

QUANTUM SIMULATION IN STRONGLY CORRELATED OPTICAL LATTICES

BY

DAVID C. MCKAY

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, 2012

Urbana, Illinois

Doctoral Committee:

Professor Paul Kwiat, Chair
Associate Professor Brian DeMarco, Co-Director of Research
Associate Professor Joseph Thywissen,
University of Toronto, Co-Director of Research
Professor David Ceperley
Professor Paul Selvin

Abstract

An outstanding problem in physics is how to understand strongly interacting quantum many-body systems such as the quark-gluon plasma, neutron stars, superfluid ^4He , and the high-temperature superconducting cuprates. The physics approach to this problem is to reduce these complex systems to minimal models that are believed to retain relevant phenomenology. For example, the Hubbard model — the focus of this thesis — describes quantum particles tunneling between sites of a lattice with on-site interactions. The Hubbard model is conjectured to describe the low-energy charge and spin properties of high-temperature superconducting cuprates. Thus far, there are no analytic solutions to the Hubbard model, and numerical calculations are difficult and even impossible in some regimes (e.g., the Fermi-Hubbard model away from half-filling). Therefore, whether the Hubbard model is a minimal model for the cuprates remains unresolved. In the face of these difficulties, a new approach has emerged — quantum simulation. The premise of quantum simulation is to perform experiments on a quantum system that is well-described by the model we are trying to study, has tunable parameters, and is easily probed. Ultracold atoms trapped in optical lattices are an ideal candidate for quantum simulation of the Hubbard models. This thesis describes work on two such systems: a ^{87}Rb (boson) optical lattice experiment in the group of Brian DeMarco at the University of Illinois to simulate Bose-Hubbard physics, and a ^{40}K (fermion) optical lattice experiment in the group of Joseph Thywissen at the University of Toronto to simulate Fermi-Hubbard physics.

My work on the ^{87}Rb apparatus focuses on three main topics: simulating the Bose-Hubbard (BH) model out of equilibrium, developing thermometry probes, and developing impurity probes using a 3D spin-dependent lattice. Theoretical techniques (e.g., QMC) are adept at describing the equilibrium properties of the BH model, but the dynamics are unknown — simulation is able to bridge this gap. We perform two experiments to simulate the BH model out of equilibrium. In the first experiment, published in Ref. [1], we measure the decay rate of the center-of-mass velocity for a Bose-Einstein condensate trapped in a cubic lattice. We explore this dissipation for different Bose-Hubbard parameters (corresponding to different lattice depths) and temperatures. We observe a decay rate that asymptotes to a finite value at zero temperature, which we interpret as evidence of intrinsic decay due to quantum tunneling of phase slips. The decay rate exponentially increases with temperature, which is consistent with a cross-over from quantum tunneling to thermal activation. While phase slips are a well-known dissipation mechanism in superconductors, numerous effects prevent unambiguous detection of quantum phase slips. Therefore, our measurement is among the strongest evidence for quantum tunneling of phase slips. In a second experiment, published in Ref. [2] with theory collaborators at Cornell University, we investigate condensate fraction evolution during fast (i.e., millisecond) ramps of the lattice potential depth. These ramps simulate the BH model with time-dependent parameters. We determine that interactions lead to significant condensate fraction redistribution during these ramps, in agreement with mean-field calculations. This result clarifies adiabatic timescales for the lattice gas and strongly constrains bandmapping as an equilibrium probe.

Another part of this thesis work involves developing thermometry techniques for the lattice gas. These techniques are important because the ability to measure temperature is required for quantum simulation and to evaluate in-lattice cooling schemes. In work published in Ref. [3], we explore measuring temperature by directly fitting the quasimomentum distribution of a thermal lattice gas. We attempt to obtain quasimomentum distributions by bandmapping, a process in which the lattice depth is reduced slowly compared

to the bandgap but fast with respect to all other timescales. We find that these temperature measurements fail when the thermal energy is comparable to the bandwidth of the lattice. This failure results from two main causes. First, the quasimomentum distribution is an insensitive probe at high temperatures because the band is occupied (i.e., additional thermal energy cannot be accommodated in the kinetic energy degrees of freedom). Second, the bandmapping process does not produce accurate quasimomentum distributions because of smoothing at the Brillouin zone edge. We determine that measuring temperature using the in-situ width overcomes these issues. The in-situ width does not asymptote to a finite value as temperature increases, and the in-situ width can be measured directly without using a mapping procedure. In a second experiment, we investigate using condensate fraction (obtained from the time-of-flight momentum distribution) as an indirect means to measure temperature in the superfluid regime of the BH model. Since no standard fitting procedure exists for the lattice time-of-flight distributions, we define and test a procedure as part of this work. We measure condensate fraction for a range of lattice depths varying from deep in the superfluid regime to lattice depths proximate to the Mott-insulator transition. We also vary the entropy per particle, which is measured in the harmonic trap before adiabatically loading into the lattice. As expected, the condensate fraction increases as entropy decreases, and the condensate fraction decreases at high lattice depths (due to quantum depletion). We compare our experimental results to condensate fraction predicted by the non-interacting, Hartree-Fock-Bogoliubov-Popov, and site-decoupled-mean-field theories. Theory and experiment disagree, which motivates several future extensions to this work, including calculating condensate fraction (and testing our fit procedure) using quantum Monte Carlo numerics, and experimentally and theoretically investigating the dynamics of the lattice load process (for the finite-temperature strongly correlated regime).

Finally, we develop impurity probes for the Bose-Hubbard model by employing a spin-dependent lattice. A primary accomplishment of this thesis work was to develop the first 3D spin-dependent lattice in the strongly correlated regime (published in Ref. [4]). The spin-dependent lattice depth is proportional to $|g_F m_F|$, enabling the creation of mixtures of atoms trapped in the lattice ($m_F \neq 0$) co-trapped with atoms that do not experience the lattice ($m_F = 0$). We use the non-lattice atoms as an impurity probe. We investigate using the impurity to probe the lattice temperature, and we determine that thermalization between the impurity and lattice gas is suppressed for larger lattice depths. Using a comparison to a Fermi's golden rule calculation of the collisional energy exchange rate, we determine that this effect is consistent with suppression of energy-exchanging collisions by a mismatch between the impurity and lattice gas dispersion. While this result invalidates the concept of an impurity thermometer, it paves the way for a unique cooling scheme that relies on inter-species thermal isolation. We also explore impurity transport through the lattice gas. In other preliminary measurements, we also identify the decay rate of the center-of-mass motion as a prospective impurity probe.

A separate aspect of this thesis work is the design and construction of a new ^{40}K apparatus for single-site imaging of atoms to simulate the 2D Fermi-Hubbard model. The main component of this apparatus is high-resolution fluorescence imaging on the $4S \rightarrow 5P$ transition of K at 404.5nm. Fluorescence imaging using this transition has two advantages over imaging on the standard D2 transition at 767nm: a smaller wavelength and therefore higher resolution, and a lower Doppler temperature limit which enables longer imaging times. To validate this approach, we demonstrate the first ^{40}K magneto-optical trap (MOT) using the 404.5nm transition.

Acknowledgements

Officially, this PhD started in August 2006 — unofficially, it started five months prior when I received a phone call (in the Toronto lab of all places). Professor Brian DeMarco was on the phone inviting me to join his young group and work on “quantum simulation.” I accepted and then later joined the group in fall 2006. As soon as I arrived, I was immersed in the experiment and Brian set out to mold me into an experimental physicist. After two intense and productive years, this thesis took on a decidedly unique transformation. In fall 2008, I joined Joseph Thywissen’s group at the University of Toronto, while still remaining with Brian’s group. This was a remarkable opportunity to work on complementary experiments in the emerging field of ultracold atom quantum simulators. Balancing the demands of two groups has been a challenge at times, but it has also been a tremendously rewarding experience. This arrangement would not have been possible without the incredible support of my co-advisors Brian and Joseph. I would like to acknowledge them both and thank them for their mentorship over the years.

One benefit of working in two groups was the ability to interact with an array of various colleagues, collaborators, co-students and friends. I would like to thank them all for their support, knowledge and friendship. In particular, I would like to acknowledge Matt White, the senior (and extremely talented) graduate student in Illinois when I first joined the group. Matt immediately involved me into the experiment and quickly become my mentor (and still is to this day). I would also like to acknowledge Dylan Jervis, my co-graduate student in Toronto. I have many fond memories of our time spent building the experiment (and occasionally talking about hockey).

I would like to thank my family for all their love and support throughout the years. Of course, this thesis would not have been possible without my wife, Jessie. I will always cherish her love, understanding, and companionship (and her oatmeal chocolate chip cookies).

Finally, I would like to acknowledge the generous financial support of the Roy J. Carver fellowship, NSERC, NSF, ARO, ONR, AFOSR, and the DARPA OLE program.

So it goes,

— Dave

July 2012

Table of Contents

List of Figures	viii
List of Abbreviations	xii
List of Symbols	xiii
Chapter 1 Introduction	1
Chapter 2 Strongly Correlated Atoms in Optical Lattices	14
2.1 Ultracold Atoms	14
2.2 Ultracold Atoms in Harmonic Traps	22
2.3 Ultracold Atoms in Optical Lattices	29
2.4 Hubbard Model	40
Chapter 3 Experimental Toolkit — Tools	49
3.1 ^{87}Rb Bose-Hubbard Apparatus	49
3.2 ^{40}K Fermi-Hubbard Apparatus	57
Chapter 4 Experimental Toolkit — Probes and Techniques	86
4.1 Probes	86
4.2 Techniques	108
Chapter 5 Theoretical Toolkit	115
5.1 Non-Interacting	116
5.2 Atomic Limit	120
5.3 Mean-Field Theories	121
5.4 Exact - Small Systems	126
5.5 Comparison	126
5.6 Beyond our Toolkit	128
Chapter 6 Simulating the Bose-Hubbard Model Out of Equilibrium	129
6.1 Phase Slip Dissipation	129
6.2 Rephasing After Lattice Ramps	137
Chapter 7 Developing Direct Thermometry Probes	145
7.1 Bandmapping	146
7.2 In-situ RMS Width	150
7.3 Lattice Condensate Fraction	152
Chapter 8 Developing Impurity Probes	167
8.1 Impurity Thermometer	168
8.2 Impurity Transport Probe	183
Chapter 9 Conclusions and Outlook	186

Appendix A	Properties of ^{87}Rb and ^{40}K	189
A.1	States	189
A.2	Collisional Properties	193
Appendix B	Atoms Interacting with Electromagnetic Fields	198
B.1	Interaction Hamiltonian	198
B.2	DC Magnetic Fields	199
B.3	Oscillating Fields	200
B.4	Rabi Oscillations and Landau-Zener Transitions	204
B.5	Dipole Potential (AC Stark Shift)	207
B.6	Optical Bloch Equations	209
B.7	Scattering	213
B.8	Mechanical Effects of Scattering	217
Appendix C	Optical Lattice and Optical-Dipole Trap	222
C.1	Scalar Optical Lattice	222
C.2	The Lin- θ -Lin Retro-Reflected Lattice Potential	226
C.3	Gaussian Beams	230
C.4	Heating	233
C.5	Calibrating the Lattice with Diffraction	237
C.6	Alignment and Setup of the Spin-Dependent Lattice	238
C.7	X-Dipole Trap Alignment and Heating	239
Appendix D	Characterizing Trap Frequencies for the Lattice System	242
D.1	Experimentally Measuring the Trap Frequency	242
D.2	Parameterizing the Fixed Trap	245
D.3	Parameterizing the Lattice Confinement	246
D.4	Trap Plus Lattice Confinement	248
Appendix E	K $4S \rightarrow 5P$ Laser System and Magnetic-Optical Trap	250
E.1	$4S \rightarrow 5P$ Calculation	251
E.2	Laser System and Setup	255
E.3	$4S \rightarrow 5P$ MOT and Laser Cooling	264
Appendix F	Imaging	271
F.1	Number Calibration of Absorption Images	271
F.2	Partial Repumping	272
F.3	Resolution	272
Appendix G	Supporting Calculations	278
G.1	K $4S$ - $5P$ Magic Wavelength for Dipole Trapping	278
G.2	Quadrupole to Dipole Transfer	281
G.3	Expansion in Condensate Mean-Field	285
G.4	Simulated Bose-Hubbard TOF Distributions	291
G.5	Deriving the Probe Integrated Quasimomentum Distribution	292
Appendix H	Circuits	294
H.1	Laser Lock	294
H.2	Experiment Monitor and Interlock	299
H.3	Mixer	299

Appendix I	Code	307
I.1	SDMFT Code	307
I.2	HFBP	311
I.3	Small System Exact Hubbard Solver	316
I.4	Wavefunction Solvers	324
I.5	Band Structure	338
I.6	Atomic Physics	341
Appendix J	Technical Drawings	356
Appendix K	References	364

List of Figures

1.1	Increasing model complexity versus the full set of equations	2
1.2	The quantum simulator feedback cycle	4
1.3	Overview of steps to prepare ultracold atoms in an optical lattice	6
2.1	Magnetic quadrupole field	16
2.2	Retro-reflected lattice potential	18
2.3	Schematic of atomic interactions	20
2.4	Thermal distribution and condensate fraction for the semi-ideal model	27
2.5	Lattice band structure	30
2.6	Wavefunctions in the lattice	31
2.7	Energy spectrum of the combined lattice-harmonic potential	35
2.8	Hubbard parameters (t, U) versus lattice depth	37
2.9	Limitations of the tight-binding approximation	37
2.10	Effective mean-field potential from the lattice atoms	39
2.11	Bose-Hubbard phase diagram	42
2.12	Fermi-Hubbard phase diagram	46
3.1	Schematic of the ^{87}Rb apparatus	50
3.2	Preparing a state mixture using Landau-Zener sweeps	52
3.3	Typical lattice timing sequence and data for the ^{87}Rb apparatus	53
3.4	Lin-Lin lattice ($\lambda = 810\text{nm}$) versus the lin- \perp -lin lattice ($\lambda = 790\text{nm}$)	54
3.5	Observation of the SF-MI transition in a spin-dependent lattice	56
3.6	Schematic of the ^{40}K apparatus	58
3.7	^{40}K vacuum system	60
3.8	^{40}K all-metal vacuum bake	62
3.9	^{40}K support structure	64
3.10	^{40}K atomic sources	66
3.11	Simplified ^{40}K state diagram for laser locking	68
3.12	Simplified ^{87}Rb state diagram for laser locking	69
3.13	^{87}Rb beat-note laser lock setup	70
3.14	Schematic of the $^{40}\text{K}/^{87}\text{Rb}$ laser system	71
3.15	$^{40}\text{K}/^{87}\text{Rb}$ MOT setup	72
3.16	$^{40}\text{K}/^{87}\text{Rb}$ MOT cell	73
3.17	^{40}K and ^{87}Rb MOT fluorescence images	74
3.18	Schematic of the magnetic transport system	75
3.19	Magnetic transport transfer efficiency	77
3.20	Schematic of the plug beam setup	79
3.21	Schematic of the dipole and lattice setup	80
3.22	^{40}K single-site imaging setup	82
3.23	Image of the MRS-5 target through the Special Optics objective	83
3.24	Schematic of the ^{40}K apparatus computer control and monitoring system	84

4.1	Schematic of fluorescence and absorption imaging	87
4.2	Maxwell-Boltzmann distributions after TOF expansion	89
4.3	Bandmapped and quasimomentum distributions	92
4.4	Condensate distributions after TOF expansion from a harmonic trap	96
4.5	Fitting TOF condensate distributions using the heuristic method	98
4.6	Testing various TOF condensate distribution fitting techniques	99
4.7	Calculated images of a 2D lattice plane with varying resolution	104
4.8	Transitions into the spin-dependent lattice from the $ 1, 0\rangle$ state	106
4.9	Adiabatic loading into the lattice	109
4.10	Schematic of common transport (collective) modes	112
5.1	Thermodynamics of the combined lattice-harmonic potential	118
5.2	HFBP theory $T = 0$ condensate fraction versus U/t	124
5.3	Site-decoupled-mean-field theory $T = 0$ condensate fraction versus U/t	125
5.4	7-site Bose-Hubbard model condensate fraction versus U/t	127
5.5	Summary of the different theoretical methods for solving the Hubbard model	127
6.1	Schematic of the phase slip process	131
6.2	Timing sequence for the phase slip experiment	132
6.3	Sample transport data from the phase slip experiment	133
6.4	Decay rate versus inverse temperature	135
6.5	Logarithm of the decay rate versus $\sqrt{t/U}$	136
6.6	Small system calculation of ramping down the lattice	139
6.7	Condensate fraction versus ramp down time for $s = 12$	141
6.8	Experimental results for τ_{-1} versus lattice depth	142
6.9	Experimental results for τ_{-1} compared to mean-field theory predictions	143
7.1	Calculated RMS quasimomentum width versus temperature	146
7.2	Bandmapping calculation of the $n = 0$ eigenstate of the combined lattice-harmonic potential	148
7.3	Temperature measured by fitting calculated bandmapped distributions	148
7.4	Temperature measured by fitting experimental bandmapped distributions	149
7.5	Calculated RMS in-situ width versus temperature	151
7.6	Experimental results using the in-situ RMS width for thermometry	152
7.7	Various thermodynamic proxies in the SF regime of the BH lattice gas	153
7.8	Applying the peak fraction fit procedure step-by-step to experimental data	158
7.9	Applying the peak fraction fit procedure to simulated data	159
7.10	Averaged experimental lattice condensate fraction data	162
7.11	Experimental peak fraction compared to theoretical condensate fraction	163
7.12	Experimental peak fraction versus ramp time into the lattice	165
8.1	In-situ separation between the $ 1, 0\rangle$ and $ 1, -1\rangle$ gas centers	170
8.2	Probing the impurity density with microwaves	171
8.3	Impurity condensate fraction versus primary lattice depth	172
8.4	Schematic of using selective heating to test thermalization	174
8.5	Overview of parametric heating (schematic and control data at $s = 4$)	176
8.6	Impurity condensate fraction versus parametric oscillation modulation depth	177
8.7	Schematic of the dephasing experiment	178
8.8	Impurity ($ 1, 0\rangle$) dephasing data for $s = 4$	179
8.9	Heating rate of the impurity after dephasing as a function of the primary lattice depth	180
8.10	Calculated energy transfer rate	182
8.11	Impurity ($ 1, 0\rangle$) and primary ($ 1, -1\rangle$) oscillations	184
8.12	Impurity damping rate versus the primary lattice depth	185
A.1	^{40}K state diagram	190

A.2	^{87}Rb state diagram	191
A.3	^{40}K and ^{87}Rb hyperfine state energies versus B	194
B.1	Multiple level decay and driving configurations for the OBEs	211
B.2	Schematic of a MOT (along one direction) for a simplified atom	219
C.1	Scalar lattice configurations	223
C.2	Lin- θ -Lin lattice setup	226
C.3	Schematic of a Gaussian beam propagating along z	230
C.4	Optimal wavelength for a ^{87}Rb lattice with spacing d	235
C.5	Calibrating lattice depth using diffraction	237
C.6	Optimizing the waveplate angle and wavelength for the spin-dependent lattice	238
C.7	Schematic of the x-dipole trap setup	240
D.1	Example of a single-frequency trap oscillation measurement in the ^{87}Rb apparatus	243
D.2	Example of a two-frequency trap oscillation measurement in the ^{87}Rb apparatus	244
D.3	Parameterizing the hybrid-dipole trap	245
D.4	Calculated principle trap frequencies versus lattice depth	249
E.1	State scheme for the $4S \rightarrow 5P$ cooling and trapping calculation	252
E.2	Force curves versus position and velocity for the $4S \rightarrow 5P$ and $4S \rightarrow 4P$ transitions.	253
E.3	Temperature versus detuning for cooling on the $4S \rightarrow 5P$ transition.	253
E.4	Ionization energy for each K state	254
E.5	Schematic of the 405nm laser system for laser cooling and trapping	256
E.6	$4S \rightarrow 5P$ saturation spectroscopy	257
E.7	Laser setup for probing the $4S \rightarrow 5P$ transition	258
E.8	Results of probing the $^{40}\text{K } 4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F'$ transition	259
E.9	Custom made 405nm scanning cavity	260
E.10	High power 405nm diode spectra	261
E.11	High power diode mount and enclosure	262
E.12	$4S \rightarrow 5P$ state diagram for the MOT	264
E.13	$4S \rightarrow 5P$ MOT setup	265
E.14	Characterizing the $4S \rightarrow 5P$ MOT	266
E.15	$4S \rightarrow 5P$ MOT atom number versus hold time	267
E.16	$4S \rightarrow 5P$ fluorescence versus intensity	268
E.17	$4S \rightarrow 5P$ MOT temperature and density	269
E.18	$4S \rightarrow 5P$ molasses	270
F.1	Number calibration for the ^{87}Rb apparatus	271
F.2	Schematic of partial repumping	273
F.3	Sample repump calibration curve for the ^{87}Rb apparatus	273
F.4	Image of two point sources as a function of the distance between the sources	275
F.5	PSF at different focus depths	277
G.1	Possible transitions from/to the $5P_{3/2}$ and $4S_{1/2}$ states in potassium	279
G.2	$4S_{1/2}$ and $5P_{3/2}$ energy shift versus wavelength	280
G.3	Differential stark shift in all cascade states at the magic wavelength	281
G.4	Stages of the quadrupole-dipole transfer	282
G.5	Final temperature after transfer versus the initial quadrupole trap temperature	283
G.6	Quadrupole to dipole transfer efficiency versus initial quadrupole trap temperature	284
G.7	Numerically calculated expansion of the condensate	287
G.8	Sample non-condensate states in-trap and after 10ms expansion	288
G.9	Thermal distributions after numerically solving the mean-field expansion	289
G.10	Cross-section of simulated lattice condensate fraction data at $s = 6$	290

G.11 Noise samples used to create simulated data	290
H.1 Laser lock panel photographs	295
H.2 Laser lock circuit schematic pg. 1	296
H.3 Laser lock circuit schematic pg. 2	297
H.4 Laser lock circuit schematic pg. 3	298
H.5 Interlock/MOT photodiode panel photographs	300
H.6 MOT Photodiode monitor circuit schematic	301
H.7 Interlock circuit schematic pg. 1	302
H.8 Interlock circuit schematic pg. 2	303
H.9 Interlock circuit schematic pg. 3	304
H.10 Relay board circuit schematic for the interlock	305
H.11 Mixer circuit schematic	306
J.1 Technical drawing: ^{40}K top flange	356
J.2 Technical drawing: ^{40}K bottom flange	357
J.3 Technical drawing: ^{40}K science chamber	358
J.4 Technical drawing: ^{40}K MOT cell	359
J.5 ^{40}K science chamber window coatings specifications	360
J.6 ^{40}K imaging window coating specifications	361
J.7 Dimensions of the Special Optics objective	362
J.8 MTF of the Special Optics objective	363

List of Abbreviations

AFM	- Antiferromagnet
AOM	- Acousto-optic modulator
BEC	- Bose-Einstein condensate
BH	- Bose-Hubbard
BI	- Band insulator
COM	- Center of mass
D	- Dimension
DOF	- Degree of freedom
DFG	- Degenerate Fermi gas
FH	- Fermi-Hubbard
HFBP	- Hartree-Fock-Bogoliubov-Popov
MFT	- Mean-field theory
MI	- Mott insulator
MOT	- Magneto-optical trap
NA	- Numerical aperture
RF	- Radiofrequency
RMS	- Root-mean square
SC	- Superconductor
SDMFT	- Site-decoupled mean-field theory
SF	- Superfluid
TF	- Thomas-Fermi
TOF	- Time of flight

List of Symbols

X	- Quantum operators are denoted as bold
\vec{x}	- Vectors are denoted by an arrow
\hat{x}	- Unit vectors are denoted by a hat
H	- Hamiltonian
\mathbf{a}_i (\mathbf{a}_i^\dagger)	- Annihilation (creation) operator of mode i
Ψ	- Field operator
\mathbf{t}	- Time
m	- Mass
p	- Momentum
q	- Quasimomentum
\vec{x}, x, y, z	- Spatial coordinates
r	- Radial coordinate
a	- S-wave scattering length
t	- Tunneling energy (Hubbard model)
U	- Interaction energy (Hubbard model)
d	- Lattice spacing
λ	- Wavelength
\mathbf{z}	- Coordination number for the lattice
I	- Electro-magnetic field intensity
\mathbf{w}	- Gaussian beam waist (e^{-2} intensity)
s	- Lattice potential depth in E_R
E_R	- Recoil energy ($\frac{(h/\lambda)^2}{2m}$)
ω	- Harmonic potential trap frequency ($\nu = \frac{\omega}{2\pi}$)
V	- Potential energy
T	- Temperature

T_C	- Bose-Einstein condensate transition temperature
μ	- Chemical potential
N	- Number of particles
F (m_F)	- Atomic hyperfine quantum number (hyperfine projection)
B	- Magnetic field
Δ	- Frequency detuning from resonance
c	- Speed of light (2.99792458×10^8 m/s) [5]
h	- Planck's constant ($6.62606957(29) \times 10^{-34}$ Js, $\hbar = \frac{h}{2\pi}$) [5]
k_B	- Boltzmann constant ($1.3806488(13) \times 10^{-23}$ JK $^{-1}$) [5]
μ_B	- Bohr magneton ($\frac{e\hbar}{2M_e} = h \times 1.399624555(31)$ MHz/G) [5]
a_0	- Bohr radius ($52.917721092(17) \times 10^{-12}$ m) [5]

Chapter 1

Introduction

Physics endeavors to provide a mathematical, and therefore predictive framework for the fundamental phenomena of the universe. One aspect of this goal is to determine the equations that describe the behavior of a system's constituent particles. In this regard, physics has been very successful. Except at extreme scales, the ordinary matter in the universe is completely described by the quantum field theories of the standard model and general relativity. Almost all low-energy earth-bound matter (nuclei and electrons), materials which are the basis of our modern technology, can be described using the non-relativistic Schrödinger equation for the quantum wavefunction $|\Psi\rangle$,

$$i\hbar \frac{\partial |\Psi\rangle}{\partial \tau} = \mathbf{H} |\Psi\rangle. \quad (1.1)$$

Here \mathbf{H} is the Breit-Pauli Hamiltonian for N_e electrons and N_Z nuclei of charge Z [6]¹,

$$\begin{aligned} \mathbf{H} = & \sum_{i=1}^{N_e} \frac{\vec{\mathbf{p}}_i^2}{2} + \sum_j^{N_Z} \frac{m_e}{m_Z} \frac{\vec{\mathbf{P}}_j^2}{2} + \sum_{i=1}^{N_e} \sum_{j<i}^{N_Z} \frac{Z^2}{|\vec{\mathbf{R}}_{ij}|^2} - \sum_{i=1}^{N_e} \sum_{j=1}^{N_Z} \frac{Z}{|\vec{\mathbf{r}}_i - \vec{\mathbf{R}}_j|^2} + \sum_{i=1}^{N_e} \sum_{j<i}^{N_e} \frac{1}{|\vec{\mathbf{r}}_{ij}|^2} \\ & + \alpha^2 \sum_{i=1}^{N_e} \left\{ \sum_{j=1}^{N_Z} \frac{Z}{2|\vec{\mathbf{r}}_i - \vec{\mathbf{R}}_j|^3} \left((\vec{\mathbf{r}}_i - \vec{\mathbf{R}}_j) \times \vec{\mathbf{p}}_i \right) \cdot \vec{\mathbf{s}}_i - \sum_{j \neq i} \frac{1}{2|\vec{\mathbf{r}}_{ij}|^3} (\vec{\mathbf{r}}_{ij} \times \vec{\mathbf{p}}_i) \cdot (\vec{\mathbf{s}}_i + 2\vec{\mathbf{s}}_j) \right. \\ & \left. + \sum_{j<i} \left[\frac{\vec{\mathbf{s}}_i \cdot \vec{\mathbf{s}}_j}{|\vec{\mathbf{r}}_{ij}|^3} - 3 \frac{(\vec{\mathbf{s}}_i \cdot \vec{\mathbf{r}}_{ij})(\vec{\mathbf{s}}_j \cdot \vec{\mathbf{r}}_{ij})}{|\vec{\mathbf{r}}_{ij}|^5} \right] \right\} \end{aligned} \quad (1.2)$$

where $\vec{\mathbf{p}}_i$, $\vec{\mathbf{r}}_i$ and $\vec{\mathbf{s}}_i$ are the momentum, position and spin operators for electron i , $\vec{\mathbf{P}}_i$ and $\vec{\mathbf{R}}_i$ are the momentum and position operators for nuclei i , and $\alpha \approx 1/137$ is the fine structure constant. The wavefunction must be antisymmetric under electron exchange to satisfy Fermi statistics and (anti)symmetric under nuclei exchange to obey (Fermi)Bose statistics if the number of nucleons (neutrons plus protons) is (odd)even.

To a very good approximation, Eqn. 1.2 describes all aspects of Earth-bound matter, including the formation of atoms and the assembly of these atoms into molecules, gases, liquids and solids. Of course, knowing Eqn. 1.2 is only the first aspect of physics; the second is solving these equations, which is often the more difficult task. As elucidated by Philip Anderson, “The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe . . . [since] the constructionist hypothesis breaks down when confronted with the twin difficulties of scale and complexity” [7]. The

¹For clarity, this equation is presented in atomic units ($\hbar = m_e = e = 4\pi\epsilon_0 = 1$). Relativistic corrections that shift, but do not split energy levels (non-fine structure terms) have been omitted. Also, hyperfine terms associated with the nuclear magnetic momentum and finite charge radius have been omitted. In general, a fully relativistic theory (Dirac equation) is required for the electrons, but in almost all materials these relativistic effects can be expressed in terms of the included order α^2 corrections to the non-relativistic Hamiltonian. Field theoretical effects, such as the Lamb shift, are also not included.




Figure 1.1: (a) The full set of equations, Eqn. 1.2, describes the properties of electrons (black) and nuclei (blue protons, red neutrons) at all energy scales. This set of diagrams gives a schematic overview of the typical properties of a metallic nuclei, e.g., ^{87}Rb (37 protons, 50 neutrons), at different energies. For simplicity only a handful of the particles are shown and the diagrams are not to scale. At the highest energy scales (top) we have a free gas of electrons and nuclei. Lowering the energy, they form atoms (middle) and finally these atoms form a crystalline solid where the outer valence electrons become delocalized throughout the solid. (b) If we are interested in only the delocalized electron properties at the lowest energy scales described by Eqn. 1.2, then we can alternatively build up a minimal model. We start with a free gas of non-interacting electrons in a fixed volume, then add a static lattice, then add phonons (lattice vibrations) and electron interactions.

complexity of Eqn. 1.2 is evident since it contains all the necessary ingredients for tightly bound shells of electrons to form around the nuclei into atoms, for these atoms to arrange into a crystalline structure, and for the outer-shell electrons to delocalize throughout the crystal (Fig. 1.1). However, if we are only interested in the motion of these delocalized electrons, which is enough to describe a plurality of material properties, then Eqn. 1.2 is overly complicated. At these energy scales many degrees of freedom are frozen out and can be represented in the equations as static terms that can be parameterized. For example, the atomic ions (nuclei plus bound core electrons) are essentially fixed in a discrete crystalline structure, so a large number of operators in Eqn. 1.2 can be replaced by a small set of parameters specifying this arrangement. These parameters can be measured experimentally using crystallographic techniques.

Therefore, the alternative approach, shown schematically in Fig. 1.1, is to build a model that increases in complexity, as necessary, to describe a relevant aspect of the system. To add in layers of complexity we parameterize the solutions of Eqn. 1.2 and determine these parameters empirically. Applying this approach to the delocalized electrons, we start with a free electron gas (Drude-Sommerfeld model), then add a lattice, then add in additional terms such as phonons, electron-phonon interactions, electron-electron interactions (e.g., see [8]). To understand what level of complexity is required so that the model is a good representation, we have to solve this model and compare to experiments on physical systems. Once a certain

model reproduces experimental results we can use it to predict new phenomena and motivate new materials. This is illustrated by the left side of the cycle in Fig. 1.2. Examples of attainable experimental data for comparison include the excitation spectra, order parameters (broken symmetries), transport properties and mass/momentum/spin densities, and correlation functions. A full comparison to experiment involves describing these properties as a function of the microscopic model parameters and macroscopic thermodynamic state variables, temperature (T) and chemical potential (μ). A common experimental construction is the phase diagram, which identifies regions of uniform properties in the parameter space bounded by abrupt changes (phase boundaries). Example phase diagrams are illustrated in Fig. 1.2.

Although reducing Eqn. 1.2 to a model minimizes the complexity of the equations we need to solve, there is still the issue of scale. Each quantum particle may be in a number of states; the simplest example is a stationary electron which has two states, spin up (\uparrow) or spin down (\downarrow). The scaling issue arises since the spin can be in a superposition of up and down, $p_\uparrow |\uparrow\rangle + p_\downarrow |\downarrow\rangle$, where $|p_\uparrow|^2$ ($|p_\downarrow|^2$) is the probability of the spin to be up (down). However, in general p_\uparrow and p_\downarrow are complex numbers and cannot be represented as probabilities. Therefore, when the system increases to N electrons, the system can be in a superposition of all possible N -spin configurations, which requires keeping track of 2^N complex numbers. This exponential scaling makes solving general quantum models intractable on classical computers. However, there are still many models which are efficiently solvable using classical resources, although identifying these models is a non-trivial task as this depends on the level of entanglement in the system [9]. Additionally, since we are solving in order to compare to experimental measurements, we do not necessarily have to compute the complex coefficients of each state. Therefore, certain sets of properties may be well-approximated through classical numerics. For example, the equilibrium ground state density distribution of an interacting many-body Bose gas is efficiently calculated using Quantum Monte Carlo (QMC) [10], which is a probabilistic path integral sampling done on a classical computer.

Nonetheless, it would be ideal to avoid these scaling issues by calculating the model properties on a system which is itself quantum, an idea first pointed out by Feynman [13]. A certain class of quantum systems, defined by the Divincenzo criteria [14], are known as quantum computers (see [15] for a review) and are able to simulate any quantum model with local interactions [16]. For this reason, quantum computers are referred to as universal [16] or digital quantum simulators [17]. However, although quantum computers avoid fundamental quantum scaling issues, they suffer from technical scaling problems and the implementation of a quantum computer with a large number of quantum particles is an enormous technical challenge; state-of-the-art is six qubits (i.e., 2-level quantum particles) [18]. Fortunately, there exists another class of quantum systems whose intrinsic Hamiltonians are known to be well-described by our model Hamiltonian. Therefore, the properties of this system are, in effect, solutions to the model. Because this system performs the simulation continuously and is not universal, this is referred to as an analog quantum simulator² [17]. Whether a system qualifies as a quantum simulator is rather ill-defined since all systems are in some sense simulating a certain model all the time. To clarify, I propose the following ingredients for a quantum simulator:

1. A system which is well-described by a quantum model with tunable parameters,

²From this point I will refer to analog quantum simulators as quantum simulations and digital quantum simulators as quantum computers.




Figure 1.2: (Top) The quantum simulator cycle. The starting point is a physical system, e.g., a crystalline solid. Although we know the full equations of motion for the system, i.e., Eqn. 1.2, these are too complex, so a model is developed. To validate the model we compare the model properties to the experimental properties. However, issues of quantum scaling mean that many models are not solvable using classical numerics. Therefore, we design a quantum simulator; a physical system a priori known to be well-described by the model. If the model is validated by comparing the experimental measurements to the quantum simulator, the simulator can be tuned to discover novel properties, e.g., “room-temperature” superconductivity, and motivate the synthesis of new materials. If the current model is not validated, then we need to iterate and add complexity to generate a new model starting the cycle over. (Bottom) Phase diagrams are useful tools for comparing experimental data to data generated via quantum simulation. On the left we show the schematic phase diagram for a common physical system, the cuprate compounds, which notably display high-temperature superconductivity when hole-doped (motivated by [11, 12]). On the right we show schematic phase diagrams for the two-component Fermi-Hubbard (FH) model, discussed in §2.4.2, which is conjectured to describe the cuprates. However, the FH model is unsolved since the hole doped phases are unknown, as shown in the first phase diagram for T versus doping at large U/t . Since the quantum simulation parameters are adjustable, we can easily explore the parameter space, a difficult task with the physical system. The phase diagram versus t/U and $\mu = \mu_{\uparrow} = \mu_{\downarrow}$ at $T = 0$ is shown at the far right.

2. Control over the initial quantum state,
3. Probes to measure the state/properties of the system,
4. Techniques to explore non-equilibrium states.

A similar list, emphasizing issues of reliability and efficiency, has been described in [19]. Often, quantum simulators are identified by ingredient (1) with no consideration for ingredients (2)-(4), but a reliable comparison to the physical systems being simulated, see Fig. 1.2, is not possible without these additional components. In the same way, simulation on a classical computer does not just consist of the computer processor, but the software (code) and various initial conditions and methodologies for running the simulation to get useful results. Together, ingredients (1)-(4) are what we term the toolkit for quantum simulation. This toolkit is roughly divided into three main groups; tools are the elements of ingredient (1) and probes/techniques are the elements of ingredients (2)-(4). Tools are generally the physical elements of the quantum simulation (the processor) and probes/techniques are the methodologies and procedures for using these physical elements (the software).

A quantum simulator allows us to compare the properties of quantum models to experimental measurements on physical systems to discover the minimal models associated with these systems. Eventually this allows us to understand how certain parameters are linked to interesting and technically useful properties leading to the synthesis of new materials. For example, there is no proven model for the high-temperature superconducting cuprates [12], because proposed models, e.g., the Fermi-Hubbard model (§2.4.2), have not been solved. The intent is to use a quantum simulator to validate one of these models, and then perform further simulations to optimize the model parameters to maximize the superconducting transition temperature. This will in turn motivate synthesizing new materials with these parameters. The quantum simulation feedback process is illustrated in Fig. 1.2, but it leaves open two questions; how do we implement a quantum simulator and what quantum models should we simulate? The general process is to look for a well-defined and easily controlled quantum system and investigate whether its intrinsic Hamiltonian corresponds to a physically relevant, unsolvable quantum model. There are a number of candidate systems [17], such as ions, photons in cavities, superconducting circuits and quantum dots. However, in this thesis I focus on one of the most promising implementations for simulations of models related to materials: neutral ultracold atoms in optical lattices.

To prepare ultracold atoms, illustrated in Fig. 1.3, we laser cool [20] a room temperature atomic gas to $T \approx 100\mu K$ and then perform further evaporative cooling [21] so that the de Broglie wavelength is comparable to the interparticle spacing, which is the condition for quantum degeneracy. Both bosonic [22, 23] and fermionic [24] atoms have been cooled to degeneracy, so ultracold atoms can simulate both types of quantum statistics. The interactions between ultracold atoms (§2.1.2) are well-characterized, short-range, spherically symmetric, and tunable, which are important characteristics for simulation. Furthermore, DC magnetic fields and AC electric fields near resonance with optical atomic transitions can be used to shift the atomic energy levels. Shaping these external fields provides a range of trapping potentials (§2.1.1). One particular potential, the optical lattice (§C), is formed by interfering two or more laser beams to produce a periodic intensity and therefore a periodic potential³ with depth sE_R where $E_R = \frac{(\hbar\pi/d)^2}{2m}$ is the recoil

³Fields with periodic polarization can also create periodic potentials, which will comprise an important part of this work.




Figure 1.3: Schematic of preparing ultracold atoms in an optical lattice for quantum simulation. (a) Collecting atoms using laser cooling and trapping (§2.1). (b) Evaporative cooling in a trapping potential (§2.1, §2.1.1). (c) Ultracold atoms loaded into an optical lattice (§C) with spacing d and height sE_R . Atoms can tunnel between adjacent sites and atoms on the same site interact.

energy (d is the lattice spacing). Near quantum degeneracy, the wavefunctions of atoms in optical lattices are localized to the potential wells in the lowest vibrational state with order unit occupancy per well. Atoms can then tunnel between lattice sites and interact strongly when two or more occupy the same site.

The intrinsic Hamiltonian for ultracold atoms in optical lattices is, in certain limits, the single-band Hubbard model ([25, 26], §2.3.4), generically written as,

$$\mathbf{H} = - \sum_{\langle i,j \rangle, \sigma} t_{ij, \sigma} \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{j, \sigma} + \sum_{i, \sigma, \sigma'} \frac{1}{2} U_{i, \sigma, \sigma'} \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{i, \sigma'}^\dagger \mathbf{a}_{i, \sigma} \mathbf{a}_{i, \sigma'} + \sum_{i, \sigma} (\epsilon_{i, \sigma} - \mu_{i, \sigma}) \mathbf{n}_{i, \sigma} \quad (1.3)$$

where $\langle i, j \rangle$ is a sum over nearest-neighbor lattice sites, σ is the spin index, $t_{ij, \sigma}$ is the tunneling energy, $U_{i, \sigma, \sigma'}$ is the interaction energy, $\epsilon_{i, \sigma}$ is the kinetic plus potential energy offset of site i , and $\mu_{i, \sigma}$ is the chemical potential which adjusts the average number of σ atoms per site. The operator $\mathbf{a}_{i, \sigma}$ ($\mathbf{a}_{i, \sigma}^\dagger$) annihilates (creates) a particle in state σ at site i and $\mathbf{n}_{i, \sigma} = \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{i, \sigma}$ is the number operator. In this notation, the operators are anti-commuting for the Fermi-Hubbard model and commuting for the Bose-Hubbard model. In this thesis I focus on two variants, the single-component Bose-Hubbard model, Eqn. 2.93, and the two-component Fermi-Hubbard model, Eqn. 2.94. In experimental realizations of the Hubbard model with ultracold atoms there is typically an external harmonic potential, so $\epsilon_i = \frac{1}{2} m \omega^2 r_i^2$. This system satisfies the tunability requirement as the model parameters t and U are easily tuned by varying the lattice depth (§2.3).

