low Δ/U data for $s = 14$.

by heating, we define the reversible fractional change in N_0/N as

$$\delta(N_0/N) = \frac{(N_0/N)_{\text{bandmap}}}{(N_0/N)_{\text{adiabatic}}} - 1, \quad (5.5)$$

which is shown in figure 5.17 for the data points in figure 5.15. If disorder results in no change in N_0/N or in a change entirely due to heating, then $\delta(N_0/N) = 0$, which is indicated by the dashed line. The reversible reduction in condensate fraction evident in figure 5.17 at all lattice depths is more pronounced at higher s and saturates above $\Delta/U = 15$ for $s = 14$.

5.9 Discussion

Theoretical work using mean-field theory [19–21, 23–25], renormalization group [26], replica theory [22], and quantum Monte Carlo algorithms [27–32] disagree on the nature of the ground state phase diagram.

The sampling of proposed phase diagrams for the DBH in figure 5.18 illustrates two of the primary points of contention. First, whether or not a direct transition between superfluid (SF) and Mott insulator (MI) is possible in the presence of disorder and second, if, as suggested in [129], disorder can in fact cause a transition from MI to SF. One physical explanation for this perhaps counterintuitive result is that the deep wells in the disordered lattice are screened by the interparticle repulsion, allowing a superfluid to exist on the remaining, shallower wells. Fisher *et al.* [19], argue that a Bose-glass (BG) phase always intervenes between the MI and SF phases, implying that disorder transforms the MI phase into a BG phase. However, recent results using, for example, stochastic mean-field [21] and replica theory [22] indicate that disorder can change the MI phase directly into a SF.

Quantitative predictions from theoretical work to date on the DBH model have mostly addressed superfluid fraction, not condensate fraction as we measure here, the difference being due to quantum depletion arising from the strong interactions in the DBH model. However, condensate fraction can be calculated easily in many theoretical approaches and, motivated by work with cold atoms, very recent studies have presented predictions for condensate fraction [21, 48].

The non-condensate component observed may include atoms in MI or BG (with short coherence length) phases in addition to non-condensed atoms resulting from quantum depletion (reversible) and thermal excitation (irreversible). The condensate component may include a contribution from a BG phase with a long coherence length, i.e., nearly equal to the diameter of the gas.

At zero temperature, the BH model with a parabolic confining potential predicts a SF phase with negligible quantum depletion at $s = 6$, a strongly depleted SF phase at $s = 12$, and coexisting MI and depleted SF phases (as nested spherical SF and unit filling MI shells surrounding a SF core) at $s = 14$. The mean-field theory calculation described in section

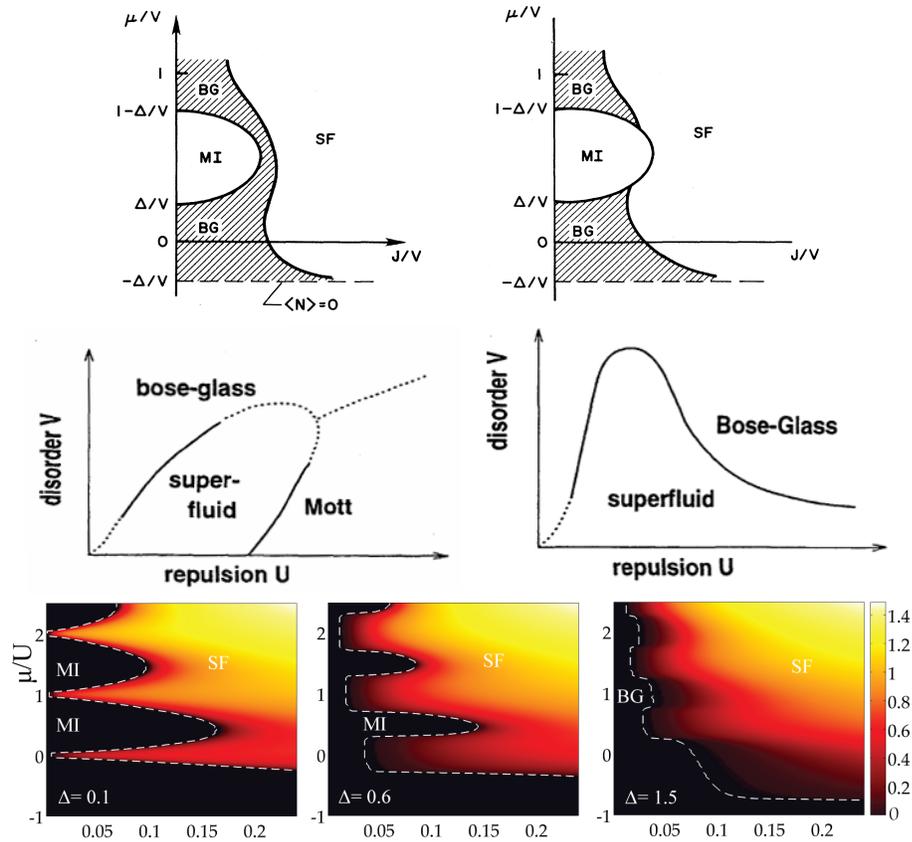


Figure 5.18 Sampling of proposed phase diagrams for the disordered Bose-Hubbard model from [19] (top), [129](center), and [21] (bottom). The authors of [19] argue that the top left diagram is the correct one (for $\Delta < U$). The central pair of figures shows phase diagrams determined by quantum Monte Carlo for commensurate (one atom per site) filling, left, and incommensurate filling, right. The lower set of phase diagrams for several different values of disorder were calculated using a stochastic mean field approach.

4.3.3 predicts $N_0/N = 0.97, 0.64,$ and 0.36 at $s = 6, 12,$ and 14 for our system; the ratio of atoms in the SF to MI phase is 0.83 at $s = 14$.

The inset to figure 5.17 shows an expanded view at low Δ for $s = 14$. We observe that for relatively weak disorder at $s = 14$, approximately a 10% reversible decrease in N_0/N is measured up to $\Delta/U = 1$ (with no change in N_0/N for adiabatic release). This result may seem to be at odds with predictions that disorder will transform the MI into SF in this regime [22, 29] and, in particular, [21] which predicts an increase in condensate fraction. Before a detailed comparison with data will be able to constrain theory, finite temperature and the distribution of DBH model parameters shown in figure 5.9 must be included in calculations of condensate fraction. Furthermore, conversion into a BG phase with a long enough correlation length that the momentum distribution is not strongly affected could mask a disorder-induced SF to insulator transition; the possibility of distinguishing the BG phase is discussed below. Finally, we note that the DBH Hamiltonian (5.4) requires corrections for multiple bands at the highest values of Δ .

5.10 Future Directions

Studies of transport in the 3D DBH model were recently performed [130] using the setup above, with the exception that the IP was replaced with the dipole trap (section 2.12). An impulse was applied to the atom cloud either in a horizontal direction, by shifting the center of the magnetic potential momentarily, or in the vertical direction by shifting the position of the dipole beam momentarily.

The DBH can be studied in 2D or 1D by making one or two of the lattice beams strong enough to completely suppress tunneling between the resulting 2D planes or 1D tubes. The phase diagram can then be explored by varying the strength of the remaining beam(s).

All the work done to date in our group, and other groups studying cold atoms in disordered potentials, has suffered from the complication of an inhomogeneous density distribution arising from the harmonic confining potential needed to hold the atoms. Red-detuned optical lattice beams necessarily add an additional confining potential. We are considering two approaches to creating a more homogenous distribution. The first is to selectively image only atoms from a small region around the center of the atom cloud, thereby restricting the measurement to smaller range of densities. This can be done by using tightly focused beams to transfer only atoms in this region to a different internal state just before release from confinement. The second is the use of a blue-detuned optical lattice, which itself produces a deconfining potential, so that the overall potential, while still confining, can be weaker.

Switching to a blue-detuned lattice enables another measurement technique which can provide insight into the phase diagram. By changing the strength of the confining potential, the compressibility of the system can be determined from the resulting change in size [51]. The superfluid and Bose-glass phases have non-zero compressibility, while the Mott insulator

phase is incompressible, so this measurement may be able to distinguish BG from MI.

Chapter 6

Simulating Quantum Magnetism

6.1 Overview

Using one quantum system to simulate the dynamics of another was first proposed by Feynman [131] as a way around the insurmountable complexity of exact calculation on classical computers: a complete quantum description of N interacting spins requires 2^N numbers. For just 300 spins, this is more numbers than there are atoms in the universe. Otherwise successful approximation methods for many-body systems encounter serious difficulties when applied to strongly interacting systems. Direct simulation of strongly-interacting models using cold atoms is a promising alternative. The understanding of weakly-interacting many-body quantum systems has yielded tremendous technological payoffs. Many anticipate that similar payoffs can be realized by achieving a better understanding of strongly interacting many-body systems, the canonical example being high-temperature superconductivity.

After a more detailed review of the atom-light interaction, we will discuss our progress toward simulation of quantum magnetism models using cold atoms in an optical lattice. The direction, polarization, and wavelength of the lattice beams can be chosen to produce a state-dependent lattice, where different states of the atom experience different potentials. The polarization of the beams can then be adjusted in real-time to move the state-dependent potentials relative to one another. This ability, which we will refer to as the shifting lattice technique, is a key tool for quantum simulation of interacting spin systems in an optical lattice. The system for implementing the shifting lattice technique was designed and built, but has yet to be installed in our apparatus. A 6.8 GHz microwave system for coherently manipulating the internal state of ^{87}Rb atoms was built and installed.

6.2 Multilevel Atom-Light Interaction

Because the wavelength of visible light is much larger than the size of an atom, it is an excellent approximation to treat an illuminated atom as experiencing an oscillating uniform electric field (the oscillating magnetic field is usually ignored since electric dipole transitions are by far the dominant process). In the so-called atomic basis, where the z axis is chosen along the magnetic field, any oscillating uniform electric field can be decomposed into

components parallel to z , right circularly polarized with respect to z , and left circularly polarized with respect to z . These components are labeled π , σ^+ , and σ^- , respectively.

From second order perturbation theory, the energy shift for a state $|i\rangle$ due to an interaction \hat{H}_1 is

$$\Delta E_i = \sum_{j \neq i} \frac{|\langle j | \hat{H}_1 | i \rangle|^2}{E_i - E_j}. \quad (6.1)$$

For light interacting with an atom, $\hat{H}_1 = -\hat{\boldsymbol{\mu}} \cdot \mathbf{E}$, where $\boldsymbol{\mu}$ is the electric dipole operator. For a two level atom with states $|e\rangle$ and $|g\rangle$ separated by an energy $\hbar\omega_0$ and light with frequency $\omega/2\pi$, this simplifies to

$$\Delta E_g = \frac{|\langle e | \hat{\boldsymbol{\mu}} | g \rangle|^2}{E_g - E_e} |E|^2 = \frac{3\pi c^2 \Gamma}{2\omega_0^3 \delta} I \quad (6.2)$$

for the ground state, where $I = \epsilon_0 c |\mathbf{E}|^2 / 2$ is the intensity, $\hat{\boldsymbol{\mu}}$ is the projection of $\boldsymbol{\mu}$ along \mathbf{E} , and we have introduced the detuning $\delta = \omega - \omega_0$ and the linewidth

$$\Gamma = \frac{\omega_0^3}{3\pi\epsilon_0 \hbar c^3} |\langle e | \hat{\boldsymbol{\mu}} | g \rangle|^2. \quad (6.3)$$

This energy shift is referred to as the light shift, AC Stark shift, or dipole potential when it is spatially dependent (due to, e.g., a spatially varying intensity). We arrived at this result by appealing to the dressed state picture, where the ground state includes a bath of n photons, so $E_g = n\hbar\omega$ and $E_e = (n-1)\hbar\omega + \hbar\omega_0$, one photon having been absorbed by the atom.