Hubbard models are believed to be relevant to physical systems and are the paradigm for describing strongly correlated quantum materials. For example, The Fermi-Hubbard (FH) model was first proposed to describe strongly correlated electrons in transition metals [27], and later the Bose-Hubbard (BH) model was proposed [28] for bosonic particles in localized potentials, such as ^4He in porous media (i.e., Vycor). In particular, the FH model in the limit $U \gg t$, also called the $t - J$ model (J is the exchange term $4t^2/U$), is thought by some to describe the physics of the cuprate high-temperature superconductors [12]. The BH model is also conjectured to describe granular superconductors [29], Josephson-junction arrays [30] and flux lattices in type-II superconductors [31]. A shared feature between these physical systems and the Hubbard model is an interaction-driven insulator transition (Mott-insulator (MI) transition [32]). As shown schematically by the process in Fig. 1.2, this model/experiment comparison lends evidence to the belief that these physical systems are well-described by the Hubbard model. Similar experimental demonstrations of

the superfluid (SF) to MI transition for bosons in optical lattices [33] and the metal to MI transition for fermions [34, 35] confirmed the theoretical predictions of [25, 26] that ultracold atoms in optical lattices are well-described by the Hubbard model. In addition to being physically relevant, the Hubbard model has no analytic solutions and the quantum scaling issues previously discussed make numerical solutions intractable for certain parameter regimes. In particular, in the regime in which superconductivity is observed in the cuprates (hole doped away from the Mott-insulator), the FH model has not been solved [26].

The concept of using ultracold atoms in optical lattices as a quantum simulator of the Hubbard model to compare to experimental observables from physical systems almost exactly achieves the goals of quantum simulation laid out in Fig. 1.2. However, in practice, ultracold atoms in optical lattices have not yet simulated the Hubbard model over parameter ranges lacking classical solutions. This is because current systems do not satisfy the full range of proposed ingredients — the toolkit is incomplete. For example, current experiments do not have arbitrary control over the initial temperature and so do not contain ingredient (2). The main impediment to this control is the inability to reach the low temperatures associated with non-trivial regimes of the Hubbard models [36]. Since the common practice is to load the atoms into the lattice after cooling in a harmonic trap, experiments are limited to the lowest achievable harmonic trap entropies. Furthermore, light scattering from the optical lattice (§C.4) adds heat to the system throughout the loading process, which must be slow to prevent exciting the gas. Light scattering is more than a temperature issue since it also excites the system out of equilibrium [37], further reducing our ability to control the initial state. The ramping process itself is problematic for state control. In principle the ramp is sufficiently slow to be adiabatic, but assessing adiabaticity is complicated since the many-body timescales are difficult to ascertain a priori. Compounding these issues is the lack of thermometry methods in strongly correlated regimes, which violates ingredient (3). The ability to measure temperature is required to evaluate cooling proposals geared towards overcoming the temperature issues. Also, thermometry is critical for comparison with both theory and with the experimental results of physical systems. The scaling of system properties with temperature is an important method for distinguishing between competing theories and for identifying novel physics, such as the linear temperature dependence of the resistivity in the pseudogap regime of the cuprates [11].

An often overlooked aspect of quantum simulation is ingredient (4); techniques to explore non-equilibrium states. Yet, equilibrium only represents a small fraction of all possible solutions to the model and so simulating non-equilibrium physics greatly increases our understanding. In fact, the physical systems we compare to are usually considered useful for their non-equilibrium properties and are labeled as such. For example, a solid is labeled as either metal, insulator, semiconductor or superconductor, depending on charge transport properties. Finally, non-equilibrium measurements have no generalized framework [38] and are among the most intractable numerical calculations. For example, while the equilibrium properties of the BH model are well-described by QMC in three-dimensions [39], QMC cannot efficiently calculate dynamical properties [40]. In contrast, for a quantum simulator, non-equilibrium and equilibrium problems are on a similar footing.

Therefore, the work in this thesis is centered around the twin themes of developing the toolkit for quantum simulation and looking at the Hubbard model out of equilibrium. Our main tool for this is a ^{87}Rb 3D optical lattice apparatus⁴ built prior to this thesis (for details see [41]) and summarized in §3.1. This apparatus creates Bose-Einstein condensates (BEC) using standard methods (see i.e., [42, 43]) and then

⁴In the group of Brian DeMarco at the University of Illinois.

applies a 3D optical lattice using laser beams generated from a Ti-Sapphire laser, which is tunable from $\lambda \approx 760 - 820\text{nm}$. The thesis work on this apparatus can be divided into three sections: simulating the BH model out-of-equilibrium (§6), developing direct thermometry probes (§7) and implementing a spin-dependent lattice (§3.1.2) to develop impurity probes (§8). The specific experiments are summarized as follows:

Phase Slip Dissipation (§6.1, [1]): We measure the decay rate of center of mass oscillations of the condensate in the lattice as a function of temperature and lattice depth (t/U) for velocities less than known Landau and dynamic instabilities [44,45]. The scaling of decay rate with $\sqrt{t/U}$ is consistent with dissipation due to phase slips [46]. Phase slips occur when a phase coherent system (i.e., a condensate) with a phase gradient, and therefore a macroscopic velocity, relaxes to a lower velocity via a phase discontinuity. For example, these are a known cause of dissipation in SF ^4He and in thin superconducting wires [47]. Phase slip dynamics are set by the phase slip activation barrier which can be overcome by either thermal activation or quantum tunneling. At high temperatures, thermal activation of phase slips in superconductors is well-understood [48], but the role of quantum phase slips is more speculative [49]. We analyze the decay rate of the condensate velocity as a function of temperature and we observe the expected exponential suppression of thermal phase slip activation at low temperatures. Additionally, the decay rate saturates to a finite value as $T \rightarrow 0$, which is evidence for quantum tunneling of phase slips. Phase slips decay into topological excitations (i.e., vortices), so another test is to image vortices in the system. Although our imaging system is not optimized, approximately 20% of our images show evidence of vortex-like features, which are not observed in the absence of motion. These results clarify our understanding of the dynamic Bose-Hubbard model. Also, they are a useful basis of comparison against the transport properties of systems thought to be described by the BH model, i.e., thin 2D superconducting films [50]. Combined with other experiments [51,52] this helps establish center of mass velocity decay in the lattice as part of quantum simulation ingredient (4).

Condensate Fraction during Fast Lattice Ramps (§6.2, [2]): In the Bose-Hubbard (BH) model, as we increase t/U (i.e., decrease the lattice depth), the condensate fraction (i.e., the fraction of atoms in the macroscopically occupied single particle state [53]) transitions from zero in the Mott-insulator regime monotonically to unity in the weakly interacting superfluid regime. To probe the dynamics of this condensate fraction growth, we adiabatically load (§4.2.1) a condensate into the lattice to depth s_i , and then linearly ramp down the lattice from s_i to $s_f = 4$ ($s_i > s_f$) in time τ and measuring the condensate fraction. In the two extreme limits, $\tau \rightarrow 0$ and $\tau \rightarrow \infty$, we measure the equilibrium condensate fraction at s_i and s_f , respectively. The timescales between these limits provide an important probe of excitations in this system. Additionally, this experiment is a careful investigation of the validity of a technique known as bandmapping for the interacting gas. Bandmapping is a probe where the lattice depth is also linearly ramped down in time τ such that τ is adiabatic with respect to the bandgap and fast compared to quasimomentum changing processes (see §4.1.3). If these conditions are met, then quasimomentum is mapped to momentum. Bandmapping times on the order of 1ms are commonly used and here we test the legitimacy of bandmapping as a condensate fraction probe. We measure condensate fraction after the ramp as a function of τ and find that the condensate fraction grows exponentially with a time constant (τ_0) of approximately 1ms. We compare τ_0 measured experimentally for $10 \leq s_i \leq 14$ to a Gutzwiller time-dependent mean-field theory [54]⁵. This theory agrees well with experiment and provides insight into the origin and scaling of

⁵Theory by collaborators Stefan Natu and Erich Mueller at Cornell University

these timescales for condensate fraction growth. These timescales are too fast for transport and are therefore entirely local and are driven by particle-hole modes which scale as U . Importantly, for bandmapping this result means that the ramp has to be slow compared to the bandgap, but fast compared to $1/U$, which is difficult because there is no separation of these timescales. Therefore, measuring condensate fraction with bandmapping may introduce significant systematic errors, and so rule it out as a quantum simulation probe satisfying ingredient (3). However, this demonstrates the utility of lattice ramps for both a probe of the excitation spectra and as a technique to explore non-equilibrium phases. Additionally, this experiment helps validate Gutzwiller time-dependent mean field theory for simulating the dynamics of the BH model.

Bandmapping and In-Situ Thermometry (§7.1, §7.2, [3]): For weakly interacting bosonic and fermionic atoms in a smooth trap, the temperature is measured by fitting the momentum distribution, which is obtained after rapidly turning off the trapping potential and letting the gas ballistically expand (§4.1.2). Typically, the momentum distribution is fitted to a semiclassical distribution (§2.2.1). The momentum distribution in the lattice cannot be described semiclassically because the rapid spatial variation of the potential means that the de Broglie wavelength is not smaller than the typical length scale of the potential [55]. However, if we write the Hamiltonian in terms of the quasimomentum and position, then there are no fast spatial terms, so the quasimomentum distribution in the lattice can be described semiclassically. The 3D momentum distribution is a periodic function of the quasimomentum distribution with a Wannier function envelope. However, once the 3D momentum distribution is integrated along the probe beam to create a 2D image, the functional dependence on quasimomentum is complicated.

We investigate using bandmapping to directly image the quasimomentum distribution. This maps quasimomentum to momentum, which we then image after expansion. We only consider weakly interacting thermal gases ($T > T_C$), so that complications associated with the separation of timescales introduced by interactions for the degenerate gas are not relevant. We fit images obtained from bandmapping to semiclassical quasimomentum distributions with temperature as a free parameter for a range of known entropies. Here entropy is determined by the temperature in the harmonic trap before adiabatically loading into the lattice. Since the gas is non-interacting, we use non-interacting thermodynamics to predict the in-lattice temperature.

We find that when the predicted temperature T_1 in the lattice is comparable (and larger) to the bandwidth $12t$, the temperature T_2 obtained from the quasimomentum fit is larger than T_1 . Through comparison to 1D numerical simulations of bandmapping using a Crank-Nicolson solver, we conclude that this failure is a result of adiabaticity breaking down at the band edge. This is exacerbated by the added complication that the quasimomentum distribution becomes increasingly insensitive to temperature as the band becomes full (i.e., $k_B T \geq 12t$). Since the band is full, any additional thermal energy must go into the spatial degrees of freedom. This motivates using the spatial distribution for thermometry. Therefore, we fit the in-trap RMS width and demonstrate, via comparison to T_1 , that it is a reliable method for thermometry when $k_B T > 12t$. Although our measurements utilize non-interacting gases, the qualitative effect of interactions is to also populate higher quasimomentum states, and equivalently position is the only sensitive temperature variable. This has been exploited in subsequent high-resolution imaging experiments [56–58] in the strongly correlated regime of the BH model. This is an important result for further developing thermometry methods for quantum simulation to satisfy ingredient (3) and for validating other thermometry methods, such as in

Ref. [59].

Condensate Fraction Thermometry (§7.3): Since the gases trapped in a lattice are a closed system, direct thermometry probes for quantum simulation using atoms in optical lattices must measure temperature via its relation to a system observable. However, these relations require theoretical input, which may not be available or is unreliable given that we are running a quantum simulation which aims to probe unknown physics. Therefore, we need to identify aspects of the system which are both sensitive to temperature and related to temperature from validated theories. In the harmonic trap, the approach is to fit the tails of the momentum distribution, where the interactions are minimal, to a non-interacting distribution. However, we previously demonstrated that the equivalent approach in the lattice, measuring the quasimomentum distribution [3], does not work. Also, in the interesting regions of the BH phase diagram, the quasimomentum distribution is strongly altered by interactions [60]. For experiments with high resolution in-situ imaging [56–58], temperature can be measured at the trap edge where the density is low enough to ignore interactions [61–63], but for strong interactions and low temperatures this region becomes quite small.

We investigate the possibility of measuring temperature in the superfluid (SF) regime using the peak fraction, an experimentally measurable proxy for condensate fraction. When the system is condensed, the distribution after expansion is bimodal [64]; condensed atoms are in narrow peaks. We define peak fraction as the fraction of atoms in these peaks after subtracting off the non-condensate atoms; we assume the non-condensate distribution under the peaks is flat. The peak fraction can thus be evaluated without requiring the exact functional form. Also, the bimodal nature is robust [64], so the peak fraction is only weakly dependent on experimental issues such as imaging resolution, imaging direction with respect to the lattice axes, and expansion time. Similar peak fractions have been used in a number of experiments to identify the SF-normal transition [65–67].

In this experiment, we establish our heuristic method for fitting lattice expansion images to extract peak fraction, and then we do a quantitative comparison of the measured peak fraction to the condensate fraction predicted assuming adiabatic loading using several theoretical approaches: non-interacting thermodynamics, Hartree-Fock-Bogoliubov-Popov theory [64], and site-decoupled-mean-field theory [68]. In developing our peak fraction fit, we compare to simulated lattice data (§G.4) and do a full numerical calculation for the harmonic gas expansion to better understand how the non-condensate distribution under the condensate peak is affected after time-of-flight (§G.3). We find that the comparison of peak fraction to predicted condensate fraction is poor over the entire parameter range. Although some heating is expected during the ramp, counterintuitively we find states which appear to have lower entropy in the lattice. For entropies where a condensate exists in the harmonic trap, but is predicted to disappear in the lattice, we find the condensate persists in the lattice. This is possible evidence for condensate metastability after loading into the lattice. This has an important impact for ingredient (2) of quantum simulation, since the initial state is not in equilibrium as expected. Fundamentally, lattice loading is a dynamic process and further understanding requires developing new theoretical tools.

3D Spin-Dependent Lattice in the Hubbard Regime (§3.1.2, [4]): The most general approach to thermometry is to have a well-understood system, in our case a harmonically trapped impurity, in thermal contact with the quantum simulator: an impurity thermometer. This technique is also central to several

cooling proposals [69, 70] that rely on shifting entropy from the simulator to a harmonically trapped reservoir of atoms. Such a system can be created by using two different atomic species and tuning the lattice wavelength [71, 72] or by using two states (“spins”) of the same species and tuning polarization. Part of this thesis work was to develop a spin-dependent lattice to realize the latter implementation.

Spin-dependent lattices are created by rotating the polarization of the retro-reflected lattice beam 90° to obtain a lin- \perp -lin lattice (§C.2.3), resulting in a lattice depth which is proportional to the atomic magnetic moment, $g_F m_F$. For ^{87}Rb , a mixture of the states $|F = 1, m_F = 0\rangle$ (or $|2, 0\rangle$) and $|1, -1\rangle$ in a lin- \perp -lin lattice realizes a system where a harmonically trapped gas ($|1, 0\rangle$) is in contact with a lattice gas ($|1, -1\rangle$). Our work represents the first time a fully spin-dependent lattice was realized in all three-dimensions for a lattice gas in the Hubbard regime⁶; previous work was in 1D [73–76]. Using a mixture, we demonstrated the SF to MI transition with the lattice state while the other state was virtually unchanged. After realizing the 3D spin-dependent lattice, we investigated probing the lattice gas with the impurity as described in the next two paragraphs.

Impurity Thermalization (§8.1): Thermal contact is a necessary requirement to realize an impurity thermometer using a spin-dependent lattice. Therefore, we investigated thermalization between the impurity and the lattice gas. Based on several consistent results, we conclude that thermalization is poor between these two systems. An obvious mechanism for poor thermalization is if the two gases are not physically in contact with each other. We rule out separation on trap length scales by direct in-situ imaging. This still leaves the possibility that the gases are separated on lattice length scales, which are difficult to spatially resolve by imaging. Therefore, we developed a microwave spectroscopy technique (§4.1.6) that measures the density distribution of the impurity by driving atoms from free states into lattice bound states (see also [76] for microwave driving in spin-dependent lattices). Using this technique, we do not observe any concrete evidence for phase separation.

In our first thermalization test (the ‘direct’ test), we create a mixture of spin states and measure the impurity temperature as we turn on the lattice. We observe no change in the impurity temperature (T_{im}) as a function of the lattice depth. Loading into the lattice adiabatically cools the lattice gas, so there is a limited temperature gradient between the systems ($\Delta T \leq T_{im}$). Therefore, we focused on methods for heating the lattice gas. In the second test, we put energy into the lattice gas by parametrically oscillating the lattice depth [77, 78] and measuring the impurity condensate fraction as a function of the drive strength. At higher lattice depths, the drive strength required to change the impurity condensate fraction by the same amount also increases, indicative of reduced thermalization. To quantify this result, we need to know the lattice temperature. However, we do not fully understand energy relaxation during parametric oscillation, which is an area of current study [79–81].

In our final test of thermalization, we first dephase the lattice gas, which fills the band, and then measure the impurity condensate fraction versus time after the dephasing. We convert condensate fraction to temperature and fit to a line to get the heating rate, which decreases monotonically with lattice depth. This test can be directly compared to theory since the initial conditions are characterized; a filled band effectively

⁶3D state-dependent lattices are created in the MOT and give rise to sub-Doppler cooling processes. However, the MOT lattice is not providing the trapping force and the atoms in the MOT are not described by a Hubbard model.

corresponds to $T_{latt} = \infty$ for a single-band uniform lattice. Therefore, we perform a Fermi's golden rule calculation of the energy-exchange rate for $T_{latt} = \infty$ and $T_{im} = 0$. The calculated energy-exchange rate follows the same trend as the experimental data, and implies that poor thermalization is caused by the mismatch between the lattice and harmonic dispersion relations. This experiment has important implications for impurity thermometry and for constraining cooling proposals that rely on thermal contact between the lattice gas and a reservoir.

Impurity Transport (§8.2): A direct method to probe interactions between two systems is to initialize relative center-of-mass (COM) motion and measure the resulting decay of the motion caused by collisions. Here we explore collisions between the impurity and the lattice gas by starting in-trap COM oscillations of the impurity with the lattice gas at rest. Similar schemes in 1D tubes [82,83] have been shown to probe the static and dynamic structure factors. In this experiment, we let the impurity oscillate over several periods through the 3D lattice gas and measure the decay rate of the oscillation. In thermal gases this decay rate is a direct measure of the thermalization time [84], and in harmonically trapped condensates damping is the result of both non-linear mean-field dynamics and binary collisions [85]. We measure that the decay rate decreases to nearly zero (normalized to the single component decay rate) as we approach the MI transition. One possibility is that this rate is probing the excitation spectrum which is becoming gapped. It has also been suggested that number fluctuations in the lattice gas, which are Poissonian at low lattice depths, but zero in the MI regime, may appear as a disordered potential for the impurity [86,87] and dephase the oscillations. This result may reveal important information regarding thermalization of the impurity and may also be an important probe for the excitation spectrum for quantum simulation to satisfy ingredients (3) and (4).

2D Fermi-Hubbard Apparatus with Single-Site Imaging Resolution (§3.2): A separate component of this thesis was developing (including design, construction, and testing) a ^{40}K 2D optical lattice apparatus⁷ (§3.2). This apparatus is specifically designed to achieve single-site imaging [57,58] and manipulation of a 2D degenerate Fermi gas in an optical lattice and will help satisfy ingredients (2) and (3) of quantum simulation for the 2D Fermi-Hubbard model. Single-site imaging will allow us to measure important particle-hole and spin statistics to identify the predicted antiferromagnetic correlations. Once achieved, single site manipulation will allow doping to explore the parameter regime where d-wave superconductivity is thought to occur. While this apparatus has not yet realized a degenerate lattice gas, progress during this thesis includes: constructing the physical apparatus (including the lasers, control, magnetic trapping, and vacuum system), realizing a ^{40}K MOT, and demonstrating transport of ^{40}K atoms from the MOT chamber to the final science chamber where the degenerate gas will be prepared. A specific experiment performed as part of this thesis work was realizing a 405nm MOT for ^{40}K .

4S \rightarrow 5P K MOT at 405nm (§E.3, [88]): A unique feature of the ^{40}K apparatus is that we plan on imaging using fluorescence from the 4S \rightarrow 5P transition at 405nm (see Fig. A.1). Compared to standard imaging on the 4S \rightarrow 4P transition at 767nm, 405nm imaging can achieve a higher resolution since the Rayleigh criterion scales as λ (Eqn. F.4). Also, for single site imaging, which requires fluorescence collection times on the order of seconds for a sufficient signal-to-noise ratio, the excitation light needs to simultaneously perform laser cooling. Since the linewidth of the 405nm transition is approximately six times less than

⁷This work was completed in the group of Joseph Thywissen at the University of Toronto.

767nm, we expect enhanced Doppler cooling.

To exploit this transition, we developed a 405nm laser system based on GaN diode lasers locked to the $4S \rightarrow 5P$ transition of ^{39}K and capable of delivering 30mW of single-mode light to the MOT chamber. Using this light we demonstrated the first free-space laser cooling on this transition for K; similar work on the $nS \rightarrow (n+1)S$ transition has been performed for He^* [89, 90] and ^6Li [91]. We observe an almost ten times increase in density and temperatures as low as $63(6)\mu\text{K}$. The combined temperature and density improvement in the 405 nm MOT corresponds to a twenty-fold increase in phase space density over our 767 nm MOT. This is an important first step towards developing our high-resolution imaging probe for simulating the 2D FH model.

In addition to the sections detailing these specific experiments, §2 is an overview and background of ultracold atoms in optical lattices, §3 and §4 give an overview of our quantum simulation toolkit (including elements developed in this thesis), and §5 discusses theoretical tools used and developed to complement our quantum simulator. Background information, technical details (i.e., calibrations, code and circuits), derivations, and supporting calculations are covered in the Appendix.

Chapter 2

Strongly Correlated Atoms in Optical Lattices

In this thesis, I use ultracold atoms trapped in optical lattices as a quantum simulator for the Hubbard model. This chapter is a background on ultracold atoms, optical lattices, and the Hubbard model as a foundation for the experimental work discussed later. As briefly mentioned in the introduction, ultracold atoms are ideal candidates for quantum simulation because they are quantum particles with configurable potentials (§2.1.1) and simple, tunable interactions (§2.1.2). To trap atoms in an optical lattice, we first create a degenerate gas in a harmonic trap. In section §2.2, we review the properties and theory of atoms in harmonic traps. Next, the atoms are adiabatically loaded into an optical lattice formed from the interference of several laser beams (§C). In the optical lattice, the atomic states form a band structure and the wavefunctions are localized to lattice wells. Combined with the atomic interactions this system realizes the Hubbard model. These topics are detailed in §2.3. Finally, we discuss the properties of this model in §2.4

2.1 Ultracold Atoms

Although there is no precise definition, ultracold atoms are approximately atoms with a center-of-mass temperature, typically $T < 1\mu\text{K}$, sufficiently low that the de Broglie wavelength is comparable to the interparticle spacing. In the regime satisfying this condition, the atoms must be described quantum mechanically and are thus quantum particles¹. A limiting factor to attaining the ultracold regime is that nearly all atomic species solidify at significantly higher temperatures. To avoid this issue, systems of ultracold atoms are very dilute, and thus are metastable to forming a solid and/or deeply bound molecules. In some cases (e.g., the strongly repulsive gas), these timescales limit current experimental progress [92]. The variety of atomic species that have been cooled to the ultracold regime continues to increase. However, the majority of experiments utilize alkali metal atoms for technical feasibility and optimal atomic properties. The work in this thesis involves two such atoms, ^{87}Rb and ^{40}K , the properties of which are detailed in §A.

The standard protocol for preparing ultracold atoms is laser cooling and collection in a magneto-optical trap (MOT), followed by evaporative cooling in a conservative trapping potential (see Fig. 1.3). In laser cooling, atoms slow down by absorbing photons from a laser beam. The atoms preferentially absorb from one of a pair of counter-propagating beams because they are Doppler shifted into resonance. The MOT provides an additional mechanism to make the absorption spatially dependent and provide a restoring force for trapping; for more theoretical details see §B.8.2 and §E. Experimental details for the MOTs in this thesis are given in §3.1, §3.2.6 and §E. This process needs to be performed in an ultrahigh vacuum system

¹Atoms are “cold” ($T \lesssim 1\text{mK}$) when 2-particle interactions must be described using quantum mechanics (i.e., quantum scattering). A broader definition of ultracold atoms are atoms cooled into the regime where quantum collisions are restricted to s-wave scattering ($T \lesssim 100\mu\text{K}$). See §2.1.2 for details.

to thermally isolate the atoms and prevent collisions with residual gases. The original source of the atoms is either a slow atomic beam [93] or room-temperature vapor [94]. In all our experiments, we load from a vapor produced by heating a small metal dispenser which activates a chemical reaction, thereby releasing the atomic gas into the vacuum system. Once $\approx 10^9$ atoms are collected in a MOT they are transferred to a conservative potential (see §2.1.1) in a vacuum region with minimal residual gas ($P < 10^{-11}$ Torr), and the temperature is further reduced by evaporative cooling. This involves selectively removing the highest energy atoms, and then collisions re-thermalize the gas to a lower temperature (see [21], §3.2.8). The atoms enter the ultracold regime after sufficient evaporative cooling.

As quantum particles for quantum simulation, ultracold atoms have a number of advantages. First, since atoms are composite particles comprised of spin- $\frac{1}{2}$ electrons, neutrons and protons, they can be either bosons (total integer spin) or fermions (total half-integer spin). Therefore, they can be used to simulate models with either type of quantum statistics. We exploit this fact in this thesis since ^{87}Rb is a boson and ^{40}K is a fermion. Another advantage is that atoms have a number of internal states and can be initialized into an arbitrary coherent or incoherent mixture of these states. Experiments with atoms in a mixture of states can simulate multi-component models. Specifically, mapping states to pseudospins simulates spin-models. This is important since at least two states are required to simulate non-trivial fermion models such as the Fermi-Hubbard model². Atoms interact with electromagnetic fields (§B), which allows us to probe the simulation results (§4.1) and create electromagnetic potentials to confine the atoms and influence the simulation model. The next section will explore this topic in more detail. Finally, ultracold atoms are interacting, which is a feature of most classically unsolvable models. These interactions are well-approximated by a contact potential characterized by a single parameter, the s-wave scattering length. Importantly, this parameter can be tuned from $-\infty$ to $+\infty$ using a single experimental knob. These tunable interactions are discussed in more detail in §2.1.2.

2.1.1 Tailored Potentials

Atoms interact with electromagnetic fields predominantly via the magnetic-dipole and electric-dipole interaction (Eqn. B.1). When these fields are resonant with the energy difference between internal states, this interaction drives atomic population coherently between these states (see §B.4). When the fields are off-resonance, the interaction causes energy shifts in the states. If these shifts are spatially dependent, then a conservative potential is formed for the atoms. Two types of fields create the strongest potentials — DC magnetic fields, and AC electric fields oscillating at optical frequencies (light).

DC magnetic fields (§B.2) create strong shifts because the atomic ground state, labeled by the quantum numbers $|F, m_F\rangle$ (see §A.1 for details), have a permanent magnetic dipole. For sufficiently slow moving atoms, the potential is

$$V(\vec{x}) = \mu_B m_F g_F |\vec{B}(\vec{x})|. \quad (2.1)$$

AC magnetic fields can also form potentials directly [95] or by dressing [96], but these are more specialized and not relevant to this work. The simplest type of magnetic field for trapping atoms is created by running current in opposite directions through a pair of coils (see Fig. 2.1). Close to the center of the coils, the field

²Fermions in the same state do not interact at the temperatures of interest; see §2.1.2.




Figure 2.1: (a) Quadrupole coil pair separated by a distance $2L$ and with coil radius R . Each coil has equal and opposite current I and N turns of wire. (b) Contour plot of the quadrupole field magnitude in Gauss and arrows illustrating the local field direction for typical experimental parameters: $N = 100$ turns, $I = 20A$, $L = 10\text{cm}$, and $R = 10\text{cm}$. There is no field coming out of the page and the field is symmetric about the y axis. (b) Field magnitude for a 1D slice along x for $y = 0$. The maximum field along this direction is $\approx 45\text{G}$. For small $|x|$, the field is well-described as linear in $|x|$ as given by Eqn. 2.3. The gradient along this direction is 13.3G/cm and twice as strong in the other direction.

is (symbols defined in Fig. 2.1),

$$\vec{B}(\vec{x}) = \frac{3}{2}NI\mu_0 \frac{R^2L}{(R^2 + L^2)^{5/2}}(x\hat{x} + y\hat{y} - 2z\hat{z}) \quad (2.2)$$

$$|\vec{B}(\vec{x})| = \frac{3}{2}NI\mu_0 \frac{R^2L}{(R^2 + L^2)^{5/2}}\sqrt{x^2 + y^2 + 4z^2} \quad (2.3)$$

which has a field zero at the center and increases linearly outwards. This is known as a magnetic quadrupole field and confines atoms with states $m_F g_F > 1$ (weak-field seeking states)³. From Fig. 2.1, the depth of a typical trap is 45G and the trap size has centimeter length scales. For $g_F m_F = 1$, this corresponds to a trap depth of $T = V/k_B = 3\text{mK}$. Therefore, magnetic traps are characterized by large depths and volumes and so are optimal for initial trapping of atom clouds out of the MOT. As we cool atoms in a quadrupole trap, they collect near the zero-field region where they have a high probability of diabatic spin transitions and subsequent ejection from the trap (see §B.2). This loss process, called Majorana loss, implies that the quadrupole trap is not suitable for ultracold atoms. Instead, we require a trap with a non-zero magnetic field minima. This type of trap, known as a Ioffe-Pritchard (IP) trap, has a field magnitude (nearby the minimum) of,

$$|B| = |B_0| + A_x x^2 + A_y y^2 + A_z z^2 \quad (2.4)$$

where A_i are the field-magnitude curvatures. IP traps are designed such that $A_i > 0$, and therefore form a harmonic trap. The relation of B_0, A_i to the physical parameters depends on the specific configuration (see [43], for example). With the exception of the fields close to patterned magnetic surfaces [97], magnetic traps for sufficiently small gases can be described as either quadrupole (Eqn. 2.3) or harmonic (Eqn. 2.4).

The other type of potential for ultracold atoms, a “dipole-potential” (§B.5), is due to light interacting with the atom via the electric-dipole interaction. The light induces an electric dipole in the atom by coupling the ground state to an electronic excited state, and this dipole interacts with the electric field of the light. The potential is therefore proportional to the intensity of the field. Additionally, to determine the full potential we must consider all possible excited states as shown in Eqn. B.58. For alkali atoms, the potential from a linearly polarized light field with intensity I and frequency ω is

$$V(\vec{x}) = \frac{\pi c^2 \Gamma}{2\omega_0^3} \left(\frac{2}{\Delta_{D2}} + \frac{1}{\Delta_{D1}} \right) I(\vec{x}), \quad (2.5)$$

where $\Delta_{D1(D2)} = \omega - \omega_{D1(D2)}$ is the detuning of the light frequency from the atomic D1 (D2) transition and Γ is the decay rate from the excited state. The values for $\omega_{D1(D2)}$ and Γ are given by Fig. A.1 for ⁴⁰K and Fig. A.2 for ⁸⁷Rb. The above equation assumes the detuning is larger than the hyperfine splitting, which is almost always the case, and ignores so-called counter-rotating terms. This potential is discussed in more detail in §B.5, and the full potential for arbitrary polarizations is given by Eqn. B.60. However, Eqn. 2.5 retains most of the essential features of the dipole potential. First, the dipole potential can be either attractive ($\Delta < 0$, “red-detuned”) or repulsive ($\Delta > 0$, “blue-detuned”), depending on the sign of the detuning. Furthermore, for linear polarizations and/or large detunings, the potential is state-independent. This expands the number of trappable states and decouples the state degree-of-freedom from the potential which allows, for example, the exploration of state-dependent interactions [98].

³There are no traps with magnetic field maxima by Earnshaw’s theorem.

Importantly, since the potential is proportional to the intensity of an *optical* field, we can create features in the intensity on the order of $\lambda = c/\omega \approx 1\mu\text{m}$ as limited by diffraction theory. The most straightforward potential is obtained by focusing a single-mode laser beam, which has an intensity near the focus of [99,100]

$$I(\vec{x}) = \frac{2Pz_0^2}{\pi w_0^2(z_0^2 + z^2)} \exp\left[-2\frac{r^2 z_0^2}{w_0^2(z_0^2 + z^2)}\right], \quad (2.6)$$

where z_0 is the Rayleigh length (Eqn. C.44), w_0 is the beam waist, P is the beam power, z is the coordinate along the beam propagation axis, and r is the radial coordinate orthogonal to z . The potential is harmonic close to the focus (see Eqn. C.48), with strong confinement perpendicular to the beam axis and weak confinement along the beam. For typical experimental parameters — $\lambda = 1064\text{nm}$, ^{87}Rb , $P = 1\text{W}$, and $w_0 = 100\mu\text{m}$ — the trap depth is $k_B \times 9.7\mu\text{K}$ (including counter-rotating terms), and the harmonic trap frequencies are $\nu_r = 97\text{Hz}$ and $\nu_z = 0.2\text{Hz}$. Therefore, the dipole potential is shallow compared to the magnetic potential. Another possibility is to combine these two types of traps. For example, a repulsive optical dipole trap focused to the center of a magnetic quadrupole — a “plugged” quadrupole [101] — prevents Majorana losses (see §3.2.9). Another hybrid potential, used in the ^{87}Rb apparatus and detailed in §3.1, is a single-beam dipole potential focused below the center of a magnetic quadrupole [102,103]. The off-center quadrupole provides harmonic confinement along the weak axis of the dipole potential and offsets gravity along the vertical axis. One disadvantage of optical traps is that they are not perfectly conservative. Photon scattering via spontaneous emission causes heating; see §B.8.1 and §C.4 for details. The heating rates scales as Δ^{-2} , and so it can be made negligible by using large detunings.




Figure 2.2: (a) A standing wave intensity pattern is created by retro-reflecting a light beam. Here the intensity is proportional to $\cos^2(kx)$, where $k = 2\pi/\lambda$ is the wavevector. The distance between intensity maxima (minima) is $d = \lambda/2$, which is the lattice spacing. The depth of the optical lattice potential is $s \times E_R$, where $E_R = (\hbar k)^2/2m$ is the recoil energy ($E_R = k_B \times 172\text{nK}$ for ^{87}Rb and $\lambda = 800\text{nm}$). If the lattice is red(blue)-detuned the atoms reside near the intensity maxima (minima). (b) Incoherently adding two orthogonal standing waves creates a 2D lattice potential.

Since the optical potential is proportional to the intensity (i.e., the squared modulus of the electric field), an interesting class of potentials arise from the coherent addition of two or more optical fields. The simplest of these is to retro-reflect a single beam, which creates a periodic standing wave in intensity and a periodic potential as illustrated in Fig. 2.2. This is the basis of the optical lattice potential described in detail in §C. For multiple fields at different frequencies, the coherent (interference) terms oscillate at the frequency difference and, if large enough, these terms will time average to zero. Then the total potential is the incoherent sum of the potentials from each individual field. For example, if we have retro-reflected beams along three orthogonal directions with slightly different frequencies then the potential is,

$$V(\vec{x}) = E_R [s_x \cos^2(kx) + s_y \cos^2(ky) + s_z \cos^2(kz)] \quad (2.7)$$

where $k = 2\pi/\lambda$ is the wavevector of the light, and s_i are the lattice depths (in E_R) along each direction. Often in experiments, $s_x = s_y = s_z = s$, and Eqn. 2.7 is a 3D cubic lattice potential. By adjusting the overall and relative strengths of s_i , this potential can also be utilized to explore physics in lower dimensions. If $s_x \gg s_y, s_z$ then the atoms are restricted to moving in a series of 2D planes and similarly if $s_x, s_y \gg s_z$ then the atoms move in 1D tubes. Beyond the cubic lattice, more advanced interference geometries can realize hexagonal [104], triangular [105], kagome [106], honeycomb [107] and superlattices [108, 109]. Additionally, by adjusting the relative polarization of the interfering beams the lattice depth can be state-dependent. A specific configuration is when forward and retro-reflected beams are linearly polarized, but rotated with respect to one another by an angle θ (i.e., a lin- θ -lin lattice; see §C.2). In this thesis work, we implement and utilize a lattice where $\theta = \pi/2$ (lin- \perp -lin), which has a fully state-dependent lattice potential depth: $s \propto g_F m_F$ (§3.1.2, §8). Finally, a special type of interference pattern, optical speckle, is generated by interfering beams from many directions with random phases. This speckle field has a very short correlation length and can be used to simulate disorder in physical materials. This is part of another research direction in the Illinois group, but not a specific aspect of this thesis (see [41, 110]).

2.1.2 Tunable Interactions

Alkali atoms interact via a van der Waals potential $V = -C_6/r^6$ at long range and a highly repulsive exchange potential at short range. This interaction is spherically symmetric in the relative coordinate between the atoms. Since the free-space eigenstates, $|F, m_F\rangle$, are mixed in the interaction region, a given pair of atoms in states $|F_1, m_{F1}\rangle, |F_2, m_{F2}\rangle$ are coupled to a number of different Born-Oppenheimer potentials, referred to as channels. This situation is illustrated in Fig. 2.3.

At sufficiently low temperature⁴, the effect of the interaction potential must be treated using the formalism of quantum scattering (see e.g., [112]). Since the potential is spherically symmetric, we can decompose the scattering into spherical waves of total angular momentum l . The strength of the interaction is characterized by the l^{th} -wave phase shift δ_l where $\delta_l = 0$ corresponds to no interaction. The collision cross-section for particles of relative wavenumber k is

$$\sigma = \frac{4\pi}{k^2} \sum_l (2l+1) \sin^2(\delta_l). \quad (2.8)$$

⁴ $T < T_{vdW} = k_B \times E_{vdW}$ (E_{vdW} is the van der Waals energy), where $T_{vdW} \approx 300\mu\text{K}$ for ^{87}Rb and 1mK for ^{40}K [111].




Figure 2.3: (a) The interaction potential between two alkali atoms. Typical length scales of the interaction region are a few nanometers (grey shaded region). For atoms in free-particle eigenstates (spin “up” or “down”), there is a particular interatomic potential (blue curve, open channel), but this is coupled to the potential for other states (red curve, closed channel) since the states are mixed in the interaction region. A schematic representation of the radial wavefunction in the open channel ($r\Psi$) is shown, which is asymptotically sinusoidal. The scattering length a is the distance from the origin to where the asymptotic $r\Psi$ is 0. (b) For an attractive square well potential, which qualitatively displays the same features as the atomic potential, the scattering length changes as the depth of the potential is adjusted. In particular, the scattering length varies through $\pm\infty$ as a bound state appears. A similar resonance (Feshbach resonance) occurs when a bound state in a (coupled) closed channel is equal to the free particle energy.

For $l \neq 0$ scattering there is a centrifugal barrier with potential energy

$$\frac{\hbar^2}{2m} \frac{l(l+1)}{r^2}. \quad (2.9)$$

If the initial kinetic energy is not sufficient to overcome this barrier, then the atoms cannot reach the interaction region, and $\delta_l \rightarrow 0$, $\sigma \rightarrow 0$. This is true at ultracold temperatures, which implies that only s-wave ($l = 0$) interactions are relevant. Since the s-wave state is symmetric under particle exchange, fermions in the same state do not have s-wave interactions.

The s-wave scattering phase is more conveniently represented as a length

$$-a = \lim_{k \rightarrow 0} \frac{\tan(\delta_0)}{k}. \quad (2.10)$$

The scattering length is where the free-particle radial wavefunction ($r\Psi$) asymptotically crosses through zero, as illustrated in Fig. 2.3. The cross-section is

$$\sigma = (f+1) \frac{4\pi a^2}{1+k^2 a^2}, \text{ and} \quad (2.11)$$

$$\sigma_{|ka| \ll 1} = (f+1)4\pi a^2, \quad (2.12)$$

where $f = \pm 1$ for identical bosons (fermions), and $f = 0$ for distinguishable particles. The scattering length is a sensitive function of the potential and typically must be determined by experimentally measured parameters. For example, if we approximate the potential as an attractive spherical well of depth V_0 (with $k_0 = \sqrt{2MV_0/\hbar^2}$) and size r_0 then [113]

$$a = r_0 \left(1 - \frac{\tan(k_0 r_0)}{k_0 r_0} \right), \quad (2.13)$$

which is plotted in (b) of Fig. 2.3.

A resonance in a occurs when $k_0 r_0 = (1+2n)\pi/2$, which corresponds to the formation of a bound state in the potential on the $a > 0$ side of resonance with energy $-1/a^2$. Therefore, we can tune the scattering length, and therefore interactions, by changing the potential depth. While we cannot directly change the depth, we can change the relative depth between different coupled channels since each may have a different magnetic moment. In this way, a bound state in the closed channel (red curve in Fig. 2.3) crosses the energy of the free-particle in the open channel, and the scattering length also goes through a resonance. This ‘‘Feshbach’’ resonance can be characterized as a magnetic field dependent scattering length [111]

$$a(B) = a_0 \left(1 + \frac{\Delta}{B - B_0} \right), \quad (2.14)$$

where a_0 is the background scattering length, Δ is the resonance width and B_0 is the resonant field. These parameters are summarized for the atoms used in this thesis in Table A.3. The availability of a Feshbach resonance for tuning the interactions is a main advantage for simulations with ultracold atoms.

It is inconvenient to use the full atomic interaction potential for calculating the resulting many-body

effects of the interaction. We can replace the true potential with a pseudo-potential [114]⁵,

$$\begin{aligned} V(r) &= \frac{2\pi a \hbar^2}{\tilde{m}} \delta(r) \frac{\partial}{\partial r} r, \\ &= g \delta(r) \frac{\partial}{\partial r} r, \end{aligned} \quad (2.15)$$

where \tilde{m} is the reduced mass ($\tilde{m} = m/2$ for particles of the same mass). If we consider changes in the wavefunction to first-order in the interaction, then we can leave out the regularization term ($\frac{\partial}{\partial r} r$), and

$$V(r) \approx g \delta(r). \quad (2.16)$$

For example, the first order energy shift due to the interaction from Eqn. 2.16 is linear in a , which cannot be obtained perturbatively from the true potential. This highlights the importance of a in characterizing interactions in the many-body system.

Despite the many advantages of alkali interactions, they are not suitable for simulating certain models with anisotropic or long-range interactions, such as $1/r^3$ dipole-dipole-type interactions⁶. However, different types of ultracold atoms with large magnetic dipole moments, such as Dy [115], Er [116] and Cr [117], and ultracold molecules with large electric dipole moments, such as Rb-K, have promise for exploring these types of interactions. A more fundamental limitation resulting from interactions are inelastic loss processes. Since alkali potentials support deeply bound states, free particles are not the true ground state and are therefore metastable to forming molecules and ultimately solids. For bosons, this rate scales as a^4 [118] and limits the ability to explore strongly interacting Bose gases. For free fermions on the repulsive branch ($a > 0$), the rate scales as a^6 [118]. Fortuitously, if fermions are paired in the bound state formed on the $a > 0$ side of the Feshbach resonance, the loss rate scales as $a^{-2.5}$ [119]. Additionally, these loss processes require a three-particle collision and so can be suppressed by using low density samples.

2.2 Ultracold Atoms in Harmonic Traps

The starting point of quantum simulation with ultracold atoms is preparing the gas in a harmonic trap. Additionally, this is where the thermodynamic state of the system is characterized before superimposing the lattice. Therefore, it is important to understand the theory of the gas in a harmonic potential, both for practical purposes and as a foundation for quantum simulation. Furthermore, the theory of the harmonic gas elucidates concepts that will also be relevant to more complicated Hubbard model physics.

2.2.1 Non-Interacting Atoms: Bose-Einstein Condensation and Fermi Degeneracy

As with any system, we start with a simplified case: no interactions. While this may seem overly trivial, non-interacting theory includes quantum statistics, which gives rise to non-trivial properties. Also, while the theory will fail for the strong interactions necessary for quantum simulation, there is usually some part of the system (i.e., a low density region at the edge of the trap) for which the non-interacting description is valid

⁵The pseudo-potential is valid outside the interaction region.

⁶Dipole-dipole interaction effects can be observed for alkali atoms, but the energy scales are very low.

and important for measuring global properties, such as temperature and chemical potential. In describing the non-interacting gas in the harmonic trap I will focus on two aspects: the single-particle wavefunctions and thermodynamics.

Solving for the eigenstates and eigenvalues of the harmonic oscillator is detailed in any standard quantum textbook (i.e., [120]) with the well-known result,

$$E_n = \hbar\omega \left(n + \frac{1}{2} \right) \quad (2.17)$$

$$a_{ho}^2 = \frac{\hbar}{m\omega} \quad (2.18)$$

$$\psi_n(x) = \left(\pi 2^{2n} a_{ho}^2 (n!)^2 \right)^{-1/4} e^{-x^2/2a_{ho}^2} H_n \left[\frac{x}{a_{ho}} \right], \quad (2.19)$$

where H_n are the Hermite polynomials, and a_{ho} is the harmonic oscillator length. The most important wavefunction $\psi_n(x)$ is for the ground state $n = 0$,

$$\psi_0(x) = \left(\pi a_{ho}^2 \right)^{-1/4} e^{-x^2/2a_{ho}^2}, \quad (2.20)$$

which is Gaussian. From these states we can construct the mean value of a general operator \mathbf{O} ,

$$\langle \mathbf{O} \rangle = \sum_n \left(e^{(E_n - \mu)/kT} \pm 1 \right)^{-1} \langle \psi_n | \mathbf{O} | \psi_n \rangle \quad (2.21)$$

and any thermodynamic quantity starting from the grand canonical potential⁷

$$\Omega = \mp k_B T \sum_n \log \left[1 \pm e^{(E_n - \mu)/kT} \right], \quad (2.22)$$

for either bosons ($-$) or fermions ($+$).

This method is completely general, but there is a simpler approach: the semiclassical approximation. In this approximation, $\vec{\mathbf{x}}$ and $\vec{\mathbf{p}}$ are treated as classical variables, so that,

$$\langle \mathbf{O} \rangle = \frac{1}{h^3} \int d^3x d^3p O(x, p) \left[e^{(p^2/2m + V(x) - \mu)/kT} \pm 1 \right]^{-1}, \quad (2.23)$$

$$\Omega = \mp k_B T \frac{1}{h^3} \int d^3x d^3p \log \left[1 \pm e^{(p^2/2m + V(x) - \mu)/kT} \right]. \quad (2.24)$$

The prefactor, h^{-3} , is the volume in phase space⁸ per minimum uncertainty quantum state. For example, we can use the semiclassical approximation to determine the mean density,

$$n(x) = \frac{1}{h^3} \int d^3p \left(e^{(p^2/2m + V(x) - \mu)/kT} \pm 1 \right)^{-1}, \quad (2.25)$$

$$= \mp \lambda_T^{-3} Li_{3/2} \left[\mp e^{-(V(x) - \mu)/kT} \right] \quad (2.26)$$

⁷The entropy is $S = -\frac{\partial \Omega}{\partial T}$.

⁸Phase space refers to the combined \mathbf{x}, \mathbf{p} coordinate system.

where $Li_n[x] = \sum_{i=1}^{\infty} \frac{x^i}{i^n}$ is the polylogarithm function and

$$\lambda_T = \sqrt{\frac{h^2}{2\pi m k T}} \quad (2.27)$$

is the thermal de Broglie wavelength, which is an important defining characteristic of the system. First, the de Broglie wavelength defines the peak phase-space density

$$\rho_0 = n_0 \lambda_T^3, \quad (2.28)$$

where n_0 is the peak spatial density. This quantity, ρ_0 , roughly determines the number of particles in the quantum ground state, which in turn determines the relative importance of quantum statistics. Also, the de Broglie wavelength is a good indication of quantum motion, and subsequently when λ_T is smaller than the characteristic length scale of the potential (i.e., a_{ho} for the harmonic trap), the semiclassical approximation is valid.

Bosons and fermions display different behavior when $\rho_0 \rightarrow 1$ and quantum statistics become important. For bosons, the ground state becomes macroscopically occupied via a phase transition known as Bose-Einstein condensation (BEC), which occurs when $T < T_C$. For the 3D harmonic potential

$$T_C = \frac{0.94 \hbar \omega N^{1/3}}{k_B} \quad (2.29)$$

where N is the total number of particles. For fermions, the Pauli exclusion principle dictates that the maximum state occupancy is 1. As $T \rightarrow 0$, the N lowest energy states become occupied with exactly one particle, and the energy of the system is finite. The energy of the highest occupied state E_F is the Fermi energy. For a 3D harmonic potential,

$$E_F = \hbar \omega (6N)^{1/3}. \quad (2.30)$$

Since state occupancies start to approach 1 when $T < T_F = E_F/k_B$, below T_F the system is referred to as a degenerate Fermi gas (DFG). The low temperature thermodynamics can be written as functions of the reduced temperature, $\tilde{T} = T/T_C$ for bosons and $\tilde{T} = T/T_F$ for fermions. Some of these thermodynamic quantities for the 3D harmonic potential are to lowest order in \tilde{T} (top bosons, bottom fermions⁹),

$$\frac{S}{k_B N} = \begin{cases} 3.6 \tilde{T} \\ \pi^2 \tilde{T} \end{cases}, \quad (2.31)$$

$$\frac{U}{N} = \begin{cases} 2.7 k_B T_C \tilde{T}^4 \\ \frac{3}{4} E_F \left(1 + \frac{2\pi^2}{3} \tilde{T}^2\right) \end{cases}, \quad (2.32)$$

$$\frac{N_0}{N} = 1 - \tilde{T}^3. \quad (2.33)$$

where S is the entropy and U is the total energy. The last property is the condensate fraction, which is the fraction of atoms in the macroscopically occupied state, which applies only to the bosonic systems. BEC and DFG are not theoretical constructs — both have been observed with ultracold atoms: BEC in 1995 [22] and DFG in 1999 [121].

⁹For fermions these are the Sommerfeld expansions.

2.2.2 Weak Interactions

Since all real systems have and need interactions to satisfy the ergodicity that validates thermodynamics, the next level of theory is to consider weak interactions. The difficulty with interactions, and why they lead to unsolvable models, is that unlike for non-interacting systems, the wavefunctions and state occupancies are coupled. In general, we need to solve for the complete N-particle wavefunctions that must be properly (anti-)symmetrized for (fermions) bosons. For this problem, it is more convenient to use operators that act directly on the many-body state instead of on individual particles. These second-quantized operators have the exchange symmetries explicitly built-in. For example, the operator \mathbf{a}_i , acting on the many-body wavefunction, $|n_1, n_2, \dots, n_i, \dots\rangle$ removes a single particle from the system in mode i ,

$$\mathbf{a}_i |n_1, n_2, \dots, n_i, \dots\rangle = \sqrt{n_i} |n_1, n_2, \dots, n_i - 1, \dots\rangle. \quad (2.34)$$

The quantum statistics of the operators are based on the commutation relations for Bose operators [122]

$$[\mathbf{a}_i, \mathbf{a}_j^\dagger] = \delta_{ij}, \quad (2.35)$$

$$[\mathbf{a}_i^\dagger, \mathbf{a}_j^\dagger] = 0, \quad (2.36)$$

$$[\mathbf{a}_i, \mathbf{a}_j] = 0, \quad (2.37)$$

and anti-commutation relations for Fermi operators

$$\{\mathbf{a}_i, \mathbf{a}_j^\dagger\} = \delta_{ij}, \quad (2.38)$$

$$\{\mathbf{a}_i^\dagger, \mathbf{a}_j^\dagger\} = 0, \quad (2.39)$$

$$\{\mathbf{a}_i, \mathbf{a}_j\} = 0. \quad (2.40)$$

In this second-quantized notation, the grand canonical Hamiltonian is [123],

$$\begin{aligned} \mathbf{H} = & \sum_{\sigma} \int d^3x \Psi_{\sigma}^{\dagger}(\vec{x}) \left(-\frac{\hbar^2}{2m} \nabla^2 + V(\vec{x}) - \mu_{\sigma} \right) \Psi_{\sigma}(\vec{x}) + \\ & \sum_{\sigma, \sigma'} \int d^3x d^3x' \Psi_{\sigma}^{\dagger}(\vec{x}) \Psi_{\sigma'}^{\dagger}(\vec{x}') \left[g_{\sigma, \sigma'} \delta(\vec{x} - \vec{x}') \frac{\partial}{\partial |\vec{x} - \vec{x}'|} |\vec{x} - \vec{x}'| \right] \Psi_{\sigma}(\vec{x}) \Psi_{\sigma'}(\vec{x}') \end{aligned} \quad (2.41)$$

where σ indexes the spin, and $\Psi_{\sigma}(\vec{x}) = \sum_i \Psi_i(\vec{x}) \mathbf{a}_i$ is the field operator which removes a particle at position \vec{x} . The term inside the square brackets $[\dots]$ is the regularized interaction pseudo-potential Eqn. 2.15 and $V(\vec{x})$ is the single-particle potential. If we drop the regularization term in the interaction,

$$\begin{aligned} \mathbf{H} = & \sum_{\sigma} \int d^3x \Psi_{\sigma}^{\dagger}(\vec{x}) \left(-\frac{\hbar^2}{2m} \nabla^2 + V(\vec{x}) - \mu_{\sigma} \right) \Psi_{\sigma}(\vec{x}) + \\ & \frac{1}{2} \sum_{\sigma, \sigma'} g_{\sigma, \sigma'} \int d^3x \Psi_{\sigma}^{\dagger}(\vec{x}) \Psi_{\sigma'}^{\dagger}(\vec{x}) \Psi_{\sigma'}(\vec{x}) \Psi_{\sigma}(\vec{x}) \end{aligned} \quad (2.42)$$

Complete solutions to Eqns. 2.41 and 2.42 are generally unattainable, and therefore we utilize a number of approximations. Before delving into specifics, we would like to highlight two important approximations. The first approximation is self-consistent mean-field theory. Here we obtain an effective single particle Hamiltonian for the σ particles

$$\mathbf{H}_\sigma = -\frac{\hbar^2}{2m}\nabla^2 + V(\vec{x}) - \mu_\sigma + (f+1)g_\sigma n_\sigma(\vec{x}) + \sum_{\sigma' \neq \sigma} g_{\sigma,\sigma'} n_{\sigma'}(\vec{x}) \quad (2.43)$$

where $f = 1$ for bosons and $f = -1$ for fermions, representing the exchange term for the interaction and resulting in no interactions for identical fermions. This Hamiltonian is inadequate for a condensed system, as discussed next. The self-consistency condition is that we must solve Eqn. 2.43 to get the single particle eigenstates and eigenvalues, and then use Eqn. 2.23 to determine the density terms appearing in Eqn. 2.43, and iterate until the density converges. Another common approximation, useful for bosons and fermions, is to treat the potential as a slowly varying chemical potential. We then solve Eqn. 2.42 for a uniform system with chemical potential $\tilde{\mu}$, for example to determine the density $n[\tilde{\mu}]$. Then the spatial density variation is

$$n(\vec{x}) = n[\mu_0 - V(\vec{x})], \quad (2.44)$$

which is known as the local density approximation (LDA). In the following, we will look at some specific cases for single component bosons and two-component fermions.

For a condensed system of single-component bosons there is macroscopic occupation of the single-particle ground state $\Phi(\vec{x})$. Because of the large number of particles in this low energy mode, the condensate density is high and interactions become important. However, the presence of a dominant mode in the wavefunction simplifies our approach to solving Eqn. 2.42. Once we solve for $\Phi(\vec{x})$, we can expand Eqn. 2.42 around low-order $\Phi(\vec{x})$ terms. This standard method, and its various approximations, have been well-established over the course of numerous theoretical treatments and are summarized in e.g., [123–125]. The starting point for these solutions is to rewrite the field operator as

$$\Psi(\vec{x}) = \Phi(\vec{x}) + \tilde{\Psi}(\vec{x}) \quad (2.45)$$

where $\tilde{\Psi}(\vec{x})$ is the field operator for all atoms not in the condensate ($\langle \tilde{\Psi}(\vec{x}) \rangle = 0$). To solve for $\Phi(\vec{x})$ we substitute Eqn. 2.45 into Eqn. 2.42, and write out the Heisenberg equations of motion for the operator,

$$i\hbar \frac{\partial \Phi(\vec{x}, \mathbf{t})}{\partial \mathbf{t}} = \langle [\mathbf{H}, \Psi(\vec{x}, \mathbf{t})] \rangle \quad (2.46)$$

If we make the Hartree-Fock-Bogoliubov-Popov approximation¹⁰ [124], then the above reduces to the Gross-Pitaevskii equation [123]

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + V(\vec{x}, \mathbf{t}) - \mu + g(|\Phi(\vec{x}, \mathbf{t})|^2 + 2n_T(x)) \right] \Phi(\vec{x}, \mathbf{t}) = i\hbar \frac{\partial \Phi(\vec{x}, \mathbf{t})}{\partial \mathbf{t}} \quad (2.47)$$

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + V(\vec{x}) + g(|\Phi(\vec{x})|^2 + 2n_T(x)) \right] \Phi(\vec{x}) = \mu \Phi(\vec{x}) \quad (2.48)$$

¹⁰In this approximation $\langle \tilde{\Psi}(\vec{x}) \tilde{\Psi}(\vec{x}) \rangle = 0$ and $\langle \tilde{\Psi}^\dagger(\vec{x}) \tilde{\Psi}(\vec{x}) \tilde{\Psi}(\vec{x}) \rangle = 0$.

where $n_T(\vec{x}) = \langle \tilde{\Psi}^\dagger(\vec{x}) \tilde{\Psi}(\vec{x}) \rangle$ is the density of atoms not in the condensate¹¹; to zero-order n_T is given by the non-interacting semi-classical expression Eqn. 2.26. When $n_T \approx 0$ ($T \rightarrow 0$ and $na^3 \ll 1$) the condensate density for a uniform system is $|\Phi|^2 = \mu/g$. To solve for Φ in a 3D harmonic trap we use the Thomas-Fermi approximation (ignore kinetic energy i.e., $\nabla^2 \rightarrow 0$) and find

$$|\Phi(x)|^2 = \begin{cases} \frac{\mu - \frac{1}{2}m\omega^2 x^2}{g} & , \quad x \leq \sqrt{\frac{2\mu}{m\omega^2}} \\ 0 & , \quad x > \sqrt{\frac{2\mu}{m\omega^2}} \end{cases}. \quad (2.49)$$

We can determine μ as a function of total number by integrating the density,

$$\mu = \frac{15^{2/5}}{2} \left(\frac{N_0 a}{a_{ho}} \right)^{2/5} \hbar\omega, \quad (2.50)$$

and the condensate size (the Thomas-Fermi radius) is,

$$r_{TF}^2 = a_{ho}^2 15^{2/5} \left(\frac{Na}{a_{ho}} \right)^{2/5}. \quad (2.51)$$

The importance of taking interactions into consideration is highlighted by Eqn. 2.51. For typical experimental values ($N = 1 \times 10^5$, $a = 100a_0$, $\omega = 2\pi \times 50\text{Hz}$, $m = m_{Rb}$), the ratio of the Thomas-Fermi (TF) radius to the harmonic oscillator length a_{ho} (the ground state size in the non-interacting case) is approximately 14.




Figure 2.4: The semi-ideal model for condensed bosons. (a) The effective single particle potential (grey curve) for the non-condensate atoms for $N = 10^5$, $N_0/N = 0.47$, $\omega = 2\pi \times 50\text{Hz}$, $a = 100a_0$ and $m = m_{Rb}$. The non-condensate density is also shown (red curve) corresponding to $T/T_{C0} = 0.7$ ($T_{C0} = 104.75\text{nK}$). (b) The condensate fraction versus reduced temperature for the same parameters using the semi-ideal model (blue circles) versus the non-interacting gas (black curve).

The next-order improvement for calculating $n_T(\vec{x})$ at finite temperatures is to include the condensate density calculated from the GP equation with n_T set to zero (Eqn. 2.47), into the mean-field Hamiltonian

¹¹ $n_T \neq 0$ when $T = 0$ for finite interactions because of quantum depletion. However, for weakly interacting alkali gases in the harmonic trap this depletion is typically less than 1%.

for the non-condensate atoms (Eqn. 2.43),

$$\mathbf{H} = -\frac{\hbar^2}{2m}\nabla^2 + V(\vec{x}) - \mu + 2g|\Phi(\vec{x})|^2. \quad (2.52)$$

This is known as the semi-ideal model [123, 126, 127] and it takes into consideration the changes to the effective potential experienced by the non-condensate atoms due to the mean-field repulsion from the condensate. It does not take into account interaction between non-condensate atoms or describe sound modes (phonons) and therefore $n_T = 0$ when $T = 0$. However, the semi-ideal model accurately describes the interacting thermodynamics for the harmonically trapped gas. In the semi-ideal model, the thermal atoms experience a potential that is a harmonic trap at the edges, but harmonically repulsive in the middle due to interactions with the condensate (for $g > 0$). This potential and the thermal atom distribution is illustrated in Fig. 2.4. Solving this model using the semi-classical approximation and assuming a TF condensate distribution (Eqn. 2.49), the thermodynamic quantities are ($\alpha = \sqrt{\mu/kT}$)

$$\xi[j] = \frac{4\sqrt{j}\alpha}{\sqrt{\pi}} - e^{-j\alpha^2} \operatorname{erfi}(\sqrt{j}\alpha) + e^{j\alpha^2} \operatorname{erfc}(\sqrt{j}\alpha), \quad (2.53)$$

$$\Omega = kT \left(\frac{kT}{\hbar\omega} \right)^3 \sum_{j=1}^{\infty} \frac{\xi[j]}{j^4}, \quad (2.54)$$

$$N_T = \left(\frac{kT}{\hbar\omega} \right)^3 \sum_{j=1}^{\infty} \frac{\xi[j]}{j^3}, \quad (2.55)$$

$$U = \left(\frac{kT}{\hbar\omega} \right)^3 \sum_{j=1}^{\infty} \frac{\xi[j]}{j^3} \left(\frac{3kT}{j} - \mu \right), \quad (2.56)$$

$$S = -\frac{\Omega}{T} + \frac{U}{T} - \frac{\mu N_T}{T}, \quad (2.57)$$

where μ is determined by the number of condensate atoms using Eqn. 2.50. These equations must be solved self-consistently for fixed N and T . A primary effect of interactions between the condensate and the thermal component is a shift in T_C

$$T_C \approx T_{C,0} \left(1 - 1.33 \frac{a}{a_{ho}} N^{1/6} \right) \quad (2.58)$$

where $T_{C,0}$ is the non-interacting T_C (Eqn. 2.29). Using the same experimental parameters for calculating the Thomas-Fermi radius, this corresponds to an approximately 3% negative shift. The semi-ideal model also describes non-negligible interaction shifts in other thermodynamic quantities, for example the condensate fraction (see Fig. 2.4) and entropy. The interacting thermodynamics is well-described by the semi-ideal model; next-order effects such as Bogoliubov modes (phonons), interactions between thermal atoms and beyond mean field T_C shifts [128] have a negligible role on harmonic trap thermodynamics for our experimental parameter range.

There are a number of new complications when we consider solving Eqn. 2.41 for fermions instead of bosons. For one, the essential nature of fermions is that the occupancy of any state is ≤ 1 (Pauli exclusion principle), so the expansion around a single mode used for the bosonic solutions is generically not applicable. Also, interacting fermions require at minimum two components, so the parameter space is greatly increased. Solutions depend on the total and relative density of the two components. Furthermore, different compo-

nents may form bound pairs, which can significantly alter the nature of the system since two fermions can form a composite boson. An overview of the entire solution space is well beyond the scope of this work, so we restrict ourselves to a weakly interacting two-component mixture with $N_\uparrow = N_\downarrow$.

When the atoms are free and $a > 0$ ($a = a_{\uparrow,\downarrow}$), this is a repulsive gas which is well-described by the mean-field Hamiltonian Eqn. 2.43. As interactions increase beyond the weak limit considered here ($k_F|a| > 1$), it is conjectured that the repulsive gas is ferromagnetic due to the Stoner instability, however, the lifetime of the gas is very short [92]. Tightly bound pairs can form for $a > 0$ due to the bound state in the inter-atomic potential (see §2.1.2). These pairs act as effective bosons with a pair scattering length related to the free-fermion scattering length, $\tilde{a} = 0.6a$ [119]. The gas of paired particles is well-described by the semi-ideal model discussed for the Bose gas where $a = \tilde{a}$ and $m = 2m_F$.

When the gas is free and weakly attractive ($a < 0, k_F|a| \ll 1$), the bulk properties are described by the mean-field Hamiltonian with the added possibility of forming weak pairs near the Fermi surface. Below the temperature T_{BCS} [129]

$$T_{BCS} \approx 0.28T_F e^{-\frac{\pi}{2|k_F a|}} \quad (2.59)$$

the system transitions to a BCS-like superfluid [130].

One of the most active areas of study is fermions in the strongly interacting limit $|k_F a| > 1$ and in particular unitarity $|k_F a| \rightarrow \infty$. In these limits, more advanced techniques, e.g., Fermi-Liquid theory, are required to solve Eqn. 2.41. In some areas of the parameter space, there are competing, unproven theoretical solutions. Therefore, experiments in this limit are also quantum simulators, but not the subject of this thesis (see i.e., [129, 131]).

2.3 Ultracold Atoms in Optical Lattices

Once we create and characterize an ultracold gas in the harmonic trap, the next step towards quantum simulation is to apply an optical lattice. Even in the case of the non-interacting gas, we will see that the lattice has a large effect on the density of states. In particular, for a large enough lattice depth, the atoms in the lattice are well-described by a basis with wavefunctions localized to each lattice well. Expanding out the field operators in this basis and substituting into Eqn. 2.41 reveals the Hubbard model. However, this description is not perfect, and I will note several limitations to this simulation. Furthermore, I will discuss state-dependent lattices, which allow extensions to the simplest Hubbard model.

2.3.1 Non-Interacting Atoms in the Lattice: Band Structure

Single, non-interacting atoms in a uniform 3D cubic lattice are described by the Schrödinger equation

$$\left[-\frac{1}{k^2} \nabla^2 + s_x \sin^2(kx) + s_y \sin^2(ky) + s_z \sin^2(kz) \right] \phi(\vec{x}) = E \phi(\vec{x}), \quad (2.60)$$

where E is in units of E_R and $k = \pi/d$ (d is the lattice spacing). Since Eqn. 2.60 is separable along each direction, we will limit ourselves to solving the 1D problem. We can apply Bloch's theorem [8] because the




Figure 2.5: (a) As the lattice is applied ($s = 1$) the free particle dispersion (red curve) splits into bands (black curves). (b) The bands are most commonly represented in the reduced zone scheme. At $s = 10$, the bands have flattened considerably.

potential is periodic, giving

$$\phi_q(x) = e^{iqx/\hbar} u_q(x), \quad (2.61)$$

where q is the quasimomentum (a good quantum number in the lattice), and $u_q(x)$ has the same periodicity as the lattice. Substituting Eqn. 2.61 into Eqn. 2.60 we find

$$\left[-\frac{1}{k^2} \left(\frac{1}{i} \frac{d}{dx} + \frac{q}{\hbar} \right)^2 + s_x \sin^2(kx) \right] u_q(x) = E_q u_q(x). \quad (2.62)$$

We write out the potential and u_q in a Fourier series as

$$s_x \sin^2(kx) = \frac{s_x}{2} \left(1 + e^{-i(2k)x} + e^{i(2k)x} \right) \quad (2.63)$$

$$u_k(x) = \sum_j b_j e^{i(2kj)x}, \quad (2.64)$$

then the matrix form of Eqn. 2.62 is

$$\mathbf{C}\mathbf{b} = E_q \mathbf{b}, \quad (2.65)$$

where

$$C_{j,l} = \delta_{j,l} \frac{1}{k^2} \left[-(2jk)^2 + q(2jk) + q^2 \right] + \frac{s_x}{2} (\delta_{j,l} + \delta_{j,l-1} + \delta_{j,l+1}). \quad (2.66)$$

To solve these equations numerically, we truncate the Fourier series at some finite j , typically $|j| \leq 10$. For a given value of q , solving these equations gives a number of discrete energy levels the we call “bands” (labeled by n); the solutions are also periodic in $2q_B$ ($q_B = \hbar\pi/d$). The single non-periodic region of q ($-q_B < q < q_B$) is known as the Brillouin zone. The dispersion relations, E vs. q , for the lattice are shown in Fig. 2.5. As the lattice depth increases, the bands flatten since the particles have little kinetic energy via tunneling. At the very bottom of the band, the dispersion is still quadratic and the atoms have an effective

mass

$$m^* = \left(\frac{\partial^2 E}{\partial q^2} \right)^{-1}. \quad (2.67)$$

When the system is predominantly at the bottom of the band i.e., at very low temperatures, the lattice thermodynamics and dynamics can be obtained from the free-particle solutions by substituting m^* for m . Eventually, in the deep-lattice limit, the bands are separated by $\hbar\omega_{\text{latt}}$ where $\omega_{\text{latt}}^2 = 2sE_Rk^2m^{-1}$ is the trap frequency if we approximate the bottom of each lattice well as a harmonic oscillator. In this limit the lattice is essentially an array of isolated harmonic traps.

The change from a free particle dispersion to a band structure is the main impetus for adding the lattice. From the perspective of quantum simulation, the most important effect is that the maximum kinetic energy for a specific band is bounded and can be adjusted downward from the free particle result. This has an effect on thermalization and energy distribution among the degrees of freedom in the system. In particular, this is a parameter that adjusts the ratio between kinetic energy and interaction energy, which is essential for realizing the strongly correlated regime.

2.3.2 Tight Binding




Figure 2.6: (a) Ground band ($n = 0$) $q = 0$ Bloch wave ($\phi_{n=0,q=0}(x)$) for $s = 6$ (blue curve). The lattice potential is shown for spatial reference (grey curve). As expected the wavefunction is peaked at each lattice well. (b) The ground band Wannier state for the lattice site centered at $x = 0$ ($w_{n=0,j=0}(x)$).

The eigenstates of Eqn. 2.60 are $\phi_{q,n}(x)$, where q is the quasimomentum (a continuous variable in the infinite lattice) and n the discrete band index. Since these so-called Bloch waves can be decomposed into plane waves (Eqn. 2.64), the normalization condition is¹²,

$$\int dx \phi_{q,n}^*(x) \phi_{q',n'}(x) = \delta_{n,n'} \delta(q - q'). \quad (2.68)$$

¹²Delta-function normalization is a standard technique employed to handle the infinitely large lattice limit, however, it introduces issues with units (i.e., the units in Eqn. 2.61 and Eqn. 2.68 do not agree). An alternate, more well-defined procedure is to assume the particles are confined to a 1D box of length L (L/d is an integer) in which case the quasimomentum is quantized to $q_j = 2\pi j\hbar/L$ where j is an integer and the Bloch waves are $\phi_{q_j}(x) = e^{iq_j x/\hbar} u_{q_j}(x) L^{-1/2}$. The normalization condition is then $\int_0^L dx \phi_{q_j}^*(x) \phi_{q_{j'}}(x) = \delta_{j,j'}$, so these states are properly normalized.

A sample $q = 0$ Bloch wave is illustrated in Fig. 2.6. The Bloch waves form a complete basis set and so any arbitrary wavefunction $\psi(x)$ can be decomposed as

$$\psi(x) = \sum_n \int_{-q_B}^{q_B} dq \tilde{\psi}_n(q) \phi_{q,n}(x), \quad (2.69)$$

$$\tilde{\psi}_n(q) = \int dx \phi_{q,n}^*(x) \psi(x), \quad (2.70)$$

where $\tilde{\psi}_n(x)$ is the wavefunction in the quasimomentum-space of band n .

If $\psi(x)$ is spatially localized then the Bloch waves are not an optimal basis set; instead we can use a complementary basis localized around individual lattice sites. This Wannier basis is defined in terms of the Bloch waves as

$$w_{j,n}(x) = \int_{-q_B}^{q_B} dq e^{iqR_j/\hbar} \phi_{q,n}(x) \quad (2.71)$$

where j is the site index, n is the band index, R_j is the distance to site j , and $w_{j,n}(x)$ is normalized to unity. The ground band Wannier function is illustrated in Fig. 2.6. The generic wavefunction $\psi(x)$ expressed in terms of the Wannier basis is

$$\psi(x) = \sum_{j,n} \alpha_{j,n} w_{j,n}(x), \quad (2.72)$$

$$\alpha_{j,n} = \int dx w_{j,n}^*(x) \psi(x), \quad (2.73)$$

where $\alpha_{j,n}$ is the Wannier coefficient in band n at site j . Writing the quasimomentum wavefunction in terms of $\alpha_{j,n}$,

$$\tilde{\psi}_n(q) = \sum_j \alpha_{j,n} e^{iqR_j/\hbar}, \quad (2.74)$$

and is therefore given by the Fourier series of the Wannier coefficients. In the deep lattice limit the Wannier functions are connected to the harmonic oscillator eigenstates at each site.

The advantage of the Wannier functions is that they are a convenient basis for expressing the lattice Hamiltonian in second-quantized form. The field operator in the Wannier basis is

$$\Psi(x) = \sum_{j,n} w_{j,n}(x) \mathbf{a}_{j,n} \quad (2.75)$$

where $\mathbf{a}_{j,n}$ is the annihilation operator in band n and site j . Substituting Eqn. 2.75 into the non-interacting part of Eqn. 2.41 (interactions will be added in §2.3.4),

$$\begin{aligned} \mathbf{H} = & \sum_{i,j \neq i,n,n'} \mathbf{a}_{i,n}^\dagger \mathbf{a}_{j,n'} \int dx w_{i,n}^* \left[-\frac{\hbar^2}{2m} \frac{d^2}{dx^2} + s_x E_R \sin^2(kx) \right] w_{j,n'} + \\ & \sum_{i,n} \left[\left(\int_{-q_B}^{q_B} dq E_{q,n} \right) - \mu \right] \mathbf{a}_{i,n}^\dagger \mathbf{a}_{i,n} \end{aligned} \quad (2.76)$$

where $E_{q,n}$ are the band energies found by solving Eqn. 2.62. The above form of the Hamiltonian is not

particularly useful until we make two important approximations. First, we restrict ourselves to the lowest band (i.e., $n = 0$ only), since higher bands are typically frozen out in the ultracold regime¹³. Second, for sufficiently deep lattices, the Wannier states are strongly localized to a particular site, and so we only retain the most dominant term with i and j as neighboring sites — the “tight-binding” approximation, which is essentially valid for $s \gtrsim 4$. The validity of these approximations is further discussed in §2.3.5.

Making these approximations, the Hamiltonian simplifies to

$$\hat{H} = -t \sum_{\langle i,j \rangle} \mathbf{a}_i^\dagger \mathbf{a}_j - \mu \sum_i \hat{a}_i^\dagger \hat{a}_i, \text{ where} \quad (2.77)$$

$$t = \int dx w_j^*(x) \left[-\frac{\hbar^2}{2m} \frac{d^2}{dx^2} + s_x E_R \sin^2(kx) \right] w_{j+1}(x), \quad (2.78)$$

where $\langle i, j \rangle$ is a sum over neighboring sites, and the band energy offset $\int dq E_q$ is absorbed into μ . This form easily lends itself to generalization, since we can model any number of lattice geometries by controlling the sum. For example, for a nD lattice we sum over $\mathbf{z} = 2n$ neighboring sites, where \mathbf{z} is known as the coordination number. We can also generalize the form of t , which can differ along various directions or be dependent on ij for an inhomogeneous lattice. Physically, t is the tunneling energy, and if we start a particle in a single site with \mathbf{z} neighbors, it will completely transfer to the other sites in time $\hbar(t\sqrt{\mathbf{z}})^{-1}$. An approximate analytic form for t (Eqn. 2.78) as a function of s , exact in the $s \rightarrow \infty$ limit, is [132],

$$\frac{t}{E_R} = \frac{4}{\sqrt{\pi}} s^{3/4} e^{-2\sqrt{s}} \quad (2.79)$$

It is easy to see that the state,

$$|q\rangle = \sum_j e^{iqR_j/\hbar} \hat{a}_j^\dagger |0\rangle \quad (2.80)$$

is an eigenstate of Eqn. 2.77 by direct substitution. The eigenvalues give the tight-binding dispersion,

$$E = 2t \left[1 - \cos \left(\frac{q}{q_B} \pi \right) \right], \quad (2.81)$$

where $4t$ is the bandwidth. In this limit the effective mass is

$$m^* = \frac{q_B^2}{\pi^2 2t}. \quad (2.82)$$

The tight-binding form of the Hamiltonian, Eqn. 2.77, emphasizes physics local to each site, a feature which will be explored in the next two sections.

2.3.3 Combined Lattice-Harmonic Potential

While the infinitely large and perfect lattice potential is important conceptually, it is not physical. The potential will always have additional terms that shift the energy of different lattice sites. In any practical cold atoms experiment the lattice must have an overall confining potential to restrict the atoms to a finite region of the lattice. Fortunately, adding in additional potential terms to Eqn. 2.77 is straightforward using

¹³In subsequent references to the Wannier function $w_j(x) = w_{j,n=0}(x)$.

the Wannier basis. For an additional potential $V(x)$, the resulting single-band Hamiltonian is,

$$\mathbf{H} = -t \sum_{\langle i,j \rangle} \mathbf{a}_i^\dagger \mathbf{a}_j - \mu \sum_i \mathbf{a}_i^\dagger \mathbf{a}_i + \sum_{i,j} \left[\int dx w_i^*(x) V(x) w_j(x) \right] \mathbf{a}_i^\dagger \mathbf{a}_j. \quad (2.83)$$

If $V(x)$ is sufficiently smooth compared to the lattice then we only keep terms for $i = j$. Furthermore the integral simplifies,

$$\left[\int dx w_i^*(x) V(x) w_i(x) \right] \approx V(x_i) \quad (2.84)$$

to the potential evaluated at the center of site i . The tight-binding lattice Hamiltonian with a smooth potential is,

$$\mathbf{H} = -t \sum_{\langle i,j \rangle} \mathbf{a}_i^\dagger \mathbf{a}_j + \sum_i [V(x_i) - \mu] \mathbf{a}_i^\dagger \mathbf{a}_i. \quad (2.85)$$

Typically the required overall confinement is provided by a harmonic potential, and at minimum there is always some harmonic confinement from the lattice beams themselves (see §C.3.2), and so

$$V(x_i) = \frac{1}{2} m \omega^2 x_i^2. \quad (2.86)$$

In [133], analytic expressions are given for the eigenstates and eigenvalues for the combined lattice-harmonic potential in terms of the dimensionless parameter, $\alpha = 8t/m\omega^2 d^2$,

$$E_n = \begin{cases} \frac{m\omega^2 d^2}{8} a_n(\alpha), & n \text{ even} \\ \frac{m\omega^2 d^2}{8} b_{n+1}(\alpha), & n \text{ odd} \end{cases} \quad (2.87)$$

where $a_n(\alpha)$ and $b_{n+1}(\alpha)$ are the Mathieu characteristic values. The spectrum of energies for $\alpha = 2500$ is shown in Fig. 2.7. At low energies ($E \ll 4t$), the spectrum is characteristic of a harmonic oscillator with frequency $\omega^* = \omega \sqrt{M/M^*}$ because of the effective mass from by the lattice potential (see Eqn. 2.82). When the energy is greater than the bandwidth, the kinetic energy scale is frozen out. In the pure lattice potential there are no states with these energies because of the bandgap. In this regime, the energy is given by $\frac{1}{2} m \omega^2 x_n^2$, and there are two degenerate energies for particles at $\pm x_n$. In §5.1, we will discuss the thermodynamics of the combined lattice-harmonic potential.

2.3.4 Adding Interactions: The Hubbard Model

Up to this point we have left out inter-particle interactions, but adding them to Eqn. 2.85 is also straightforward using the Wannier basis. Using the field operators written in the terms of the Wannier basis¹⁴

¹⁴The interaction term is not separable, so we need to use the full 3D field operator.




Figure 2.7: Eigenvalues for the combined lattice-harmonic potential for $\alpha = 2500$ ($s \approx 6$, $\omega = 2\pi \times 20\text{Hz}$, m_{Rb}) in units of $\hbar\omega$ using the formulas from [133]. The grey dashed line indicates where $E = 4t$. For $E \lesssim 4t$, the system behaves as a harmonic oscillator. The energies in this region are highlighted in inset (a), and fit to $E \approx 0.68n\hbar\omega$ (without the lattice $E = n\hbar\omega$). The harmonic oscillator frequency in the lattice is reduced because the effective mass in the lattice is approximately twice the bare mass. For $E \gtrsim 4t$, the system behaves as particles isolated to individual wells sequentially further from the center of the trap. The energies in this region are highlighted in inset (b) and fit to $E = \frac{1}{2}m\omega^2 d^2(n/2)^2 + E_0$ where E_0 is an offset. At each value of n there are two degenerate values corresponding to particles at distances $\pm dn/2$ away from the center of the trap.