For multilevel atoms, we can add the light shifts from each transition. For an alkali atom at detunings large compared to the excited state hyperfine structure, the dipole potential for an atomic state $|F, m_F\rangle$ (with Lande g-factor g_F) becomes

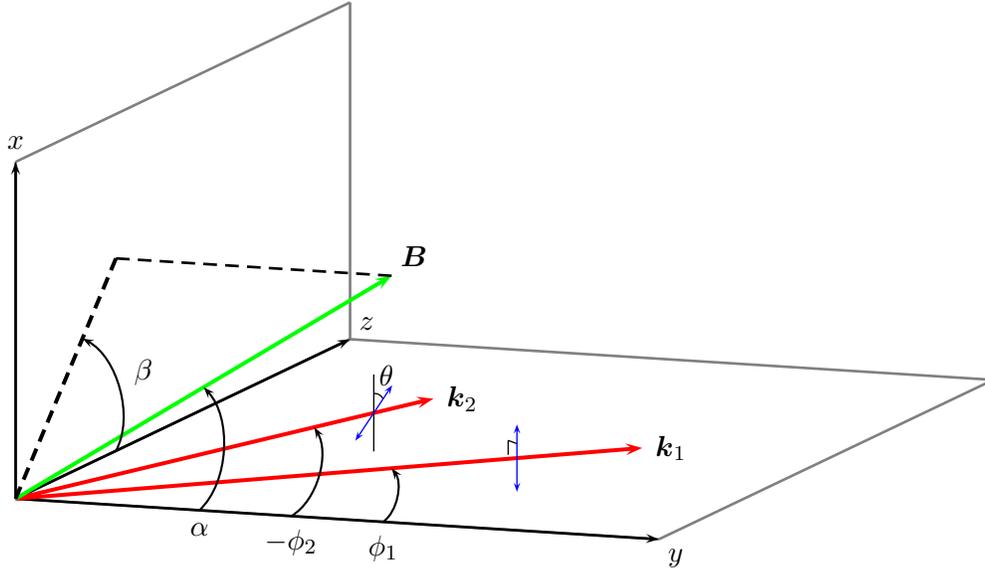
$$U_{\text{dip}}(\mathbf{r}) = \sum_P \frac{\pi c^2}{2} \left(\frac{\Gamma_{D2}}{\omega_{D2}^3} \frac{2 + P g_F m_F}{\delta_{D2}} + \frac{\Gamma_{D1}}{\omega_{D1}^3} \frac{1 - P g_F m_F}{\delta_{D1}} \right) I_P(\mathbf{r}), \quad (6.4)$$

where P denotes the polarization in the atomic basis: $P = 0$ for π and $P = \pm 1$ for σ^\pm . $I_P = c\epsilon_0 |\mathbf{E}_P(\mathbf{r})|^2 / 2$ is the intensity of the P component [132]. The detunings from the $^2P_{3/2}$ (labeled D2) and $^2P_{1/2}$ (D1) excited states are $\delta_{D2} = \omega - \omega_{D2}$ and $\delta_{D1} = \omega - \omega_{D1}$. The linewidths for the D2 and D1 transitions are Γ_{D2} and Γ_{D1} .

The scattering rate is

$$R(\mathbf{r}) = \sum_P \frac{\pi c^2}{2\hbar} \left(\frac{\Gamma_{D2}}{\omega_{D2}^3} \frac{2 + P g_F m_F}{\delta_{D2}^2} + \frac{\Gamma_{D1}}{\omega_{D1}^3} \frac{1 - P g_F m_F}{\delta_{D1}^2} \right) I_P(\mathbf{r}). \quad (6.5)$$

Notice that the scattering rate goes as δ^{-2} vs. the δ^{-1} scaling of the dipole potential.



beam	ϕ	θ	α	β
1	$\pi/2$	0	$\pi/4$	$\pi/4$
2	$\pi/2$	0	$-\pi/4$	$\pi/4$
3	$\pi/2$	0	$\pi/4$	0

Figure 6.1 Lattice beam and magnetic field geometry for equations (6.7a) and (6.8a). The two traveling wave beams creating the lattice are shown with \mathbf{k} vectors \mathbf{k}_1 and \mathbf{k}_2 in red and electric fields \mathbf{E}_1 and \mathbf{E}_2 in blue.

Since scattering of a single photon results in significant heating or loss of a trapped atom, it is best to work at the largest detuning possible for the desired potential depth, the limiting factor being the available optical power. We also remark that for a blue-detuned lattice ($U_{\text{dip}} > 0$), the atoms are trapped at the nodes of the standing wave, greatly reducing the intensity seen by the atoms and hence the scattering rate for a given lattice depth. The overall dipole force from a blue-detuned Gaussian beam is repulsive, necessitating additional external confinement or beam shaping to prevent atoms from being ejected from the lattice as it is turned on.

6.3 Optical Lattices for Multilevel Atoms

Here, we will present the potential for a 1D state-dependent optical lattice. The potential for a 3D cubic lattice can be obtained by simple addition. Referring to figure 6.1, we take one beam to be linearly polarized in the x -direction and propagating in the $y - z$ plane at an angle ϕ_1 to the y -axis. The second beam propagates in the same plane at an angle ϕ_2

to the y -axis and with linear polarization at an angle θ to the plane normal. The resulting electric field is

$$\mathbf{E}(\mathbf{r}) = \mathbf{E}_1(\mathbf{r}) + \mathbf{E}_2(\mathbf{r}), \quad (6.6)$$

where

$$\mathbf{E}_1(\mathbf{r}) = \sqrt{\frac{2I_0}{c\varepsilon_0}} e^{ik(-y \cos \phi_1 + z \sin \phi_1)} (1, 0, 0), \quad (6.7a)$$

$$\mathbf{E}_2(\mathbf{r}) = \sqrt{\frac{2I_0}{c\varepsilon_0}} e^{ik(-y \cos \phi_2 - z \sin \phi_2)} (\cos \theta, -\sin \theta \sin \phi_2, \sin \theta \cos \phi_2) \quad (6.7b)$$

Note that we have ignored the $e^{i\omega t}$ time dependence since it will drop out below. For a retro-reflected lattice, we take $\phi_1 = \phi_2 = \pi/2$ to represent beams propagating the $\pm z$ direction. If $\theta = 0$, we have so-called lin||lin lattice.

The intensity components for the field (6.6) in the atomic basis, π , σ^- , and σ^+ are, respectively,

$$I_0(\mathbf{r}) = \frac{c\varepsilon_0}{2} |\mathbf{E}(\mathbf{r}) \cdot (\sin \alpha \sin \beta, \cos \alpha, \sin \alpha \cos \beta)|^2 \quad (6.8a)$$

$$I_-(\mathbf{r}) = \frac{c\varepsilon_0}{2} |\mathbf{E}(\mathbf{r}) \cdot (\cos \beta + i \sin \beta \cos \alpha, -i \sin \alpha, i \cos \beta \cos \alpha - \sin \beta)|^2 \quad (6.8b)$$

$$I_+(\mathbf{r}) = \frac{c\varepsilon_0}{2} |\mathbf{E}(\mathbf{r}) \cdot (\cos \beta - i \sin \beta \cos \alpha, i \sin \alpha, -i \cos \beta \cos \alpha - \sin \beta)|^2 \quad (6.8c)$$

Plugging these into (6.4) yields the optical lattice potential. Scattering rates can be obtained from (6.5). The values of α and β for our system are given in figure 6.1. It can be shown that for any linear polarization and any propagation direction relative to the magnetic field, the σ^+ and σ^- components will always have equal intensities. Because of this and the fact that $U_{\text{dip},\sigma^+} + U_{\text{dip},\sigma^-} = 2U_{\text{dip},\pi}$, the depth of a lin||lin lattice is independent of polarization angle and propagation direction. Note that the lattice potential seen by the atoms is always sinusoidal.

6.4 Stroboscopic Quantum Simulation

Following Jane *et. al.* [133], we consider the simulation of a Hamiltonian for a system of many spin-1/2 particles,

$$H_N = \sum_a H^{(a)} + \sum_{a \neq b} H^{(ab)}, \quad (6.9)$$

where $H^{(a)}$ are one-body terms and $H^{(ab)}$ are two-body interactions, with a and b labeling lattice sites. We begin by noting that this Hamiltonian can be decomposed into two body terms using the Trotter formula [134]. For a Hamiltonian which can be written as $H =$

$\sum_{i=1}^m H_i$, the Trotter formula allows us to write the propagator as

$$e^{-iH\tau/\hbar} = \lim_{n \rightarrow \infty} \left[e^{(-iH_1\tau/\hbar)/n} \dots e^{(-iH_m\tau/\hbar)/n} \right]^n, \quad (6.10)$$

In words, the evolution of the system with Hamiltonian H can be approximated with a series of short time steps in which the system evolves according to each of the pieces H_i . Thus the many body Hamiltonian (6.9) can be simulated by composing short evolutions generated by one and two body terms $H^{(a)}$ and $H^{(ab)}$. This is the foundation of the so-called stroboscopic method of quantum simulation.

Let us now consider the problem of obtaining the desired two-body interaction $H^{(ab)}$, or equivalently, $U^{(ab)}(t) = \exp(iH^{(ab)}t/\hbar)$, from a given interaction $H_0^{(ab)}$. We will assume that we can perform arbitrary one-body unitary transformations. Note such unitary transformations immediately allow us to generate the evolution due to the one body term $H^{(a)}$. Let us consider an alternating sequence of short evolutions according to $H_0^{(ab)}$ and unitary transformations V_i to give a total evolution

$$\begin{aligned} U(t) &= \prod_{i=1}^n V_i \exp\left(-iH_0^{(ab)}t_i/\hbar\right) V_i^\dagger \\ &= \prod_{i=1}^n \exp\left(-iV_i H_0^{(ab)} V_i^\dagger t_i/\hbar\right) \end{aligned} \quad (6.11)$$

where $t = \sum_i t_i$ and all time steps t_i are short compared to the characteristic timescale of $H_0^{(ab)}$. The second expression can be verified by expanding the exponential. Note that the product is time-ordered. The two-body unitary transformations can be written in terms of single body operators as $V_i = u_i^{(a)} \otimes v_i^{(b)}$.

For t small compared to the eigenvalues of $H_0^{(ab)}/\hbar$, we can expand the exponentials to obtain

$$U(t) = 1 - \frac{it}{\hbar} \sum_{i=1}^n p_i V_i H_0^{(ab)} V_i^\dagger + O(t^2), \quad (6.12)$$

where $p_i = t_i/t$. This propagator produces evolution according to an effective Hamiltonian

$$H_{\text{eff}}^{(ab)} = \sum_{i=1}^n p_i V_i H_0^{(ab)} V_i^\dagger + O(t). \quad (6.13)$$

We of course require that the local operations V_i be performed on a time scale fast compared to that of $H_0^{(ab)}$.

We have now reduced the problem of simulating the evolution of $H^{(ab)}$ for a short time t to finding a set of unitary transformations V_i and time steps $p_i t$ such that $H_{\text{eff}}^{(ab)} = H^{(ab)}$. Repeated application of $U(t) = e^{-iH_{\text{eff}}^{(ab)}t/\hbar} + O(t^2)$ allows us to simulate $H^{(ab)}$ for longer times.