(Eqn. 2.75), we substitute into the interaction term of the second-quantized Hamiltonian (Eqn. 2.42),

$$\frac{1}{2} \sum_{\sigma, \sigma'} g_{\sigma, \sigma'} \sum_{\vec{i}, \vec{j}, \vec{k}, \vec{l}} \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{j, \sigma'}^\dagger \mathbf{a}_{\vec{k}, \sigma'} \mathbf{a}_{\vec{l}, \sigma} \left[\int d^3x w_{i, \sigma}^*(\vec{x}) w_{j, \sigma'}^*(\vec{x}) w_{\vec{k}, \sigma'}(\vec{x}) w_{\vec{l}, \sigma}(\vec{x}) \right] \quad (2.88)$$

$$w_{i, \sigma}(\vec{x}) = w_{i_1, \sigma}(x) w_{i_2, \sigma}(y) w_{i_3, \sigma}(z) \quad (2.89)$$

where we have kept the spin index σ and the site indices are vectors because the interaction integral is 3D. Since the Wannier function, $w_i(x)$, is localized to site i , the tight-binding approximation for interactions is to only consider interactions on the same lattice site i.e., $i = j = k = l$. Therefore the interaction term simplifies to

$$\frac{1}{2} \sum_{\sigma, \sigma', i} U_{\sigma, \sigma', i} \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{i, \sigma'}^\dagger \mathbf{a}_{i, \sigma'} \mathbf{a}_{i, \sigma}, \quad (2.90)$$

where the interaction energy is

$$U_{\sigma, \sigma', i} = g_{\sigma, \sigma'} \left[\int dx |w_{i, \sigma}(x)|^2 |w_{i, \sigma'}(x)|^2 \right]^3. \quad (2.91)$$

In the deep lattice limit, $U = \sqrt{8\pi}(a/d)s^{3/4}$ [132]. Adding the interaction term to Eqn. 2.85 (and generalizing for several possible spin-components),

$$\hat{H} = - \sum_{\langle i, j \rangle, \sigma} t_\sigma \hat{a}_{i, \sigma}^\dagger \hat{a}_{j, \sigma} + \sum_{i, \sigma} [V_\sigma(x_i) - \mu_\sigma] \hat{a}_{i, \sigma}^\dagger \hat{a}_{i, \sigma} + \frac{1}{2} \sum_{i, \sigma, \sigma'} U_{\sigma, \sigma'} \hat{a}_{i, \sigma}^\dagger \hat{a}_{i, \sigma'}^\dagger \hat{a}_{i, \sigma'} \hat{a}_{i, \sigma} \quad (2.92)$$

This is the Hubbard model that we introduced in Eqn. 1.3. In deriving this model, we assume particles are ultracold atoms. Therefore, as first discussed in [25], this demonstrates the possibility that ultracold atoms in an optical lattice are well-described by, and thus should simulate, the Hubbard model. In Fig. 2.8 we plot the Hubbard parameters t and U and the ratio U/t versus the lattice depth. As the lattice depth increases, the Wannier functions become more strongly localized to individual lattice sites. The tunneling energy t , which requires overlap between Wannier functions on adjacent sites, decreases exponentially and the interaction energy U increases. This plot illustrates the second key element to quantum simulation — the tunability of the model parameters.

In the rest of the thesis I focus on two specific instances of the Hubbard model. The first is the single-component Bose-Hubbard model,

$$\mathbf{H} = -t \sum_{\langle i, j \rangle} \mathbf{a}_i^\dagger \mathbf{a}_j + \sum_i [V(x_i) - \mu] \mathbf{n}_i + \frac{U}{2} \sum_i \mathbf{n}_i (\mathbf{n}_i - 1) \quad (2.93)$$

and the two-component Fermi-Hubbard model (labeled \uparrow and \downarrow),

$$\mathbf{H} = - \sum_{\langle i, j \rangle, \sigma} t_\sigma \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{j, \sigma} + \sum_{i, \sigma} [V_\sigma(x_i) - \mu_\sigma] \mathbf{n}_{i, \sigma} + U_{\uparrow, \downarrow} \sum_i \hat{n}_{i, \uparrow} \mathbf{n}_{i, \downarrow} \quad (2.94)$$




Figure 2.8: (a) Hubbard parameters, t and U in units of E_R for $d = 400\text{nm}$, $a = 100.4a_0$, and $m = m_{Rb}$ versus lattice depth. These were calculated numerically using Wannier functions created by solving Eqn. 2.65 and truncating at the 7th Fourier term. (b) Ratio of U/t versus lattice depth. The large increase in U/t is mainly due to the exponential decrease in t as the lattice depth increases.




Figure 2.9: Analyzing the limits of the tight-binding approximation. (a) Ratio of next-nearest neighbor tunneling, t' , to nearest-neighbor tunneling, t , versus lattice depth. Typically, the tight-binding limit is given as $s > 4$, at which point t' is -7% of t . (b) Ratio of nearest-neighbor U' to on-site interactions U versus lattice depth. Ignoring the U' term is a good approximation once interactions are non-negligible, $s \gtrsim 8$.

2.3.5 Limitations

Before discussing the Hubbard model in more detail, it is important to understand the limits of the approximations made during the derivation. These approximations mainly involved neglecting terms in the full Hamiltonian. If these terms are large, they can lead to new phases not described by Eqn. 2.92. For example, next-nearest-neighbor tunneling terms can realize frustrated magnetism for square lattices [134], and a model with nearest-neighbor interactions (the extended Bose-Hubbard model) exhibits a charge-density-wave phase [135]. First, we evaluate the relative strength of the next-nearest-neighbor tunneling term t' ,

$$t' = \int dx w_j^*(x) \left[-\frac{\hbar^2}{2m} \frac{d^2}{dx^2} + s_x E_R \sin^2(kx) \right] w_{j+2}(x). \quad (2.95)$$

A simple method to determine t' is to look at the band structure calculated numerically in §2.3.1. The dispersion with next-nearest neighbor tunneling is

$$E(q) = 2t \left[1 - \cos\left(\frac{q}{q_B} \pi\right) \right] + 2t' \left[1 - \cos\left(2\frac{q}{q_B} \pi\right) \right], \quad (2.96)$$

so

$$8t' = 2E\left(\frac{q_B}{2}\right) - E(q_B). \quad (2.97)$$

In Fig. 2.9 we plot t'/t versus the lattice depth. At the lowest lattice depth that we perform experiments $s = 2$, $|t'/t| = 14\%$. However, most of our experiments are performed with $s \geq 6$ where $|t'/t| < 4\%$. The next approximation was to only keep on-site, ground band terms of the confining potential. For a harmonic potential we can estimate the neglected terms, and they are always negligible.

For interactions the first approximation is that only on-site terms are important. From Eqn. 2.88 the first-order off-site term is,

$$U' = \frac{1}{2}g \left[\int dx w_{i+1}^*(x) w_i^*(x) w_i(x) w_i(x) \right] \left[\int dx |w_i(x)|^4 \right]^2. \quad (2.98)$$

In Fig. 2.9 we plot U'/U versus the lattice depth. At low lattice depths the ratio is non-negligible, however, at these depths the overall importance of interactions is small. When interactions start to influence the system for $s \gtrsim 8$, this term can safely be ignored. The other approximation we make for the interactions is that we do not consider the influence of higher bands, nor do we use the fully regularized pseudo-potential. These approximations are analyzed in detail in [136] and they find two important corrections to the Hubbard interaction term. First, the value for U is 10% less than expected from Eqn. 2.91 for the parameters of their experiment, which are similar to the parameters in our ^{87}Rb apparatus (§3.1). Second, there is an additional interaction term, $\frac{U_3}{6}n(n-1)(n-2)$, where U_3 is less than 10% of U . These corrections are important for precision comparisons with theory, but should not affect the qualitative properties of the model. Finally, we have also ignored the weak, long-range $(1/r^3)$ magnetic dipole-dipole interaction between atoms. This energy is on the order of $10^{-8}E_R$ for atoms in adjacent sites, and therefore has no effect on the system at the energy scales of interest.

One of the biggest constraints for simulating the Hubbard model is the validity of using a conservative Hamiltonian (Eqn. 2.42), to describe an optical potential, which has fundamental non-Hermitian processes

arising from spontaneous emission. In general, this necessitates a master equation approach as in [37]. However, the simpler approach is to treat spontaneous emission as a heat source on top of Eqn. 2.92. The heating rate is described in detail in §B.8.1 and §C.4. Whether this is a limitation depends on the energy and time scale of the phenomenon being simulated, however, in many cases these scales are not known. For known phases, e.g., the Mott insulator (see §2.4.1), we can estimate that the energy scale is U and the time scale is \hbar/t . This type of analysis is performed in detail in [36]. The general conclusion of this analysis is that heating is not a limitation to achieving the MI phase, as proven by its experimental realization. However, realizing magnetic phases will require very low heating rates and therefore large detunings from electronic resonances of at least several hundred nm. This is an important consideration for the ^{40}K experiment presented in this thesis (§3.2).

2.3.6 State-Dependent Lattices




Figure 2.10: 2D slice through the effective potential seen by harmonically trapped atoms in the state dependent lattice due to a mean-field interaction with the lattice atoms. Here the lattice atoms have a mean density of one atom per lattice site, $s = 10$, $d = 400\text{nm}$, $m = m_{Rb}$, and the scattering length between them is $100a_0$. The depth of the potential is plotted in E_R . Note that unlike the cubic optical lattice potential, there are no local minimum of the periodic interaction potential.

Inherent in Eqn. 2.92 is that the lattice and trapping potentials may be state-dependent. As mentioned in §2.1.1 and detailed in §C.2, such a lattice can be realized using optical lattices with polarization gradients. A limiting case is when one state experiences no lattice, which is experimentally explored in this thesis. Clearly this state cannot be described in the tight-binding formalism, so we need to generalize Eqn. 2.92. For the case of one component in the lattice (α) and one component in a harmonic trap (β),

$$\mathbf{H} = \mathbf{H}_{\text{Hubb}} + \mathbf{H}_{\text{HO}} + \mathbf{H}_{\text{int}} \quad (2.99)$$

where \mathbf{H}_{Hubb} is the Bose-Hubbard Hamiltonian for the α component (Eqn. 2.93), \mathbf{H}_{HO} is the second-quantized, interacting harmonic oscillator Hamiltonian for the β component, and \mathbf{H}_{int} is the mutual interaction between the α and β particles. To write out \mathbf{H}_{HO} and \mathbf{H}_{int} we first construct the 3D harmonic

oscillator field operator,

$$\Psi_{\text{HO}}(\vec{x}) = \sum_{\vec{n}} \psi_{n_1}(x) \psi_{n_2}(y) \psi_{n_3}(z) \mathbf{b}_{\vec{n}} \quad (2.100)$$

where n is a sum over the harmonic oscillator eigenstates $\psi_n(x)$ given by Eqn. 2.19, and $\mathbf{b}_{\vec{n}}$ is the annihilation operator for state $\vec{n} = \{n_1, n_2, n_3\}$. We substitute the harmonic oscillator field operator and the lattice field operator (Eqn. 2.75) into Eqn. 2.42 to get

$$\begin{aligned} \mathbf{H}_{\text{HO}} = & \sum_{\vec{n}} [\hbar\omega_{\beta}(n_1 + n_2 + n_3) - \mu_{\beta}] \mathbf{b}_{\vec{n}}^{\dagger} \mathbf{b}_{\vec{n}} + \\ & \frac{g_{\beta}}{2} \sum_{\vec{n}, \vec{\sigma}, \vec{p}, \vec{q}} \left[\int d^3x \psi_{\vec{n}}^*(\vec{x}) \psi_{\vec{\sigma}}^*(\vec{x}) \psi_{\vec{p}}(\vec{x}) \psi_{\vec{q}}(\vec{x}) \right] \mathbf{b}_{\vec{n}}^{\dagger} \mathbf{b}_{\vec{\sigma}}^{\dagger} \mathbf{b}_{\vec{p}} \mathbf{b}_{\vec{q}}, \end{aligned} \quad (2.101)$$

$$\mathbf{H}_{\text{int}} = g_{\alpha, \beta} \sum_{\vec{n}, \vec{\sigma}, i, j} \mathbf{b}_{\vec{n}}^{\dagger} \mathbf{b}_{\vec{\sigma}}^{\dagger} \mathbf{a}_i^{\dagger} \mathbf{a}_j \left[\int d^3x \psi_{\vec{n}}^*(\vec{x}) \psi_{\vec{\sigma}}^*(\vec{x}) w_i^*(\vec{x}) w_j(\vec{x}) \right] \quad (2.102)$$

where i, j is a sum over lattice sites, $w_i(\vec{x})$ is the 3D Wannier function at site i , and $g_{\alpha, \beta}$ is the contact interaction between particles α and β (Eqn. 2.15). The last term \mathbf{H}_{int} is the mutual interaction between the gases and can be simplified if we treat the lattice gas as a mean-field with site occupancies $\mathbf{n}_i = \langle \mathbf{n}_i \rangle = n_i$. In this case, the harmonically trapped atoms experience an effective lattice potential,

$$V_{\text{eff}}(\vec{x}) = g_{\alpha, \beta} \sum_i \langle n_i \rangle \left| w_0(\vec{x} - \vec{r}_i) \right|^2 \quad (2.103)$$

which is illustrated in Fig. 2.10 for a uniformly filled lattice.

The model given by Eqn. 2.99 has many interesting inherent features. However, our main goal is to use the harmonically trapped gas as a probe of the gas simulating the Hubbard model. See §3.1.2, §4.1.7, and §8 for further details.

2.4 Hubbard Model

The previous sections, §2.1–§2.3, illustrated on theoretical grounds that ultracold atoms in optical lattices are well-described by the Hubbard model with tunable parameters. However, to experimentally validate this conjecture we must compare the results of optical lattice experiments to the Hubbard model in parameter regimes where the model has known solutions. Therefore, in this section we will summarize these regimes and their properties.

Understanding the known properties of the Hubbard model is also important to identify the Hubbard model as a candidate to describe a relevant physical system. However, it is not sufficient that a system be just described by the Hubbard model, it must be described in an unsolved regime. These motivations will also be discussed in this section.

2.4.1 Bose-Hubbard Model

In the simplest version of the Hubbard model, the uniform single-component Bose-Hubbard model (Eqn. 2.93), there are three dimensionless parameters, t/U , μ/U and T/U . In the non-interacting limit ($U/t \rightarrow 0$ (§5.1)),

the system is qualitatively similar to a harmonically trapped non-interacting gas (§2.2.1) and is condensed for $T < T_C$ ($T_C = 5.591t$ for one-particle per site in 3D [39]). For a uniform lattice, the ground state is the zero quasimomentum state. In the Fock basis this state is represented as (for N particles and M sites)

$$|\Psi\rangle_{U=0} \propto \prod_j^N \left(\sum_i^M \mathbf{a}_i^\dagger \right) |0\rangle. \quad (2.104)$$

Each individual atom is maximally delocalized throughout the lattice to minimize kinetic energy. For large N, M this is nearly equivalent to writing a product of coherent states on each site ($\bar{n} = N/M$, the average number of particles per site),

$$|\Psi\rangle_{U=0} = \prod_i^M \left(e^{-\bar{n}/2} \sum_{j=0}^{\infty} \frac{\bar{n}^{j/2}}{\sqrt{j!}} |j\rangle_i \right) \quad (2.105)$$

This is a convenient form for calculation. Also, it clearly shows that the number fluctuations on a given site in the non-interacting state are Poissonian.

In the opposite limit $t/U \rightarrow 0$ (i.e., the atomic limit, §5.2), the Hamiltonian is separable on each site, and the ground state has an integer number of atoms per site

$$|\Psi\rangle_{t=0} = \prod_i |n\rangle_i \quad , \quad \text{where } \mu/U < n < \mu/U + 1. \quad (2.106)$$

There are no number fluctuations in this state and each atom is localized to minimize interaction energy. Importantly, these two states have very different momentum distributions

$$n(p) = |w(p)|^2 \sum_{j,k} e^{ipd/\hbar(j-k)} \langle \mathbf{a}_j^\dagger \mathbf{a}_k \rangle, \quad (2.107)$$

where $w(p)$ is the Fourier transform of the Wannier function, and j, k is a sum over lattice sites. For the non-interacting state, $\mathbf{a}_i |\Psi\rangle_{U=0} = \sqrt{\bar{n}} |\Psi\rangle_{U=0}$, and

$$n(p)_{U=0} = |w(p)|^2 \bar{n}^2 \sum_{j,k} e^{ipd/\hbar(j-k)}, \quad (2.108)$$

$$= |w(p)|^2 \bar{n}^2 M^2 \sum_n \delta_{p, 2q_B n}, \quad (2.109)$$

which has sharp diffraction peaks at $p = 2q_B n$ (n is an integer), where the peak heights decrease as the Wannier function envelope. In the atomic limit,

$$n(p)_{t=0} = |w(p)|^2 N, \quad (2.110)$$

and the momentum distribution is a featureless Wannier function. Since the sum over all sites in the $t = 0$ ground state adds coherently to create sharp diffraction peaks in the momentum distribution, this state is often referred to as having “phase coherence.”

The transition between these two limits characterizes the main feature of Hubbard models, the SF to Mott-insulator (MI) transition [28]. At $T = 0$, as t/U is tuned from $U = 0 \rightarrow t = 0$, there is a quantum phase




Figure 2.11: Phase diagram of the single component Bose-Hubbard model from site-decoupled mean-field theory (§5.3.2). (a) Condensate fraction $\langle \mathbf{a} \rangle^2 / \bar{n}$, versus μ/U and t/U calculated numerically using the SDMFT (§5.3.2). The lattice depth for $d = 400\text{nm}$, $a = 100a_0$, $m = m_{Rb}$ is shown on the top axis for reference. There are two main regimes: the superfluid (SF) and MI. The MI regime forms characteristic lobes corresponding to different integer number of atoms per site. The critical point for the transition with $\langle n \rangle = 1$ from QMC is shown as a white circle [39]. An inhomogeneous trapped system is well-described by the LDA and the effective chemical potential follows the grey short-dashed line from the trap center to the trap edge. (b) Condensate fraction versus T/U and t/U at $\mu = U/2$ calculated numerically using the SDMFT (§5.3.2). The critical temperature decreases as the system approaches the MI regime. For finite T in the MI regime, the system can be characterized by the number variance, $\langle \mathbf{n}^2 - \langle \mathbf{n} \rangle^2 \rangle / \langle \mathbf{n} \rangle$ (plotted in inset at $t/U = 0$), which increases sharply with temperature. Based on this, in [137] the melting temperature of the MI is given as $T = 0.2U$.

transition at finite t/U to a MI, a state with strongly reduced number fluctuations, that is incompressible ($\partial n/\partial \mu = 0$), and which has an excitation gap. The qualitative features of this transition, in 3D, are shown in Fig. 2.11 and generated using a mean-field approximation. Evident in Fig. 2.11 are two ways to approach the MI region from the SF. The first is to start in the SF with $\bar{n} = 1$ and then decrease t/U until the MI transition occurs at $(t/U)_C$. This transition belongs to the $(d+1)$ dimensional XY model universality class [28]. The other transition is by adding or subtracting atoms at fixed t/U . This transition is more generic and well-described by mean-field [39]. The MI transition in 3D was observed in a system of bosonic ultracold atoms in a lattice by measuring the reversible loss of sharp diffraction features in the momentum distribution and the emergence of an excitation gap [33]. This observation strongly confirms the conjecture that ultracold atoms simulate the Hubbard model. The MI transition has also been observed in 2D [138, 139] and 1D [139].

A defining characteristic of the MI region is reduced number fluctuations, and an integer number of atoms per site at $t/U = 0$ and $T = 0$. Reduced number fluctuations were detected indirectly through spin-changing collisions [140] and directly using high-resolution in-situ imaging [56–58]. Number fluctuations are also reduced in the SF state proximate to the MI regime. Equivalently, the condensate fraction at $T = 0$ decreases due to quantum depletion. Interactions also drive down T_C as $t/U \rightarrow 0$, until $T_C = 0$ at the MI boundary. In Fig. 2.11, we plot condensate fraction versus t/U and T/U . The reduction of condensate fraction was measured in 2D [66] and the T_C boundary in 3D [67]. In the MI region there is no defined phase transition at finite temperature; however, particle-hole excitations essentially destroy the MI domains at $t/U = 0$ for $T \gtrsim 0.2U$ [137].

For experiments with harmonic trapping, the system can be described using the LDA, which was introduced in §2.2.2. In the LDA, the local phase at radial distance r in the trap is given by the uniform phase for chemical potential $\mu = \mu_0 - \frac{1}{2}m\omega^2 r^2$ where μ_0 is the global chemical potential that determines the overall number of atoms. In Fig. 2.11, the LDA is represented by the vertical line through the phase diagram. The harmonic trap confers some advantages for quantum simulation. For example, the trap allows for a one-shot realization of numerous chemical potentials at fixed t/U . This change in μ vs. r also gives access to thermodynamic derivatives which can be used to construct equations of state [62]. However, the trap has disadvantages since measurements in momentum are a convolution of all trapped phases, which makes it difficult to identify individual phase boundaries. Additionally, measurements of specific critical points are impeded since only a small fraction of the system is at any particular chemical potential. At $t/U \approx 0$, the LDA predicts sharp steps in number per site going outward from the center of the trap. These “Mott shells” were measured indirectly from level shifts [141] and spin collisions [140], and directly using high-resolution imaging [56–58].

A number of theoretical techniques exist that describe the properties of the BH model with high precision. In 1D, exact numerics are possible using the density matrix renormalization group (DMRG) technique, which gives $(U/t)_C = 3.61(13)$ [142]. In higher dimensions, Quantum Monte Carlo (QMC) is numerically exact and can simulate the equilibrium properties of up to 10^6 particles in 3D [39] using current computational resources. In 2D, QMC calculations give $(U/t)_C = 16.75$ [39], and in 3D $(U/t)_C = 29.34(2)$ [143]. For most quantities, dynamical mean-field theory (DMFT) [144, 145] requires less computational resources than QMC and is nearly as accurate, particularly when fluctuations are not important. An even simpler technique is

static mean field theory, which is exact as $z \rightarrow \infty$ [28, 68] and predicts $(U/t)_C = 5.8z$. In 3D ($z = 6$) this is only 20% less than the QMC value (see Fig. 2.11). Close to and in the MI region, one can perform strong-coupling expansions [31, 146]. However, these expansions fail near the critical point [31]. In the SF regime, there is the established Bogoliubov theory; however, this clearly breaks down at stronger interactions as it fails to predict the SF-MI transition [147]. The biggest gap between theory and experiment, and where BH quantum simulation can make the biggest impact, is for time-dependent and out-of-equilibrium many-body physics. These properties cannot be computed using QMC. Instead, the only reliable theory is time-dependent Gutzwiller mean-field [54] and possible extensions to DMFT.

2.4.2 Fermi-Hubbard Model

There are two important differences when comparing the two-component Fermi-Hubbard model (Eqn. 2.94) to the single component Bose-Hubbard model discussed in the previous section. The first difference is that the added spin degree-of-freedom opens up an additional axis in parameter space and allows for a number of new “magnetic” phases. Second, due to Fermi statistics, the basis per site is restricted to

$$|0\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle \quad (2.111)$$

in the single-band limit. This is most evident in the non-interacting case. Here the spins are decoupled, but there is no straightforward ground state as there is for bosons (Eqn. 2.104) because the particles fill up quasimomentum states to the Fermi energy. For a single spin component, if there are N particles and M sites ($N \leq M$), then the ground state is

$$|\Psi\rangle \propto \sum_{j=1}^{M!/(M-N)!N!} |f_j\rangle e^{i\phi_j} \prod_k^M (\mathbf{a}_k^\dagger)^{n_{k,j}} |0\rangle, \quad (2.112)$$

where the index j sums over the set $\{n_{j,1}, \dots, n_{j,M}\}$ of all site occupancy permutations adding up to N particles, and k is the site index. For each arrangement j there is an amplitude $|f_j|$ and a phase ϕ_j , which are complicated functions of the specific permutation. For example, the ground state for $N = 2$ and $M = 3$ (in a 1D chain) is (for one spin-component)¹⁵

$$|\Psi\rangle = \frac{1}{2} \left(|110\rangle + |011\rangle + \sqrt{2}|101\rangle \right). \quad (2.113)$$

The Fermi statistics are captured by the basis and unlike the boson ground state for $N = 2$ and $M = 3$, there are no terms where a site is doubly occupied (e.g., $|200\rangle$).

If $N = M$ there is only one possible arrangement,

$$|\Psi\rangle = \prod_k^M \hat{a}_k^\dagger |0\rangle \quad (2.114)$$

This is known as a band insulator (BI), since the state is insulating by virtue of all quasimomentum states being filled. In our grand-canonical model, this occurs when $\mu > 4Dt$. It follows that number fluctuations

¹⁵Using the convention that $|n_1, n_2, n_3\rangle = (\mathbf{a}_1^\dagger)^{n_1} (\mathbf{a}_2^\dagger)^{n_2} (\mathbf{a}_3^\dagger)^{n_3} |0\rangle$.

are also reduced as the state approaches the BI. For fermions $\hat{n}^2 = \hat{n}$, so

$$\frac{\langle \mathbf{n}^2 - \langle \mathbf{n} \rangle^2 \rangle}{\langle \mathbf{n} \rangle} = 1 - \langle \mathbf{n} \rangle, \quad (2.115)$$

$$= 1 - \frac{N}{M}. \quad (2.116)$$

For low fillings the number fluctuations are Poissonian, identical to the ground state of the BH model, but for high fillings the number fluctuations vanish as the system becomes a BI. This is driven by Fermi statistics and not interactions. In the $t = 0$ limit, the FH model separates into individual sites and the ground state is an integer number of atoms per site. Unlike the BH model, the number of states is truncated by Fermi statistics (Eqn. 2.111). The phase diagram (for $t = 0$) in the parameter space μ_\uparrow/U and μ_\downarrow/U at $T = 0$ is shown in Fig. 2.12. There are two phases — the BI and the MI. Along the line $\mu_\uparrow = \mu_\downarrow = \mu, \mu < U$ the ground state is degenerate. This is known as half-filling because $\langle \mathbf{n}_\uparrow \rangle = \langle \mathbf{n}_\downarrow \rangle = 1/2$.

Similar to the BH model, going between the non-interacting and atomic limits involves a metal to insulator transition at (in mean-field) $U/t \approx \sqrt{2}5$ [148], characterized by a vanishing compressibility and a reduction in the number of doubly occupied sites. However, unlike the Bose system there is not an abrupt change in the phase coherence since the system starts in a Fermi-liquid. Using measurements of compressibility [34] and double occupancy [35], the MI phase has been observed for fermions in optical lattices with $n_\uparrow = n_\downarrow$. These measurements confirm that the Hubbard model description extends to fermionic ultracold atoms in optical lattices. To get a qualitative sense of the phase diagram, results of an exactly diagonalized 5-site system (§5.4) are shown in Fig. 2.12.

The spin degree-of-freedom in the FH model means that there can also be spin correlations (i.e., magnetization) in the system. In particular, as interactions reduce number fluctuations deep in the MI phase, spin is the only free parameter. At half-filling and $\mu_\uparrow = \mu_\downarrow$, the FH model reduces to a spin-1/2 antiferromagnetic Heisenberg model (see, e.g., [149])

$$\mathbf{H} = \frac{J}{2} \sum_{\langle i,j \rangle} \vec{\mathbf{S}}_i \cdot \vec{\mathbf{S}}_j, \quad (2.117)$$

where $J = 4t^2/U$ is the superexchange energy. The origin of this term is that a small admixture of doubly occupied sites can reduce the kinetic energy. However, no tunneling is possible unless opposite spins are on adjacent sites. A similar reduction can be made for the two-component Bose-Hubbard model with one atom per site. In that case, $J < 0$, which is the ferromagnetic Heisenberg model. Eqn. 2.117 means that the absolute ground state as $t/U \rightarrow 0$ is putatively antiferromagnetic. As the system is cooled, the onset of magnetic ordering occurs at the Néel temperature T_N , which is $T_N = zJ/4k_B$ in the simplest mean-field theory. Although Eqn. 2.117 is only valid deep in the MI regime when $t/U \ll 1$, antiferromagnetic magnetic correlations develop for all t/U , however, T_N is very small when t/U is large [150]. Magnetic correlations are plotted at $T = 0$ for the 5-site model in Fig. 2.12. It is instructive to look at the ground state of the two-site antiferromagnetic Heisenberg model,

$$|\Psi\rangle = \frac{1}{\sqrt{2}} (|\uparrow, \downarrow\rangle - |\downarrow, \uparrow\rangle) \quad (2.118)$$

which is a singlet. Detecting these correlations is an outstanding experimental issue [151]. The entropy




Figure 2.12: FH Phase diagram at $T = 0$. (a) Atomic limit ($t = 0$) phase diagram as a function of the relative chemical potentials. (b-d) Qualitative 2D FH phase diagrams for $\mu_\uparrow = \mu_\downarrow = \mu$ via exact diagonalization of one site connected to 4 adjacent sites. (b) Average number in the middle site versus t/U and μ/U . (c) Number fluctuations $\langle n^2 - \langle n \rangle^2 \rangle / \langle n \rangle$. There are two main regions with no number fluctuations. At the top is the band insulator (BI), which is a trivial non-interacting phase. At the lower-left edge is the Mott insulator (MI), where number fluctuations are reduced due to interactions. (d) Magnetic correlations $\sum_{i=2}^5 \langle \vec{S}_1 \cdot \vec{S}_i \rangle$, where $\vec{S} = 1/2 \left[\left(\mathbf{a}_\uparrow^\dagger \mathbf{a}_\downarrow + \mathbf{a}_\downarrow^\dagger \mathbf{a}_\uparrow \right), -i \left(\mathbf{a}_\uparrow^\dagger \mathbf{a}_\downarrow - \mathbf{a}_\downarrow^\dagger \mathbf{a}_\uparrow \right), \left(\mathbf{a}_\uparrow^\dagger \mathbf{a}_\uparrow - \mathbf{a}_\downarrow^\dagger \mathbf{a}_\downarrow \right) \right]$. Antiferromagnetic (AFM) correlations develop on the approach to the MI phase.

required for Néel ordering is also very low, $S/Nk_B \approx \ln(2)/2$, which is around a factor of two lower than current experiments [152].

Theoretically, the FH model is less tractable to solve than the BH model because of the fermion sign problem [153], which greatly limits most QMC protocols. Exceptions to this include QMC at half-filling [154], determinantal QMC on finite lattices [155], and proof-of-principle calculations using diagrammatic Monte Carlo [156]. Instead, there are a number of approximate methods, such as dynamical mean-field theory (DMFT) [153], which is exact as $t \rightarrow 0$, $U \rightarrow 0$, or $z \rightarrow \infty$ and is accurate for temperatures $T \gtrsim t$. To better account for fluctuations, DMFT can be extended using a cluster approximation (e.g., [157]). Another technique, which is highly accurate so long as it converges, are high-temperature series expansions [158]. DMFT and high-temperature series expansions are compared in [159]. Exact diagonalization is possible on 16 site clusters [160]. Together with the LDA, these methods can describe a number of properties of trapped systems.

For the uniform 1D system there is an exact solution [161], which predicts a MI for $U \neq 0$. However, for the harmonically trapped case [162], QMC calculations indicate a local MI phase transition at $U/t \approx 3$. In higher dimensions the transition is determined to be $U/t \approx 15$ [163] (3D uniform), $U/t \approx 4.8\sqrt{z}$ [148] (uniform) and $U/t \approx 12$ [159] (3D trapped). There are a number of calculations of magnetic ordering. In the Heisenberg limit in 3D ($U/t \gtrsim 12$), high temperature expansions give $k_B T_N = 0.957J$ [154]. Simulating the Néel temperature for all t/U in 3D using QMC (at half-filling) there is a maximum in $k_B T_N = 0.33t$ at intermediate coupling $U/t = 8$ [152, 154]. The more relevant thermodynamic quantity for experiments is the entropy at T_N , which is $S/Nk_B = \ln(2)$ for mean-field, but quantum fluctuations reduce this to $\ln(2)/2$ [152]. In 2D, QMC gives a similar value, $S/Nk_B \approx 0.4$ [155]. Experimentally, the whole trapped system does not necessarily need to be as low as S_N since phases at the trap edge can hold entropy.

Away from half-filling and at low temperatures, there is considerable uncertainty regarding the FH phase diagram. Providing definitive answers in this region is one of the main goals of quantum simulation using optical lattices. The results of doping are hard to predict. For example, although the half-filled ground state is always antiferromagnetic, removing just one particle causes the system to be ferromagnetic as $U \rightarrow \infty$ [164]. The most controversial conjecture is that competition between antiferromagnetism and delocalization leads to the emergence of a d-wave superfluid [165] at very low temperatures. Some temperature estimates are $0.01t/k_B$ [40] and $0.02t/k_B$ for 10% doping in 2D and $U/t = 4$ [166]. The existence of this state in the FH model has not been proven [160].

2.4.3 Physical Motivations

The main goal of quantum simulation as laid out in Fig. 1.2 is to determine the minimal model describing important physical systems and to “solve” that model. Therefore, ultracold atoms in optical lattices only qualify as a true quantum simulator if the Hubbard model is a candidate model for a class of physical systems. To be a candidate, the known properties of the Hubbard model, discussed in the last two sections, must agree with experimental data from these systems. For the case of the Bose-Hubbard model, this would entail a system with bosonic quantum particles that undergoes a superfluid to insulator transition. Examples include granular superconductors [29], Josephson-junction arrays [30], flux lattices in type-II superconduc-

tors [31] and ^4He in porous media [28].

The Fermi-Hubbard model is more widely applicable since the quantum particles of almost all materials are electrons. Therefore, the FH model is the paradigm for strongly correlated electron systems where electrons are thought to be highly localized, e.g., due to the particular band structure of the material. For example, DMFT solutions of the Hubbard model match well to experimental data from transition metal oxides, such as V_2O_3 [153]. The strongest motivation is that the FH model is a candidate model for cuprate (high T_C) superconductors, such as $La_{2-x}Sr_xCuO_4$ [12]. The phase diagram of these materials, e.g., Fig. 1 of [12], is an insulating antiferromagnet at half-filling and a d-wave superconductor when doped. In the known regimes of the FH model, these phase diagrams are very similar, and the emergence of d-wave superconductivity agrees with some, but not all, calculations of the Hubbard model. This is a clear motivation for pursuing quantum simulation of the FH model.

Another motivation for simulating these models is to validate different theoretical techniques and establish their regimes of validity. Since the Hubbard model is among the simplest non-trivial models with only a few parameters, it is an ideal point of comparison between experiment and theory.

Chapter 3

Experimental Toolkit — Tools

To realize a Hubbard model quantum simulator using ultracold atoms in optical lattices (see overview §2.1), we need to possess an experimental toolkit. This toolkit is comprised of three main parts: tools, probes and techniques. Tools are the set of physical elements, “the apparatus”, required to create an ultracold gas in an optical lattice. Probes and techniques are the elements and procedures that then initialize, measure and perturb the gas to obtain relevant simulation results. These are detailed in §4. Using a computer analogy — tools are the physical components inside the box (e.g., the motherboard, memory, and processor), and probes and techniques are the interface devices (e.g., the monitor, mouse, and keyboard) and software.

In this chapter, I will detail the tools involved in this thesis work. Indeed, a significant portion of the time required to complete this thesis was spent developing a number of these tools. Of course, many tools were developed by prior and current colleagues. Where necessary, I will include these tools, with appropriate credit, for completeness. The work done in this thesis covers apparatuses in two separate labs: a Bose-Hubbard simulator at the University of Illinois using ^{87}Rb atoms, and a Fermi-Hubbard simulator at the University of Toronto using ^{40}K atoms.

3.1 ^{87}Rb Bose-Hubbard Apparatus

The first apparatus, at the University of Illinois, is designed to create ultracold bosonic gases of ^{87}Rb in an optical lattice in order to simulate Bose-Hubbard physics. This apparatus was primarily constructed by post-doctoral associate Hong Gao and graduate student Matt White. A detailed overview of the system is found in Matt White’s thesis [41]; a schematic of the apparatus is shown in Fig. 3.1. To summarize Ref. [41], atoms are collected from a Rb vapor in an ultrahigh vacuum glass cell using a magneto-optical trap (MOT). The MOT loads $> 10^9$ ^{87}Rb atoms in 30 seconds. The atoms are transferred into a magnetic quadrupole trap (see §2.1.1) in the state $|F = 1, m_F = -1\rangle = |1, -1\rangle$ (see Fig. A.2). The “cart quadrupole” trap is on a moveable cart which mechanically carries the atoms through a differential pumping tube into another glass cell that has orders of magnitude better vacuum. This is the science cell shown in Fig. 3.1.

Once in the science cell, the specific details differ depending on the chronological order of the projects, since the experiment was partly reconfigured in Fall 2007. Prior to this date, atoms were transferred from the cart quadrupole to a Ioffe-Pritchard (IP) magnetic trap (§2.1.1), which was fixed in reference to the vacuum system. After transfer, the cart would return to the MOT cell. Evaporative cooling, using radiofrequency (RF) magnetic Zeeman transitions, was used to cool the gas in the IP trap to the ultracold regime. In Fall 2007, the IP trap irreversibly broke, so a new procedure was implemented by adding a dipole trap. Now




Figure 3.1: (a) Illustration of the ^{87}Rb Bose-Hubbard apparatus. Left to right: the cart quadrupole magnetic coils (i), the MOT cell (ii), the transport tube (iii), and the science cell (iv). The cart quadrupole coils are on a mechanical cart and move from the MOT cell to the science cell. The science cell region is shown in detail illustrating the pinch quadrupole coils (v), the dipole trap beams (black), and the lattice beams (red). Also indicated are the coordinate system which will be used throughout. The crossing beam of the dipole trap is not always used. A set of bias coils provide a uniform field along the y axis (not shown). The main imaging beam travels along the x axis. (b) A schematic of the lattice path is shown with an optional quarter-waveplate for generating spin-dependent lattices. An EOM may be added in the future for dynamic polarization adjustment.

experiments use partial RF evaporation in the cart quadrupole, then transfer the atoms to a fixed quadrupole (“pinch quadrupole”). Once atoms are in the pinch quadrupole the cart returns to the MOT. Further RF evaporation is done in the pinch quadrupole until Majorana losses become severe. At this point, the atoms are transferred into a hybrid optical-magnetic trap. This trap consists of a laser beam¹ focused to a waist of $120\mu\text{m}$ located $100\mu\text{m}$ below the center of the quadrupole. This beam forms a dipole trap (§2.1.1, §C.3.1) with strong radial confinement. The pinch quadrupole provides a linear potential along \hat{z} and additional harmonic confinement along y and x . Efficient evaporation is achieved by increasing the quadrupole field and pushing atoms out of the dipole trap [102, 103], as shown in Fig. 2.26 of [41]. After this stage of evaporation the gas is close to T_C .

The final stage of evaporation depends on the particular type of experiment. For experiments that require a single spin component ($|1, -1\rangle$), the quadrupole remains on, but is reduced until the force from the magnetic vertical gradient is equal and opposite to the gravitational force². Axial confinement is provided by the magnetic quadrupole. Evaporation proceeds to low temperatures by decreasing the dipole power; high energy atoms escape along the vertical direction. For experiments requiring multiple spin states there can be no magnetic field gradients present once we create a state mixture. Although we create mixtures at the end of evaporation (§3.1.1), we find it optimal to complete the last stage of evaporation in the final trap configuration. Therefore, the pinch quadrupole is turned completely off before the last evaporation stage, and a 10G uniform magnetic field along \hat{y} is provided by a set of bias coils. Since there is no longer axial quadrupole confinement, the dipole beam is recycled back into the experiment (see Fig. 3.1 and Fig. C.7) and crossed with the initial beam to form a crossed-dipole (‘x-dipole’) trap; see §C.7 for the alignment procedure. Although the crossing beam is only needed in this last stage, there is no separate control of the crossing beam power, so in these experiments the dipole trap is always crossed. To remove the quadrupole field, it is ramped down in 1s and the dipole power is simultaneously ramped up to 4W. Evaporation proceeds by reducing the dipole power after which the state mixture is created. From MOT to the ultracold regime, the cycle takes $\approx 70\text{s}$. Both traps, the hybrid-dipole and x-dipole, are characterized by three independent trap frequencies that are measured empirically (§D for the measurement procedure). Typically, the mean trap frequency is $\bar{\nu} = (40 - 50)\text{Hz}$.

3.1.1 Microwaves for State Preparation

For experiments requiring a state mixture (i.e., spin-dependent lattice experiments), the next step is to prepare this mixture starting from the single component $|1, -1\rangle$ gas. Since loss rates are higher for atoms in $F = 2$ (due to relaxation back to $F = 1$), we almost always utilize mixtures of $|1, -1\rangle$ and $|1, 0\rangle$. In a 10G field, the separation between these states is $\Delta E_{|1, -1\rangle, |1, 0\rangle} = h \times 7\text{MHz}$. However, in the low B-field regime we cannot use RF magnetic fields to transfer directly to $|1, 0\rangle$ because the energy splitting between $|1, 0\rangle$ and $|1, 1\rangle$ ($\Delta E_{|1, 0\rangle, |1, 1\rangle}$), is the same as $\Delta E_{|1, -1\rangle, |1, 0\rangle}$. Therefore, there are two main methods for creating the mixture: RF transitions at high-field or 2-step transitions using microwaves. A resolvable RF transition

¹The laser source is an IPG Photonics YLR-10-1064-LP fiber laser (1064nm, linearly polarized, 10W maximum power). The beam optics are illustrated in Fig. 2.25 of [41].

²The magnetic quadrupole gradient is slightly less than gravity so that evaporating atoms may escape along the vertical direction.




Figure 3.2: (a) State diagram illustrating the two microwave transitions we use to prepare a $|1, -1\rangle/|1, 0\rangle$ mixture. (b) Probability to transfer atoms from $|1, -1\rangle$ to $|1, 0\rangle$ versus the microwave power attenuation during the first sweep from $|1, -1\rangle$ to $|2, 0\rangle$ (blue, dashed arrow in (a)). This sweep is 20ms long and 50kHz wide with a center frequency of 6827.65MHz. The second sweep (red, dashed arrow in (a)) is a full power sweep 5ms long and 60kHz wide sweep with a center frequency of 6834.7MHz. The line is a free-parameter fit using the Landau-Zener formula (Eqn. B.38). To measure the transfer probability we image atoms in $|1, 0\rangle$ and $|1, -1\rangle$ after time-of-flight Stern-Gerlach separation.

can be made at higher B-fields due to the quadratic Zeeman shift (Eqn. A.13)

$$\Delta E_{|1,-1\rangle,|1,0\rangle} - \Delta E_{|1,0\rangle,|1,1\rangle} \approx h \times (0.144B^2) \text{ kHz/G}^2; \quad (3.1)$$

at $B = 40 \text{ G}$ the difference in transition frequencies is 230kHz. This allows population transfer using RF to $|1, 0\rangle$ without also transferring to $|1, 1\rangle$. However, we prefer to use a 2-step transition using microwaves (see §6.7 of [41] for details of the source), as shown in Fig. 3.2. These transitions are accomplished by sweeping the frequency of the source through resonance; for details see §B.4. To control the transfer probability, we change the microwave power for the $|1, -1\rangle \rightarrow |2, 0\rangle$ transition (see Fig. 3.2). Maximum power is used for the $|2, 0\rangle \rightarrow |1, 0\rangle$ transition so that all the atoms are transferred out of $|2, 0\rangle$ and the result is a mixture of only $|1, 0\rangle$ and $|1, -1\rangle$ atoms.

Before creating a state mixture, the cloud of $|1, -1\rangle$ atoms is condensed and in thermal equilibrium with N atoms at temperature T . The transfer process results in two clouds with atom numbers $N_{|1,-1\rangle}, N_{|1,0\rangle} < N$ where condensate fraction (N_0/N) and T are approximately preserved. Since T_C is proportional to $N^{1/3}$, the T_C of each component is less than $T_{C,i}$ (the transition temperature before the transfer), and so the system after transfer is out-of-equilibrium. Therefore, to allow the system to rethermalize we wait 100ms after the transfer before proceeding with the experiment. Since our motivation is to use the $|1, 0\rangle$ component as an impurity, $N_{|1,0\rangle}$ is often much less than N and so $T_{C,|1,0\rangle}$ can be small and sometimes even less than T . In this case, we have observed the impurity condensate “melt”; data is shown in Fig. 5 of Ref. [4].

To simultaneously image different spin components at separate locations in the camera plane, we apply a magnetic field gradient along the \hat{z} axis during time-of-flight expansion (“Stern-Gerlach field”). This

gradient ($\partial B_z/\partial z$) provides a spin-dependent force

$$\vec{F}_{SG} = (\mu_B g_F m_F) \frac{\partial B_z}{\partial z} \hat{z}, \quad (3.2)$$

which separates the spin components after a sufficiently long time-of-flight. To generate the Stern-Gerlach field we have a single coil directly underneath the science cell (“Stern-Gerlach coil”). This coil is not shown in Fig. 3.1.

3.1.2 Lattice




Figure 3.3: (a) Typical experimental sequence for the optical lattice potential depth (Eqn. 3.3) with $\tau = 200\text{ms}$ and $t_0 = 100\text{ms}$. At the end of the experiment we can turn the lattice off as fast as possible ($< 100\text{ns}$, “snap-off”, blue curve), linearly in $100\mu\text{s}$ – 1ms (“bandmapping”, red curve), or linearly in 10s of ms (“adiabatic”, grey curve). (b) Schematic of atomic diffraction from a “snapped-off” lattice (grey spheres). The imaging beam (red) projects the atomic density integrated along \hat{x} onto the camera plane (y - z), resulting in overlapping diffraction peaks. An experimental absorption image for an $s = 10$ lattice is shown.

Once the ultracold gas has been created and is in the appropriate state mixture the lattice is superimposed on the gas. The lattice depth along each direction $s(\mathbf{t})E_R$, is increased according to (see Fig. 3.3)

$$s(\mathbf{t}) = \begin{cases} 0, & \mathbf{t} < 0 \\ s_0 \frac{e^{\mathbf{t}/\tau} - 1}{e^{t_0/\tau} - 1}, & 0 \leq \mathbf{t} < t_0 \\ s_0, & \mathbf{t} \geq t_0 \end{cases} \quad (3.3)$$

where typically $\tau = 200\text{ms}$ and $t_0 = 100\text{ms}$. These values were selected heuristically by optimizing condensate fraction after turning on and then slowly off the lattice. In general, the issue of optimal ramping is unresolved, but our timescales agree well with those used in similar experiments (see, e.g., [167]).

The lattice itself is comprised of three retro-reflected laser beams (see Fig. 3.1) with orthogonal wavevectors. All the beams are split off from the same laser source³ and coupled into three separate fibers to ensure sufficient mode quality and pointing stability (see Fig. 4.5 of [41]). The direction of each forward traveling

³The source is a Verdi pumped Ti:Sapphire laser with $\approx 1\text{W}$ output.




Figure 3.4: (a) Simplified electronic state diagram of ^{87}Rb illustrating the relevant detunings Δ for the lattice. (b) Intensity of the retro-reflected lin-lin and lin- \perp -lin lattice for a forward beam intensity of 50W/cm^2 (i.e., approximately 5mW focused to an $80\mu\text{m}$ beam waist). (c) Potential depth for the retro-reflected lin-lin lattice at $\lambda = 812\text{nm}$ and lin- \perp -lin lattice at $\lambda = 790\text{nm}$ for $|1, -1\rangle$ and $|1, 0\rangle$ and for forward beam intensity 50W/cm^2 . The lin-lin potential is purely attractive and state-independent, whereas the lin- \perp -lin lattice for $|1, -1\rangle$ (with $\hat{k} \cdot \hat{B} = 1$) is a combination of attractive and repulsive. The $|1, 0\rangle$ state experiences no lin- \perp -lin potential.

beam, using the coordinate system defined in Fig. 3.1, is

$$\begin{aligned}\hat{k}_1 &= -\frac{1}{\sqrt{2}}\hat{x} + \frac{1}{2}\hat{y} + \frac{1}{2}\hat{z} \\ \hat{k}_2 &= \frac{1}{\sqrt{2}}\hat{x} + \frac{1}{2}\hat{y} + \frac{1}{2}\hat{z} \\ \hat{k}_3 &= -\frac{1}{\sqrt{2}}\hat{y} + \frac{1}{\sqrt{2}}\hat{z}\end{aligned}\quad (3.4)$$

The imaging laser beam travels along \hat{x} , so absorption images are formed from the integrated density profile along \hat{x} with \hat{y} and \hat{z} as the camera axes (see §4.1.1). This projects atomic diffraction peaks from lattice beams 1 and 2 on top of each other, as shown in Fig. 3.3.

The polarizations of the forward lattice beams are all linear and mutually orthogonal. The polarization of the retro-reflected beam is also linear, but may be rotated by an angle θ with respect to the forward beam polarization using a quarter-wave plate⁴. The full lattice potential for this configuration (lin- θ -lin) is

$$\begin{aligned}V(\vec{x}) &= \frac{I_0 \pi c^2 \Gamma}{\omega_0^3} \left\{ \left(\frac{2}{\Delta_{D2}} + \frac{1}{\Delta_{D1}} \right) [1 + \cos(\theta) \cos(2\vec{k} \cdot \vec{x})] + \right. \\ &\quad \left. g_F m_F \left(\frac{1}{\Delta_{D2}} - \frac{1}{\Delta_{D1}} \right) (\hat{k} \cdot \hat{B}) \sin(\theta) \sin(2\vec{k} \cdot \vec{x}) \right\},\end{aligned}\quad (3.5)$$

⁴Space has been left in the apparatus to put an electric-optical modulator (EOM) for dynamic polarization control [73]; for more details see §6 of [41].

where I_0 is the peak intensity of the single beam, Δ_{D1} and Δ_{D2} are the detunings from the D1 and D2 lines of ^{87}Rb (see the simplified state structure in Fig. 3.4), and \hat{B} is the direction of the quantizing magnetic field. For the derivation of Eqn. 3.5 see §C.2. We have dropped counter-rotating terms, which are important for large detunings; these are included in the full potential as given in Eqn. C.31. In this thesis work, we realize the two limits of this potential: the lin-lin lattice ($\theta = 0$) and the lin- \perp -lin spin-dependent lattice ($\theta = \pi/2$).

The lin-lin lattice is created when $\theta = 0$ (i.e., the waveplate is simply omitted), and the potential results from the intensity standing wave created by the forward-going and retro-reflected beams. The potential is given by dropping the second term of Eqn. 3.5, and so this lattice is state-independent. For our typical wavelength ($\lambda = 812\text{nm}$), the lattice depth (in E_R) is $s = 2.8 \times 10^{-4} I_0 [\text{mW}/\text{cm}^2]$, and the intensity and potential are plotted in Fig. 3.4.

When $\theta = \pi/2$, the second term of Eqn. 3.5 gives rise to a fully state-dependent lattice (lin- \perp -lin) with a potential depth proportional to the magnetic moment $g_F m_F$. The first term in Eqn. 3.5 becomes a constant offset. Since the retro-reflected beam polarization is linear and orthogonal to the forward beam, there is no intensity standing wave, and the lattice potential is fully due to polarization effects. In theory, a lin- \perp -lin lattice exists for all λ , but we find it optimal to run at $\lambda = 790\text{nm}$ ⁵ where the first term of Eqn. 3.5 (i.e., the scalar light shift) goes to zero. This eliminates residual lin-lin lattices that arise due to imperfect polarization rotation of the retro-reflected beam and from polarization impurities in the forward-going beam. As described in §C.6, these effects can be demonstrated using pulses of the lattice light. At $\lambda = 790\text{nm}$, the lattice depth (in E_R) for $|1, -1\rangle$ is $s = 3.2 \times 10^{-4} I_0 [\text{mW}/\text{cm}^2] (\hat{k} \cdot \hat{B})$, and the intensity and potential are plotted in Fig. 3.4.

Our apparatus represents the first fully three-dimensional spin-dependent lattice realized in the strongly correlated regime. A number of groups have implemented 1D spin-dependent lattices [73, 75, 76]. One of the main differences is that in 1D the prefactor $\hat{k} \cdot \hat{B}$ can be set to unity by aligning the lattice and magnetic field directions. However, in 3D if $\hat{k}_i \cdot \hat{B} = 1$ for direction i then it will be zero for the other two directions. Therefore, we must ensure that \hat{B} has a sizeable projection along all of the lattice directions (Eqn. 3.4). In our case, $\hat{B} = \hat{y}$, which has a near optimal projection along each beam. To demonstrate the 3D strongly correlated regime, we realize the SF-MI transition (§2.4.1) for the lin- \perp -lin lattice in Fig. 3.5 using a mixture of $|1, -1\rangle/|2, 0\rangle$ and $|1, -1\rangle/|2, -2\rangle$ atoms.

The alignment procedure for the lin-lin lattice is detailed in §4.9 of [41] and the calibration procedure in §C.5. To summarize — calibration is performed by pulsing one direction of the lattice and measuring the atomic diffraction ratio versus pulse time, which is periodic for short ($< 100\mu\text{s}$) pulses. The periodicity is related to the gap between the ground and second excited bands (see Fig. 2.5). This method typically calibrates the lattice to within a 5% uncertainty. There are some small differences when aligning the lin- \perp -lin lattice, which are detailed in §C.6.

Since the lattice beams have a Gaussian profile, each beam also adds harmonic confinement, which is

⁵The value calculated using the known atomic structure of ^{87}Rb is $\lambda = 790.028\text{nm}$.




Figure 3.5: Demonstrating the SF-MI transition for the spin-dependent lin-⊥-lin lattice. Lattice depths are given for each spin component. Images at the same $|1, -1\rangle$ depth, in the same black outline are from a single experimental run, and the spin components were separated during expansion using a Stern-Gerlach field. For the purposes of this figure, empty regions of the image have been cropped out. Black, dashed lines indicate the expected mean-field SF-MI transition depth. (a) Diffraction images from a mixture of $|1, -1\rangle$ and $|2, 0\rangle$ atoms for increasing lattice depth. (b) Diffraction images from a mixture of $|1, -1\rangle$ and $|2, -2\rangle$ atoms. Since $g_F m_F$ for $|2, -2\rangle$ is twice as large as for $|1, -1\rangle$ for the same lattice realization, the $|2, -2\rangle$ atoms experience twice the lattice depth as given by Eqn. 3.5. In the mixture, the $|2, -2\rangle$ atoms undergo the SF-MI transition at half the lattice *intensity* as the $|1, -1\rangle$ atoms, but at the same depth. Mutual interaction effects are small since the two spins are not in overlapping lattices; the lattice minima are shifted from each other by half a period. Figure adapted from Ref. [4] (© 2010 IOP Science).

approximately,

$$\omega_{latt}^2 = \xi \frac{4sE_R}{m\mathbf{w}^2} \left(1 - \frac{1}{2\xi\sqrt{s}} \right), \quad (3.6)$$

where $\xi = 1$ for red-detuning, and $\xi = 0.5$ for the lin- \perp -lin lattice at 790nm. The beam waist \mathbf{w} is $120\mu\text{m}$. For more details, see §C.3.2. The total harmonic confinement is a combination of all individual harmonic confinements: lattice, dipole, and magnetic. To combine these is generally complicated, since the axes of all individual traps are not aligned. In §D we solve the full problem in detail. The straightforward quadrature sum of the geometric mean trap frequencies

$$\bar{\omega} = \sqrt{\bar{\omega}_0^2 + \omega_{\text{lattice}}^2} \quad (3.7)$$

is accurate to within a few percent.

An unavoidable consequence of applying an optical lattice to the gas is heating due to momentum diffusion from spontaneous emission [168]. For the lin-lin lattice, the ratio of lattice depth to heating can be minimized by using large, red detunings. This same strategy does not work for the lin- \perp -lin lattice, since this ratio has an upper bound. For the lin-lin lattice at $\lambda = 812\text{nm}$, the theoretical heating rate per unit lattice depth is 1.8nK/s . For the lin- \perp -lin lattice at $\lambda = 790\text{nm}$, the theoretical heating rate per $|1, -1\rangle$ unit lattice depth is 19.4nK/s . The lin- \perp -lin heating rate is independent of state and field direction, which implies that the lattice beams equally heat atoms in $|1, 0\rangle$ even though there is no $|1, 0\rangle$ potential. These heating rates give the average energy input per particle (in units of k_B) into the gas due to the occurrence of many scattering events. However, the energy input from each individual scattering event is a complicated function of the lattice arrangement, detuning, and the density matrix of the gas. Therefore, on experimental timescales, when the average scattering number of scattering events per particle is much less than one, the effect of these events on a particular quantum simulation measurement is an open question [37, 169]. For example, if the energy of a scattering event does not thermalize back into the bulk gas over the course of the experiment, the effective heating rate may be much less than calculated assuming the standard formulas [168]. For a more detailed discussion see §C.4.

3.2 ^{40}K Fermi-Hubbard Apparatus

The second apparatus involved in this thesis, located at the University of Toronto, is designed to create an ultracold degenerate Fermi gas comprised of ^{40}K atoms in an optical lattice to simulate the Fermi-Hubbard model. In particular, the goal is to develop a new tool — single-site imaging resolution of a 2D lattice plane — to simulate aspects of the 2D Fermi-Hubbard model believed to be relevant to cuprate superconductors (§2.4.3).

The thesis work on this experiment involved the many facets of designing and building an apparatus. The work is not complete, however, almost all the main components have been built and integrated together. For those elements not already built, there is at least a specific design or idea and most are currently under construction. Therefore, we will give a detailed overview of the apparatus as it exists at the time of writing (i.e., spring 2012). As is the case for any system still in the construction phase, the plan today may not be the reality tomorrow. In fact, many elements of the system we will describe have already gone through a

number of iterations in response to the issues that have thus far preventing completing the apparatus. Where relevant, we will point out aspects and designs that have changed and why. Not only is this documentation important for contextualizing the current apparatus, but also to hopefully avoid similar mistakes in the future. A detailed overview of the system will also be available in a forthcoming thesis by Dylan Jervis.

3.2.1 Overview




Figure 3.6: Schematic of the two-chamber ^{40}K apparatus as described in the text. Labels: (i) MOT chamber, (ii) transport system, (iii) science chamber, (iv) quadrupole coils, (v) RF/ microwave antenna, (vi) dipole beam (two sets), (vii) Feshbach coils, (viii) vertical lattice beam, (ix) in-plane lattice beam (two sets of retro-reflected beams), (x) imaging window, and (xi) custom objective (NA=0.6 at 404.5nm). Not all laser beams are illustrated.

Our system, illustrated in Fig. 3.6, consists of a two-chamber vacuum system connected by a differential pumping tube. In section (i), atoms are collected from a high pressure vapor in a magneto-optical trap (MOT) and then transported to section (iii) (the ‘science chamber’), which is a low pressure region where the atoms are evaporated to the ultracold regime and loaded into a lattice. The main feature of this system is a 200 μm thick, 5mm diameter sapphire window for high resolution imaging with a custom objective optimized for 404.5nm light. In terms of physical construction, only two elements, the lattice and imaging system, are incomplete. At present, the biggest impediment to progress is integrating all the components

together — specifically attaining enough atoms in the MOT and then transporting those atoms to the science chamber with high enough efficiency. While some of these problems are issues of optimization, mostly they are hardware issues related to the atomic sources, which I will discuss in more detail in §3.2.6.

There are several routes to attaining an ultracold gas of ^{40}K , and we have iterated through a number of these. For reasons detailed in §3.2.3, we are currently pursuing sympathetic cooling of ^{40}K using ^{87}Rb . Therefore, in our current setup, we collect ^{87}Rb atoms and ^{40}K atoms in the MOT. The ^{87}Rb vapor is created using a commercial alkali metal dispenser from SAES, and the ^{40}K vapor is created using a home-built (see [121] for design) dispenser containing enriched potassium (10.8% ^{40}K).

Once in the MOT, atoms are loaded into a magnetic quadrupole trap and transported (ii) to the science chamber (iii) by ramping the current sequentially on and off in sets of overlapping quadrupole coil pairs (§3.2.7). The pressure in the science chamber is sufficiently low, so the lifetime of trapped atoms is on the order of several minutes. This allows enough time for efficient evaporation to the ultracold regime. In past iterations, when we attempted ^{40}K only evaporation, too few atoms were transported to the science chamber to start evaporation. This was caused by a combination of poor transport efficiency and low initial MOT number (see §3.2.7 and §3.2.6). Our belief is that ^{87}Rb - ^{40}K sympathetic cooling will be a more robust solution, since only large numbers of ^{87}Rb atoms are required. Therefore, the remaining overview is the *planned* sequence of events. Although these steps have not been demonstrated, the physical elements required have been, or are well towards being, constructed.

In the science chamber, the atoms are held in a magnetic quadrupole trap (iv). At low temperatures, atoms can flip their spin in the zero-field region near the center of the trap (i.e., Majorana losses). To prevent these losses, we plan on utilizing a “plugged” quadrupole [23, 101] by focusing a blue-detuned (repulsive) optical beam through the quadrupole center. The source of this $\lambda \approx 760\text{nm}$ laser beam is a Ti:Sapphire laser coupled into an optical fiber. In the plugged quadrupole trap, ^{87}Rb atoms will be evaporated via transitions to untrapped states induced by an RF frequency magnetic field created by a current loop close to the imaging window (v). ^{40}K atoms are not evaporated as much because the RF frequency removes ^{40}K atoms which are 9/4 times more energetic. As the ^{87}Rb gas cools, the ^{40}K gas cools via ^{87}Rb - ^{40}K collisions. Once sufficiently cold, the ^{87}Rb gas will be ejected from the trap, and the ^{40}K atoms will be loaded into a crossed-dipole trap (vi) at $\lambda = 1053.57\text{nm}$ ⁶, located close to the imaging window (x).

A hyperfine state mixture consisting of $|9/2, -9/2\rangle$ and $|9/2, -7/2\rangle$ atoms will be prepared using microwave and/or RF transitions. In this mixture we can access the Feshbach resonance (§2.1.2) at 202.1G using a pair of bias coils (vii) to increase the collision rate. Further evaporation will be performed by lowering the dipole trap depth until the gas enters the ultracold regime. At this point, atoms will be loaded into 2D planes by retro-reflecting light (from the same source as the dipole laser) from the imaging window, which has a high-reflection coating at 1052nm (viii). Finally, a square lattice is applied in these planes (ix). For imaging, the lattice is ramped to full power, and beams at 404.5nm cause scattering (and cooling) on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition (see §E). This light is collected by a custom objective (xi) with a numerical aperture of 0.6 (resolution $\approx 400\text{nm}$) located on the air side of the imaging window for single-site imaging. Not pictured in Fig. 3.6 is the support structure for the objective, consisting of a 3-axis translation stage and

⁶This is a “magic” wavelength for the $4S \rightarrow 5P_{3/2}$ transition of ^{40}K ; see §3.2.10 and §G.1.

two-axis tilt control as well as the tube lens and CCD camera. This hardware is currently being developed. More detail on each of these experimental steps is provided in the following sections.

3.2.2 Vacuum System and Support Structure




Figure 3.7: ^{40}K vacuum system including all pumps. The MOT side is pumped with a 20L/s ion pump and the science side with a 60L/s ion pump and titanium-sublimation pump (TSP). A residual gas analyzer (RGA) is used to detect leaks and diagnose dispenser performance issues. (Inset) Expanded view of the three elements of the science chamber. In the middle is a custom built stainless steel, 12-sided chamber from Kimball Physics. One of the side ports has an increased bore for high conductance to the vacuum pumps. The top and bottom recessed flanges are custom built by UKAEA.

A necessary component of any ultracold atom experiment is an ultrahigh vacuum (UHV) system. Because of the conflicting demands of high pressure for atom loading and low ($< 10^{-11}$ Torr) pressure for evaporation, simulation, and imaging, we have a two-chamber system as illustrated in Fig. 3.7. The MOT side, which is at relatively high pressure ($\approx 10^{-9}$ Torr), is pumped using a single 20L/s ion pump (Duniway RVA-20-DD-M). The science side is pumped using a 60L/s ion pump (Duniway RVA-60-DD-M) and a titanium-sublimation pump (TSP). The two chambers are connected by a narrow tube with a right-angle for differential pumping. At the narrowest point, the tube has a 0.5" OD (0.43" ID). The tube that connects vertically into the science chamber is 149.44mm long and has a 0.75" OD (0.68" ID). The whole length of the differential tube is best illustrated in Fig. 3.18. A rough estimate of the total conductance [170] for H_2 is 5L/s through the tube. If the H_2 pumping speed in the science chamber is 200L/s (a conservative estimate

for the TSP), then the H_2 pressure is 40 times lower in the science chamber. Differential pumping for the alkalis is better, both because they are heavier and because they are pumped by all surfaces. In practice, the differential pumping is sufficient as the lifetime has been measured to be greater than 200s.

The vacuum system has three valves. A VAT 48124 all-metal gate valve separates the MOT side from the science side, and it can be used to limit alkali contamination when the experiment is not in operation and to independently bring either side of the vacuum to atmospheric pressure. This is mostly useful for replacing the atomic sources in the MOT. Two all-metal right angle valves (VG Scienta ZCR40R) provide air access to either side, which is used to attach a turbomolecular pump during initial pump down and bake (until the ion pump takes over).

The centerpiece of the vacuum system is the science chamber; an expanded view is provided in Fig. 3.7. The main element is a stainless-steel chamber from Kimball physics. This chamber has 12 2.75" ports which are used for windows; one of the ports has an expanded bore (4.5" CF) to increase conductance from the chamber to the ion pump and TSP. All 12 ports have 7056 glass vacuum windows (Duniway VP-275-150), which are broadband anti-reflection coated for 700-1100nm. Exact coating specifications are given in Fig. J.5. The top and bottom sides of the chamber accept 8" flanges. We attach two custom recessed titanium flanges built by UKAEA to these ports. The vertical transport tube (connecting to the MOT side) is welded to the bottom flange. The top flange has a bottom insert with a specially brazed 200 μ m, 5mm diameter sapphire window for high resolution imaging (more details in §3.2.10). Before constructing this flange, the window strength under vacuum was tested in Ref. [171]. The brazing process slightly curved the window in towards the vacuum side. The radius of curvature, ≈ 30 cm, is not expected to effect the resolution. The window assembly was anti-reflection coated at 404.5nm and high-reflection coated at 1000-1060nm; see the coating specifications in Fig. J.6. Technical drawings of these parts are included in §J.

The entire system was assembled and evacuated according to standard UHV protocol. All surfaces, except the inside of the valve, RGA, and MOT cell, were thoroughly cleaned with hot water and alconox, and then liberally rinsed with hot water, acetone, methanol, and distilled water (in that order). Parts were left to dry in HEPA-filtered air. Most metal parts, in particular the stainless steel chamber, were vacuum baked at 400°C for two months to reduce hydrogen outgassing (see Fig. 3.8). Despite using silver plated gaskets and screws, significant corrosion occurred during the bake, and the flanges had to be chiseled off, thereby destroying several parts in the process. Consequently, residual corroded gasket material on conflat knife edges was a concern. This was primarily removed by scrubbing with a 1M nitric acid solution. After the acid scrub and a cleaning cycle, these parts were air baked at 450°C to build an oxide layer on the vacuum side to further reduce hydrogen outgassing. The TSP was separately vacuum baked and outgassed at 360°C.

After the system was assembled as shown in Fig. 3.7, it was baked in three stages. First, the system was baked while being pumped by a turbomolecular pump only. Next, the ion pumps were reduced in temperature and turned on. Finally, the all-metal valve was cooled down, the valve was closed, and we baked with only the ion pump running. The total bake time was typically 1-2 weeks. For the MOT chamber, we split the first bake stage into two sub-stages. First we baked for several days to drive water out of the atomic sources. Then we turned on the sources at operational current for outgassing. This will be discussed in more detail in §3.2.4.




Figure 3.8: ^{40}K system all-metal bake. All the metal parts were assembled (in no particular order) and baked under vacuum for two months at 400°C . (a) Images from the bake. (i) System in the oven before the bake. The walls of the oven are aluminum-foil-covered firebricks (Thermal Ceramic K23). (ii) The oven enclosed during the bake. A turbomolecular pump outside the oven pumps on the system through a 1m-long bellows. (iii) The main chamber after the bake and the chisel required to remove the blank flanges on the side ports. Note the oxidation on the outside versus the inside of the system, which had not yet been baked in air. (b) H_2 partial pressure measured with an RGA versus the bake time in days. The pressure continued to decrease as we baked. The bake temperature is shown as an inset. The dip before day 20 was the result of a broken heater element.

Where possible, and particularly for the MOT chamber, which has many glass-metal transitions, we built a firebrick oven for the bake to ensure uniform heating. For some parts of the system this is not possible, so they were wrapped in Al foil and insulation. In all cases, heat was supplied by lengths of “heater tape” (Omega STH series) plugged into variable-voltage transformers (“variatics”) running off 120V AC. Temperature was monitored by ≈ 10 thermocouples, and the voltage on each transformer was adjusted manually to maintain a uniform temperature. The science chamber was baked to a maximum temperature of 200°C; a 180°C temperature was maintained at the imaging window. The MOT chamber was baked to 150°C. Due to our system design, several of the transfer quadrupole coils (see Fig. 3.6) were put on before the vacuum was fully assembled and cannot be removed. Therefore, we needed to arrange our bake around these coils.

A residual gas analyser (SRS RGA100 with Option 1) with an electron multiplier on the MOT-chamber side allows for extremely sensitive He leak testing (10^{-13} Torr He noise floor). We detected several leaks this way. Twice, there was a leak in the TSP feedthrough, which required replacement. The imaging window developed a leak during the first vacuum bake, likely due to an accidental temperature spike. Ultimately, we replaced the entire flange and successfully rebaked. Initially, we used a nickel (versus copper) gasket on the bottom flange of the science chamber to reduce eddy currents. After the bake, this gasket badly leaked, and so it was replaced with a standard copper gasket. This problem may have been caused by thermal expansion mismatches between the titanium flange, stainless steel chamber, and nickel gasket, exacerbated by nickel’s poorer sealing qualities. For example, conflat the knife-edge indentation into the nickel was very shallow compared to the deep indentation typically formed in a copper gasket. Even the copper gasket developed a very small leak. We injected a mixture of Vacseal and acetone around the gasket edge (in the gap between the flanges) using a syringe, which has provided a seal for several years. Since Vacseal has been known to fail, this should be checked in the event that the vacuum performance degrades.

To support the system and create an interferometrically stable platform for the lattice and imaging, the chamber is mounted on an anodized aluminum collar set on 4”x4” aluminum posts. These posts also form the base for three aluminum breadboards (2-Thorlabs PBG12104, 1-Thorlabs PBG12111), which is an interferometrically stable platform. A schematic and photograph is shown in Fig. 3.9.

3.2.3 Assessing Different Routes to Ultracold ^{40}K Atoms

In comparison to ^{87}Rb , preparing ultracold ^{40}K faces three main challenges. First, as a fermionic isotope it requires two distinguishable particles for evaporative cooling. Second, the natural abundance of ^{40}K is very low (0.01%). Lastly, the sub-Doppler laser cooling processes for ^{40}K are weak [172] or non-existent [173], which degrades the initial condition for evaporative cooling. There are two main approaches to overcome these problems. One is to sympathetically cool ^{40}K using another species or isotope (almost always ^{87}Rb), at the cost of significantly complicating the system. The second option is to proceed with evaporation using two spin-states of ^{40}K . However, because of poor laser cooling, this approach requires a large initial atom number, which is difficult to produce given the low natural abundance. Surveying 17 groups with ultracold ^{40}K gases⁷, 12 used sympathetic cooling (10 using ^{87}Rb [174–183], one using ^6Li [184], and one using ^{41}K [185]), and 5 cooled using ^{40}K [173, 186–189] only. The large number of groups using ^{87}Rb for

⁷Based on our best effort, this includes every group with ultracold ^{40}K gases as of August 2011.




Figure 3.9: ^{40}K vacuum system support structure and science chamber optics deck. (Top) schematic with one of the breadboards removed for viewability. (Bottom) Photograph of the system before the vacuum system was fully constructed. The base is set of 8 4"x4" aluminum posts. Optics are mounted on aluminum breadboards from Thorlabs.

sympathetic cooling is not surprising. ^{87}Rb MOTs are large ($N > 10^9$), with efficient sub-Doppler cooling, and ^{87}Rb has a natural abundance of 27.8% with readily available sources. From a technical perspective, all the optics for a ^{40}K and ^{87}Rb MOT are interchangeable, since the wavelengths for the D2 transitions for these species (766.7nm and 780.2nm), are sufficiently close together. These two wavelengths can be easily combined using a polarizing beam-splitter and a multi-order waveplate.

Initially, we decided to pursue an all ^{40}K evaporation scheme, based on the simplicity of the approach. We were foiled in two regards. First, achieving a large MOT number ($> 10^9$ atoms) depends on a quality source of enriched K, which was a problematic issue (see §3.2.4). Also, this large MOT must be transferred into the magnetic trap and then transported to the evaporation region without significant heating or loss. No group has succeeded in cooling ultracold ^{40}K gases using a fixed coil transport scheme, and indeed, we found the transport efficiency too low (see §3.2.7). Transport efficiency issues also rule out sympathetic cooling with ^{41}K , for which the transport efficiency was even worse than ^{40}K . Based on assessing the success of various groups, in addition to our own experience, our current approach is sympathetic cooling using ^{87}Rb atoms.

3.2.4 Atomic Sources

Atomic sources that produce a high purity atomic gas into a UHV system are essential to the operation of our apparatus. The source is the first in a long chain of elements required to reach the ultracold regime, and so small source issues are compounded. Typically, experiments using alkali atoms employ either a ‘lump’ of pure metal ($> 1\text{g}$) or a commercial alkali metal dispenser. A dispenser is a small metal vessel containing compounds that react to produce the alkali gas when heated by passing current through the dispenser via external leads. A widely used commercial dispenser is produced by SAES. The low natural abundance (0.01%) of ^{40}K means that using a gram-sized source is not feasible. Instead, we must use very expensive enriched potassium⁸, which is only available in mg quantities. Small amounts of pure enriched metal plated into a glass ampoule are available. These can be placed into the vacuum system as long as the vapor pressure is well-controlled so that the metal is not rapidly pumped away. For example, at $T=300\text{K}$ (27°C), the vapor pressure of K is $1.8 \times 10^{-8}\text{Torr}$ [190], which results in 0.06mg/day lost to a 20L/s pump. For small quantities of material, a dispenser-type source is desirable because the release of the material is easy to regulate — when there is no current through the dispenser, no K is released.

In the 1990s, commercially enriched dispensers were not available, so a solution was developed at JILA [121] that uses enriched KCl and ground calcium, which undergoes the following reaction when heated



These ‘homebuilt’ dispensers are used in 8 of the 17 experiments surveyed, but they are time-consuming to make and require a dry environment, since calcium oxidizes in ≈ 10 minutes in 40% relative humidity. Several years ago, commercially enriched dispensers produced by Alvacat appeared on the market. These operate differently than the SAES and ‘homebuilt’ dispensers as no chemical reaction is used. Instead, the enriched K metal is alloyed with barium, which decreases the vapor pressure so that an increased temperature is

⁸K (in the form of KCl) enriched to 10.8% ^{40}K is available for \$132/mg from Trace Sciences (July, 2011).




Figure 3.10: (a) Our first dispenser mounting scheme, and our (b) current scheme. (c) Residual gas analyzer (RGA) traces firing the dispensers for the first time under vacuum. Alvatec, SAES and homebuilt sources are shown. Dispensers are considered clean when they are hydrogen-limited and when initial water outgassing is complete. The data for the Alvatec dispensers was taken from one system with ion pumps on, and the data for the homebuilt and SAES dispensers were taken in another system with only a turbomolecular pump (so that we could access higher pressures). For the Alvatec dispensers, a large variation in the behavior between dispensers was noted, as shown in (c). Whereas one dispenser was almost clean to begin with, another dispenser had such a high water content that the ion pump controller shut off automatically. For the homebuilt and SAES dispensers these traces are representative of all dispensers (i.e. there is much less variation). The SAES dispenser produces a short carbon dioxide and water spike, but is quickly hydrogen dominated after being turned on. The homebuilt dispensers initially outgas a large volume of water. After performing approximately 5 current spikes, the water outgassing rate decreases and eventually disappears. After outgassing, water is the least dominant partial pressure when operating at 8A .

required to release a sufficient quantity of K. Alvatec’s claim was that these dispensers have a lower overall vacuum pressure during operation than other sources. Because the alloyed metal can oxidize in air, the dispensers are filled with argon and sealed with indium. Once in the vacuum system, the indium seal opens after the dispenser is heated to 160°C.

Our first two MOT cells used Alvatec dispensers. To mount the dispensers, we used a standard technique previously employed with SAES dispensers and carried out by Ron Bihler at Technical glass (where the MOT cells were constructed). The dispensers are spot welded to a “pin press”, which consists of metal rods sealed directly to the glass wall (see Fig. 3.10). When done correctly, this method creates robust electric connections, since there are no mechanical connections that can loosen. Additionally, this method uses the minimal amount of metal located inside the vacuum chamber. Minimizing metal surface area is important, since K bonds to metal; our early tests with all-metal chambers could only produce a small MOT with a high-flux source. This bonding technique had been perfected for SAES dispensers, and there were some issues transitioning to the Alvatec dispensers. The Alvatec sources in our first MOT cell (#1) outgassed a large amount of water, had visual evidence of oxidation, and some of the spot welds failed during the bake. We believe that during the bonding process, these dispensers were heated above the melting point of indium and exposed to air for a long period of time. These sources could only produce MOTs consisting of 10^7 atoms.

In our next cell (#2), we used the same mounting technique, but the pin press was longer to keep the dispensers further from the glass during the bonding process (of the pin press to the cell). Additionally, they were sunk to a cold finger during bonding. The RGA traces we measured when we fired these dispensers for the first time are shown in Fig. 3.10. A standard technique used to test atomic sources is to look for fluorescence from a resonant laser beam sent through the cell while the sources are activated. We were unable to observe fluorescence from the Alvatec dispensers in cell #2, despite operating them up to 13A of current, which created sufficient dispenser temperature to visually glow orange⁹. At these high currents a black film, presumably barium, was deposited on the inside of the glass cell. Fortunately, the dispensers were located in a separate tube projecting from the cell, and the film did not coat any of the MOT windows. A separate experimental group in our lab also coated their cell with a black film emitted from Alvatec dispensers. This film could only be removed with sulfuric acid [191]. Ultimately, we generated sufficient vapor pressure for MOT operation at $N \approx 10^8 - 10^9$. While we cannot say with certainty that these sources did not run to their specification, the lack of fluorescence remains troubling. Strong fluorescence has always been seen using our homebuilt dispensers. Cell #2 included SAES Rb dispensers on the same side of the MOT cell as the Alvatec source. When we turned these “Rb dispensers” on after a year of running just the K dispensers, mostly K vapor and only a small amount of Rb vapor was emitted. It may be that K dispensers were put into the system by accident instead of Rb dispensers. At that time, SAES did not mark the dispensers; now all dispensers are marked with their alkali metal content. Another possibility is that the K vapor poisoned the reactants in the SAES dispensers.

For our third and current cell (#3), we have 3 homebuilt K dispensers each with 7.5mg of K at 10.8% enrichment and 3 commercial SAES Rb dispensers. Additionally, we used an electrical feedthrough on a

⁹Due to variations in how dispensers are mounted, the same current in different vacuum systems may produce different dispenser temperatures. Therefore, we found it imperative to visually estimate the temperature using the color of light emitted by the dispenser. We observed that dispensers emitted alkali when emitting enough IR to be visible to our IR viewer (FJW Optical Systems).

1.33" CF flange instead of a pin press. The dispensers are mounted using mechanical crimp connectors, as shown in Fig. 3.10. This is a necessary change because of the difficulties associated with shipping homebuilt dispensers to the glassblower; we mount the dispenses ourselves on site. This design also eliminates thermal stresses on the dispensers and allows us greater flexibility to change them at a later date (without making a new cell). Guided by the possible poisoning in cell #2, we located the Rb and K sources on opposite sides of the cell (see Fig. 3.16). The RGA scans measured while firing the homebuilt and SAES dispensers are shown in Fig. 3.10. During these scans, we detected strong fluorescence from all dispensers, which is an encouraging indication of quality. Ultimately, poor dispenser performance ended up delaying this project by close to two years, which hopefully serves as a cautionary tale for those building future experiments.

3.2.5 MOT Laser System




Figure 3.11: Simplified ^{40}K state diagram illustrating the required laser frequencies for a standard $4S_{1/2} \rightarrow 4P_{3/2}$ MOT at 766.7nm. Also shown is how the ^{39}K states are used for locking. Frequencies are referenced to the bare $4S_{1/2} \rightarrow 4P_{3/2}$ transition [192].

As shown in Fig. 3.11, to create a magneto-optical trap, we need several laser beams of various powers which are narrow ($<1\text{MHz}$ linewidth) and accurately tuned to an absolute frequency within a part in 10^8 . The ‘trap’ beam performs laser cooling and trapping on the cycling transition $4S_{1/2}, F = 9/2 \rightarrow 4P_{3/2}, F' = 11/2$ for ^{40}K ($5S_{1/2}, F = 2 \rightarrow 5P_{3/2}, F' = 3$ for ^{87}Rb). The trap beam is 10–50MHz red-detuned from this transition and requires $\approx 200\text{mW}$ of total power. Occasionally atoms are depumped by this beam to the ground hyperfine state $F = 7/2$ for ^{40}K ($F = 1$ for ^{87}Rb), which is not excited by the trap beam. Therefore, we also need a ‘repump’ beam from $4S_{1/2}, F = 7/2 \rightarrow 4P_{3/2}, F' = 9/2$ for ^{40}K ($5S_{1/2}, F = 1 \rightarrow 5P_{3/2}, F' = 2$ for ^{87}Rb). For ^{87}Rb , this beam only requires $\approx 10\text{mW}$, but ^{40}K atoms are more frequently depumped because the excited state hyperfine splitting is small, so the repump beam is essentially an additional cooling and trapping beam and requires $\approx 100\text{--}200\text{mW}$ of power. Auxiliary low power beams are required for state preparation and imaging. These include the ‘probe’ beam on resonance with the cycling transition and an optical




Figure 3.12: Simplified ^{87}Rb state diagram [193] illustrating the required laser frequencies for a standard $5S_{1/2} \rightarrow 5P_{3/2}$ MOT at 780.2nm. Frequencies are referenced to the bare $5S_{1/2} \rightarrow 5P_{3/2}$ transition.

pumping (OP) beam resonant with $4S_{1/2}, F = 9/2 \rightarrow 4P_{3/2}, F' = 9/2$ for ^{40}K ($5S_{1/2}, F = 2 \rightarrow 5P_{3/2}, F' = 2$ for ^{87}Rb). All four of these beams are shown with respect to the state diagram for ^{40}K in Fig. 3.11 and for ^{87}Rb in Fig. 3.12.

The laser system which generates these beams can be divided into two main sub-systems. The first part is designed to create narrow, frequency stabilized beams from which all the other beams will be derived. For frequency stability these lasers are locked to an atomic transition. These source lasers (‘master’ lasers) are all external-cavity diode lasers in with frequency-selective feedback. For ^{40}K we have one master laser (Laser #1), which is a linear Fabry-Perot cavity with an intra-cavity interference filter [194–196]. Rotating the filter performs coarse wavelength tuning. A piezo-electric transducer controls the cavity length, which, along with the diode current, allows for fine wavelength tuning. Master laser #1 is referenced to the $4S_{1/2} \rightarrow 4P_{3/2}, F = 1/2$ crossover peak of ^{39}K saturation spectroscopy (the excited state is not resolved) as shown in Fig. 3.11. A direct reference to ^{40}K is not possible using a natural abundance vapor cell. To generate the error signal, we frequency modulate the saturation spectroscopy probe beam using an electro-optical modulator (EOM) at 15MHz and demodulate and filter the photodiode signal; see [194] for details and layout. We feedback to the laser using homebuilt servo electronics (see §H.1 for circuit diagrams).