We now specialize to spin-1/2 systems, for which we choose two hyperfine states of the atom to serve as $|\alpha\rangle$ and $|\beta\rangle$. By using more atomic states, larger spins can be simulated. Below, we will show that the shifting lattice technique allows us to produce nearest-neighbor interactions (interactions between more distant neighbors can be produced by simply shifting the lattice further) of the form

$$H_1 = \sum_a H_0^{(a,a+1)}, \quad (6.14)$$

where

$$H_0^{(ab)} = \gamma \sigma_z^{(a)} \otimes \sigma_z^{(b)}, \quad (6.15)$$

or equivalently,

$$U_1 = \prod_a \exp[-i\theta_1 \sigma_z^{(a)} \otimes \sigma_z^{(a+1)}/\hbar], \quad (6.16)$$

where θ_1 can be adjusted by changing the interaction time. Since any spin-1/2 operator can be written in the form $u\sigma_z u^\dagger$ where u is a unitary operator, arbitrary inhomogeneous unitary operations, $V_\mu \equiv u_\mu^{(a)} \otimes v_\mu^{(b)}$, would allow us to achieve any spin-1/2 two-body Hamiltonian,

$$H^{(ab)} = \sum_{\mu,\nu=x,y,z} M_{\mu\nu} \sigma_\mu^{(a)} \otimes \sigma_\nu^{(b)}. \quad (6.17)$$

Such operations would require the addressing of individual atoms in the lattice. Microwave transitions can be used to apply an arbitrary unitary transformation to all atoms in the lattice at once, i.e., $V_\mu = u_\mu^{(a)} \otimes u_\mu^{(b)}$. It is easy to show [135] that this introduces the restriction in (6.17) that the matrix M be real and symmetric.

Many Hamiltonians of interest satisfy this requirement. For example, consider the interaction term for the antiferromagnetic (AF) Heisenberg Hamiltonian,

$$H^{(ab)} = -J \sum_{\mu=x,y,z} \sigma_\mu^{(a)} \otimes \sigma_\mu^{(b)}, \quad (6.18)$$

where $J < 0$. An evolution according to this Hamiltonian can be simulated by homogenous unitary operations alternated with short evolutions under $H_0^{(ab)}$ (6.15), if $\gamma > 0$. The unitary operations and time steps in (6.11) are easily calculated to be:

$$\begin{aligned} p_1 &= \frac{1}{3}, & V_1 &= 1 \otimes 1, \\ p_2 &= \frac{1}{3}, & V_2 &= \frac{1 - i\sigma_x}{\sqrt{2}} \otimes \frac{1 - i\sigma_x}{\sqrt{2}}, \\ p_3 &= \frac{1}{3}, & V_3 &= \frac{1 - i\sigma_y}{\sqrt{2}} \otimes \frac{1 - i\sigma_y}{\sqrt{2}}. \end{aligned} \quad (6.19a)$$

Recall that $p_i = t_i/t$ and notice that V_2 and V_3 are simply rotations.

6.5 Systems of Interest

Simulation of the 1D and 2D ferromagnetic Ising and 2D AF Heisenberg models can be used to validate and demonstrate the stroboscopic methods described above. These are special cases of the anisotropic Heisenberg Hamiltonian,

$$-\sum_{\langle a,b \rangle} \sum_{\mu=x,y,z} J_{\mu} \sigma_{\mu}^{(a)} \otimes \sigma_{\mu}^{(b)} - \mu_M \sum_a \sum_{\mu=x,y,z} H_{\mu} \sigma_{\mu}^{(a)}, \quad (6.20)$$

where \mathbf{H} is an external magnetic field and μ_M is the magnetic moment. The first summation is over nearest neighbors on the lattice. For the Ising model, $J_x = J_y = 0$, $H_x = H_y = 0$, the exact solution is known in 1D and 2D. Comparing simulation results to these solutions can be used to debug, optimize, and validate the simulation procedure. A next step could be implementation of the 2D AF Heisenberg Hamiltonian, $J_x = J_y = J_z \equiv J < 0$ which is mostly understood but for which no exact solution exists. For example, by looking at density-density correlations 4.4.2 in state sensitive images we could verify the presence of AF (Néel) ordering by observing a lattice period for each state twice that of the underlying optical lattice. The 2D AF Heisenberg model is of particular interest because it is relevant to the 2D Hubbard model (which reduces to the AF Heisenberg in the large U/J limit) and cuprate superconductors.

Another set of interesting systems are spin liquids in frustrated geometries, a square lattice with next-nearest-neighbor interactions, for example. More exotically, Santos *et. al.* [136] have shown that it is possible to create an ideal Kagome lattice with tunable couplings using three lattice beams.

6.6 Implementation

6.6.1 Overview

We choose two states of ^{87}Rb to serve as the effective spin up and spin down states, for example, $|F = 1, m_F = -1\rangle \equiv |\alpha\rangle$ and $|F = 2, m_F = -2\rangle \equiv |\beta\rangle$. This choice has the advantage that the two states can be coupled by a single microwave photon. The procedure for implementing stroboscopic quantum simulation is as follows. First, the system is prepared with one atom on each site of the optical lattice in a known state. We then use the ability to independently control the lattices seen by the two states to produce state-selective controlled collisions which generate an effective spin-spin interaction between neighboring sites. Homogenous unitary operations, i.e., rotations of the effective spins, are produced by microwave pulses. Stimulated Raman transitions, driven by two phase-locked lasers, can be used to perform spin rotations with regional addressing in the lattice. Major challenges in implementation include magnetic field stability, spontaneous scattering from the lattice beams, and heating during the lattice shifts.

6.6.2 State-Dependent Lattices

For a lin||lin lattice, where the forward and retro-reflected beams have the same linear polarization, we get a simple standing wave and the lattice potential arises from the spatially varying intensity. If the retro-reflected polarization is rotated by 90° , the lin \perp lin case, there is no longer any spatial variation in intensity. Instead, the polarization varies from σ^+ to linear to σ^- to linear and finally back to σ^+ over a distance $\lambda/2$. While for large detunings relative to the fine structure, so that $\delta_{D1} \approx \delta_{D2}$, the polarization dependence disappears from the light shift (6.4), at small detunings the polarization dependence of the light shift will produce a lattice potential.

Figure 6.2 shows the potentials seen by atoms in two different states, $|F = 2, m_F = -2\rangle$ and $|F = 1, m_F = -1\rangle$ as the polarization is rotated through 180° degrees (this corresponds to a half-wavelength shift of the retro-reflected beam, which is why the potential wells shift by a quarter wavelength, rather than return to their original positions). Notice that the lattice shifts in opposite directions for the two states so that at the completion of the 180° degree rotation the $|F = 1, m_F = -1\rangle$ component of the state in each well has been brought together with the $|F = 2, m_F = -2\rangle$ component of the state in the well to its right. The change in lattice depth as the polarization is rotated can be reduced by working at a different wavelength or compensated for by adjusting the optical power as the polarization is rotated. Figure 6.3 shows the potentials for the other two lattice beams in our system for two different polarizations, demonstrating that the shifting lattice technique works for all three lattice beams even with a fixed magnetic field direction.

The shifting lattice technique has been demonstrated in one dimension by Mandel *et al.* [101]. For simulating 2D magnetic systems, independent control along two different lattice directions is required. To implement the technique, the retro-reflection path for each lattice beam must be modified, as shown in figure 6.4. A quarter-wave plate (QWP) and an electro-optic modulator (EOM), with its axes at 45° to the QWP axes are placed just before the retro-reflection mirror in the lattice beam path to control the polarization of the retro-reflected light (see figure 4.6). The Jones matrix for this configuration is

$$\begin{pmatrix} \cos bV & \sin bV \\ \sin bV & -\cos bV \end{pmatrix}, \quad (6.21)$$

where V is the voltage on the modulator and b is a constant with units of radians per volt. This is just the Jones matrix for a polarization rotation so that if linearly polarized light enters the system, the retro-reflected light is also linear and rotated by bV . The half-wave plate rotates the incoming polarization to align it with an axis of the QWP (this is done since it is difficult to mount the EOM at an arbitrary angle). To achieve the full 180° rotation needed to shift the lattice by one site, the EOM must be driven with several hundred volts. For each EOM, a ± 500 V supply and a high voltage operational amplifier

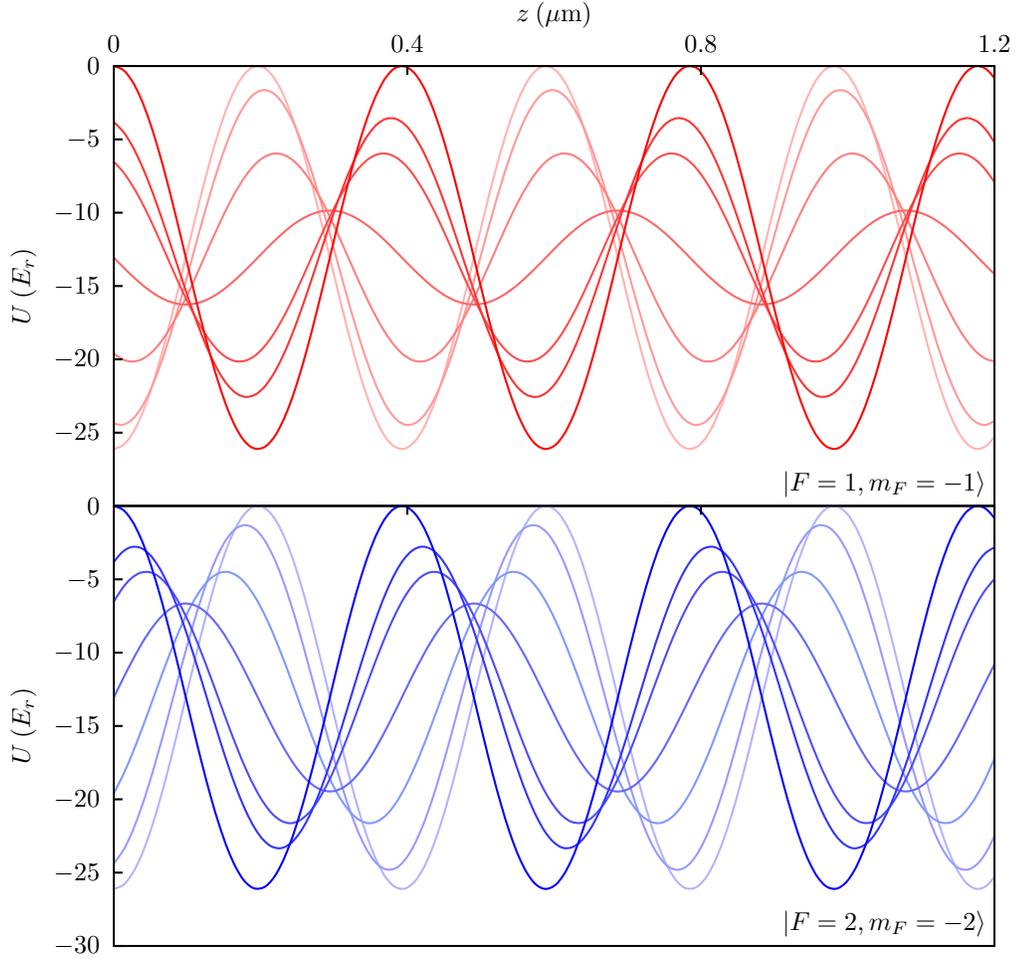


Figure 6.2 State dependent potentials for a 785 nm lattice as the polarization of the retro-reflected beam is rotated through 180° in 30° increments (increasing color). For this calculation, the intensity of all forward and retro beams is 0.4 W/mm^2 at 785 nm, corresponding to a lattice depth of $26.1 E_R$ for $\theta = 0$. In practice, the lattice beam power can be adjusted during the shift process to minimize the change in lattice depth.