For ^{87}Rb we have two master lasers: lasers #2 and #3. Laser #3 is an identical design to #1 and Laser #2 is a commercial New Focus Vortex laser which uses a Littman-Metcalf laser cavity, although we expect to replace this with an interference filter laser in the near future. Laser #3 is locked to the $^{87}\text{Rb} F = 1 \rightarrow F' = 1/2$ crossover using the same FM modulation scheme as used for laser #1. Laser #2 is not referenced to an atomic transition, but instead referenced directly to #3. Light from both lasers is mixed on a fast photodiode (EOT ET-4000, 10GHz unamplified GaAs photodiode). The photodiode signal is




Figure 3.13: Optical and electronic setup for beat-note stabilization of Laser #3 to Laser #2. Beams with the same polarization from the two lasers are combined using a non-polarizing beam splitter cube. The combined light is focused onto a fast photodiode, resulting in a DC signal plus a time-dependent component at the frequency difference between the lasers. This photodiode signal is DC blocked, amplified, divided down, and compared to a reference signal (VCO) using the ADF4007 phase-frequency detector. The error signal is zero when the absolute frequency difference between the lasers is equal to 32 times the VCO frequency. The sign of the frequency difference determines the sign of the error signal slope. A snapshot of the beat note on a spectrum analyzer is shown (averaged). The FWHM is approximately 1MHz, indicating that further servo optimization is possible.

amplified, divided down by a factor of 64 in frequency, and compared to a VCO (Mini-circuits ZX95-310A, 160-360MHz tuning range) which is frequency-divided by a factor of 2 using a phase-frequency detector (Analog devices ADF4007). This scheme produces an error signal which is used to lock the difference between the frequencies of lasers #2 and #3 to 32 times the VCO frequency. This is a convenient method for generating a laser beam which is shifted in frequency with respect to a reference beam up to 10GHz. In addition, laser #2 is widely tunable, over hundreds of MHz, by simply changing the VCO frequency. The beat-note setup is shown in Fig. 3.13.

The next section of the laser system splits, amplifies, and frequency shifts the three reference beams to produce the beams required for the MOT. The main component we use for frequency control are acousto-optical modulators (AOMs), which use acoustic waves excited by RF (75-400MHz, depending on the specific model) to diffract the beam. The ± 1 diffraction order is frequency shifted \pm the RF drive frequency. By adjusting the relative angle between the AOM and incoming beam, up to 80-90% of the optical power will diffract into either the ± 1 order. The fraction of diffracted optical power is a function of the RF power, which makes AOMs also useful for power control and as a fast (μ s timescale) switch. The diffraction angle changes with the RF frequency, so a single AOM pass cannot be used for dynamic frequency control. For dynamic control, this angle can be canceled in a special ‘cats-eye’ double-pass configuration (see 3.14).

To power-amplify the beams we either injection lock a higher power diode or pass the beam through a tapered amplifier. For injection locking we use a Faraday isolator to send a small amount of the reference light into a higher power diode laser that is not in an external cavity (i.e., ‘free running’). Under the right




Figure 3.14: Schematic of the $^{40}\text{K}/^{87}\text{Rb}$ laser system. Beam paths are not exactly as they appear on the optics table, and some elements have been omitted for simplicity. All beam splitters are polarized. Acousto-Optic Modulators (AOMs) are listed with the typical frequency. The square projection on the side of the AOM represents the RF input. Beams angled away from/toward the RF increase/decrease in frequency. AOMs in a double-passed “cat’s-eye” configuration are shown followed by a lens, waveplate and mirror. Here the beam picks up twice the frequency shift and exits from the opposite port of the beam splitter. For both ^{40}K and ^{87}Rb , the probe and optical pumping utilizes the same fiber-coupled beam, which is dynamically adjusted in frequency. The ^{87}Rb repump beam is shifted into resonance with a single-pass AOM since no dynamic frequency tuning is required. The ^{87}Rb trap, probe, and optical pumping beams also only use single-pass AOMs since the frequency of Laser #3 is tuned by adjusting the reference VCO in the beat-lock setup (Fig. 3.13); the AOMs are only for power control and fast switching. When the frequency of Laser #3 is set so that the ^{87}Rb trap beam is at the correct frequency, the ^{87}Rb probe/OP beam is close to resonance for optical pumping.

conditions, injection causes lasing with the same frequency characteristics as the input light. This is a highly non-linear amplifier, because the output power is fixed to the output power of the free-running laser. Below a threshold input power, injection locking will fail, and the output is completely unusable. Therefore, above threshold, injection locking removes amplitude noise in the injection light. For example, injection locking has some tolerance to angular deflections of the input beam. Therefore, as shown in Fig. 3.14, we typically inject after double passing through an AOM. As we dynamically tune the AOM frequency (to tune the optical frequency) there is some change in both the diffraction beam power and direction. As long as the laser stays injection locked, these power and direction changes do not impact the light used in the experiment. Injection is limited to about 100mW of output power. For beams requiring higher powers (e.g., the ^{87}Rb trap, ^{40}K trap, and ^{40}K repump beams), we pass the beams through tapered amplifiers (TAs). TAs are solid-state linear amplifiers, with a maximum output of $\approx 1\text{W}$, depending on the specific device (we use the Eagleyard EYP-TPA series). The downside of using TAs is that the spatial mode of the output is poor. We send the beam through a single-mode fiber, which improves the mode quality at the cost of 50-70% less power. Up to 300mW of single-mode light for the MOT is available at the output of the fiber. The overall setup is illustrated in Fig. 3.14.

3.2.6 MOT




Figure 3.15: Optical layout of the $^{40}\text{K}/^{87}\text{Rb}$ MOT. Beams for ^{40}K and ^{87}Rb are first combined using a polarizing beam splitter cube and a dichroic waveplate (λ for 780.2nm and $\lambda/2$ for 766.6nm from LensOptics). After the first 4x telescope, an optional dark spot [197] is added to the repump beam. The repump and trap beams are overlapped using a non-polarizing beamsplitter (Thorlabs BSW-08), which operates as a 60:40 beamsplitter for polarized light. The resulting two beams are expanded and then split into the three beam paths for the MOT. The top beam path, which is the vertical MOT beam, has the most power since it does not need to be split into two beams. There is the option to discard some of the light if better beam balance is required, but in practice we use the maximum amount of power in all beams. (Inset) One of the MOT beams going through the cell. Our MOT beams are in a retro-reflected configuration. A dichroic mirror (DM) allows 404.5nm MOT beams to be added (see §E).

Once on the MOT table, the ^{40}K and ^{87}Rb beams are combined into one dual-wavelength repump beam and one dual-wavelength trap beam using polarizing beam splitters and dichroic waveplates (LensOptics




Figure 3.16: Schematic of the $^{40}\text{K}/^{87}\text{Rb}$ MOT Cell; see Fig. J.4 for technical drawings sent to the glassmaker. (Bottom right) Cell shown with beams (red) and quadrupole coil. (Top Right) Actual cell produced by Technical Glass.

27-order waveplates with λ retardance for 780.25nm and $\lambda/2$ retardance for 766.7nm). Next, the trap and repump beams are combined, split into three beams and expanded. These three MOT beams then pass through windows into the glass cell in a retro-reflected configuration as shown in Fig. 3.15. The MOT Cell is shown in Fig. 3.16. The MOT beams are apertured by the windows; the horizontal windows are 1.75" in diameter and the vertical windows are 1.5" in diameter. The magnetic quadrupole field for the MOT is provided by a set of coils on the top and bottom of the glass cell. These coils form the first quadrupole pair of the magnetic transport (see §3.2.7). We use a set of three magnetic field shims (not shown) to null magnetic fields along all axes and to provide a quantization axis for optical pumping. Because the vapor pressure of K at room temperature is low, we heat the source tube to 60°C and the main cell to 50°C. The MOT can be monitored by collecting fluorescence through the side or front viewing windows (0.75" diameter). We send a probe beam down the axis of the transport tube to take absorption images and for optical pumping.




Figure 3.17: Fluorescence images of the (a) ^{40}K and (b) ^{87}Rb MOTs. The image of the ^{87}Rb MOT is from the current setup acquired in 0.7ms; the ^{40}K MOT is from our previous cell acquired in 3ms.

In typical operation, the MOT takes 10-30s to achieve a steady-state atom number. After loading the MOT, we optimize for transfer into the magnetic trap. We first do a compression stage (CMOT) by decreasing the repump power and increasing the trap detuning, then a molasses stage (involving laser cooling only and no magnetic field gradients), and finally optical pumping along the transport axis with a quantization field applied by one of the shims. We have observed a ^{87}Rb MOT in our current system and a ^{40}K MOT in our previous cell. Fluorescence images of these MOTs are shown in Fig. 3.17. Since we are still optimizing our MOT, we will not specify exact parameters; these will be available in Dylan Jervis's future thesis. We also performed laser cooling and trapping on the $4S \rightarrow 5P$ (404.5nm) transition of ^{40}K (see §E for details).

3.2.7 Magnetic Transport

After CMOT, molasses, and optical pumping, we load the atoms into a magnetic quadrupole trap by increasing the current in the MOT coils to a 100G/cm gradient. To move the atoms to the science chamber for evaporation, we use a fixed-coil magnetic transport scheme that employs dynamically adjusted currents in a series of overlapping coils. The fixed-coil method has no mechanically moving parts, but the setup and electronics are significantly more complicated than the moving coil system used in the ^{87}Rb experiment. As a result, achieving operational transport occupied a good portion of our time. Although we eventually succeeded in transporting atoms, the efficiency is low and optimization is still an ongoing project. One of the




Figure 3.18: Schematic of the magnetic transport path illustrating the transport coils and differential pumping tube. The horizontal section of the transport utilizes 12 quadrupole coil pairs that require unipolar current; three pairs are on at any one time. Because of the large spacing between the MOT quadrupole coil pair and the next coil pair, a “push” coil is used to provide the axial field required to move the MOT quadrupole zero across this gap. At the end of the horizontal section, the atoms are transferred to a series of vertically stacked coils. Using bipolar current, we transfer the atoms vertically through the center of these coils. (Inset) Photograph of the horizontal transport coils and mounts. The mounts are made from brass with a teflon spacer to prevent induced currents. The coils are glued to the mounts using thermally conductive epoxy, and the mounts are water cooled via soldered on pipes.

complications of our setup is that, unlike most other transport systems, we transport the atoms horizontally 365mm and vertically 174mm. The system is shown schematically in Fig. 3.18. In particular, getting the vertical transport to work with high efficiency has been challenging. The initial design and construction are detailed in [198, 199]. The definitive account of this transport system will be in Dylan Jervis’s thesis, so here I will only summarize and highlight our progress to date and note some important considerations.

The horizontal transport section is a standard design based on [200]. In our system, 12 quadrupole coil pairs are overlapped perpendicular to the transport axis, as shown in Fig. 3.18. In the middle of the transport system, a unipolar current in three adjacent coils creates a quadrupole trap with a gradient of 100G/cm along z with a 1.72 aspect ratio¹⁰. The atoms start and finish in a single quadrupole coil pair, so the aspect ratio has to ramp up and then back down to 1. There is a ‘push’ coil at the beginning which provides extra bias and gradient along x to push the trap from the first pair of coils (the MOT coils) into the second pair of coils. This is because the overlap between the MOT and second pair is large to allow the vertical MOT beam into the MOT cell. At the end of the horizontal transport, an extra pair (horizontal-vertical transfer pair) has been added to maximize the overlap into the vertical transport. This pair was added after the initial design and construction described in [198, 199]. The horizontal coils are driven from a single 30V, 200A supply with each coil pair regulated by a single low-side FET.

In the vertical transport section, the atoms are transported vertically along the coil axis as shown in Fig. 3.18; by symmetry, the aspect ratio is 1. Unlike in the horizontal transport section, each individual coil has an adjustable current, and four of the coils are bipolar. Therefore, these coils are driven by FETS in an H-bridge configuration. The two top coils form the final quadrupole coil pair for evaporation. They are in a special H-bridge configuration so that they end up electronically in series to reduce common-mode current noise. The circuit details and datasheets are listed at <http://www.physics.utoronto.ca/~astummer/>. Power handling of the FETs is an important consideration; the power dissipated in each FET is

$$P_{FET} = I(V_{supply} - IR_{coil}). \quad (3.9)$$

If we are only driving a single coil we can adjust V_{supply} to minimize this power. However, since we are driving several coils from a single supply we do not have independent control, and a large amount of power may be inadvertently dissipated in the FET.

To estimate the transfer efficiency to a certain distance x_{trans} , we transport from the MOT to x_{trans} and back, and then measure the number of atoms recaptured by the MOT. Data for ^{87}Rb is shown in Fig. 3.19. If we assume that the loss rate is the same in each direction, the data indicates a 60% efficiency in the horizontal transport section and a 33% efficiency in the vertical transport section. There are three main loss processes. The first is due to collisions with the vapor in the MOT cell during the time it takes to enter the low pressure region of the pumping tube. This can be improved by operating with low vapor pressure and loading the MOT longer, or by increasing the transport speed during this stage. However, the latter is not a desirable solution because it leads to heating of the trapped gas. The second loss mechanism is “clipping”,

¹⁰The aspect ratio is defined as the gradient along y to the gradient along x . In a three-coil transport scheme, the aspect ratio and gradient are free parameters. However, if we constrain the aspect ratio to be constant as the trap zero moves and that the quadrupole coil pairs are unipolar, then the aspect ratio is fixed. For the geometry of our system, this fixed aspect ratio is 1.72.




Figure 3.19: (a) Fraction of ^{87}Rb atoms transported from the MOT cell to the end of the vertical transport section and to the science chamber. Overall, approximately 17% of the ^{87}Rb atoms in the MOT are transported to the science chamber. (b) Roundtrip efficiency to the end of the horizontal transport system (normalized to holding in the MOT cell) versus the initial width (controlled by temperature) for ^{41}K . To heat the gas, we varied the repump power during molasses. This is strong evidence that clipping against the differential pumping tube is a main loss process during transport.

which occurs when the RMS size of the cloud along direction i (in the direction of the differential pumping tube wall),

$$\bar{x}_i^2 = 2 \left(\frac{k_B T}{g_F m_F \mu_B \frac{\partial B_i}{\partial x_i}} \right)^2, \quad (3.10)$$

is comparable to the radial size of the tube. To minimize clipping we need to load a cold gas, which emphasizes the importance of proper cooling in the MOT and loading into the trap without heating. Eqn. 3.10 is somewhat deceiving because it seems to suggest that increasing the gradient will linearly decrease the size. However, we must increase the gradient adiabatically, and then the size only scales as the ratio of gradients to the $\frac{1}{3}$ power. Clipping is exacerbated in the vertical transport section because the weak gradients are in the radial tube direction. For example, in the horizontal transport section the weakest gradient is along x , but that direction does not clip on the tube. The weakest gradient along a radial tube direction is y , for which $\partial B_y / \partial y = 63 \text{ G/cm}$. On the other hand, in the vertical transport section, the weakest gradients are $\partial B_x / \partial x = \partial B_y / \partial y = 50 \text{ G/cm}$, and both are in radial tube directions. To give typical number, for $g_F m_F = 1$, $T = 300 \mu\text{K}$, and $\partial B_i / \partial x_i = 50 \text{ G/cm}$, the cloud diameter (twice the RMS radius) is $\approx 2.5 \text{ mm}$, and the smallest diameter of the transport tube is 11 mm . Clipping can also occur if the transport system and the tube are not properly aligned.

The main uncertainty in predicting clipping is the initial temperature in the magnetic trap. We observe that the transport efficiencies for three atomic species, ^{87}Rb , ^{40}K , and ^{41}K , are increasingly worse in that order, which is correlated with the fact that the MOT temperature also increases in that order. As a quantitative test, we heated the atoms in the magnetic trap before transport, and the resulting efficiency curve (Fig. 3.19) supports the conclusion that clipping increases as the temperature increases. The decreased efficiency of ^{40}K versus ^{87}Rb transport due to temperature issues may explain why all experiments utilizing fixed coil transport systems use sympathetic cooling by ^{87}Rb (§3.2.3). This is a major reason for our switch. A final (and large) loss process is due to the coil electronics, which is specifically an issue in the vertical

transport section. Small current spikes and dead zones in the crossover region of the bipolar channels was strongly correlated to atom loss. Ultimately these can be minimized with judicious and time-consuming optimization of the electronics and coil currents.

3.2.8 Science Chamber Evaporation

After the transport, the atoms are in the science chamber and held in a quadrupole trap (see Fig. 3.6 and Fig. 3.7). The science chamber vacuum is excellent ($P < 10^{-11}$ Torr), so the lifetime is long enough for evaporative cooling to the ultracold regime. Inside the top flange, centered on the imaging window, we have installed a 1.3GHz coil for ^{40}K evaporation (see [201] for details), which also doubles as an RF coil for m_F -changing transitions, e.g., for ^{87}Rb evaporation. As mentioned in §3.2.3, our original plan was to carry out an all ^{40}K evaporation scheme. Initially, the idea was to evaporate with a polarized sample (all ^{40}K atoms in $|9/2, 9/2\rangle$) in a pure magnetic quadrupole (QP) trap and then transfer to a dipole trap. This is the technique used in the ^{87}Rb Bose-Hubbard apparatus (see §3.1). Although the gas is polarized, fermions can still thermalize due to p-wave elastic collisions until they are suppressed when $T < 30\mu\text{K}$ [202]. Calculations (§G.2) indicated that this approach was feasible and that we could transfer a high fraction of the atoms to the dipole trap before Majorana losses become severe.

However, in practice our attempts to evaporative ^{40}K were unsuccessful. Although we observed atoms ejected from the quadrupole trap by applying microwaves, we did not see evidence of decreasing temperature and increasing phase-space-density. One issue is that we started with a small atom number ($N < 5 \times 10^7$) at high temperature ($T > 200\mu\text{K}$). However, the approach may be flawed. Another group [203] also observed very inefficient ^{40}K evaporation in a pure QP trap even though the standard Majorana rate formula (Eqn. B.11) does not predict a problem, and number loss is not observed. The survey of ^{40}K experiments (§3.2.3) revealed that a pure QP evaporation has not been implemented by any group. One possibility is that the large spin of ^{40}K ($F=9/2$) means that small spin rotations occur at a faster rate because the gap between m_F states is proportional to $1/F$ (see the adiabaticity criterion in §B.2). This does not cause large atom loss because these states are still trapped, but they are trapped in successively weaker potentials. This reduces the density and therefore the elastic collision rate.

To overcome this problem we decided to optically plug the QP trap [23, 101]. Here a repulsive optical potential (‘plug beam’) is sent through the low-field region of the QP trap. Atoms are repelled from this region due to the plug beam and therefore cannot undergo Majorana spin flips. To create our plug beam, we use a Ti:Sapphire laser tuned to 760nm with up to 400mW after passing through a single-mode optical fiber. We focus the beam through center of the QP trap to a $40\mu\text{m}$ waist. A schematic of the optical setup and the plug beam in the science chamber is shown in Fig. 3.20. Coarse alignment of the plug is typically done by evaporating the atomic gas in the presence of the plug beam and observing the deformation of the density profile. Fine alignment is done by minimizing Majorana losses (i.e., maximizing atom number) when evaporating to a low temperature. With only ^{40}K we are unable to perform any evaporation, so alignment was not possible. In our new $^{87}\text{Rb}/^{40}\text{K}$ setup, we can use the ^{87}Rb signal to align the plug. This works better than ^{40}K because the ^{87}Rb number is higher and ^{87}Rb evaporation in an unplugged QP is efficient. After the plug is aligned we will sympathetically cool the ^{40}K by evaporating ^{87}Rb .




Figure 3.20: (a) Layout of the plug beam optics designed to focus the beam to a $40\mu\text{m}$ waist at the center of the science chamber. Active feedback to the AOM from a photodiode signal after the fiber intensity stabilizes the plug beam. (b) Schematic of the plug beam going through the gas in the center of the science chamber.

3.2.9 Dipole Trap and Lattice Beams

After evaporation in the plugged QP trap, we will transfer to a crossed-dipole trap at $\lambda = 1054\text{nm}$ ($\lambda = 1053.57\text{nm}$). Once in the dipole trap we can further evaporate using a ^{40}K $|9/2, -9/2\rangle/|9/2, -7/2\rangle$ mixture, which has a Feshbach resonance at 202.1G (see Table. A.3). The ultracold ^{40}K gas in the dipole trap near the imaging window will be loaded into a series of 2D planes formed by a $\lambda = 1054\text{nm}$ (from the same source as the dipole trap) beam retro-reflected off the imaging window; the window has a high-reflection coating at that wavelength (see Fig. J.6). Next, a 2D cubic lattice will be applied in these planes by retro-reflecting 1054nm beams through the science chamber side viewports (Fig. 3.6). Images illustrating the atoms in the dipole trap and the lattice are shown in Fig. 3.21. Exact beam waists and ellipticity are still open issues. Some main considerations are having a high enough lattice depth during the imaging phase (§3.2.10), keeping the Fermi energy below the bandgap, and avoiding heating during transfer between the dipole trap and lattice.

The dipole and lattice beams all originate from the same source: an NP Photonics “Rock” fiber laser with a $< 10\text{kHz}$ linewidth, 150mW power, $\lambda = 1053.57\text{nm}$, and a 30GHz thermal tuning range. This light is sent directly into a NuFern 40W fiber amplifier, which adds $< 5\text{kHz}$ frequency noise (in quadrature). Our source wavelength is not arbitrary — it is a “magic” wavelength for the $4S_{1/2} \rightarrow 5P_{3/2}$ transition in ^{40}K (see calculation, §G.1). “Magic” in this case refers to the fact that the dipole potential formed in the $4S_{1/2}$ state, which couples predominantly to the $4P$ states, is identical to the shift in the $5P_{3/2}$ state (due to coupling to the $12D$ state). Since 1054nm is far-detuned from the $4S \rightarrow 4P$ transitions, the heating rate is small. Also, when imaging at 404.5nm, the 1054nm photons cannot ionize the $5P_{3/2}$ state. For reference, the ^{40}K $4S_{1/2}$ shift is $(0.137\mu\text{K} \times k_B)\text{mm}^2\text{W}^{-1}$, and the ^{87}Rb $5S_{1/2}$ shift is $(0.156\mu\text{K} \times k_B)\text{mm}^2\text{W}^{-1}$ ($\lambda = 1054\text{nm}$).

An outline of the dipole/lattice optical layout for beam splitting and stabilization is shown in Fig. 3.21. We use a series of AOMs to control the intensity of the beams. Photodiodes close to the chamber and after the fibers (for the lattice) are used to feedback to the AOM to provide intensity stabilization. Since the




Figure 3.21: (a) Dipole trap and lattice beam optical layout for splitting the source light and intensity stabilizing each beam. (b) Illustration of the atoms in the dipole trap close to the imaging window. (c) Illustration of the atoms in the lattice. The vertical lattice is much stronger than the in-plane lattice, so that the atoms are described as a 2D gas in each plane.

lattice beams are all retro-reflected, optical isolators after the fibers prevent the retro light from damaging or causing high intensity instabilities in the fibers. These optics are still in the process of being set up, so the layout in Fig. 3.21 is not final.

3.2.10 Imaging

After the ^{40}K gas is loaded into the lattice, our goal is to image the gas in-situ with single-site resolution. Our imaging system, shown in Fig. 3.22, is oriented so that the imaging plane is aligned with the 2D planes formed by the vertical lattice beam. Therefore, we will first need to select atoms in a single horizontal plane, otherwise the image will contain contributions from out-of-focus atoms. To isolate a single plane, we plan on applying a magnetic field gradient so that a Landau-Zener microwave sweep only transfers atoms in the imaging plane from $F = 9/2 \rightarrow F = 7/2$ [57, 204]. Then, we will apply resonant light on the $4S_{1/2}, F = 9/2 \rightarrow 4P_{3/2}, F' = 11/2$ cycling transition at 767nm to remove atoms in all other planes; atoms transferred to $F = 7/2$ will be unaffected. Once a single plane of atoms remains, we will ramp the lattice depth as high as possible, limited by total optical power¹¹ and fiber power handling. This procedure freezes the atoms to individual lattice sites.

Next, to realize single-site imaging requires two main ingredients: high signal-to-noise and high imaging resolution (i.e. comparable to the lattice spacing $d = 527\text{nm}$). To fulfill these demands we plan on performing fluorescence imaging (§4.1.1) by exciting the atoms on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition at $\lambda = 404.5\text{nm}$ (see Fig. 3.22). This type of imaging is unable to determine an arbitrary number of atoms per site as pairs of atoms on lattice sites will quickly photoassociate and become lost, leaving behind lattice sites with either zero or one atom. To achieve high signal-to-noise we must collect a large number of photons per atom ($> 10^3$); the number of photons collected is the product of the collection efficiency, scattering rate, and total imaging time. The collection efficiency is

$$\frac{1 - \cos(\theta)}{2}, \quad (3.11)$$

where θ is the maximum angle for rays from a point source to be accepted by the lens ($NA = \sin(\theta)$ is the numerical aperture; see Fig. 4.1). To realize large θ (high NA), we have obtained a custom 0.6NA objective designed and built by Special Optics which is diffraction limited at $\lambda = 405\text{nm}$ ($NA = 0.45$ at $\lambda = 767\text{nm}$) (see Fig. J.7 for dimensions). Our imaging system will therefore collect 10% of the scattered photons.

Long imaging times allow for large improvements in signal. This will be possible in our setup because the atoms are trapped to individual lattice sites. Furthermore, by red-detuning our excitation beams, we will laser cool these single atoms as they fluoresce. Laser cooling allows for long imaging times because the temperature does not continuously increase as it would for resonant imaging. However, for the atoms to remain trapped we need to ensure that the steady state temperature is much lower than the lattice depth. This is an advantage of 404.5nm imaging — although the scattering rate is lower when imaging at 404.5nm ($\Gamma = 2\pi \times 185\text{kHz}$) versus 767nm imaging on the $4S_{1/2} \rightarrow 4P_{3/2}$ transition ($\Gamma = 2\pi \times 6.04\text{MHz}$), the Doppler temperature is 5 times smaller. This significantly reduces the power requirements for the lattice during the imaging stage. If we excite on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition, we can also collect 767nm photons scattered at a rate of $\Gamma \approx 2\pi \times 700\text{kHz}$ from the multi-photon decay path. Collecting these photons allows for

¹¹For fixed power, the lattice depth can also be increased by dynamically decreasing the lattice beam waist.

background-free imaging (i.e., there is no excitation beam at 767nm). A further reduction in temperature is possible using sub-Doppler cooling, which is the procedure for single-site imaging in ^{87}Rb [57, 58]. These processes are not typically observed in free space cooling of ^{40}K , although this may be different in a lattice. Currently we have developed the laser system at 404.5nm and demonstrated free space cooling of ^{40}K on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition (§E). Cooling in the lattice is complicated by the lattice potential shift between the $4S_{1/2}$ and $5P_{3/2}$ states. Therefore, our lattice will be at the “magic wavelength” $\lambda = 1053.57\text{nm}$ where the dipole shift is the same in both states (§3.2.9).




Figure 3.22: (a) Simplified state structure for fluorescence imaging of atoms on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition at 404.5nm (see Table A.2 for the transition properties). For simplicity the intermediate states $5S_{1/2}$, $3D_{3/2}$, and $3D_{5/3}$ have been combined; see Fig. A.1 for the full state structure. Atoms are excited by a beam slightly red-detuned from the $4S_{1/2} \rightarrow 5P_{3/2}$ transition (blue, solid line). Once in the excited state, atoms can decay (dashed, grey lines). The wavelength of the scattered photon and the transition rate A are shown. The rate to decay straight back to $4S_{1/2}$ and emit a 404.5nm photon is $2\pi \times 185\text{kHz}$. There is also a decay path that emits three photons of approximate wavelengths, $3\mu\text{m}$, $1.2\mu\text{m}$, and 768nm , which occurs at a rate of $2\pi \times 1\text{MHz}$. The total linewidth of the excited $5P_{3/2}$ state is $2\pi \times 1.19\text{MHz}$. Because the transition can decay to the dark $F = 7/2$ ground state, a repump beam (red, solid line) is applied on the $4S_{1/2}, F = 7/2 \rightarrow 4P_{3/2}$ transition. (b) Schematic of the imaging setup. Atoms in the lattice (red beams) are excited from the side by the 404.5nm excitation beam (blue). The fluorescence is collected by a 0.6 NA objective through a $200\mu\text{m}$ thick sapphire window. (c) Photograph of the custom-built imaging window (vacuum side up). This piece is attached to the bottom of the top recessed flange.

The other ingredient of single-site imaging is high-resolution. Specifically, our imaging resolution must

be smaller than the lattice spacing¹². The standard resolution limit according to the Rayleigh criterion for incoherent light is [205]

$$\sigma = \frac{1.22\lambda}{2NA}. \quad (3.12)$$

Therefore, to achieve high resolution, we need high NA , small wavelength, and minimal aberrations (i.e., deviations from ideal diffractive imaging). Imaging at 404.5nm improves our resolution limit by almost a factor of two over imaging on the D2 transition at 767nm. For our 0.6NA, diffraction-limited objective the resolution is $\sigma = 410\text{nm}$ at $\lambda = 405\text{nm}$ ($\sigma = 1040\text{nm}$ for diffraction-limited $NA = 0.45$ at $\lambda = 767\text{nm}$) (see Fig. J.8 for objective MTF). To achieve a high NA , the working distance is 2.43mm. Since the objective is in air and the atoms are in ultra-high vacuum, we need to use an extremely thin, low aberration window. This window consists of a $200\mu\text{m}$ thick, 12mm diameter sapphire disk (Meller Optics) custom brazed to a vacuum flange (UKAEA) with a 5mm diameter window aperture, as shown in Fig. 3.22.




Figure 3.23: Sample image of the MRS-5 through the Special Optics objective. The zoomed in region has a 600nm pitch (i.e., each line is 300nm wide).

We have investigated a number of methods to test ex-situ the objective resolution and window aberrations. For coarse testing, we use a 1951 USAF resolution target. Since the USAF target does not have fine enough features to reach the diffraction limit, we used the MRS-5 (Geller MicroAnalytical Laboratory) target, which has features as small as 80nm (see Fig. 3.23). Although useful, illuminating this target is difficult because it is tungsten layered on silicon and therefore it must be illuminated through the objective. Other tests include imaging the slanted edge of a cleaved silicon wafer [206] and imaging a grid of 200nm diameter holes spaced 500nm apart in a 100nm thick layer of silver (fabricated by Sang-Hyun Oh group at the University of Minnesota [207]). To test the window aberrations, we have performed some of these resolution measurements with and without the window and we also tested the window phase in a Twyman-Green interferometer [208]. We measured the RMS phase roughness to be 0.06λ .

To take site-resolved images, we need to rigidly mount the objective with respect to the lattice and have three-axis translation combined with two-axis tilt control. Planning and initial construction of this mount is underway. Another future consideration is using the objective to send beams into the experiment for

¹²The maximum resolution for which one can claim single-site imaging is a complicated function of the number of photons collected, noise, coherence of the fluorescence light and the atomic density in the lattice (see §4.1.5). However, the single-site criteria $\sigma \leq d$ is approximately correct.

high-resolution manipulation of atoms trapped in the lattice.

3.2.11 Computer Control




Figure 3.24: Schematic of computer control and monitoring in the experiment.

To run a typical experimental sequence requires controlling all the components of the apparatus with μs timing and then acquiring and analyzing the resulting image. The core of this process is a set of output analog and digital signals controlled by an FPGA-based computer ('ADWIN' Pro II from Jäger). The ADWIN timing array is loaded via TCP/IP from a standard Windows computer ('Control Computer'). A MATLAB program ('Lattice Sequencer') on the control computer compiles user generated MATLAB scripts into the timing array. The Lattice Sequencer also sends information about the current sequence, i.e., the value of the parameter being changed during each run, to a second computer ('Imaging Computer'), which is connected directly to the camera. At the end of the experimental sequence, a MATLAB program ('Lattice Acquisition') loads images from the camera and processes them using information sent by the Lattice Sequencer. For example, a typical set of experimental runs might involve cycling through a parameter such as a given laser power P , which is controlled via an analog voltage from the ADWIN. For each sequence, the control computer, based on the instructions in a MATLAB file, determines the ADWIN voltage required for laser power P . The control computer loads that information into the ADWIN and sends a file with the variable P to the imaging computer. At the end of the experimental sequence, the imaging computer saves the image with the value of P in the filename.

At the hardware level, the ADWIN is the direct interface with the experiment. The analog voltages are $\pm 10\text{V}$ and the digital outputs are standard TTL. The ADWIN has a 300MHz clock, but we only update every $5\mu\text{s}$ based on our requirements. To prevent noise, ground loops, and damage to the ADWIN, all the digital channels are isolated at the ADWIN, and analog signals are isolated at the experimental component if necessary.

It is also important to be able to monitor the apparatus during the cycle to ensure the ADWIN sequence is implemented as expected. For passive monitoring we have a National Instruments 32 channel analog input card (NI-PCI-6224) connected to the experiment to monitor, for example, the MOT level, coil currents, and laser powers. This is based on the same device in the ^{87}Rb apparatus [41]. For critical components it is also necessary to have active monitoring. Most of the magnetic field coils require large currents during operation (in excess of 100A), and therefore they are water cooled. If the water cooling fails and/or the coils are left on too long, the coils could quickly overheat and cause irreparable damage to themselves and the apparatus. If the water flow (Proteus 08004BN4 flow monitor) falls below a set value or the coil temperatures (monitored with 30k thermistors, DigiKey Part#490-4662-ND) exceed set values, the current supplies are automatically shut off via a relay. The monitor circuit diagram is given in §H.2. An overview of these various components and how they connect is shown in Fig. 3.24.

Chapter 4

Experimental Toolkit — Probes and Techniques

Using the tools described in the last chapter (§3), we prepare an ultracold gas in a lattice which is well-described by the models we are simulating. However, this is only the first aspect of the full quantum simulation. We also need probes to read out the results and techniques to manipulate the initial state of the system.

A number of standard probes and techniques have been developed over a number of years for ultracold gas experiments, and more specifically, for optical lattice experiments. In this work, we extend some of these existing probes/techniques and also develop new ones of our own. In this chapter, I summarize all the probes/techniques we use, with particular detail given to our new work. Our experimental work on thermometry probes is given in §7, on impurity probes in §8, and non-equilibrium techniques in §6.

4.1 Probes

Probes are the collection of physical elements, procedures, and analyses we use to detect the atoms and measure the state of the system. All our probes have one common theme — using near resonant light to image the atoms. However, how we manipulate the atoms just before the image and the image analysis set different probes apart. Since imaging is common to all probes, I will describe it first as a separate section.

The probes described in this section are certainly not a complete list for ultracold atom experiments. For instance, a different class of probes ionize the atoms with light [209] or electrons [210] and directly detect the ions. Another class use resonant light, but do not form an image and instead measure the diffraction pattern of the scattered light in the far-field [211, 212] to perform crystallography. This is one proposed method for measuring spin correlations in the Hubbard model [213].

4.1.1 Imaging - Absorption and Fluorescence

Atoms interact strongly with electric fields (i.e., light) via the electric-dipole interaction (Eqn. B.1). In §2.1.1, we discussed the case when the field frequency is far-detuned from an atomic resonance frequency. For this situation, the atom-field interaction shifts the energy of the atomic states proportional to the field intensity, which can be exploited to make atomic potentials. In this far-detuned limit, changes to the electric field from the interaction are small and typically neglected. However, close to an atomic resonance, atoms have a dispersive and absorptive effect on the field, both of which can be used for imaging. For our purposes we will ignore the dispersive effects by imaging on resonance where these effects vanish. On resonance, the atoms absorb photons from an incoming beam and re-radiate them in a dipole pattern (see §B.7.3). For




Figure 4.1: Schematics of (a) fluorescence and (b) absorption imaging setups in $f - 2f - f$ configuration. The collection angle of the imaging is θ . For fluorescence, an excitation beam is incident from an angle greater than θ and only the scattered light is imaged. In absorption imaging, the excitation beam is sent straight into the imaging system and the shadow created by light scattered out of the beam (black lines) is imaged.

dense enough samples we need to worry about re-absorption and coherent effects such as superradiance, which we will ignore in the following discussion.

To use light scattering for imaging we need to transmit the electric field at the atoms to a CCD detector using an imaging system. This is performed by optical elements (i.e., lenses) which collect and refocus the field. An imaging system has a maximum angle of acceptance (see Fig. 4.1) which defines the numerical aperture, $NA = n \sin(\theta)$, where n is the index of refraction of the medium between the emitter and lens. If we image a point emitter, the fraction of light collected is $\frac{1 - \cos(\theta)}{2}$, and the resolution is $\approx 0.61\lambda(NA)^{-1}$ (see §4.1.5 and §F for more details).

There are two complementary types of imaging - absorption and fluorescence. Fluorescence imaging is conceptually more straightforward. For this type we send a beam ('excitation beam') through the atoms, and a lens collects light scattered from the atoms perpendicular to the excitation beam (see Fig. 4.1). For resonant light and no re-absorption, the detected signal W at position (x, y) on the detector is

$$W(x, y) = \frac{1 - \cos(\theta)}{2} \gamma \int_0^{\tau_0} d\tau \int dz n(x, y, z, \tau) \frac{\Gamma}{2} \frac{I(x, y, z, \tau)}{I_{sat} + I(x, y, z, \tau)}, \quad (4.1)$$

where γ is the detector efficiency, τ_0 is the imaging time and I_{sat} is the saturation intensity (Eqn. B.70). $I(x, y, z, \tau)$ is the local intensity of the excitation beam, noting that I is itself a function of n and the beam direction since the beam is attenuated as it travels through the gas. Therefore, fluorescence imaging is difficult to calibrate. However, it does have two distinct advantages. Since the detector signal is directly proportional to τ_0 , we can increase the signal by collecting fluorescence for long times ($\tau_0 \rightarrow \infty$), which is a feature we exploit for single-site imaging (see §4.1.5). Also, there is no signal for zero atoms (assuming low detector dark counts), so fluorescence imaging is useful for discriminating between zero and ≥ 1 atoms.

In absorption imaging, the imaging lens is placed in the path of the laser beam ('probe beam') we send through the atoms as illustrated in Fig. 4.1. Each atom scatters light out of the probe beam, which is therefore attenuated. For a sufficiently weak probe beam ($I \ll I_{sat}$) the attenuation is linear, so the detector signal is,

$$W(x, y) = \gamma \int_0^{\tau_0} d\tau I_0(x, y, \tau) e^{-\sigma \int n(x, y, z, \tau) dz} \quad (4.2)$$

assuming the depth of field of the imaging system is larger than the atom cloud, σ is the absorption cross-section (Eqn. B.77), and γ and \mathfrak{t}_0 are defined as above for fluorescence. If we take an image a time $\tilde{\mathfrak{t}}$ later after the atoms are gone,

$$W_2(x, y, \tilde{\mathfrak{t}}) = \gamma \int_{\tilde{\mathfrak{t}}}^{\tilde{\mathfrak{t}}+\mathfrak{t}_0} d\mathfrak{t} I_0(x, y, \mathfrak{t}). \quad (4.3)$$

If n and I are time-independent we can take the logarithm of the ratio of these signals,

$$\ln \left(\frac{W_2(x, y)}{W(x, y)} \right) = \sigma \int n(x, y, z) dz \quad (4.4)$$

to get a signal which is proportional to the atomic density with no calibrations. Importantly, the final signal does not depend on inhomogeneities in the probe beam. Also, the final signal is not a function of the imaging time and therefore we can take fast images of the atoms before the distribution changes, limited only by photon shot noise. On the other hand, we cannot increase our signal by imaging for arbitrarily long times as we can in fluorescence¹. In practical realizations of fluorescence and absorption imaging there are technical and fundamental noise sources (i.e., photon shot noise). These issues will be discussed in more detail in §F.

Although atoms will scatter close to any resonance, optimal imaging occurs when scattering on a cycling transition. A cycling transition is a ground state-excited state pair coupled by the excitation/probe beam, where the excited state can only decay to this one ground state. In alkali atoms this is the $nS_{1/2}|F, \pm F\rangle \rightarrow nP_{3/2}|F+1, \pm(F+1)\rangle$ transition driven using $\hat{\sigma}^\pm$ polarized light² (Eqn. B.22), where $F = I + 1/2$ (I is the nuclear spin) and $nS_{1/2}$ is the ground state (see state diagrams Figs. A.1 and A.2). If our excitation/probe beam is on resonance with that transition, it will not excite atoms in the other hyperfine ground state $\tilde{F} = I - 1/2$. Therefore, we can control the effective density by putting (or keeping) some atoms in \tilde{F} . For example, in the ^{87}Rb apparatus we prepare the atoms in the $F = 1$ state, but the cycling transition is $F = 2 \rightarrow F' = 3$. To image, we transfer a controlled fraction of atoms from $F = 1 \rightarrow F = 2$ (“partial repumping”); see §F.2 for a more technical discussion. Partial repumping is particularly important for absorption imaging due to saturation. When n is too large, the slope $\frac{\partial V}{\partial n}$ becomes too small and the signal is dominated by noise and background light (e.g., off-resonant light in the probe beam path). Empirically for the ^{87}Rb apparatus, this upper limit is $e^{-\sigma \int n} = e^{-2}$.