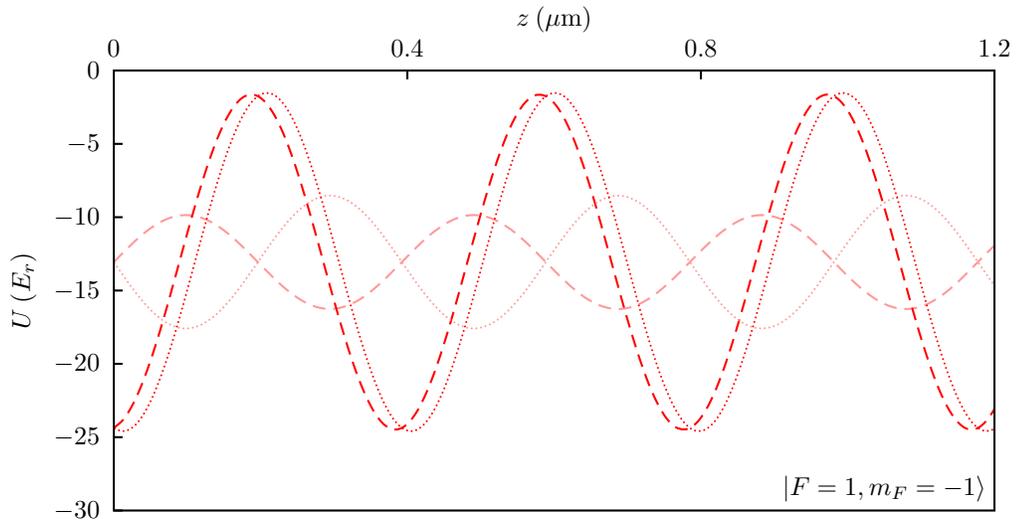


Figure 6.3 Potentials for $|F = 1, m_F = -1\rangle$ state for the two other lattice beams in our system (dashed and dotted) with retro-reflection polarizations of 30° and 90° (lighter color). The optical power and wavelength are the same as in figure 6.2.

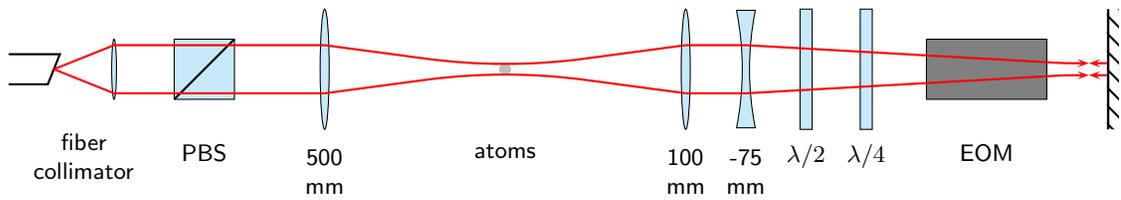


Figure 6.4 Optical system for implementing shifting lattice technique.

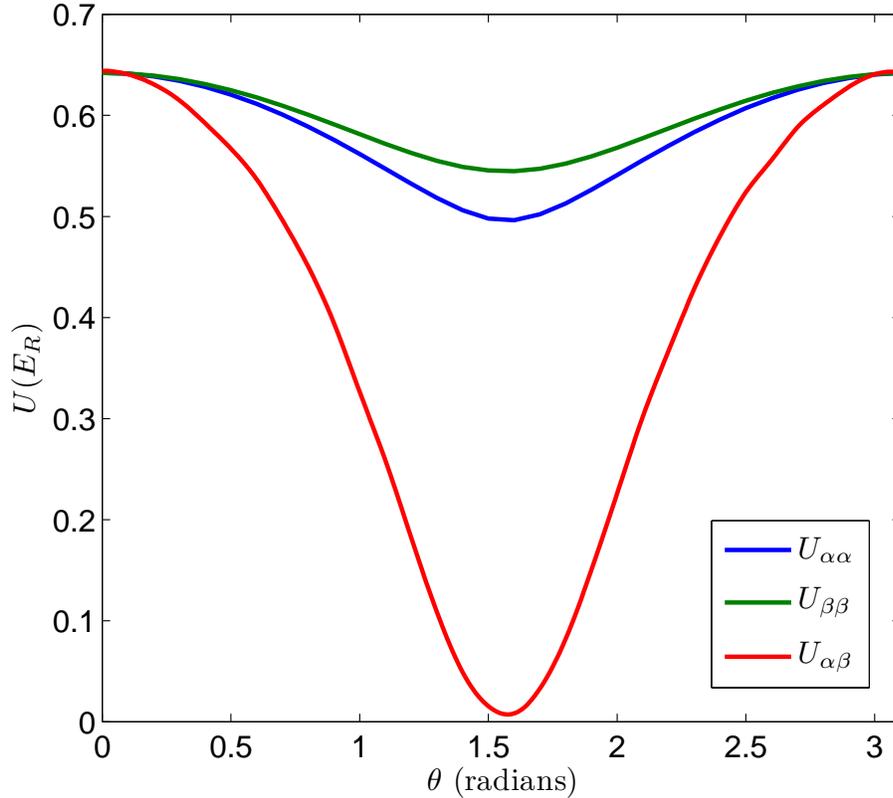


Figure 6.5 Interaction energies $U_{\alpha\alpha}$, $U_{\beta\beta}$, and $U_{\alpha\beta}$ vs. retro beam polarization angle θ for lattice beam 1; α is $|F = 1, m_F = -1\rangle$ and β is $|F = 2, m_F = -2\rangle$. The decrease in $U_{\alpha\alpha}$ and $U_{\beta\beta}$ results from the decrease in lattice depth for $\theta \neq 0$. The same parameters are used as in figure 6.2, except that the lattice beam power is adjusted to partially compensate for the change in lattice depth.

circuit (using a Cirrus/Apex PA94 IC) were constructed. The ± 500 V supply consists of four linear OEM voltage supplies (250 V, 0.1 A) connected in series (note that this is not recommended by the manufacturer).

6.6.3 State Preparation

Preparation of the initial state of uniform singly-occupied sites can be accomplished according to the proposal by Rabl *et. al.* [137]. The first step is loading a BEC into the lattice, then turning up the lattice through the superfluid/Mott insulator phase transition to localize the atoms on lattice sites. This is done with a sufficiently dense BEC to give a central core of mostly multiply occupied sites with a very low probability of vacant sites (section 4.3.5) [138].

We now consider the single-site Hamiltonian including two different internal atomic states. We will show that it is possible to transfer exactly one atom on each site of the optical lattice to the other state (i.e., flip the spin) via adiabatic rapid passage. Specifically,

the detuning of a coupling between the states, here, a microwave signal, is swept through a single avoided crossing.

If the potential minima for states α and β are offset by a distance Δx , say in the \hat{x} direction, without otherwise changing the potentials (so that we can use the same ground state wavefunctions), the interaction energy between states α and β is

$$U_{\alpha\beta} = \int |w_\alpha(\mathbf{r})|^2 |w_\beta(\mathbf{r} - \Delta x \hat{x})|^2 d^3r. \quad (6.22)$$

This value is plotted in figure 6.5, along with $U_{\alpha\alpha}$ and $U_{\beta\beta}$. Note that $U_{\alpha\beta}$ will always be smaller than $U_{\alpha\alpha}$, which is equal (or very nearly equal) to $U_{\beta\beta}$ for the same lattice depth. Both these restrictions can be overcome by using other methods to tuning the atomic interactions, such as Feshbach resonances [10]. If we now apply an RF field coupling states α and β with frequency ω and spatially dependent Rabi rate $\Omega(\mathbf{r})$, the full single site Hamiltonian is

$$\hat{H} = \frac{U_{\alpha\alpha}}{2} \hat{n}_\alpha(\hat{n}_\alpha - 1) + \frac{U_{\beta\beta}}{2} \hat{n}_\beta(\hat{n}_\beta - 1) + U_{\alpha\beta} \hat{n}_\alpha \hat{n}_\beta - \frac{W}{2} (\hat{\alpha}^\dagger \hat{\beta} + \hat{\beta}^\dagger \hat{\alpha}) - \hbar \delta \hat{n}_\beta, \quad (6.23)$$

where $\delta = \omega - \omega_0$ and

$$W = \int w_\alpha^*(\mathbf{r}) \Omega(\mathbf{r}) w_\beta(\mathbf{r} - \Delta x \hat{x}) d^3r. \quad (6.24)$$

The interaction terms follow from the usual Hamiltonian in the rotating wave approximation [59]:

$$\hat{H}_{\text{int}} = \frac{\hbar}{2} \begin{bmatrix} 0 & \Omega \\ \Omega & -2\delta \end{bmatrix} \quad (6.25)$$

Assuming all atoms start in state α , corresponding to the state $|N, 0\rangle$, we want to determine the range over which to sweep δ to move adiabatically to the state $|N - 1, 1\rangle$ for some range of N values, while avoiding the crossing between $|N - 1, 1\rangle$ and $|N - 2, 2\rangle$. From (6.22), we see that the $|N, 0\rangle, |N - 1, 1\rangle$ crossing occurs at $\hbar\delta_1(N) = (U_{\alpha\beta} - U_{\alpha\alpha})(N - 1)$ and the $|N - 1, 1\rangle, |N - 2, 2\rangle$ crossing occurs at $\hbar\delta_2(N) = U_{\alpha\beta}(N - 3) - U_{\alpha\alpha}(N - 2) + U_{\beta\beta}$ —this is the value that we must avoid after sweeping through δ_1 . Note that we neglect higher-order transitions which transfer more than one atom between states α and β .

Given the restrictions on how the interactions may be tuned in our system, satisfying these constraints for the case considered in [137] of sweeping from smaller to larger detuning, $\delta_i < \delta_f$, is difficult. If instead, $\delta_i > \delta_f$, $\delta_2(N)$ is encountered before $\delta_1(N)$ and the problem becomes easy. If we wish to transfer a single atom from α to β on sites with between N_{\min} and N_{\max} atoms, we require $\delta_i > (U_{\alpha\beta} - U_{\alpha\alpha})(N_{\min} - 1)/\hbar$ and $\delta_f < (U_{\alpha\beta} - U_{\alpha\alpha})(N_{\max} - 1)/\hbar$ (recall that $U_{\alpha\beta} - U_{\alpha\alpha} < 0$).

6.6.4 Controlled Interaction

Spin-spin interactions between neighboring atoms (not necessarily nearest-neighbors) can be simulated by shifting the lattices for the two states relative to one another to bring atoms in opposite states into proximity for a controlled amount of time—the interaction between the atoms produces a phase shift [101, 139]. The process is shown schematically in figure 6.6. Notice that, if the lattice for $|\alpha\rangle$ is shifted to the right, the interaction only occurs between the $|\alpha\rangle$ component of the state on a site and the $|\beta\rangle$ component of the state on the site to the right. We have:

$$|\alpha\rangle_1|\alpha\rangle_2 \rightarrow e^{-i2\phi_\alpha}|\alpha\rangle_1|\alpha\rangle_2 \quad (6.26)$$

$$|\alpha\rangle_1|\beta\rangle_2 \rightarrow e^{-i(\phi_\alpha+\phi_\beta+\phi_{\alpha\beta})}|\alpha\rangle_1|\beta\rangle_2 \quad (6.27)$$

$$|\beta\rangle_1|\alpha\rangle_2 \rightarrow e^{-i(\phi_\beta+\phi_\alpha)}|\beta\rangle_1|\alpha\rangle_2 \quad (6.28)$$

$$|\beta\rangle_1|\beta\rangle_2 \rightarrow e^{-i2\phi_\beta}|\beta\rangle_1|\beta\rangle_2 \quad (6.29)$$

where the subscripts on the state vectors label lattice sites. Theoretically, the collisional phase shift is just

$$\phi_{\alpha\beta} = \frac{1}{\hbar} \int U_{\alpha\beta} dt, \quad (6.30)$$

where the integral is over the interaction time. In reality, the collisional phase shift must be determined experimentally [101]. The additional kinetic phases, ϕ_α , ϕ_β , arise from the shifting process and can be removed by unitary transformations or simply incorporated into the definitions of $|\alpha\rangle$ and $|\beta\rangle$. It is then easy to see that the effective Hamiltonian for this evolution is proportional to

$$\frac{1}{2}(\sigma_z^{(1)} + 1) \otimes \frac{1}{2}(\sigma_z^{(2)} - 1) = -|\alpha\rangle_1|\beta\rangle_2\langle\alpha|_1\langle\beta|_2. \quad (6.31)$$

Since all sites along a lattice direction are shifted together, we have

$$\begin{aligned} & \sum_{i=1}^n \frac{1}{2}(\sigma_z^{(i)} + 1) \otimes \frac{1}{2}(\sigma_z^{(i+1)} - 1) \\ &= \frac{1}{4} \sum_i (\sigma_z^{(i)} \otimes \sigma_z^{(i+1)} - \sigma_z^{(i)} + \sigma_z^{(i+1)} - 1) \\ &= \frac{1}{4} \sum_i \sigma_z^{(i)} \otimes \sigma_z^{(i+1)} + \frac{1}{4}(-\sigma_z^{(1)} + \sigma_z^{(n)} - n) \\ &\rightarrow \frac{1}{4} \sum_i \sigma_z^{(i)} \otimes \sigma_z^{(i+1)} \end{aligned} \quad (6.32)$$

To arrive at the final expression, the desired spin-spin interaction, we have neglected the two sites at the edge of the lattice, which is reasonable for large n , and discarded the constant term $-n/4$.

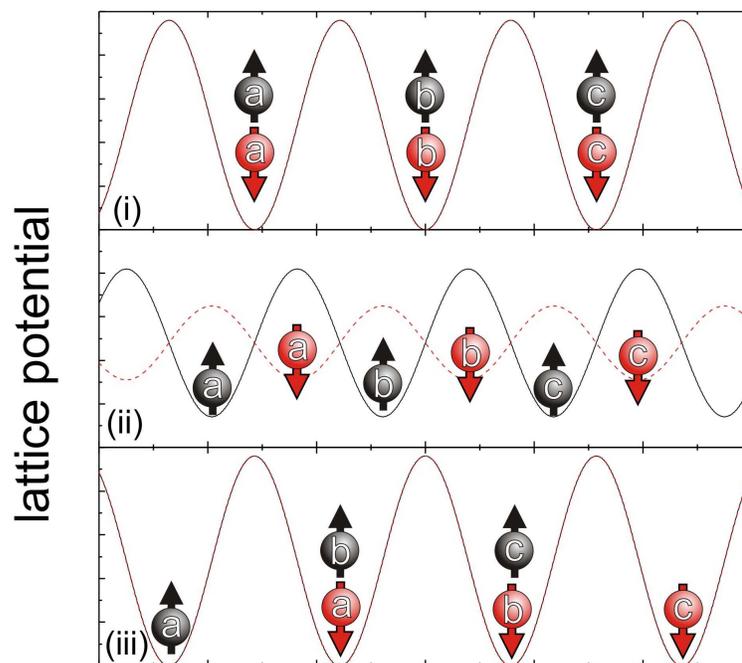


Figure 6.6 Effective spin-spin interactions using the shifting lattice technique. The potential experienced by different atomic states is shown as the lattice is shifted. Here, the $|\alpha\rangle$ (\uparrow) components of atoms a and b interact with the $|\beta\rangle$ (\downarrow) components of atoms b and c , respectively.

Adiabaticity requires that the shifting process proceed slowly enough to satisfy $|\dot{\bar{x}}(t)| \ll a_0\omega$ [140], where \bar{x} is the distance between antinodes (equivalently, nodes) of the lattices seen by the two states, ω is the oscillator frequency for a lattice site, and $a_0 = \sqrt{\hbar/(m\omega)}$ is the corresponding ground state width. Deep in the Mott insulator regime, $\omega/2\pi \sim 50$ kHz, requiring shift times $\gg 20 \mu\text{s}$.

6.6.5 Measurement

The final step is readout of the simulation. Measurement of atom shot noise correlations in the momentum distribution is expected to be a useful technique. Many images of atoms after sufficient expansion to reveal the momentum distribution are averaged together and the autocorrelation function for resulting image is calculated. This technique has been used to reveal the ordering of the Mott insulator state [117].

The dispersion relation for elementary excitations can be measured via Bragg spectroscopy [141]. This technique involves measuring the amount of Bragg scattering (by measuring the amount of heating of the atoms) for a given energy and momentum difference between two laser beams shining on the atoms (i.e., the beams have different frequencies and approach from different directions). A final measurement technique of interest is state-selective Bragg scattering [142], in which a probe beam resonant with only one of the spin states is Bragg reflected at a specific angle when anti-ferromagnetic ordering is present.

6.7 Microwave System

6.7.1 Microwave Source

A 2–8 GHz yttrium iron garnet (YIG) tuned oscillator (Microsource MCO-0207-500-01), henceforth called the VCO, is used to generate a 6.8 GHz microwave signal to drive transitions between $F = 1$ and $F = 2$ states in ^{87}Rb , figure 6.7. The frequency of the VCO is adjusted by the magnetic field from two coils acting on a YIG sphere. A high inductance coil is used for coarse tuning while a lower inductance coil is used for fast, fine tuning. The VCO is phase-locked, ultimately to the LORAN-C frequency reference (section 3.2.4), by the signal chain shown in figure 6.8. Computer control of the DDS provides tuning of the VCO frequency and phase, limited by the phase-locked loop (PLL) bandwidth. A voltage-controlled attenuator (Pulsar Microwave AAT-22V) is used to control the amplitude of the microwave signal and a microwave power amplifier drives a coil mounted near the science cell. The amplifier is protected from power reflected by the coil using a circulator. The power amplifier is switched off when the microwave signal is not in use to provide further attenuation and eliminate the need for active cooling.

Phase noise in the microwave signal results in reduced fidelity for spin rotations. Therefore, significant effort was devoted to the PLL design. Minimizing the factor by which the

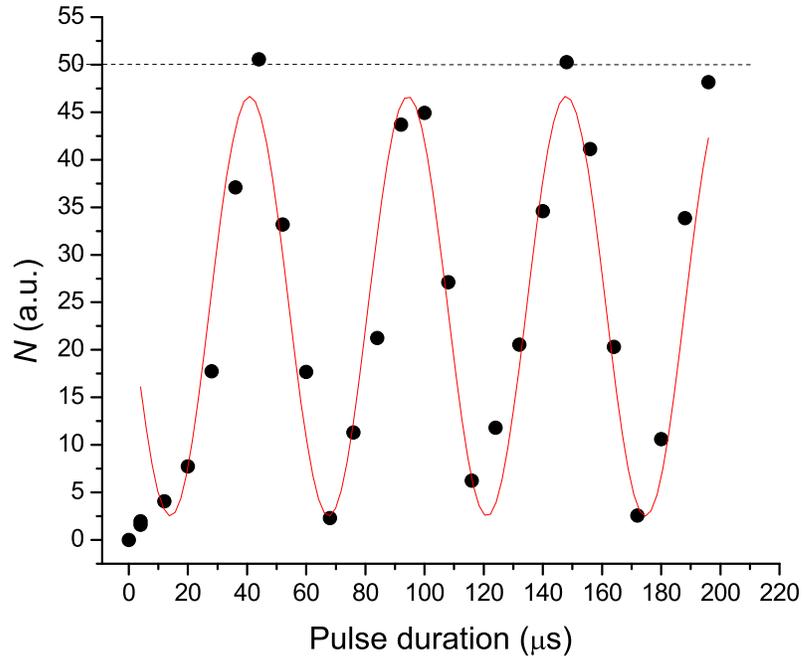


Figure 6.7 Rabi oscillations between $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = -2\rangle$ driven by a microwave signal at 6826.267 MHz and 3 dBm. The dashed line indicates the average number of atoms determined by optical repumping. The red trace is sinusoidal fit giving a π time of 26.7 μs . The initial delay is due to the turn-on time of the microwave attenuator.

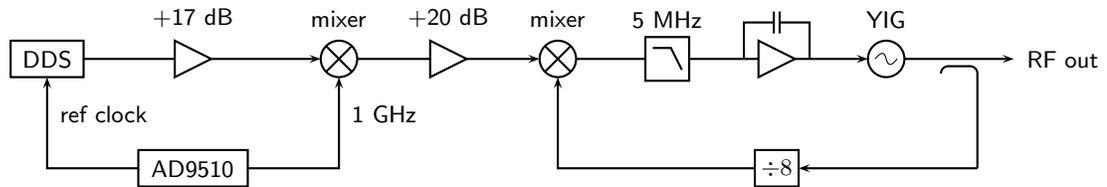


Figure 6.8 Phase-locked loop for YIG oscillator. The AD9510 is part of the frequency reference system described in section 3.2.4. The DDS is a AD9959 (500 MHz ref clock). Using the second mixer as a phase detector, the frequency-divided YIG output is locked to the lower sideband (at $1 \text{ GHz} - f_{\text{DDS}}$) produced by the first mixer. The frequency divider is a Hittite Microwave HMC434.