Ultimately, absorption and fluorescence imaging give us a method to convert the integrated density distribution of the atoms into a measureable signal. Although the equilibrium density in the trap is an important quantity to measure (§4.1.5), we are not limited to only this probe. By precisely manipulating the atoms before imaging the density, we can perform a number of probes which are described in the following sections.




Figure 4.2: TOF Expansion of a $T = 10\text{nK}$ gas from a $\nu = 50\text{Hz}$ harmonic trap using Maxwell-Boltzmann statistics. (a) Distributions after different expansion times obtained by summing the expanded harmonic oscillator eigenstates (Eqn. 4.6) with Maxwell-Boltzmann factors. Comparison of the exact solution to the semiclassical expansion and semiclassical momentum distributions after (b) 5ms expansion and (c) 20ms expansion. The semiclassical expansion is a very accurate description of the exact expansion. After 20ms the distribution is well-described by the momentum distribution with $p = mx/\tau$.

4.1.2 Momentum Distribution from Time-of-Flight (TOF) Expansion

To probe the momentum distribution we use time-of-flight (TOF) imaging. In this probe, the atoms are quickly released from their trapping potentials, and then allowed to expand for some variable amount of time, typically 20ms, before we take an absorption image. To see that the density distribution after expansion is the momentum distribution, we use the propagator for free particles (assuming no interactions) [112]

$$G(x, x', \tau) = \sqrt{\frac{m}{2\pi i \hbar \tau}} e^{im(x-x')^2/2\hbar\tau}, \quad (4.5)$$

so that,

$$\Psi(x', \tau) = \int dx \Psi(x, \tau=0) G(x, x', \tau). \quad (4.6)$$

After sufficient time³,

$$\Psi(x', \tau) = \sqrt{\frac{m}{2\pi i \hbar \tau}} e^{im(x')^2/2\hbar\tau} \int dx \Psi(x, \tau=0) e^{-imx'x/\hbar\tau}, \quad (4.7)$$

¹Absorption imaging for longer times increases the signal-to-noise ratio by decreasing the relative photon shot noise, but this improvement is limited (e.g., by CCD depth). The optimal imaging technique — fluorescence or absorption — is situational and depends on the details of the setup (e.g., the number of atoms), the noise sources, and the information we want to obtain from the image.

²In practice, it is difficult to eliminate all light in other polarizations, and so there is always some probability of off-resonant excitation on a non-cycling transition.

³For a harmonic trap $\tau \gg \omega^{-1}$

which is the Fourier transform with $\hbar k = mx'/\mathfrak{t}$. This means that phase inhomogeneities in the trap map to density inhomogeneities after expansion, which can then be measured using our imaging techniques. For example, when we look at lattice diffraction (i.e., for calibrating the lattice) we pulse the lattice on which imparts a phase on the atoms in the trap. If we were to take an in-trap image we would see no evidence of that phase, however, after expansion the phase results in diffraction peaks. Although the above presentation is focused on single-particle expansion, for many-body states we similarly expand the field operators. Information about these states can be extracted from higher order correlations in the momentum distribution [214].

For finite temperatures we have to take the thermal average of Eqn. 4.6. However, we can also use a semiclassical approximation where the initial density and momentum (phase-space) distribution, $f(x, p)$, is semiclassical (Eqn. 2.23), and we assume the atoms ballistically expand so that,

$$n(x', \mathfrak{t}) = \int dx f \left[x, \frac{m}{\mathfrak{t}}(x' - x) \right] \frac{1}{h^3} \int dx \left[e^{\left(\frac{m}{2\mathfrak{t}^2}(x' - x)^2 + V(x) - \mu \right)/kT} \pm 1 \right]^{-1}. \quad (4.8)$$

For harmonic traps, the final distribution is a polylogarithm (i.e., Eqn. 2.26), where (after integrating over the probe direction),

$$n(x', \mathfrak{t}) \propto \mp Li_2 \left[\mp e^{-\frac{(x')^2}{2\sigma^2} + \frac{\mu}{kT}} \right] \quad (4.9)$$

$$\sigma^2 = \frac{kT}{m} \left(\mathfrak{t}^2 + \frac{1}{\omega^2} \right) \quad (4.10)$$

In the high-temperature limit, where Maxwell-Boltzmann statistics are valid, generically,

$$\langle (x')^2 \rangle = \frac{\mathfrak{t}^2}{m^2} \langle p^2 \rangle + \langle x_0^2 \rangle \quad (4.11)$$

which for a harmonic trap is,

$$\langle (x')^2 \rangle = \frac{kT}{m} \left(\mathfrak{t}^2 + \frac{1}{\omega^2} \right), \quad (4.12)$$

which is a standard method for using the width after expansion as a temperature probe. A comparison between the full quantum expansion and the semiclassical expansion is illustrated in Fig. 4.2, which demonstrates that the semiclassical approximation is very accurate. An advantage of the TOF method is evident from Eqn. 4.11 which shows that the size increases linearly with expansion time (at long times). This greatly reduces the required imaging resolution and imaging complications due to high densities. The main assumption of this probe is that there are no interactions. If there are, then Eqn. 4.6 is not valid. While there are nearly always some interactions, it is typically the case where some part of the gas (i.e., the low-density edges) is well-described by Eqn. 4.6. Also, after a short period of expansion, interactions are negligible because the density has decreased. However, for strongly interacting systems the dynamics immediately following release are a current area of research [215].

During TOF expansion, the atoms experience a constant downward acceleration due to gravity, and, so

absent any other forces, the center of mass position is $x = x_0 + \frac{1}{2}gt^2$. If we add a magnetic field gradient along gravity during expansion we can increase or decrease the effective g depending on the magnetic moment,

$$\tilde{g} = g + \frac{\mu_B}{m} g_F m_F \frac{\partial |B|}{\partial z}. \quad (4.13)$$

There are two advantages to adding a magnetic field gradient. First, we can completely cancel gravity, $\tilde{g} = 0$, and allow for a very long expansion, which is important for revealing very small features of the gas (i.e., vortices). Second, the gradient separates different $|F, m_F\rangle$ states (as in the classic Stern-Gerlach experiment), so that we can gain access to the distribution of each individual state in a multi-state experiment.

Although we have focused on the momentum distribution, TOF imaging also reveals the center of mass momentum. This is important for experiments on transport (§4.2.3). The total center of mass momentum is not changed by interactions because of translational symmetry.

4.1.3 Quasimomentum Distribution from Bandmapping and TOF

Using TOF expansion to probe the momentum distribution is useful for harmonically trapped systems, but there are additional complications when the system is in a lattice. These are evident from the tight-binding (§2.3.2) form of the lattice momentum distribution,

$$n(p) = |w(p)|^2 \sum_{j,k} e^{ipd/\hbar(j-k)} \langle \mathbf{a}_j^\dagger \mathbf{a}_k \rangle \quad (4.14)$$

First, the envelope of the distribution is the Fourier transform of the Wannier function (Eqn. 2.71) reflecting the zero-point momentum of the atoms in the ground state of each lattice site. Second, because of the lattice periodicity, the momentum is also periodic. Since both these terms do not add thermodynamic state information, the momentum distribution is redundant. Instead, the simplest distribution is the quasimomentum distribution, $f(q)$, defined for $|q| < q_B$ and which is zero outside this region. We can write Eqn. 4.14 in terms of $f(q)$,

$$n(p) = |w(p)|^2 \sum_j f(q + P_j) \quad (4.15)$$

where P_j is a reciprocal lattice vector. Therefore, it would be useful if we could directly probe the quasimomentum distribution. In theory, $f(q)$ can be obtained by measuring the momentum distribution and then using Eqn. 4.15. However, since the imaging process integrates along the probe direction, this transformation can be non-trivial, especially if the lattice and probe axes are not aligned.

Fortunately, the quasimomentum distribution can be directly imaged by bandmapping immediately before TOF expansion. Bandmapping is a technique in which the lattice is ramped down slowly compared to the bandgap, but fast compared to all other timescales [216, 217]. The effects of this procedure are illustrated in Fig. 4.3. The main effect is to greatly suppress the lattice modulations in the density. This maps quasimomentum in the lattice to momentum in the distribution with the lattice off ($q \rightarrow p$), and then the TOF expansion maps the momentum distribution to density (as discussed in §4.1.2). Although we are only considering this as a probe for atoms in the lowest band, another benefit of bandmapping is that after




Figure 4.3: (a) In the top red curve, we plot the distribution of atoms in a lattice ($s = 6, d = 400\text{nm}$) plus harmonic trap ($\nu = 54.5\text{Hz}$) at $T = 15\text{nK}$ ($k_B T/t \approx 1.7$). The black curve is the distribution after linearly ramping the lattice down in $750\mu\text{s}$ (bandmapping). The main effect is to remove the density modulation due to the lattice, but the overall envelope remains the same. The bottom curves are after a 20ms TOF expansion. Bandmapping greatly reduces the atomic population outside the first Brillouin zone. (b) 2D quasimomentum distributions after imaging using the probe geometry of the ^{87}Rb experiment as given by Eqn. 4.19. The top image is at $k_B T/t = 1.7$, which is the same as the data in (a) and the bottom image is for $k_B T/t = 5$ to demonstrate how the band fills.

TOF it spatially separates the populations in the different bands, which is used in a number of experiments (e.g., [218]).

To measure the thermodynamic state from bandmapping, we need an expression for the quasimomentum distribution in terms of T and μ . In the same way we used the semiclassical approximation for momentum in §4.1.2, we can also write down an expression for the quasimomentum distribution (see 5.1 for details),

$$n(q_1, q_2, q_3) = \frac{1}{(2\pi)^3} (4\pi) \int_0^\infty dr r^2 \left[\zeta^{-1} e^{2t\beta \sum_i [1 - \cos(q_i \pi / q_B)] + \frac{1}{2} \beta m \omega^2 r^2} \pm 1 \right]^{-1} \quad (4.16)$$

$$= \mp \frac{1}{8\pi^{3/2}} \left(\frac{\beta m \omega^2}{2} \right)^{3/2} \text{Li}_{3/2} \left[\mp \zeta e^{-2t\beta \sum_i [1 - \cos(q_i \pi / q_B)]} \right] \quad (4.17)$$

where $\beta = (k_B T)^{-1}$ and $\zeta = e^{\mu\beta}$ is the fugacity. The final probe integrated distribution depends sensitively on the imaging setup since Eqn. 4.17 lacks rotational symmetry. If we image along a lattice direction, then

$$n(q_1, q_2) \propto \sum_{n=1}^{\infty} \frac{(\mp)^n \zeta^n e^{-2nt\beta [3 - \sum_{i=1}^2 \cos(q_i \pi / q_B)]} I_0 [2nt\beta]}{n^{3/2}}, \quad (4.18)$$

where $I_0(x)$ is the modified Bessel function of the first kind. If we integrate along \hat{x} and the lattice vectors as given by Eqn. 3.4, which is the setup for the ^{87}Rb apparatus, then we can express the distribution as a doubly infinite sum

$$n(\alpha_1, \alpha_2) \propto \sum_{n=1}^{\infty} (\mp)^n \frac{e^{-2nt\beta [3 - \cos(b)]} \zeta^n}{n^{3/2}} \left\{ 2(\pi - a) I_0 [\cos(a) 2nt\beta] + \sum_{j=1}^{\infty} 4I_j [\cos(a) 2nt\beta] \frac{\sin[j(\pi - a)]}{j} \right\}, \quad (4.19)$$

where α_1, α_2 are coordinates along \hat{x} and \hat{y} , $a = |\alpha_1 + \alpha_2|/\sqrt{2}$, and $b = |\alpha_1 - \alpha_2|/\sqrt{2}$. In §G.5 we derive Eqn. 4.19 from Eqn. 4.17.

The central component of this probe is that the bandmapping procedure maps quasimomentum to momentum. However, this is only exact for non-interacting systems and when the ramp-down is very long with respect to the bandgap. Therefore, one aspect of this thesis work is to understand the limitations of this probe. In §7.1 we look at using this probe in more detail for thermometry of weakly interacting bosons in the lattice. There we find that bandmapping does not reproduce the quasimomentum distribution at the level required for thermometry when the temperature is comparable and greater than t/k_B . In §6.2, we explore the bandmapping timescales for an interacting system by looking at the condensate fraction as a function of the bandmapping time. We find that interactions quickly redistribute quasimomentum making it impossible to be both slow compared to the bandgap and fast compared to these processes. Ultimately, bandmapping as a quasimomentum probe needs to be used cautiously after considering these limitations.

4.1.4 Condensate Fraction from TOF

The probes in the last two sections focused on measuring the state of the system from the low order moments of the distributions after TOF expansion (i.e., the RMS size). However, when the system is Bose condensed (§2.2.1), a new probe — condensate fraction — emerges. Condensate fraction (N_0/N) is the fraction of

atoms in the macroscopically occupied single particle ground state⁴, and, together with the total number, uniquely determines T and μ . For a weakly interacting harmonically trapped gas, the relation between N_0/N , N and T , μ is analytic (Eqn. 2.57). In the lattice, numerical methods are required, but these are very accurate (see Fig. 2.11).

When the system is condensed, the distribution of particles is bimodal in an appropriate basis⁵. This is the essential property that makes condensate fraction a useful probe. Instead of determining the moments of the distribution, which is susceptible to technical noise (i.e., distortions during expansion) and interactions during TOF, condensate fraction only requires identifying the condensate mode and taking the ratio with the total number. As the modes (condensate and non-condensate) become more separated, the probe is less sensitive to the functional form of the fit. In the limit that the modes are completely separate, condensate fraction can be measured by number counting.

The optimal distribution (in-trap or momentum) for probing condensate fraction is a non-trivial property of the model we are probing. For non-interacting atoms in a harmonic trap, the spatial and momentum distribution are equivalent, except for the caveat that we need to expand to image the momentum distribution (§4.1.2). If the system has strong repulsive interactions, then the condensate and non-condensate are spatially well-separated, although the boundaries between the two may not be easy to identify. Repulsive interactions also counteract the effect of the trap, which increases the size of the condensate in trap and decreases the size in momentum, but there is nothing preventing overlap in momentum between the modes. Additionally, these interactions cannot be ignored during expansion unless they can be turned off with a Feshbach resonance [219]. Therefore, there has been work on both probing in-trap [220] and after expansion. Our work focuses on the latter, since the issues created by interactions during expansion are mitigated by the technical issues which are improved by expansion, such as resolution, stray fields, and re-absorption. A technical absorption imaging issue for condensate probes, which is present both in-trap and after expansion, is caused by the disparate densities of the condensate and non-condensate modes. Because of the saturation issue discussed in §4.1.1, we use two absorption images from back-to-back experimental cycles to get a high signal-to-noise distribution for fitting. First we take a partially repumped image (see §F.2 for calibration details) to get an unsaturated image of the condensate. Next, we take a fully repumped imaging where the condensate peak is saturated to get high signal for the non-condensate region.

For the weakly repulsive gas in a harmonic trap (§2.2.2), the effect of interactions during expansion is well-described by the GP equation (Eqn. 2.47). The largest effect is on the condensate itself due to its high density. In the trap, interactions result in a Thomas-Fermi (TF) profile for the condensate (Eqn. 2.49), which is significantly larger than the non-interacting condensate (i.e., the ground state of the harmonic oscillator). To calculate the expansion of the condensate, we solve the time-dependent GP equation with no trapping

⁴The macroscopically occupied single particle ground state (i.e., $\Phi(x)$ in Eqn. 2.47) is not necessarily the same as the non-interacting ground state (see Ref. [53] for a more rigorous discussion).

⁵For example, non-interacting particles in a box always have a uniform average density, so only the momentum distribution is bimodal.

potential. It has been shown [221] that starting from a TF profile, the expansion just rescales the widths as

$$\begin{aligned} r_{TF,i}(\mathbf{t}) &= \xi_i(\mathbf{t})r_{TF,i}(0), \\ \frac{d^2\xi_i}{d\mathbf{t}^2} &= \frac{\omega_i^2}{\xi_i\xi_1\xi_2\xi_3}. \end{aligned} \quad (4.20)$$

Therefore, the common zero-order fit for the TOF expansion of the condensate gas is,

$$n(x', \mathbf{t}) = n_{th}(x', \mathbf{t}) + n_{TF}(x'/\xi(\mathbf{t})) \quad (4.21)$$

where $n_{th}(x')$ is the probe integrated semiclassical expansion of the thermal gas given by Eqn. 4.9 (with $\mu = 0$), and $n_{TF}(x'/\xi(\mathbf{t}))$ is the probe integrated TF profile (assuming integration along z),

$$n_{TF}(x, y) = \begin{cases} \frac{4}{3}\mu r_{TF,z} \left(1 - \frac{x^2}{r_{TF,x}^2} - \frac{y^2}{r_{TF,y}^2}\right)^{3/2}, & \frac{x^2}{r_{TF,x}^2} + \frac{y^2}{r_{TF,y}^2} \leq 1 \\ 0, & \text{else} \end{cases} \quad (4.22)$$

This distribution is shown in Fig. 4.4. In practice we leave the widths and amplitudes of these two functions as free-parameters in the fit, so

$$n_{\text{fit}}(x, y) = n_{th} Li_2 \left[e^{-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)} \right] + n_0 \left(1 - \frac{x^2}{r_{TF,x}^2} - \frac{y^2}{r_{TF,y}^2}\right)^{3/2}, \quad (4.23)$$

and the condensate fraction from these parameters is,

$$\frac{N_0}{N} = \frac{\frac{2\pi}{5} n_0 r_{TF,x} r_{TF,y}}{\frac{2\pi}{5} n_0 r_{TF,x} r_{TF,y} + 2\pi n_{th} \sigma_x \sigma_y Li_3[1]/Li_2[1]}. \quad (4.24)$$

Using Eqn. 4.9 and Eqn. 2.51 with Eqn. 4.20 for a spherical trap, we can estimate the separation between the non-condensate and condensate modes by looking at the ratio of the widths after expansion,

$$\frac{\sigma^2}{r_{TF}^2} \approx 0.44 \left(\frac{a_{ho}}{a}\right)^{2/5} \frac{\left(1 - \frac{N_0}{N}\right)^{1/3}}{\left(\frac{N_0}{N}\right)^{2/5}} N^{-1/15}. \quad (4.25)$$

For $\nu = 50\text{Hz}$, $N = 10^5$, $a = 5\text{nm}$ and $m = m_{Rb}$, this ratio is < 1 when $N_0/N \approx 0.9$. Eqn. 4.25 scales weakly with our adjustable experimental parameters (ω and N), so the mode overlap is essentially a fixed function of N_0/N . Some techniques have been explored to fully separate the modes using time-dependent potentials such as Bragg scattering [126] and by pinning the non-condensate with a lattice [222]⁶.

Eqn. 4.21 is essentially exact in the limits $N_0/N \rightarrow 0$, $N_0/N \rightarrow 1$, but is approximate in between these limits, particularly near the edge of the condensate. The next order approximation is to treat the condensate as a time-dependent mean-field potential for the non-condensate atoms,

$$V(x, t) = 2gn_{TF}(x, t) \quad (4.26)$$

where $n_{TF}(x, t)$ is the expanded TF profile given by solving Eqn. 4.20 and g is the contact parameter

⁶We see a similar effect in our lattice transport measurements described in §6.1.




Figure 4.4: 1D slices through the 2D distribution of a condensed, weakly interacting harmonically trapped Bose gas after TOF expansion and imaging. (a) Experimental ^{87}Rb data (7 shot average) after 20ms expansion and $\nu = 35.76(8)\text{Hz}$. The actual data was imaged at lower peak OD using partial repumping and has been converted to the on-resonance OD using a calibration factor. These data correspond to roughly $N_0/N = 0.5$. (b) The distribution given by Eqn. 4.21 for 10ms expansion, $T = 90\text{nK}$, and $\nu = 50\text{Hz}$ (^{87}Rb) corresponding to $N_0/N = 0.505$ and $N_0 = 10^5$ atoms. (c) Calculated expansion data for the same parameters as (b), and then integrated along the probe direction. Details for the calculation are given in §G.3. (d) Comparing the distributions from (b) and (c). The main discrepancies are at the condensate edge and the peak value of the distribution.

(Eqn. 2.15). The condensate mean-field potential is twice as strong for the non-condensate atoms versus the condensate atoms themselves (i.e., the non-linear term in Eqn. 4.20) because of the exchange term in the interaction. The non-condensate expansion with the potential in Eqn. 4.26 was solved by Monte Carlo sampling the classical particle trajectories in [223]. Here we do a full calculation by first determining the eigenstates for the harmonic oscillator with the condensate mean-field (§2.2.2) and then propagating each state using the time-dependent Schrödinger equation with the potential given by Eqn. 4.26. Finally, we add the states together using Bose-Einstein statistics to get the finite temperature expanded distribution. The full details of the calculation are given in §G.3 and the resulting distribution, compared to experimental data and Eqn. 4.21, is given in Fig. 4.4.

The calculated data deviates from the zero-order fit (Eqn. 4.21) near the edge of the condensate, whereas near the tails of the distribution the two are almost identical. This motivates a hybrid ‘multi-step’ approach that applies Eqn. 4.21 in piecewise fashion to different areas of the distribution. The general idea is to avoid directly fitting the distribution near the condensate edge. First, we fit the entire image using the full equation to determine the condensate location and size. Next, we remove (mask out) a region 10% larger than the condensate and fit the remaining distribution to the non-condensate part of Eqn. 4.21. Then we subtract this fit from the initial image, including the masked out region, and fit the residual to just the condensate part of the fit. This approach was formalized and quantified in [224]. The multi-step approach relies on some theoretical input for the functional form of the distribution.

To alleviate the need for theoretical input to the fit, another fitting approach is the ‘heuristic’ method. The goal of this method is to separate out the two modes, condensate and non-condensate, and sum to get the number of non-condensate atoms. First, we do a Gaussian fit, with the condensate peak manually masked out, to account for most of the non-condensate atoms (the number of atoms under the Gaussian is N_{Gauss}). Next, we subtract the Gaussian from the distribution, zoom into a small region around the condensate, and fit the distribution to the TF function (Eqn. 4.22). The number of atoms underneath the TF peak N_{peak} is a combination of condensate and non-condensate atoms. We use an empirical estimate of the number of non-condensate atoms under the TF peak by assuming the non-condensate distribution is flat. This assumption is motivated by the results of our mean-field expansion calculation; see (c) of Fig. 4.4. We sum around the perimeter of the condensate to get the average density, and multiply this average density by the TF area to determine the number of non-condensate atoms under the peak N_{under} . The remaining atoms N_{sum} (i.e., atoms not under the Gaussian or the TF peak) are counted by summing over the entire image with the condensate masked out. The condensate fraction from this procedure is

$$\frac{N_0}{N} = \frac{N_{\text{peak}} - N_{\text{under}}}{N_{\text{sum}} + N_{\text{Gauss}} + N_{\text{peak}}}. \quad (4.27)$$

This fit procedure is demonstrated in Fig. 4.5. The heuristic method is important for situations with little theoretic input, such as in the lattice or for general quantum simulations. Using the calculated data we quantify the different approaches by fitting for the condensate fraction and comparing to the known value. The results are summarized in Fig. 4.6. Although the heuristic method is the best in terms of relative error, in practice we use the multi-step approach since it is more robust against noise. The precision is also sensitive to the exact size of the masks used, which must be selected a priori.




Figure 4.5: Demonstrating the heuristic method for fitting the condensate fraction on the $T = 70\text{nK}$ ($N_0/N = 0.66$) calculated data. (a) The data (black curve) is fit to a Gaussian (red curve) with the condensate peak masked out. (b) The data with the Gaussian subtracted off (black curve) is fit to a TF peak (red curve). (c) The non-condensate is all atoms outside the peak plus a flat distribution underneath the peak (blue). (d) The total non-condensate distribution from the fit is the non-condensate from step (c) added to the Gaussian (blue). The non-condensate distribution from the data is given by the red curve. The fit is very accurate except at the center of the distribution where it over-estimates the non-condensate. The Gaussian fit from step (a) is illustrated as a dashed, white curve. As expected, this does not accurately fit the data near the condensate edge.




Figure 4.6: Testing the three fitting methods against calculated expansion data (§G.3) for different temperatures. In the zero-order method we fit directly using Eqn. 4.21 (square, red points). The non-condensate function in this fit is strongly peaked under the condensate which overestimates the non-condensate number and therefore underestimates the condensate fraction by about 10%. The multi-step method applies Eqn. 4.21 in steps to different regions of the distribution (described in text). This improves the error, but it still underestimates the condensate fraction because of the peak in the non-condensate fitting function. Finally, the heuristic method (described in text and Fig. 4.5) has the best error. However, in practice this type of fit is the most susceptible to noise.

In the lattice, all these same advantages and considerations still apply for using condensate fraction as a probe⁷, except we have less theoretical insight into the functional form of the expanded distribution. Although this makes it more challenging to determine an appropriate fitting procedure, one main advantage of using condensate fraction is that we can use fitting models with minimal theoretical input, such as the heuristic approach. One part of this thesis work is to extend the heuristic approach to the lattice gas and test its efficacy by fitting experimental data. These experiments are detailed in §7.3.

4.1.5 In-Situ Imaging

A complementary approach to imaging the system after expansion is to image in-trap (in-situ). However, this approach is often not utilized because it presents a large number of technical barriers; we list a few major issues here. To start, the trapping potentials, which are either on or just recently turned off, can lead to spatially dependent light scattering cross-sections, and therefore make it difficult to interpret images. Also, without expansion, the system has a very high density, which can result in high losses from light-assisted collisions and skewed distributions from re-absorption pressure. Finally, depending on the desired measurements, in-situ imaging requires an order of magnitude better imaging resolution than is typically needed to image after expansion. Nevertheless, in spite of these challenges, in-situ imaging has a number of advantages and solves fundamental problems associated with probing expanded distributions (discussed below). Therefore, one part of this thesis work is developing a site-resolved imaging probe for the ⁴⁰K apparatus (§3.2.10). We also use low-resolution in-situ images in the ⁸⁷Rb apparatus for specific measurements (§7.2, §8.1.1).

Technical issues aside, in almost all cases in-situ imaging is intrinsically superior to imaging expanded distributions for measuring properties in thermodynamic equilibrium. For in-situ imaging, the measured distribution is directly comparable to the thermodynamic distribution. In contrast, distributions measured after expansion have to be compared to the thermodynamic distribution obtained after solving the time-dependent expansion problem. At best, for weak interactions, solving this problem is time-consuming, and at worst, for strong interactions, the time-dependent solutions are incompatible with our theoretical approaches, such as QMC.

In-situ imaging is also advantageous for Hamiltonians where the specific form of the potential (e.g., a lattice) and/or strong local interactions skew the relative importance of spatial to momentum degrees of freedom. This is typically the case for the models we are simulating. In these types of systems, the effect of the smoothly varying confining potential is well-described by the local density approximation (LDA, §2.2.2), which states that at position \vec{x} the local properties are given by the properties of a uniform system with an effective chemical potential $\tilde{\mu} = \mu - V(\vec{x})$, where μ is the thermodynamic chemical potential and $V(\vec{x})$ is the smooth confining potential. These uniform system properties are the basis of comparison with both theory and experiments on the physical systems we are simulating as embodied by the simulator cycle in Fig. 1.2. With in-situ imaging, we can directly probe this locally uniform physics, whereas expanded images are a complicated average over the entire trap. This is a powerful tool, because in a single image there is information about a range of precisely known chemical potentials. For example, this can be exploited to

⁷As long as the gas is in the superfluid region of the Bose-Hubbard model.

measure the equation of state since in a harmonic trap [62],

$$P(\tilde{\mu}) = \frac{m\omega_y\omega_z}{2\pi} \int \int dy \tilde{n} \left(\sqrt{\frac{2(\mu - \tilde{\mu})}{m\omega_x^2}}, y, z \right) \quad (4.28)$$

based on the Gibbs-Duhem relation, where \tilde{n} is the density measured by in-situ imaging (already integrated along the imaging direction). This technique has been successfully used for quantum simulation of the Fermi gas at unitarity [225, 226].

Another implication of the LDA is that the trapped system may contain different phases in several regions of the trap. For example, in the Bose-Hubbard model (see Fig. 2.11) there can be a superfluid in the center of the trap, a Mott-insulating shell, and then superfluid on the outside. In-situ imaging is required to probe these phase boundaries, which is essential for constructing the phase diagram for the model. Associated with these phase transitions is quantum criticality leading to universal scaling of the in-situ density [227]. In particular, for the BH model, it is possible that co-existing phase boundaries in the trap correspond to different universality classes and therefore have different critical exponents. The ability to locally image these phase boundaries is essential for measuring these exponents. Also, if the system undergoes a phase transition to a degenerate or near-degenerate state (e.g. a ferromagnet at zero field), then domains will form. The trap averaged order parameter may go to zero, so the transition must be probed locally in each individual domain. For a ferromagnet, for example, there may be no total magnetization after expansion, but a local measure of spin correlations will reveal non-zero magnetization.

For some Hamiltonians, the imbalance between spatial and momentum degrees of freedom (DOF) is strong enough that the momentum DOFs are frozen out. In this case, it is not just an advantage, but necessary to use in-situ imaging to determine thermodynamic-state information. An extreme example is the atomic limit of the Hubbard model ($t/U \rightarrow 0$) when the Hamiltonian becomes diagonal in the site-occupation basis. In this limit, the temperature is encoded only in the occupation at each site (for two-component fermions)

$$\langle n_i \rangle = \frac{e^{-\mu_{\uparrow,i}/k_B T} + e^{-\mu_{\downarrow,i}/k_B T} + 2e^{-(U+\mu_{\uparrow,i}+\mu_{\downarrow,i})/k_B T}}{1 + e^{-\mu_{\uparrow,i}/k_B T} + e^{-\mu_{\downarrow,i}/k_B T} + e^{-(U+\mu_{\uparrow,i}+\mu_{\downarrow,i})/k_B T}}, \quad (4.29)$$

where $\mu_i = \mu - \epsilon_i$ (ϵ_i is the potential energy shift of site i). Even for non-zero, but small t , the density is essentially still Eqn. 4.29, and temperature only affects the expanded (momentum) distribution to higher orders. Even in the non-interacting regime of the single-band Hubbard model, there is a finite quasimomentum bandwidth $\epsilon = 4Dt$. When $T > \epsilon$, the band is mostly full and increasing thermal energy goes exclusively to populating higher energy modes of the trap. At this point, the momentum (or quasimomentum) distribution is no longer sensitive to temperature. Our experimental work exploring this crossover is detailed in §7.1.

Finally, a technical advantage of in-situ imaging is that since the atoms start trapped we can, in theory, image for long times by continuing to hold them in the trap. However, this is only beneficial for a certain limited set of conditions. First, as mentioned in §4.1.1, we must be imaging with fluorescence. Next, each atom must be trapped to a region smaller than the desired resolution for the imaging duration. This condition is difficult to satisfy since the light scattering heats the atoms. Therefore, the atoms must be confined to tight traps, for example, individual lattice wells. Fortunately, the excitation beam can simultaneously

laser cool [228], such that temperature will reach a steady-state. Still, this temperature (10-100 μ K) is much higher than before imaging ($\ll 1\mu$ K), so very large confining potentials are required. The final condition is that the number of atoms per trapping region is small. During fluorescence, pairs of atoms in the same region are quickly (< 2 ms) lost due to light-assisted inelastic collisions [229], so that the final density distribution is determined by the parity of the atom number in each region.

A technical challenge of in-situ imaging is high resolution, but this is only required for specific measurements. In general, the imaging resolution depends on what aspect of the system we are trying to probe. To measure the cloud width, which can be used for thermometry, approximately $5 - 10\mu$ m resolution is required. One experimental aspect of this thesis was developing a thermometry probe based on the cloud width, see §7.2. To extract the equation of state or to examine the boundary between phases (e.g., image the MI plateaus), then, based on current experiments [56, 230], a sufficient resolution is approximately $2 - 4\mu$ m. However, to measure and/or address [231] the occupancy of each site in a Hubbard lattice, we need to resolve single sites of the lattice ($< 1\mu$ m). Experiments that benefit from single-site resolution include measuring antiferromagnetic spin ordering and observing correlated particle/hole pairs in the Mott insulator [232].

To achieve single-site imaging requires resolution comparable to (or smaller than) the lattice spacing d . To ease the demands of high resolution we can increase d , however, there is an upper-bound on d set by the physics we are trying to simulate. The physics of strong correlations (i.e., the Hubbard model) requires tunneling between sites, which decreases exponentially with d (at fixed potential depth). As tunneling decreases, the simulator timescales increase and heating dominates. To put a lower bound on tunneling, we need to know the sources of heat. Fundamentally, there is always heating due to scattering from lattice light. In [36] the optimal lattice wavelength (λ_L), and therefore lattice spacing ($d = \lambda_L/2$), is characterized based on a figure of merit that compares the absolute heating rate to the energies and timescales of the Hubbard model. This analysis suggests that λ_L should be as large as possible. Although the energy scales decrease and the timescales increase with large λ_L , the heating rate also decreases. However, at some point other heating mechanisms (e.g., intensity noise) start to dominate. Empirically, $\lambda_L \approx 1\mu$ m ($d \approx 500$ nm) has been used to achieve strongly correlated Hubbard phases. In our experiments we use $395 \leq d \leq 405$ nm (^{87}Rb) and $d = 526.8$ nm (^{40}K). Therefore, for single site imaging the imaging resolution (σ) needs to be ≈ 500 nm.

In an aberration-free (see §F.3) imaging system that uses incoherent scattering, the resolution is set by the diffraction limit

$$\sigma = \frac{0.61\lambda}{nNA}, \quad (4.30)$$

as discussed in §4.1.1 and in more detail in §F.3. The diffraction-limit given by Eqn. 4.30 is not insurmountable; certain scanning and non-linear techniques can achieve sub-diffraction resolution. However, the tradeoff is that imaging times are longer, multiple images are required, and/or the data must be post-processed (see [233] for a review). Therefore, restricting ourselves to Eqn. 4.30, to achieve 500nm resolution we need to have a high numerical aperture (NA) and small λ . Of course, for atoms, λ cannot be arbitrary and is limited to a discrete set of values corresponding to atomic transitions. The lower bound on λ is the inverse energy required to ionize the ground state. However, in practice, only the first few lowest energy transitions are sufficiently strong for effective fluorescence. For our ^{40}K single-site probe (under-development), we plan to image on the second-excited $4S_{1/2} \rightarrow 5P_{3/2}$ transition at $\lambda = 404.5$ nm using a custom 0.6NA objective

($\sigma = 411\text{nm}$). The technical details of this probe are discussed in the apparatus section, §3.2.10.

To best visualize the effects of resolution on single-site imaging, we simulate images for different numerical apertures assuming a single plane of atoms in a $d = 527\text{nm}$ lattice (the spacing used in the ^{40}K apparatus). Each lattice site has only zero or one atom, and the atoms scatter light incoherently. In practice, the scattered light is some combination of coherent and incoherent light [212] because of inelastic scattering in the laser cooling process and saturation (§B.7.3). The detector (a CCD camera) consists of discrete pixels, which each collect a mean number of photons. The actual number of detected photons in a given shot is given by a Poissonian distribution around the mean (because of photon-detector shot noise). There are also additional technical sources of noise, e.g., detector dark counts, but these are left out of this analysis since they are system specific. Simulated images are shown in Fig. 4.7 for a mean of 10^4 photons collected per atom⁸. Qualitatively, the ability to determine site occupancy is consistent with the resolution being smaller than the lattice spacing. Another issue when imaging the lattice is fluorescence from atoms in other planes. Since the depth of focus is [205]

$$\Delta z = \frac{1}{2} \frac{\lambda}{NA^2}, \quad (4.31)$$

adjacent planes are still in focus, which significantly complicates our interpretation. Even atoms in planes far out of focus will contribute an overall offset which reduces the imaging contrast. The simplest solution is to load the atoms in a single plane or remove atoms from all other planes before imaging, which is discussed in §3.2.10.

Ultimately, the goal of single-site imaging is to determine the state-dependent occupancy of each site in the lattice. There are two main limitations to achieving this goal. First, as mentioned above, pairs of atoms on a site are quickly lost, so the measured image is the parity of the number of atoms on each site. The second limitation is that the imaging is not state sensitive. This is manifestly the case for imaging on the $4S \rightarrow 5P$ transition since it only takes approximately five scattering events to decay into the dark hyperfine ground state. However, even if we image on the cycling transition, off-resonant decay to the dark hyperfine state will still occur before enough fluorescence signal has been collected for an effective image. For the two-component Fermi-Hubbard model there are four possible states per site ($|0\rangle, |\uparrow\rangle, |\downarrow\rangle$, and $|\uparrow\downarrow\rangle$) and two possible measured states ($|0\rangle$, and $|1\rangle$) with the following mapping:

$$|0\rangle, |\uparrow\downarrow\rangle \rightarrow |0\rangle, \quad (4.32)$$

$$|\uparrow\rangle, |\downarrow\rangle \rightarrow |1\rangle. \quad (4.33)$$

The problem is more severe when imaging bosonic gases, as the total number of states is much higher since the site occupancy is unbounded.

By performing an operation before imaging that is density and/or spin dependent, we may be able to overcome these issues. For example, a density-dependent technique, experimentally demonstrated in [234], uses interaction shifts to change the resonance for excitations to higher bands. By modulating the lattice over a range of frequencies this technique maps sites with n particles to 1 particle and all other states to

⁸The number of photons collected per atoms should go down with NA as well, but here it is kept fixed to only focus on the changing resolution.




Figure 4.7: Single-site imaging of a 2D lattice plane. (a) Fluorescence from atoms in a lattice (not to scale) is collected by a microscope objective. The distribution of atoms in the lattice ($d = 527\text{nm}$) is illustrated. The atomic width in each lattice well is set by the temperature, here set to be 20% of the lattice depth. (b) Simulated in-focus images of the atomic distribution for different numerical apertures and imaging with $\lambda = 405\text{nm}$ light. The number of collected photons per atom is 10^4 and the pixel size is 50nm (in the object plane). The $d \times d$ region for each lattice well is indicated by the grid lines in each image. These images include photon-detector shot noise and are convolved with the initial atomic distribution. Here the Rayleigh criterion ($\sigma = d$) corresponds to $NA \approx 0.46$. (c) Out-of-focus images for $NA = 0.6$. Single-site resolution refers to our ability to determine with a high degree of confidence whether a site is either empty or full. These images give a qualitative idea of the imaging resolution and focal stability required; in practice, the ability to resolve single-sites is a complicated interplay between signal, noise and analysis methodology.

0 particles by selectively exciting atoms to higher bands and out of the lattice. The total average density can be reconstructed by taking n_{max} images in consecutive experiments. In theory this technique could be sensitive to spin if the interaction is also spin-dependent. Another possible method is to reduce the lattice depth in the out-of-plane direction (creating a series of 1D tubes) and then apply a magnetic field gradient which causes different spin states to travel in different directions. After some time, the lattice is ramped back up to lock them in a new imaging plane. Two subsequent images can be then taken by focusing on the different planes. From Fig. 4.7, we see that the separation would have to be at least four lattice planes. Another proposal is to use RF or microwave fields to selectively move some atoms to the dark hyperfine state, based on spin or density, and then remove the remaining atoms using resonant light on the cycling transition. If we reduce the lattice depth during the light pulse, only a few recoils are required to eject the atoms and so the possibility of decaying to the dark state is small. After ejecting the atoms, we pump the remaining atoms back to the bright state and perform single-site imaging. This technique is especially advantageous if the atoms are deep in the Mott-insulating regime of the half-filled Fermi-Hubbard model, where the basis is essentially truncated to $|\uparrow\rangle, |\downarrow\rangle$. This operation would allow us to measure these two states with only one image since $|\uparrow\rangle$ is mapped to $|1\rangle$ and $|\downarrow\rangle$ is mapped to $|0\rangle$ (or vice-versa). To verify that the basis really is only these two single-occupied spin states, and not empty or doubly occupied sites, we take an image without doing the removal operation, which should reveal an image with one atom per site.

4.1.6 Density Probe using a Spin-Dependent Lattice

Although the discussion of §4.1.5 centered on in-situ imaging, many of the same advantages can be realized with an in-situ probe that lacks the imaging requirement. In this type of probe, the signal is measured in time-of-flight imaging (§4.1.2), but the signal itself directly depends on in-situ properties. A familiar example is a trap-averaged site occupancy probe using the clock-shift [140, 141]. In this probe, atoms are transferred from the initial state $|1\rangle$ to an unoccupied state $|2\rangle$ using RF (or microwaves depending on the exact states $|1\rangle$ and $|2\rangle$). The number of atoms transferred (N_2) is measured after expansion and state separation using a Stern-Gerlach field. N_2 is probed as a function of the RF frequency ω ; N_2 peaks when $\hbar\omega$ equals the energy separating states $|1\rangle$ and $|2\rangle$ (ΔE). Due to the differential interaction shift between states $|1\rangle$ and $|2\rangle$ (the clock shift), ΔE is a local function of the density and is shifted from the single-particle value. Therefore, this shift directly measures the trap-averaged in-situ density. Deep in the Mott-insulator region, the N_2 versus ω signal has discrete peaks corresponding to integer site-occupations. The relative weight of these peaks indicates the total number of sites with occupation N .

In a similar way, we have developed an in-situ probe where the number of atoms transferred depends on the trap-averaged density. Instead of using density-dependent interaction shifts, our probe uses spatially varying state-dependent potentials. Specifically, we use a spin-dependent lattice (§3.1.2), where the lattice potential is proportional to the magnetic moment of the state $g_F m_F$. We start in the state $|F = 1, m_F = 0\rangle$, which does not experience the lattice, and transfer to state $|2, -1\rangle$, which is lattice-bound.