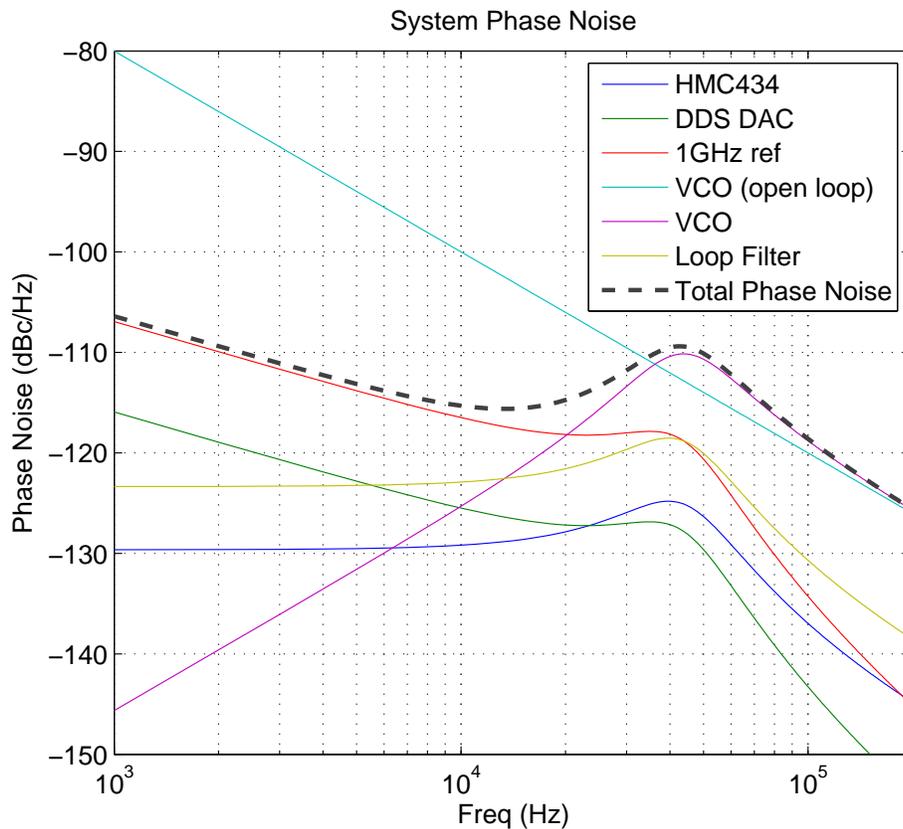


Figure 6.9 Phase noise calculated for PLL shown in figure 6.8. The contributions for each component are shown for closed loop operation (VCO refers to the YIG oscillator, HMC434 to the frequency divider). The open loop VCO phase noise is shown for comparison. The loop design bandwidth (3 dB) is 50 kHz, with a phase margin of 40° .

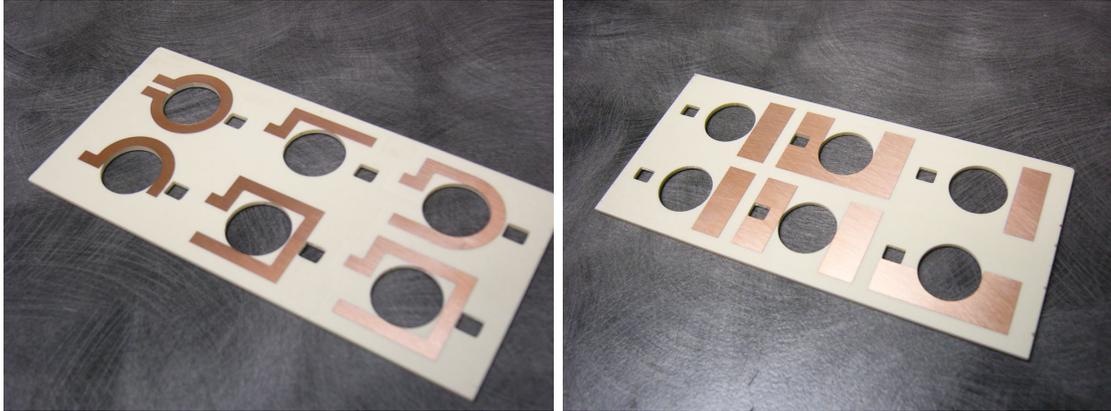


Figure 6.10 Microwave coils as fabricated. The board was subsequently diced and each coil tested separately.

VCO frequency is divided before the phase detector is essential, since the noise contributions from most of the loop components are multiplied by this factor. Upconverting the DDS output by mixing it with a very low phase noise 1 GHz signal from the AD9510 reduces the frequency division needed for the PLL signal to 8. Thus, phase noise from the DDS and reference are scaled up by only 18 dB. The ultimate lower bound on the VCO phase noise power inside the PLL feedback bandwidth is the reference phase noise power multiplied by the magnitude of the closed loop gain squared (which, inside the loop bandwidth, is equal to the frequency division factor squared). In effect, the VCO cannot be less noisy than the reference. Outside the feedback bandwidth, the phase noise is the same as for the free-running VCO.

We choose the loop bandwidth to make the phase noise everywhere less than or equal to that of the free-running VCO. This is slightly easier to follow than minimizing the integrated VCO phase noise—at some point the (often approximately white) noise of a loop component will cross the plot of free-running VCO phase noise. The first such crossing is chosen as the cut-off frequency of the loop filter. During testing, very good agreement was found between the phase noise predicted by the MATLAB calculations and the VCO phase noise measured using a spectrum analyzer.

Some difficulties were encountered with the home-built microwave amplifier (using an M/A-COM MAAPGM0060), which can be traced back to a failure to adhere to what could be deemed the first rule of amateur microwave design: each component should be placed in its own solid metal box. We subsequently purchased a 15 W commercial microwave amplifier (KU 682 XH-UM from Kuhne Electronic GmbH).

6.7.2 Microwave Coil

Six different coil designs were fabricated on high-frequency circuit boards (Rogers 4350B, distinguished from standard board by the use of lower loss dielectric), figure 6.10. The

performance of each near 6.8 GHz was characterized by measuring their reflection coefficient with a network analyzer. Rudimentary stub tuning was performed in situ by soldering sections of copper foil to the input end of the antenna trace and then slowly cutting the sections shorter while monitoring the reflection coefficient. In this way, a coil with roughly 50% coupling at 6.8 GHz was created.

Stub tuning an element with very low resistance amounts to building a high-Q resonator, not an easy task. Thus, impedance matching a loop antenna is much more sensitive to small errors in stub length and distance than, for example, matching a 25 ohm resistor.

References

- [1] Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. and Cornell, E. A. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science* **269**, 198 (1995).
- [2] Davis, K. B., Mewes, M. O., Andrews, M. R., van Druten, N. J., Durfee, D. S., Kurn, D. M. and Ketterle, W. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.* **75**, 3969–3973 (1995).
- [3] DeMarco, B. and Jin, D. Onset of Fermi degeneracy in a trapped atomic gas. *Science* **285**, 1703 (1999).
- [4] Matthews, M. R., Anderson, B. P., Haljan, P. C., Hall, D. S., Wieman, C. E. and Cornell, E. A. Vortices in a Bose-Einstein condensate. *Phys. Rev. Lett.* **83**, 2498–2501 (1999).
- [5] Andrews, M. R., Townsend, C. G., Miesner, H.-J., Durfee, D. S., Kurn, D. M. and Ketterle, W. Observation of interference between two Bose condensates. *Science* **275**, 637–641 (1997).
- [6] Walz, J., Fendel, P., Herrmann, M., König, M., Pahl, A., Pittner, H., Schatz, B. and Hänsch, T. W. Towards laser spectroscopy of antihydrogen. *J. Phys. B: At. Mol. Opt. Phys.* **36**, 649–654 (2003).
- [7] Pospelov, M. and Ritz, A. Electric dipole moments as probes of new physics. *Annals of Physics* **318**, 119 – 169 (2005).
- [8] Guest, J. R., Scielzo, N. D., Ahmad, I., Bailey, K., Greene, J. P., Holt, R. J., Lu, Z.-T., O'Connor, T. P. and Potterveld, D. H. Laser trapping of ^{225}Ra and ^{226}Ra with repumping by room-temperature blackbody radiation. *Phys. Rev. Lett.* **98**, 093001 (2007).
- [9] Greiner, M., Mandel, O., Esslinger, T., Hänsch, T. W. and Bloch, I. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature* **415**, 39 (2002).
- [10] Pethick, C. J. and Smith, H. *Bose-Einstein Condensation in Dilute Gases* (Cambridge University Press, Cambridge, 2008), 2nd edn.
- [11] Reppy, J. D. Superfluid helium in porous media. *J. Low Temp. Phys.* **87**, 205 (1992).
- [12] Kim, E. and Chan, M. H. W. Probable observation of a supersolid helium phase. *Nature* **427**, 225 (2004).

- [13] Kim, E. and Chan, M. H. W. Supersolid helium at high pressure. *Phys. Rev. Lett.* **97**, 115302 (2006).
- [14] Rittner, A. S. C. and Reppy, J. D. Disorder and the supersolid state of solid ^4He . *Phys. Rev. Lett.* **98**, 175302 (2007).
- [15] Caupin, F., Sasaki, S. and Balibar, S. Supersolidity and disorder in solid helium 4. *J. Low Temp. Phys.* **150**, 267 (2008).
- [16] C. Bruder, G. S., R. Fazio. The Bose-Hubbard model: from Josephson junction arrays to optical lattices. *Annalen der Physik* **14**, 566 (2005).
- [17] Kosterlitz, J. M. and Granato, E. Disorder in Josephson junction arrays. *Physica B* **152**, 62 (1988).
- [18] Lang, K. M., Madhavan, V., Hoffman, J. E., Hudson, E. W., Eisaki, H., Uchida, S. and Davis, J. C. Imaging the granular structure of high- T_c superconductivity in underdoped $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+\delta}$. *Nature* **415**, 412 (2002).
- [19] Fisher, M. P. A., Weichman, P. B., Grinstein, G. and Fisher, D. S. Boson localization and the superfluid-insulator transition. *Phys. Rev. B* **40**, 546–570 (1989).
- [20] Sheshadri, K., Krishnamurthy, H. R., Pandit, R. and Ramakrishnan, T. V. Percolation-enhanced localization in the disordered bosonic Hubbard model. *Phys. Rev. Lett.* **75**, 4075–4078 (1995).
- [21] Bissbort, U. and Hofstetter, W. Stochastic mean-field theory for the disordered Bose-Hubbard model (2008). [arXiv:0804.0007](https://arxiv.org/abs/0804.0007).
- [22] Wu, J. and Phillips, P. Minimal model for disorder-induced missing moment of inertia in solid ^4He . *Phys. Rev. B* **78**, 014515 (2008).
- [23] Buonsante, P., Penna, V., Vezzani, A. and Blakie, P. B. Mean-field phase diagram of cold lattice bosons in disordered potentials. *Phys. Rev. A* **76**, 011602 (2007).
- [24] Damski, B., Zakrzewski, J., Santos, L., Zoller, P. and Lewenstein, M. Atomic Bose and Anderson glasses in optical lattices. *Phys. Rev. Lett.* **91**, 080403 (2003).
- [25] Krutitsky, K. V., Pelster, A. and Graham, R. Mean-field phase diagram of disordered bosons in a lattice at nonzero temperature. *New J. Phys.* **8**, 187 (2006).
- [26] Singh, K. G. and Rokhsar, D. S. Real-space renormalization study of disordered interacting bosons. *Phys. Rev. B* **46**, 3002–3008 (1992).
- [27] Balabanyan, K. G., Prokof'ev, N. and Svistunov, B. Superfluid-insulator transition in a commensurate one-dimensional bosonic system with off-diagonal disorder. *Phys. Rev. Lett.* **95**, 055701 (2005).
- [28] Prokof'ev, N. and Svistunov, B. Superfluid-insulator transition in commensurate disordered bosonic systems: Large-scale worm algorithm simulations. *Phys. Rev. Lett.* **92**, 015703 (2004).

- [29] Krauth, W., Trivedi, N. and Ceperley, D. Superfluid-insulator transition in disordered boson systems. *Phys. Rev. Lett.* **67**, 2307–2310 (1991).
- [30] Kisker, J. and Rieger, H. Bose-glass and Mott-insulator phase in the disordered boson Hubbard model. *Phys. Rev. B* **55**, R11981–R11984 (1997).
- [31] Sengupta, P., Raghavan, A. and Haas, S. Disorder-enhanced phase coherence in trapped bosons on optical lattices. *New J. Phys.* **9**, 103 (2007).
- [32] Scalettar, R. T., Batrouni, G. G. and Zimanyi, G. T. Localization in interacting, disordered, Bose systems. *Phys. Rev. Lett.* **66**, 3144–3147 (1991).
- [33] Jaksch, D., Bruder, C., Cirac, J. I., Gardiner, C. W. and Zoller, P. Cold bosonic atoms in optical lattices. *Phys. Rev. Lett.* **81**, 3108 (1998).
- [34] Fallani, L., Lye, J. E., Guarrera, V., Fort, C. and Inguscio, M. Ultracold atoms in a disordered crystal of light: Towards a Bose glass. *Phys. Rev. Lett.* **98**, 130404 (2007).
- [35] Lye, J. E., Fallani, L., Fort, C., Guarrera, V., Modugno, M., Wiersma, D. S. and Inguscio, M. Effect of interactions on the localization of a Bose-Einstein condensate in a quasiperiodic lattice. *Phys. Rev. A* **75**, 061603 (2007).
- [36] Edwards, E. E., Beeler, M., Hong, T. and Rolston, S. L. Adiabaticity and localization in one-dimensional incommensurate lattices (2008). [arXiv:0808.4143](https://arxiv.org/abs/0808.4143).
- [37] Lye, J. E., Fallani, L., Modugno, M., Wiersma, D. S., Fort, C. and Inguscio, M. Bose-Einstein condensate in a random potential. *Phys. Rev. Lett.* **95**, 070401 (2005).
- [38] Clement, D., Varon, A. F., Retter, J. A., Sanchez-Palencia, L., Aspect, A. and Bouyer, P. Experimental study of the transport of coherent interacting matter-waves in a 1D random potential induced by laser speckle. *New J. Phys.* **8**, 165 (2006).
- [39] Clement, D., Bouyer, P., Aspect, A. and Sanchez-Palencia, L. Density modulations in an elongated Bose-Einstein condensate released from a disordered potential. *Phys. Rev. A* **77**, 033631 (2008).
- [40] Chen, Y. P., Hitchcock, J., Dries, D., Junker, M., Welford, C. and Hulet, R. G. Phase coherence and superfluid-insulator transition in a disordered Bose-Einstein condensate. *Phys. Rev. A* **77**, 033632 (2008).
- [41] Schulte, T., Drenkelforth, S., Kruse, J., Ertmer, W., Arlt, J., Sacha, K., Zakrzewski, J. and Lewenstein, M. Routes towards Anderson-like localization of Bose-Einstein condensates in disordered optical lattices. *Phys. Rev. Lett.* **95**, 170411 (2005).
- [42] Ospelkaus, S., Ospelkaus, C., Wille, O., Succo, M., Ernst, P., Sengstock, K. and Bongs, K. Localization of bosonic atoms by fermionic impurities in a three-dimensional optical lattice. *Phys. Rev. Lett.* **96**, 180403 (2006).
- [43] Anderson, P. W. Absence of diffusion in certain random lattices. *Phys. Rev.* **109**, 1492–1505 (1958).
- [44] Schwartz, T., Bartal, G., Fishman, S. and Segev, M. Transport and Anderson localization in disordered two-dimensional photonic lattices. *Nature* **446**, 52 (2007).

- [45] Weaver, R. L. Anderson localization of ultrasound. *Wave Motion* **12**, 129 (1990).
- [46] Billy, J., Josse, V., Zuo, Z., Bernard, A., Hambrecht, B., Lugan, P., Clément, D., Sanchez-Palencia, L., Bouyer, P. and Aspect, A. Direct observation of Anderson localization of matter waves in a controlled disorder. *Nature* **453**, 891 (2008).
- [47] Roati, G., D’Errico, C., Fallani, L., Fattori, M., Fort, C., Zaccanti, M., Modugno, G., Modugno, M. and Inguscio, M. Anderson localization of a non-interacting Bose-Einstein condensate. *Nature* **453**, 895 (2008).
- [48] Louis, P. J. Y. and Tsubota, M. Bosons in disordered optical potentials. *J. Low Temp. Phys* **148**, 351 (2007).
- [49] Bissbort, U. private communication (2007).
- [50] Hofstедder, W. private communication (2007).
- [51] Schneider, U., Hackermuller, L., Will, S., Best, T., Bloch, I., Costi, T. A., Helmes, R. W., Rasch, D. and Rosch, A. Metallic and insulating phases of repulsively interacting fermions in a 3D optical lattice. *Science* **322**, 1520 (2008).
- [52] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
- [53] Somaroo, S., Tseng, C. H., Havel, T. F., Laflamme, R. and Cory, D. G. Quantum simulations on a quantum computer. *Phys. Rev. Lett.* **82**, 5381–5384 (1999).
- [54] Leibfried, D., DeMarco, B., Meyer, V., Rowe, M., Ben-Kish, A., Britton, J., Itano, W. M., Jelenković, B., Langer, C., Rosenband, T. and Wineland, D. J. Trapped-ion quantum simulator: Experimental application to nonlinear interferometers. *Phys. Rev. Lett.* **89**, 247901 (2002).
- [55] Gillen, J. I., Bakr, W. S., Peng, A., Unterwaditzer, P., Fölling, S. and Greiner, M. Hybrid 2D surface trap for quantum simulation (2008). [arXiv:0812.3630](https://arxiv.org/abs/0812.3630).
- [56] Thywissen, J. private communication (2008).
- [57] Nelson, K. D., Li, X. and Weiss, D. S. Imaging single atoms in a three-dimensional array. *Nature Phys.* **3**, 556 (2007).
- [58] Lewandowski, H. J., Harber, D. M., Whitaker, D. L. and Cornell, E. A. Simplified system for creating a Bose-Einstein condensate. *J. Low Temp. Phys.* **132**, 309 (2003).
- [59] Metcalf, H. J. and van der Straten, P. *Laser Cooling and Trapping* (Springer, Berlin, 1999).
- [60] Foot, C. J. *Atomic Physics* (Oxford University Press, Oxford, 2005).
- [61] Saleh, B. E. A. and Teich, M. C. *Fundamentals of Photonics* (John Wiley, New York, 1991).
- [62] Siegman, A. E. *Lasers* (University Science Books, Mill Valley, 1986).
- [63] Horowitz, P. and Hill, W. *The Art of Electronics* (Cambridge University Press, Cambridge, 1989), 2nd edn.

- [64] Hobbs, P. C. D. *Building Electro-Optical Systems* (John Wiley, New York, 2000).
- [65] O’Hanlon, J. F. *A User’s Guide to Vacuum Technology* (John Wiley, New York, 1989), 2nd edn.
- [66] Steck, D. A. Rubidium 87 D line data (2008). <http://steck.us/alkalidata>.
- [67] Raab, E. L., Prentiss, M., Cable, A., Chu, S. and Pritchard, D. E. Trapping of neutral sodium atoms with radiation pressure. *Phys. Rev. Lett.* **59**, 2631–2634 (1987).
- [68] Wieman, C. E. and Hollberg, L. Using diode lasers for atomic physics. *Rev. Sci. Inst.* **62**, 1–20 (1991).
- [69] Hänsch, T. W. Repetitively pulsed tunable dye laser for high resolution spectroscopy. *Appl. Opt.* **11**, 895 (1972).
- [70] Littman, M. G. and Metcalf, H. J. Spectrally narrow pulsed dye laser without beam expander. *Appl. Opt.* **17**, 2224 (1978).
- [71] Demtroeder, W. *Laser Spectroscopy* (Springer, Berlin, 2003), 3rd edn.
- [72] Griffin, P. F. *Laser Cooling and Loading of Rb into a Large Period, Quasi-Electrostatic, Optical Lattice*. Ph.D. thesis, Durham University, Durham UK (2005).
- [73] MacAdam, K. B., Steinbach, A. and Wieman, C. A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb. *American Journal of Physics* **60**, 1098–1111 (1992).
- [74] DeMarco, B. L. *Quantum Behavior of an Atomic Fermi Gas*. Ph.D. thesis, University of Colorado, Boulder (2001).
- [75] Myatt, C. J., Newbury, N. R. and Wieman, C. E. Simplified atom trap by using direct microwave modulation of a diode laser. *Opt. Lett.* **18**, 649 (1993).
- [76] Dalibard, J. and Cohen-Tannoudji, C. Laser cooling below the doppler limit by polarization gradients: simple theoretical-models. *J. Opt. Soc. Am. B* **6**, 2023 (1989).
- [77] Bergeman, T., Erez, G. and Metcalf, H. J. Magnetostatic trapping fields for neutral atoms. *Phys. Rev. A* **35**, 1535–1546 (1987).
- [78] Ketterle, W. and van Druten, N. Evaporative cooling of atoms. In Bederson, B. and Walther, H. (eds.) *Advances in Atomic, Molecular, and Optical Physics*, vol. 37, 181 (1996).
- [79] White, M., Gao, H., Pasienski, M. and DeMarco, B. Bose-Einstein condensates in RF-dressed adiabatic potentials. *Phys. Rev. A* **74**, 023616 (2006).
- [80] Colombe, Y., Knyazchyan, E., Morizot, O., Mercier, B., Lorent, V. and Perrin, H. Ultracold atoms confined in rf-induced two-dimensional trapping potentials. *Europhys. Lett.* **67**, 593–599 (2004).
- [81] Castin, Y. and Dum, R. Bose-Einstein condensates in time dependent traps. *Phys. Rev. Lett.* **77**, 5315–5319 (1996).

- [82] Heinzen, D. private communication (2007).
- [83] Hung, C.-L., Zhang, X., Gemelke, N. and Chin, C. Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps. *Phys. Rev. A* **78**, 011604 (2008).
- [84] Petrich, W., Anderson, M. H., Ensher, J. R. and Cornell, E. A. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.* **74**, 3352–3355 (1995).
- [85] Naik, D. S. and Raman, C. Optically plugged quadrupole trap for Bose-Einstein condensates. *Phys. Rev. A* **71**, 033617 (2005).
- [86] Bloch, I. private communication (2007).
- [87] Stöferle, T. *Exploring Atomic Quantum Gases in Optical Lattices*. Ph.D. thesis, Swiss Federal Institute of Technology, Zurich (2005).
- [88] Wall, L., Christiansen, T. and Orwant, J. *Programming Perl* (O'Reilly, Sebastopol, 2000), 3rd edn.
- [89] Mackall, M. <http://www.selenic.com/mercurial/>.
- [90] Ott, H. W. *Noise Reduction Techniques in Electronic Systems* (John Wiley, New York, 1988), 2nd edn.
- [91] Menie, G. <http://www.menie.org/georges/embedded/tftpnaive/>.
- [92] Gould, P. L., Ruff, G. A. and Pritchard, D. E. Diffraction of atoms by light: The near-resonant Kapitza-Dirac effect. *Phys. Rev. Lett.* **56**, 827–830 (1986).
- [93] Martin, P. J., Oldaker, B. G., Miklich, A. H. and Pritchard, D. E. Bragg scattering of atoms from a standing light wave. *Phys. Rev. Lett.* **60**, 515–518 (1988).
- [94] Adams, C. S., Sigel, M. and Mlynek, J. Atom optics. *Phys. Rep.* **240**, 143 (1994).
- [95] Hemmerich, A. and Hänsch, T. W. Two-dimensional atomic crystal bound by light. *Phys. Rev. Lett.* **70**, 410–413 (1993).
- [96] Grynberg, G., Lounis, B., Verkerk, P., Courtois, J.-Y. and Salomon, C. Quantized motion of cold cesium atoms in two- and three-dimensional optical potentials. *Phys. Rev. Lett.* **70**, 2249–2252 (1993).
- [97] Kerman, A. J., Vuletić, V., Chin, C. and Chu, S. Beyond optical molasses: 3D Raman sideband cooling of atomic cesium to high phase-space density. *Phys. Rev. Lett.* **84**, 439–442 (2000).
- [98] Anderson, B. and Kasevich, M. Macroscopic quantum interference from atomic tunnel arrays. *Science* **282**, 1686 (1998).
- [99] Burger, S., Cataliotti, F. S., Fort, C., Minardi, F., Inguscio, M., Chiofalo, M. L. and Tosi, M. P. Superfluid and dissipative dynamics of a Bose-Einstein condensate in a periodic optical potential. *Phys. Rev. Lett.* **86**, 4447–4450 (2001).

- [100] Morsch, O., Müller, J. H., Cristiani, M., Ciampini, D. and Arimondo, E. Bloch oscillations and mean-field effects of Bose-Einstein condensates in 1D optical lattices. *Phys. Rev. Lett.* **87**, 140402 (2001).
- [101] Mandel, O., Greiner, M., Widera, A., Rom, T., Hänsch, T. W. and Bloch, I. Controlled collisions for multi-particle entanglement of optically trapped atoms. *Nature* **425**, 937 (2003).
- [102] Petsas, K. I., Coates, A. B. and Grynberg, G. Crystallography of optical lattices. *Phys. Rev. A* **50**, 5173–5189 (1994).
- [103] Taylor, P. L. and Heinonen, O. *A Quantum Approach to Condensed Matter Physics* (Cambridge University Press, Cambridge, 2002).
- [104] Ashcroft, N. W. and Mermin, D. N. *Solid State Physics* (Harcourt, Fort Worth, 1976).
- [105] Denschlag, J. H., Simsarian, J. E., Häffner, H., McKenzie, C., Browaeys, A., Cho, D., Helmerson, K., Rolston, S. L. and Phillips, W. D. A Bose-Einstein condensate in an optical lattice. *J. Phys. B: At. Mol. Opt. Phys.* **35**, 3095 (2002).
- [106] Kapitza, L. and Dirac, P. A. M. The reflection of electrons from standing light waves. *Proc. Cambridge Phil. Soc.* **29**, 297 (1933).
- [107] Dalitz, R. H. (ed.) *The Collected Works of P. A. M. Dirac: Volume 1: 1924-1948* (Cambridge University Press, Cambridge, 1995).
- [108] Zwerger, W. Mott-Hubbard transition of cold atoms in optical lattices. *J. Opt. B: Quantum Semiclass. Opt.* **5**, S9 (2003).
- [109] Freericks, J. K. and Monien, H. Strong-coupling expansions for the pure and disordered Bose-Hubbard model. *Phys. Rev. B* **53**, 2691–2700 (1996).
- [110] Sheshadri, K., Krishnamurthy, H. R., Pandit, R. and Ramakrishnan, T. V. Superfluid and insulating phases in an interacting-boson model: Mean-field theory and the RPA. *Europhys. Lett.* **22**, 257 (1993).
- [111] Spielman, I. B., Phillips, W. D. and Porto, J. V. Condensate fraction in a 2D Bose gas measured across the Mott-insulator transition. *Phys. Rev. Lett.* **100**, 120402 (2008).
- [112] Xu, K., Liu, Y., Miller, D. E., Chin, J. K., Setiawan, W. and Ketterle, W. Observation of strong quantum depletion in a gaseous Bose-Einstein condensate. *Phys. Rev. Lett.* **96**, 180405 (2006).
- [113] Penrose, O. and Onsager, L. Bose-Einstein condensation and liquid helium. *Phys. Rev.* **104**, 576–584 (1956).
- [114] Leggett, A. J. *Quantum Liquids* (Oxford University Press, Oxford, 2006).
- [115] Carruthers, P. and Nieto, M. Phase and angle variables in quantum mechanics. *Rev. Mod. Phys.* **40**, 411–440 (1968).
- [116] Shankar, R. *Principles of Quantum Mechanics* (Plenum, New York, 1994), 2nd edn.

- [117] Fölling, S., Gerbier, F., Widera, A., Mandel, O., Gericke, T. and Bloch, I. Spatial quantum noise interferometry in expanding ultracold atom clouds. *Nature* **434**, 481 (2005).
- [118] Greiner, M., Regal, C. A., Stewart, J. T. and Jin, D. S. Probing pair-correlated fermionic atoms through correlations in atom shot noise. *Phys. Rev. Lett.* **94**, 110401 (2005).
- [119] Fölling, S. *Probing Strongly Correlated States of Ultracold Atoms in Optical Lattices*. Ph.D. thesis, University of Mainz, Mainz (2008).
- [120] Maguire, L. P., Szilagy, S. and Scholten, R. E. High performance laser shutter using a hard disk drive voice-coil actuator. *Rev. Sci. Inst.* **75**, 3077–3079 (2004).
- [121] Scholten, R. E. Enhanced laser shutter using a hard disk drive rotary voice-coil actuator. *Rev. Sci. Inst.* **78**, 026101 (2007).
- [122] Morsch, O. and Oberthaler, M. Dynamics of Bose-Einstein condensates in optical lattices. *Rev. Mod. Phys.* **78**, 179 (2006).
- [123] Gericke, T., Gerbier, F., Widera, A., Fölling, S., Mandel, O. and Bloch, I. Adiabatic loading of a Bose-Einstein condensate in a 3D optical lattice. *J. Mod. Opt.* **54**, 735 (2007).
- [124] Campbell, G. K. *⁸⁷Rubidium Bose-Einstein Condensates in Optical Lattices*. Ph.D. thesis, MIT, Cambridge, MA (2006).
- [125] McKay, D., White, M. and DeMarco, B. Lattice thermodynamics for ultra-cold atoms (2009). [arxiv:0902.1017](https://arxiv.org/abs/0902.1017).
- [126] Goodman, J. W. *Speckle Phenomena in Optics* (Roberts and Company, Englewood, 2007).
- [127] Fort, C., Fallani, L., Guarrera, V., Lye, J. E., Modugno, M., Wiersma, D. S. and Inguscio, M. Effect of optical disorder and single defects on the expansion of a Bose-Einstein condensate in a one-dimensional waveguide. *Phys. Rev. Lett.* **95**, 170410 (2005).
- [128] Zhou, S. Q. and Ceperley, D. M. Construction of localized wave functions in a disordered lattice (2009). In preparation.
- [129] Trivedi, N. Quantum phase transition in disordered systems: What are the issues? In *Proceedings of the 20th International Workshop on Condensed Matter Theories*, vol. 12, 141 (1997).
- [130] Pasienski, M., White, M., McKay, D. and DeMarco, B. Transport in the disordered Bose-Hubbard model (2009). In preparation.
- [131] Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467 (1982).
- [132] Grimm, R., Weidemüller, M. and Ovchinnikov, Y. B. Optical dipole traps for neutral atoms. *Adv. At. Mol. Opt. Phys.* **42**, 95 (2000).

- [133] Jané, E., Vidal, G., Dur, W., Zoller, P. and Cirac, J. I. Simulation of quantum dynamics with quantum optical systems. *Quant. Inf. and Comp.* **3**, 15 (2003).
- [134] Nielsen, M. A. and Chuang, I. L. *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
- [135] Masanes, L., Vidal, G. and Latorre, J. I. Time-optimal Hamiltonian simulation and gate synthesis using homogeneous local unitaries. *Quant. Inf. and Comp.* **2**, 285 (2002).
- [136] Santos, L., Baranov, M. A., Cirac, J. I., Everts, H.-U., Fehrmann, H. and Lewenstein, M. Atomic quantum gases in Kagomé lattices. *Phys. Rev. Lett.* **93**, 030601 (2004).
- [137] Rabl, P., Daley, A. J., Fedichev, P. O., Cirac, J. I. and Zoller, P. Defect-suppressed atomic crystals in an optical lattice. *Phys. Rev. Lett.* **91**, 110403 (2003).
- [138] DeMarco, B., Lannert, C., Vishveshwara, S. and Wei, T.-C. Structure and stability of Mott-insulator shells of bosons trapped in an optical lattice. *Phys. Rev. A* **71**, 063601 (2005).
- [139] Jaksch, D., Briegel, H.-J., Cirac, J. I., Gardiner, C. W. and Zoller, P. Entanglement of atoms via cold controlled collisions. *Phys. Rev. Lett.* **82**, 1975–1978 (1999).
- [140] Mandel, O., Greiner, M., Widera, A., Rom, T., Hänsch, T. W. and Bloch, I. Coherent transport of neutral atoms in spin-dependent optical lattice potentials. *Phys. Rev. Lett.* **91**, 010407 (2003).
- [141] Kozuma, M., Deng, L., Hagley, E. W., Wen, J., Lutwak, R., Helmerson, K., Rolston, S. L. and Phillips, W. D. Coherent splitting of Bose-Einstein condensed atoms with optically induced Bragg diffraction. *Phys. Rev. Lett.* **82**, 871–875 (1999).
- [142] Birkel, G., Gatzke, M., Deutsch, I. H., Rolston, S. L. and Phillips, W. D. Bragg scattering from atoms in optical lattices. *Phys. Rev. Lett.* **75**, 2823–2826 (1995).

Vita

Matthew R. White

707 W. Vermont Ave.
Urbana, IL 61801
☎ 217.621.3806
✉ mwhite@illinois.edu

Education

- 2003–2009 **Ph.D., Physics (expected)**, *University of Illinois at Urbana-Champaign*.
thesis *Ultracold Atoms in a Disordered Optical Lattice*
advisor Brian DeMarco
- 1999–2003 **B.S., Physics**, *University of California, Santa Barbara*.

Experience

- 2003–2009 **Graduate Research Assistant**, *Dept. of Physics, UIUC*.
Senior graduate student in first cold-atom research group at Illinois. The group successfully and rapidly built an apparatus for studying Bose-Einstein condensates in optical lattices.
- 2001–2003 **Undergraduate Research Assistant**, *Dept. of Chemical Engineering, UCSB*.
Participated in the development of a new type of dye-sensitized photovoltaic device; successfully developed methods for electrochemical fabrication of the device.
- 2000 **Intern**, *IBM Storage Systems Division, San Jose, CA*.
Solely responsible for a series of measurements using thermal anemometry to study air temperature and flow within hard disk drives.

Publications

- M. Pasienski, M. White, D. McKay, and B. DeMarco. Transport in the disordered Bose-Hubbard model. In preparation, 2009.
- D. McKay, M. White, and B. DeMarco. Lattice thermodynamics using ultra-cold atoms. 2009.
- M. White, M. Pasienski, D. McKay, S. Zhou, D. Ceperley, and B. DeMarco. Strongly interacting bosons in a disordered optical lattice. *Phys. Rev. Lett.*, 102:055301, 2009.
- D. McKay, M. White, M. Pasienski, and B. DeMarco. Phase-slip induced dissipation in an atomic Bose-Hubbard system. *Nature*, 453:76, 2008.
- M. White, H. Gao, M. Pasienski, and B. DeMarco. Bose-Einstein condensates in RF-dressed adiabatic potentials. *Phys. Rev. A*, 74:023616, 2006.
- J. Tang, M. White, G.D. Stucky, and E.W. McFarland. Electrochemical fabrication of large-area Au/TiO₂ junctions. *Electrochem. Commun.*, 5:497, 2003.