The essential aspect of this probe is the effect of state-dependent potentials on state transfers. For standard transitions between different hyperfine and/or Zeeman states using a time-dependent magnetic field, $B(x, t) = B(x)\hat{x}e^{i\omega t}$, the on-resonance Rabi rate (§B.4) for driving between spatial states $\Psi_i(\vec{x})$ and

$\Phi_f(\vec{x})$ is,

$$\Omega_{if} \propto \langle F, m_F | \hat{S}_x | F', m_{F'} \rangle \langle \Psi_i(\vec{x}) | B(\vec{x}) | \Phi_f(\vec{x}) \rangle \quad (4.34)$$

where Ψ_i, Φ_f are eigenstates of the Hamiltonian for the initial and final states. Since the magnetic field is essentially uniform with respect to the size of Ψ and Φ , the Rabi rate is⁹

$$\Omega_{if} \propto \langle F, m_F | \hat{S}_x | F', m_{F'} \rangle \langle \Psi_i(\vec{x}) | \Phi_f(\vec{x}) \rangle B(\vec{x}). \quad (4.35)$$

If the potential is state-independent then the eigenstates are the same ($\Phi_f(\vec{x}) = \Psi_f(\vec{x})$), then $\langle \Psi_i(\vec{x}) | \Psi_f(\vec{x}) \rangle = \delta_{if}$. Thus, the transition rate does not depend on the spatial distribution of the atom. However, if the Hamiltonian is state-dependent, then the wavefunctions can change as long as $\langle \Psi_i(\vec{x}) | \Phi_f(\vec{x}) \rangle \neq 0$. In general this means that Ψ_i will transfer to a number of states depending on how well each state is spectrally resolved. For example, short pulses which are spectrally broad result in transitions to all Φ_f weighted by the wavefunction overlap and so the final and initial spatial wavefunctions are identical. This is a manifestation of the fact that the spatial wavefunction cannot change arbitrarily fast.




Figure 4.8: Making state transitions from $|1,0\rangle$, which does not experience the lattice, to $|2,1\rangle$, which is lattice-bound. Depending on the frequency of the drive pulse, the atoms make transitions to different bands of the lattice. The drive frequency is shown with respect to the bare driving frequency next to images taken after transfer and time-of-flight. The larger weight of the diffraction peaks for increasing drive frequency is consistent with driving in the higher bands of the lattice. When the drive was tuned between these frequencies there was no state transfer.

We implement this probe in the ^{87}Rb apparatus using the spin-dependent lattice (lin- \perp -lin) described in §3.1.2. The spin-dependent lattice potential is proportional to $g_F m_F$ (the full potential is given by Eqn. 3.5).

⁹A possible exception is when the atoms are very close to a current carrying wire.

To probe the density of atoms in state $|1, 0\rangle$, we drive to a lattice-trapped state ($m_F \neq 0$). As long as the frequency width of our pulse is less than the band gap (i.e., the pulse time is less than $50\mu\text{s}$), we can selectively drive into the different lattice bands by tuning the pulse frequency. The images resulting from this process are illustrated in Fig. 4.8.

The Rabi rate for driving to the n^{th} lattice band is obtained by substituting the $q = 0$ eigenstate (in the tight-binding form) into Eqn. 4.35

$$\Omega_n \propto \sum_i \int dx \Psi(x) w_{i,n}(x) \quad (4.36)$$

where $w_{i,n}(x)$ are the Wannier functions at site i in the n^{th} band (§2.3.2), and Ψ is the wavefunction of the free particle. The Wannier functions are sharply peaked with a width much less than the lattice spacing. Therefore, Eqn. 4.36 is a probe of the density $\Psi(x)$ on the length scale of the lattice spacing and averaged over the trap. To measure the Rabi rate, we detect the number of atoms transferred to the second state by applying a Stern-Gerlach field during TOF. The fraction of atoms transferred, N_2/N , is directly related to the Rabi rate by the relation (on-resonance) $N_2/N = \sin^2\left(\frac{\Omega\tau}{2}\right)$ (Eqn. B.36). This probe was used as part of this thesis work to verify the overlap between a $|1, 0\rangle$ and $|1, -1\rangle$ gas as described in §8.1. Experimental data for the probe can be found in that section.

We are not limited to transferring from free ($m_F = 0$) to lattice states ($m_F \neq 0$). Transitions are also allowed between states that experience the lattice because the lattice potential differs in strength (depth) and/or position for different states; only states with identical $g_F m_F$ experience the same lattice potential. This means that we can drive band-band transitions between the states because eigenfunctions corresponding to different bands (for the different states) are not orthogonal. In fact, these transitions have been demonstrated as part of a sideband cooling scheme [76, 235] and as a means to transport atoms between adjacent lattice sites [76]. Analogously to the $m_F = 0$ density probe, we can use these transitions to spectrally resolve the different band populations for atoms in a given lattice state. For example, to measure the 2nd excited band population of state 1, one could drive transitions to the ground band of state 2. This transition is resonant at a lower frequency than the bare transition and so atoms in the ground band of state 1 are not resonant on any transitions to state 2. This could complement the standard bandmapping approach (§4.1.3).

4.1.7 Impurity Probe

Another type of in-situ probe, developed in this thesis, utilizes interactions between one atomic system (i.e., an impurity) and the atomic system we want to study (“primary”). The interactions imprint the properties of the primary system onto the impurity. The most direct consequence is that these interactions lead to thermal contact and so the impurity and primary are at the same temperature. The advantage of this setup is that we can specifically select the impurity so that its properties are straightforward to measure. We still need to use a more standard probe (e.g., TOF imaging) on the impurity, but in a regime where these probes are well-understood. Impurity probes are ubiquitous throughout physics, particularly for measuring temperature. For example, one can measure temperature by putting a resistor in thermal contact with the sample under study. The temperature is a well-known function of the easily measured resistance. The main tenet of impurity probing is that the impurity cannot fundamentally alter the physics being studied.

For cold atom quantum simulators, a separate sample of cold atoms is an ideal impurity for many of the reasons discussed in §2 (e.g., interactions between the primary and simulator are well-characterized). Additionally, a cold atom impurity is already at the low energy scales of the simulator and so minimally perturbs the simulator physics. One realization is to use another state/species that has weak interactions and/or different quantum statistics than the primary, both trapped in a similar potential. For example, it is common to use an impurity of weakly interacting bosonic ^7Li to measure the temperature of a strongly interacting ^6Li Fermi gas [225]. Another method, which we pursue here, is to use another state/species that experiences a much weaker potential than the primary.

In our implementation we use the spin-dependent lattice (§3.1.2) of our ^{87}Rb apparatus. Since the lattice depth is proportional to $g_F m_F$ (Eqn. 3.5), atoms with $m_F = 0$ experience no lattice potential and are an ideal candidate as an impurity. These atoms are well-described by the theory of the weakly interacting Bose gas (§2.2.2). The primary systems (atoms in states where $m_F \neq 0$) are described by the Bose-Hubbard model (§2.3.4), which is the focus of our quantum simulation. Therefore, if we transfer a small fraction of atoms into $m_F = 0$ they should form an impurity probe for our Bose-Hubbard system. However, a main question is whether the impurity and primary interact on timescales faster than heating rates in the overall system, which is required for the probe to be useful. Developing and characterizing this impurity probe is a main part of this thesis work, so this subject is detailed in a dedicated chapter, §8.

4.2 Techniques

To this point, we have outlined the tools and probes of our experimental quantum simulation toolkit. Tools (§3) are the physical elements which allow us to prepare a system of ultracold atoms which are well-described by the models that we want to simulate (ingredient 1 of quantum simulation). Probes (§4.1) measure the state of that system (ingredient 3). However, the toolkit is lacking one last element that connects tools and probes. These are the techniques of quantum simulation — experimental protocols to perform useful quantum simulations (i.e., experiments that generate the “comparison observables” required to close the simulation loop in Fig. 1.2). Techniques include our protocols to generate well-controlled initial states in the lattice (ingredient 2) and to generate out-of-equilibrium states (ingredient 4). One view is that techniques are the ‘software’ of the quantum simulator.

4.2.1 Adiabatic Loading

To prepare an equilibrium state in the lattice, the customary technique is to adiabatically load atoms from the harmonic trap where the ultracold gas is initially created (see Fig. 4.9). The load is accomplished by ramping up the lattice; the Hubbard model parameters (t and U) are precisely controlled by servoing the lattice beam power which is calibrated to the lattice depth with high accuracy (§C.5). In the ideal case, this means the system is in thermal equilibrium in the lattice at depth $s \times E_R$ with temperature T and chemical potential μ such that the lattice entropy, $S_{\text{lattice}}(T, \mu; s)$, is the same as the entropy in the harmonic trap, S_{harmonic} . The harmonic trap entropy is measured using time-of-flight probes (§4.1.2, §4.1.4). Therefore, we have a well-characterized initial state of our simulator, satisfying ingredient 2 of quantum simulation. We can then probe the lattice state and compare to theory. A slight caveat is that the thermodynamic state




Figure 4.9: Illustrating the adiabatic ramping process. Starting from a condensate (ground state) in the harmonic trap, we load into the lattice over 100ms to $s = 10$. Because the ramp is adiabatic we load into the ground state of the lattice, which at $s = 10$ is also condensed (superfluid region of the BH model). Images are taken after time-of-flight.

in the lattice is characterized by the entropy, S , and total number, N . Although S and N uniquely defines the thermodynamic state, since T and μ are the primary inputs for theory, any theory that compares to the experiment also needs to be able to calculate the entropy.

Adiabaticity implies that if the system is in the ground state of the harmonic trap then it will simulate the ground state of the Hubbard model after loading into the lattice. Since it is almost impossible to be completely in the ground state, a common assumption is that if we are essentially in the ground state of the harmonic trap¹⁰, that this will also hold in the lattice. This assumption is validated if ground state properties predicted from theory are observed in the simulator. However, this validation violates the spirit of the simulation and is not possible when we simulate models that are theoretically intractable.

Adiabaticity is satisfied if the lattice depth is ramped up infinitely slow. In practice we must use a finite ramp time, which is assumed to be slow compared to the timescales of the model. To know these a priori requires solving the model, which is again contrary to the goals of quantum simulation. Experimentally there are two empirical methods to evaluate the timescales. The first method is to measure properties for different ramp times τ , and determine when these properties are no longer a function of τ . The second method is to load into the lattice and then ramp the lattice back down to zero depth. Since entropy is easily measured in the harmonic trap, we can observe whether entropy was generated in the ramping process. One has to be careful with the last method because if the ramp is too fast then no change will be observed. There is also an upper bound on τ set by non-ideal heating processes, such as spontaneous scattering from the lattice (discussed below), intensity noise, 3-body loss, and ultimately the vacuum lifetime. In our experiments (^{87}Rb apparatus) we use the second method to empirically determine that an exponential ramp of time constant 200ms and ramp time of 100ms is optimal. This ramp time is the same order as those used in other Hubbard simulators [167]. An additional complication is that there can be both local and global

¹⁰For example, with bosonic systems it is common to claim that a visually pure condensate is at ‘zero temperature’. In reality, the experimentally achievable upper bound on condensate fraction is no better than $N_0/N \approx 0.9$ because of technical noise.

adiabaticity timescales that differ by orders of magnitude. For example, work in this thesis demonstrated that the timescale for local coherence in the Bose-Hubbard model is $\approx 1\text{ms}$ (§6.2), but mass distribution in the trap can take 10^3 times longer [236].

A fundamental optical lattice issue is heating due to scattering from the lattice beams (see §2.3.5), which is not contained in the Hamiltonian evolution of the Hubbard model. The exact entropy increase from scattering is non-trivial when we are simulating strongly correlated phases of the Hubbard models [37]. A rough approximation is that the entropy increases monotonically with the time spent in the lattice potential. Therefore, we cannot satisfy adiabaticity by going to arbitrarily long lattice load times. Instead, the load time will have a global optima. It has been shown experimentally [67] that heating due to light scattering is the dominant source of entropy generation during the lattice loading process.

Because of entropy generation during the lattice loading process, any characterization of the thermodynamic state (T, μ) in the lattice using the harmonic trap entropy has some amount of error. This error is our motivation for developing thermodynamic probes directly in the lattice, which is part of the work in this thesis (§7, §8.1). Heating during the ramp limits the lowest possible entropies that can be obtained for quantum simulation. Therefore, unless in-lattice cooling methods are developed, we will be unable to simulate low-entropy phases [36].

4.2.2 Finite Temperature

The adiabatic loading technique loads equilibrium states with well-defined entropy into our simulator. Commonly, the entropy is minimized as much as possible so that we are simulating the “ground state” of our model. However, we also have techniques to control the exact value of the temperature (i.e., entropy) to simulate the model at finite-temperature (T). Ignoring temperature leaves out a rich set of physics which is important for the quantum simulator cycle illustrating in Fig. 1.2. The most straightforward approach to finite- T is to control evaporation to achieve a target entropy in the harmonic trap. Alternatively, we can create a degenerate gas and then allow 3-body loss processes to input energy. After preparing a finite entropy gas in the harmonic trap, we adiabatically load into the lattice as described in the previous section. Another class of finite- T techniques generate entropy directly in the lattice. For example, if we load into the lattice and then hold for a variable amount of time the entropy will increase due to light scattering. Also, we can dynamically change the lattice depth to add heat. One such method is to do a quick series of pulses to a very high lattice depth. Another similar approach is to parametrically oscillate the lattice depth, which excites the system to higher bands. When the atoms decay back to the ground back they thermalize the band energy, corresponding to an increase in temperature.

Although finite temperatures are sometimes seen as parasitic because thermal fluctuations can destroy quantum coherences, temperature is an important axis on any phase diagram. If we are to consider any model ‘solved’, we must understand the finite temperature properties. Finite temperature is a natural way to explore the excitation spectrum of a model. In superfluid ^4He , Landau was able to deduce the excitation spectrum simply from measurements of how thermodynamic quantities (i.e., the specific heat) scaled with temperature. Additionally, the temperature scales at which new phenomena emerge can reveal energy scales in the system. For example, the transition to a superconducting state below a certain temperature,

T_C , in metals indicated the emergence of a many-body energy gap (Cooper pairing). The ratio of the zero temperature gap to T_C and the temperature dependence of the gap helped to confirm the BCS theory of superconductivity [48]. Measurements of T_C also helped to reveal a new class of superconductors not described by BCS theory — the high- T_C superconductors. Many of the unexplainable phenomena associated with high- T_C materials is related to properties at finite temperature. For example, in the overdoped region at $T > T_C$, high- T_C materials have a resistivity which scales linearly with T , unlike the expected T^2 scaling of a Fermi liquid [12].

Parasitic thermal fluctuations may also be an advantage. Thermal fluctuations remove long range coherences, so that we can focus on the local physics that is the precursor to eventual global many-body physics. Again, one of the most interesting regimes of the high- T_C materials is for $T > T_C$ when the system is gapped, but not globally coherent (pseudo-gap phase). Similarly, at $T > 0$ above a quantum critical point, thermal effects actually increase the region over which the physics of the critical point play a large role [237]. Because of the interplay between thermal and quantum fluctuations, the only energy scale in this quantum critical regime is temperature. Therefore, universal scaling relations with temperature reveal the location of the quantum critical point [227], although such a point is only defined at $T = 0$.

Another area where finite temperature techniques are important is when investigating quantum-classical crossovers. For example, if there is an activation barrier to a process, then classically thermal energy is required to go over the barrier, but quantum mechanically the system can tunnel through the barrier. To differentiate between these two processes requires looking at how the rate associated with particles traversing the barrier scales with temperature. This technique was used to verify quantum versus classical phase slips as part of the experimental work in this thesis (§6.1).

4.2.3 Transport

To satisfy ingredient 4 of quantum simulation, techniques are required to investigate the system out of equilibrium. The first of these techniques is transport, which looks at the system response to a macroscopic mass current. To start transport we dynamically change an existing trapping potential or introduce a new potential which applies a force to the atoms. This force varies smoothly over the system size and in the simplest case is uniform. The different transport modes that are commonly excited are illustrated in Fig. 4.10. One method to start a dipole mode is to quickly offset the harmonic trap. This will initially cause the atoms to experience a constant force back towards the new harmonic trap center. Similarly, if we suddenly increase the harmonic confinement, an inward force is created that will cause the system to uniformly compress, thereby exciting a monopole mode. In a purely harmonic trap, the dipole mode is a sensitive probe of the trap frequency, but is completely unaffected by pair-wise interactions (generalized Kohn’s theorem [124]). However, interactions affect the other modes. For a BEC in a harmonic trap, the monopole frequency goes from 2ω in the non-interacting limit and $\sqrt{5}\omega$ in the strongly interacting mean-field limit. Similarly, the quadrupole mode frequency goes from 2ω (non-interacting) to $\sqrt{2}\omega$ (strongly interacting) [238]. However, the generalized Kohn’s theorem is not valid in a lattice, so in this case the dipole mode is a probe of many-body physics. The dipole mode is particularly easy to measure because we just need to identify the center-of-mass velocity after time-of-flight expansion.




Figure 4.10: Schematic of some of the transport modes that can be excited. In a dipole mode the cloud position changes, but the shape does not. For monopole and quadrupole modes, the position is stationary, but the cloud shape changes. Here the frequencies are shown for a non-interacting gas in a $\nu = 50\text{Hz}$ trap.

Transport studies are required to fully characterize the models we are studying because many of the phases are classified in terms of transport properties. Common examples include superfluid/superconductor (mass/charge transport without dissipation), insulator (no transport) and liquid/metal (mass/charge transport with dissipation). For the two-component Fermi-Hubbard model, the primary outstanding question is whether there is a d-wave superfluid state; transport is necessary to answer this question. Transport is also a sensitive indicator of many-body physics. As an example, consider the fractional quantum Hall effect, which is due to the correlated motion of 2D electrons in a very high magnetic field [239]. The main signature of this state is obtained by transport measurements of the Hall conductance, which plateaus at quantized rational fractions of e^2/h . Transport is also sensitive to many-body excitations that have low thermodynamic weight. As discussed in the previous section, thermodynamically superfluid ^4He is well-described by a phonon-roton spectrum. From this spectrum, the critical velocity is the ratio of the energy to momentum at the roton minimum (i.e., the Landau critical velocity). This predicted critical velocity is much larger than what is observed experimentally, because the superflow is sensitive to the nucleation of vortices [240].

Two of the experimental results of this thesis focus on using transport techniques to simulate the Bose-Hubbard model. In §6.1, we look at the role of dissipation in the dipole mode of the lattice and in §8.2, we look at the dipole mode of an impurity gas through the lattice gas.

4.2.4 Dynamic Model Parameters

The other out-of-equilibrium technique in our toolkit is the ability to vary the model parameters dynamically. This ability is quite unique to quantum simulators versus physical systems where, at best, the parameters can be tuned slowly with pressure. Unlike transport, the forces acting on the atoms are short range so this technique creates local excitations. Whereas transport tends to probe a model locally in equilibrium, but globally out-of-equilibrium, this technique allows us to access highly out-of-equilibrium states. For example, by instantaneously changing the parameters through a phase transition (i.e., a quench), we can use the equilibrium state of one phase as the initial conditions in a completely different phase. How the system relaxes and the timescales for that relaxation tell us important information about the model and the phase transition.

Another useful aspect of this technique is to evolve in regimes of the parameter space that are not well-understood using the initial ground state wavefunction of a regime that is well-understood. To do this we start in the ground state with one set of parameters and then instantaneously change to another set of parameters. This is the opposite limit of the adiabatic load technique (§4.2.1). Instead of ramping slowly so that we preserve the thermodynamic state, we ramp as fast as possible so the wavefunction is unchanged. This has been demonstrated to show phase collapse and revivals of coherent states on each lattice site in the deep Mott insulator regime [241]. We can also do the opposite; evolve starting with an unknown wavefunction in a parameter space with known solutions.

An interesting use of this technique is to look at how certain properties depend on the speed at which we vary the parameters. In the slow limit we recover the adiabatic load technique. This is useful for elucidating the timescales of local physics. One of the experiments performed in this thesis is a detailed analysis of condensate fraction as a function of ramping speed from the strong to weakly interacting superfluid regime of the Bose-Hubbard model (§6.2). For quenches there are various predictions for the generation of excitations

going through a phase transition. We can constrain these different theories by looking at the number of excitations after a quench as a function of the speed which we vary the parameters [242].

By modulating the parameters at a certain frequency we can couple to single-particle (i.e., band-to-band) and many-body excitations of the system. Measuring the system response as a function of the modulation frequency is a method to map out the excitation spectrum [139]. The decay and thermalization of these excitations is one way to generate a controlled finite temperature sample, as mentioned in §4.2.2.

Chapter 5

Theoretical Toolkit

The central idea of quantum simulation is to experimentally simulate unsolved quantum models using quantum systems. To achieve this we have assembled an experimental toolkit to simulate the Hubbard model described in detail in §3 and §4. However, the emphasis on experiment does not preclude the use of theory. In fact, theory is a key component of any experimental quantum simulation. Although quantum simulation is motivated by the failure of theory, there are many regimes where we can either exactly solve or approximately solve Hubbard models with a high degree of confidence. These theoretical solutions are critical to our overall success because they motivate, validate, and help to interpret our quantum simulator.

Theoretical solutions to the Hubbard model, although limited to certain regimes, provide an important basis of comparison. Any system, such as ultracold atoms trapped in an optical lattice, that is purportedly well-described by a Hubbard model must have overlapping properties with these theories. Once we have built our experimental quantum simulator this set of known model properties can be used to validate our system. Finally, when we run our experiment, theory helps us to interpret the results. Although we may be in a regime where theory cannot predict the exact phase of the model, theory can predict what a specific probe will measure for different quantum phases. Experimentally, this allows us to work backwards and construct a phase diagram. For example, in the Fermi-Hubbard model different theories disagree whether the ground state of a certain parameter regime is a d-wave superfluid. Yet, assuming the state is a d-wave superfluid, theory helps to predict the experimental observations.

The relationship between theory and experiment can also be reversed. Although quantum simulation is focused on modeling physical systems, it is also important for refining theory. In most regimes, theoretical solutions have to make some type of approximation. Once an approximation has been made they become tremendously powerful because of the speed and versatility of their solutions. Although they are less accurate than the experiment, which makes no approximations, the theory can solve for properties that are hard to probe experimentally, such as high-order correlation functions. Validating the approximations made in the theory is difficult theoretically, but straightforward by comparing to the results of the quantum simulation.

Therefore, any experimental toolkit for quantum simulation must have a corresponding theoretical toolkit. In §2.4 some of these theories were discussed to describe known solutions to the Hubbard model. In this section we will describe some of these theories in more detail. The focus will be theory tools directly associated with the experimental work in this thesis. Since the experimental work is almost exclusively on the Bose-Hubbard model, these theories are geared towards that model.

5.1 Non-Interacting

In the $t/U \rightarrow \infty$ limit, we can use the exact non-interacting theory discussed in §2.2.1 for the harmonically trapped gas to solve for the thermodynamics of atoms in a single-band of the harmonic plus lattice potential described by the tight-binding Hamiltonian (see also Eqn. 2.85)

$$\mathbf{H} = -t \sum_{\langle i,j \rangle, \sigma} \mathbf{a}_{i,\sigma} \mathbf{a}_{j,\sigma}^\dagger + \sum_{i,\sigma} \left(\frac{1}{2} m \omega_\sigma^2 r_i^2 - \mu_\sigma \right) \mathbf{a}_{i,\sigma}^\dagger \mathbf{a}_{i,\sigma}. \quad (5.1)$$

Exchange statistics are encapsulated in the operators of Eqn. 5.1, which (anti)-commute for (fermions) bosons. Fortunately, in the non-interacting theory there is a straightforward procedure for separating out the issues related to quantum statistics. First, we need to solve for the eigenstates and eigenenergies of the single particle problem, which for Eqn. 5.1 reduces the problem to determining the coefficients f_j of $|\Psi\rangle = (\sum_{j=-\infty}^{\infty} f_j \hat{a}_j^\dagger) |0\rangle$ (j is the site index), given [133]

$$E_n f_j = -t(f_{j-1} + f_{j+1}) + \frac{1}{2} m \omega^2 d^2 j^2 f_j. \quad (5.2)$$

The eigenvalues of Eqn. 5.2 are given by Eqn. 2.87 and illustrated in Fig. 2.7. The coefficients of eigenvalue n are [133],

$$f_j^{(n)} = \begin{cases} \frac{1}{\pi} \int_0^{2\pi} dx \, ce_n(x, -\alpha) \cos(2jx), & n \text{ even} \\ \frac{1}{\pi} \int_0^{2\pi} dx \, se_{n+1}(x, -\alpha) \sin(2jx), & n \text{ odd} \end{cases} \quad (5.3)$$

where ce and se are the even and odd periodic π solutions of the Mathieu equations with parameter α (as defined for Eqn. 2.87). Once we have the single particle energies we can construct thermodynamic functions from the grand canonical potential (Eqn. 2.22) which take into account the quantum statistics. We can also take the thermal expectation value of operators (Eqn. 2.21).

If $k_B T \gg E_1 - E_0$, or (equivalently) if the de Broglie wavelength λ_T (Eqn. 2.27) is smaller than the characteristic length scale of the potential, then we can use the semiclassical approximation, which includes quantum statistics, but uses classical variables to calculate the energy. For the harmonically trapped gas, the semiclassical approximation is naturally applied in the x, p basis (i.e., Eqn. 2.24), and the condition that $k_B T \gg \hbar \omega$ is almost always satisfied. However, in the harmonically trapped lattice gas, the two potentials set disparate length scales. In the ultracold regime, λ_T is on the order of the lattice spacing, so the semiclassical condition is not met. This is equivalent to the fact that $k_B T$ is much less than the bandgap, so that we are simulating single-band physics. Fortunately, the lattice length scale is incorporated directly into the quasimomentum operator. Rewriting the Hamiltonian in terms of \mathbf{q} and \mathbf{x} ¹,

$$\mathbf{H} = 2t \sum_i \left[1 - \cos \left(\frac{\mathbf{q}_i}{q_B} \pi \right) \right] + \sum_i \left[\frac{1}{2} m \omega^2 \mathbf{x}_i^2 - \mu \right] \quad (5.4)$$

over the range $-\infty \leq x_i \leq \infty$ and $-q_B \leq q_i \leq q_B$. In this representation of the Hamiltonian, we have projected out the bandgap and the only energy scale is given by the slowly varying harmonic potential. The

¹ \mathbf{x} in the (\mathbf{q}, \mathbf{x}) representation is a coarse-grained version of the original \mathbf{x} operator in the (\mathbf{p}, \mathbf{x}) representation. It is defined with respect to $\mathbf{q} = \hbar \mathbf{k}$ as $\mathbf{x} = i \partial / \partial \mathbf{k} + \mathbf{\Omega}$ where $\mathbf{\Omega}$ commutes with \mathbf{k} [243].

semiclassical approximation will hold in this basis, so,

$$\Omega = \mp \frac{1}{\beta(2\pi)^3} \int_{-\pi/d}^{\pi/d} d^3k \int d^3x \log \left[1 \pm \zeta e^{-\beta(\frac{1}{2}m\omega^2 x^2 + 2t(3 - \cos(k_x d) - \cos(k_y d) - \cos(k_z d)))} \right], \quad (5.5)$$

$$= \mp \left(\frac{2\pi}{d^2} \right)^{3/2} \beta^{-1} (\beta m \omega)^{-3/2} \sum_{n=1}^{\infty} \frac{(\pm)^n \zeta^n e^{-6t\beta n} I_0^3(2t\beta n)}{n^{5/2}} \quad (5.6)$$

where $\beta = (k_B T)^{-1}$, $\zeta = e^{\beta\mu}$ is the fugacity, and I_0 is the modified Bessel function of the first kind. The (non-condensate) number is

$$N = \pm (2\pi)^{3/2} (\beta m \omega^2 d^2)^{-3/2} \sum_{n=1}^{\infty} \frac{(\pm)^n \zeta^n e^{-6t\beta n} I_0^3(2t\beta n)}{n^{3/2}}, \quad (5.7)$$

and the entropy is

$$\begin{aligned} S &= -\frac{\partial \Omega}{\partial T} \\ &= \frac{5}{2} \frac{\Omega}{T} + \frac{(\mu - 6t)N}{T} \mp \frac{6t}{T} \left(\frac{2\pi}{d^2} \right)^{3/2} \beta^{-3/2} (m\omega^2)^{-3/2} \sum_{n=1}^{\infty} \frac{(\pm)^n \zeta^n e^{-6t\beta n} I_0^2(2t\beta n) I_1(2t\beta n)}{n^{3/2}}. \end{aligned} \quad (5.8)$$

In Fig. 5.1 we compare the semiclassical entropy to the exact entropy calculated using the eigenvalues (Eqn. 2.87) of Eqn. 5.1.

For bosons T_C can be determined by inverting Eqn. 5.7 when $\zeta = 1$. There are analytic solutions in two limits. When $k_B T \gg t$, then we set $e^{-6t\beta n} I_0^3(2t\beta n) \approx 1$ and,

$$k_B T_C = \left(\frac{N}{\text{Li}_{3/2}(1)} \right)^{2/3} \frac{m\omega^2 d^2}{2\pi} \quad (5.9)$$

The limit corresponds to the band being filled and the harmonic trap is the only energy scale in the problem. This has the same N scaling as the 2D free gas since there are the same degrees of freedom. From Fig. 5.1, we see that this scaling is poor unless $k_B T \gg t$. Another approximation is to sum the first few terms in the series given by Eqn. 5.7, which from Fig. 5.1 is more accurate. In the opposite limit, $t \gg T$, we can use the large z expansion of I_0

$$I_0(z) \approx \frac{e^z}{(2\pi z)^{1/2}} \left(1 + \frac{1}{8z} \right), \quad (5.10)$$

and so,

$$k_B T_C = \left(\frac{N}{\text{Li}_3(1)} \right)^{1/3} \hbar \omega \left(\frac{m}{m^*} \right)^{1/2}, \quad (5.11)$$

where m^* is the effective mass given by Eqn. 2.82. Therefore, in this limit the only effect of the lattice is to rescale the trap frequency. The full range of T_C/t versus $(N^{1/3} \hbar \omega)^2/t$ is plotted in Fig. 5.1.

To get T_F for fermions we need to invert Eqn. 5.7 for μ as $\beta \rightarrow \infty$. In this limit the argument of the




Figure 5.1: Thermodynamics of the non-interacting harmonically trapped lattice gas. (a),(b) Comparing the exact solutions using the eigenvalues from Eqn. 2.87 versus the semiclassical approximation for entropy versus temperature (Eqn. 5.8) at fixed number and $\bar{T}_C = k_B T_C / t$ (normalized temperature) versus $\bar{N} = N(m\omega^2 d^2 / t)^{3/2}$ (normalized number). In these normalized units the plots are essentially universal. For these calculations $N = 10^5$ (for entropy plot), $m = m_{Rb}$, $d = 400\text{nm}$, $t = 1.2 \times 10^{-31}\text{J}$ (tunneling value for $s = 6$ lattice), and $\nu = 50\text{Hz}$. (c) Comparing \bar{T}_C calculated using exact sums versus $\bar{N}^{2/3}$ for several approximations discussed in the text. The $k_B T_C \ll t$ curve (red) is given by Eqn. 5.11, and the $k_B T_C \gg t$ curve (grey) is given by Eqn. 5.10. The blue curve (series approximation) takes the first few terms of the series in Eqn. 5.7. (d) Comparing $\bar{T}_F = k_B T_F / t$ versus $\bar{N}^{2/3}$ for several approximations. The $k_B T_F \ll t$ curve (red) is given by Eqn. 5.15 and the $k_B T_F \gg t$ curve (grey) is given by Eqn. 5.19. The blue curve (series approximation) is given by Eqn. 5.18. (c) and (d) use the same parameters as (a) and (b).

Bessel function $(2t\beta n)$ also goes to infinity, and so we can expand using Eqn. 5.10,

$$\begin{aligned}
N &= -(2\pi)^{3/2}(\beta m\omega^2 d^2)^{-3/2} \sum_{n=1}^{\infty} \frac{(-)^n \zeta^n e^{-6t\beta n} e^{6t\beta n}}{n^{3/2} (4\pi t\beta n)^{3/2}} \\
&= -\beta^{-3} (2tm\omega^2 d^2)^{-3/2} \sum_{n=1}^{\infty} \frac{(-)^n \zeta^n}{n^3} \\
&= -\beta^{-3} (2tm\omega^2 d^2)^{-3/2} \text{Li}_3(-\zeta)
\end{aligned} \tag{5.12}$$

As $\beta \rightarrow \infty$, ζ goes to ∞ , so we can use a polylogarithm expansion

$$\text{Li}_n(-z) \approx -\frac{1}{n!} \log(z)^n, \tag{5.13}$$

and so

$$N = \frac{1}{6} \beta^{-3} (2tm\omega^2 d^2)^{-3/2} (k_B T_F)^3 \beta^3 \tag{5.14}$$

$$k_B T_F = (6N)^{1/3} \hbar \omega \left(\frac{m}{m^*} \right)^{1/2}, \tag{5.15}$$

which is only valid if $k_B T_F \ll t$ because Eqn. 5.14 is the lowest order term in a T_F/t series. Therefore, in the opposite limit ($k_B T_F \gg t$) we cannot use the prior approach. Instead we need to make an approximation to the N equation

$$N = -(2\pi)^{3/2}(\beta m\omega^2 d^2)^{-3/2} \int_{-\pi}^{\pi} d^3 k \text{Li}_{3/2} \left[-\zeta e^{-\beta 2t(3-\cos(k_x)-\cos(k_y)-\cos(k_z))} \right]. \tag{5.16}$$

We expand using Eqn. 5.13

$$\begin{aligned}
N &\approx (2\pi)^{-3/2}(\beta m\omega^2 d^2)^{-3/2} \int_{-\pi}^{\pi} d^3 k \frac{1}{\Gamma[5/2]} \log \left[\zeta e^{-\beta 2t(3-\cos(k_x)-\cos(k_y)-\cos(k_z))} \right]^{3/2} \\
&= (2\pi)^{-3/2}(\beta m\omega^2 d^2)^{-3/2} \frac{\beta^{3/2} (k_B T_F)^{3/2}}{\Gamma[5/2]} \dots \\
&\quad \int_{-\pi}^{\pi} d^3 k \left[1 - 2 \frac{t}{k_B T_F} (3 - \cos(k_x) - \cos(k_y) - \cos(k_z)) \right]^{3/2}
\end{aligned} \tag{5.17}$$

$$= (2\pi)^{3/2} (m\omega^2 d^2)^{-3/2} \frac{(k_B T_F)^{3/2}}{\Gamma[5/2]} \left[1 - 6 \frac{t}{k_B T_F} \right]^{3/2} \tag{5.18}$$

Since $k_B T_F \gg t$, we ignore the second term in the brackets, and

$$k_B T_F = (\Gamma[5/2] N)^{2/3} (2\pi)^{-1} m\omega^2 d^2, \tag{5.19}$$

consistent with a similar expression in [244]. The full range of T_F/t versus $(N^{1/3} \hbar \omega)^2/t$ is plotted in Fig. 5.1.

Although the real experiment always has interactions (unless we are near a Feshbach resonance), the non-interacting theory is still a useful tool. It is applicable when either kinetic or potential energies dominate, such as at high temperature or when the density is very low. Because we always have a trap, there will be a dilute region at the edge where we can use a fit to the non-interacting theory to measure temperature.

Furthermore, it is useful for calculating thermodynamic functions, for example, to determine the temperature if we adiabatically load into the lattice (§4.2.1). Depending on what quantity we are calculating, the weight of the non-interacting modes can be high. Energy and entropy, for instance, are dominated by high energy states where interactions are less important.

5.2 Atomic Limit

The opposite of the non-interacting gas is the atomic limit ($t/U \rightarrow 0$). Rewriting Eqn. 2.92 with $t = 0$,

$$\mathbf{H} = \sum_i \left\{ \frac{1}{2} \sum_{\sigma, \sigma'} U_{\sigma, \sigma'} \mathbf{a}_{i, \sigma}^\dagger \mathbf{a}_{i, \sigma'}^\dagger \mathbf{a}_{i, \sigma} \mathbf{a}_{i, \sigma'} + \sum_{\sigma} (\epsilon_{i, \sigma} - \mu_{\sigma}) \mathbf{n}_{i, \sigma} \right\}. \quad (5.20)$$

This Hamiltonian is diagonal in the Fock basis at each site, and the LDA is exact. Therefore, we can extract analytic expressions for thermodynamic quantities of both uniform and trapped systems [245]. Since particle-hole and trap excitations are contained in Eqn. 5.20, this is a reasonable model for describing these degrees of freedom for $t/U \neq 0$, but $t/U < 1$. For example, this $t = 0$ theory exhibits the main qualitative features of the $t/U \neq 0$ in-trap density profile [137], which for single-component bosons is

$$n_i = \frac{\sum_{n=1}^{\infty} n e^{-[\beta(U/2(n-1) + \epsilon_i - \mu)]n}}{\sum_{n=0}^{\infty} e^{-[\beta(U/2(n-1) + \epsilon_i - \mu)]n}}. \quad (5.21)$$

See Eqn. 4.29 for the equivalent expression with two-component fermions. To treat the features of the momentum distribution, we cannot use Eqn. 5.20, since there is not a definite phase relation between different sites. Similarly, at high temperature ($k_B T \gg t$), particle-hole and trap excitations dominate because the band is full, so the thermodynamics is well-described by Eqn. 5.20.

The atomic limit is the natural starting point for perturbation theory in orders of t/U . To zero-order, the ground state wavefunction for single-component bosons with $\langle n \rangle = 1$ and M lattice sites is

$$|\Psi\rangle = \prod_i^M \mathbf{a}_i^\dagger |0\rangle \quad (5.22)$$

To first-order in the tunneling term, the state in Eqn. 5.22 is connected to any state with a single particle-hole pair (there are $6M$ such states). The wavefunction to first-order in perturbation theory is

$$|\Psi\rangle \propto \prod_i^M \mathbf{a}_i^\dagger |0\rangle + \frac{t}{U} \sum_{\langle j, k \rangle} \mathbf{a}_k^\dagger \mathbf{a}_j \left(\prod_i^M \mathbf{a}_i^\dagger |0\rangle \right) \quad (5.23)$$

From Eqn. 5.23 we can extract first-order corrections to the momentum distribution. Continuing to higher orders is more suited to strong-coupling expansions of the Green's function (e.g., [146]). These types of expansions are not considered in this thesis.

5.3 Mean-Field Theories

The next level of theory must bridge the gap between the non-interacting and atomic limits. Since there is no exact solution, a common technique is to search for a mean-field solution. The vital component of such a theory is determining the best aspect of the model to treat as the mean-field. Starting from the non-interacting limit, it is natural to treat interactions as the mean-field. On the contrary, starting from the atomic limit it is more natural to treat tunneling as a mean-field. We will therefore look at an interaction mean-field theory (§5.3.1) and a tunneling mean-field theory (§5.3.2).

Before going over details, it is helpful to summarize the essential properties and approach to a mean-field theory by discussing the ferromagnetic spin-1/2 Ising model

$$\mathbf{H} = J \sum_{\langle i,j \rangle} \mathbf{S}_{z,i} \mathbf{S}_{z,j}. \quad (5.24)$$

From the perspective of a spin on site i , it interacts with quantum spins on neighboring sites. The (static) mean-field approach replaces those quantum spins with the average spin from all the neighbors

$$\mathbf{S}_{z,i} \mathbf{S}_{z,j} = \langle \mathbf{S}_{z,i} \rangle \mathbf{S}_{z,j} + \mathbf{S}_{z,i} \langle \mathbf{S}_{z,j} \rangle - \langle \mathbf{S}_{z,i} \rangle \langle \mathbf{S}_{z,j} \rangle, \quad (5.25)$$

and the model becomes

$$\mathbf{H} = zJ \sum_i (\langle \mathbf{S}_z \rangle \mathbf{S}_{z,i} - \langle \mathbf{S}_z \rangle^2), \quad (5.26)$$

where z is the number of neighboring spins (coordination number), and $\langle \mathbf{S}_z \rangle$ is given self-consistently ($\xi = \frac{zJ}{k_B T}$) as

$$\langle \mathbf{S}_z \rangle = \frac{e^{-\xi \langle \mathbf{S}_z \rangle} - e^{\xi \langle \mathbf{S}_z \rangle}}{e^{-\xi \langle \mathbf{S}_z \rangle} + e^{\xi \langle \mathbf{S}_z \rangle}}. \quad (5.27)$$

When we solve Eqn. 5.27, we get that $\langle \mathbf{S}_z \rangle = 0$ for $\xi < 1$ and $\langle \mathbf{S}_z \rangle = \pm 1$ for $\xi \geq 1$, corresponding to a ferromagnetic transition. Therefore, the mean-field approach is able to describe interesting features of the interacting model without an exponential increase in computational complexity. The consequence, however, is that the approximation will breakdown in some regime, and it is not always straightforward to understand where this will occur. An indication of how the theory will breakdown is that the mean-field neglects fluctuations. For a spin with $z = 2$ (1D Ising model), these fluctuations are important since a single spin-flip can change the total spin-spin interaction energy by a factor of two. In fact, the mean-field theory completely breaks down in 1D, and if we exactly solve Eqn. 5.24 there is no ferromagnetic transition. However, for $z \rightarrow \infty$ these individual spin-flips are negligible and the mean-field theory is exact. Another situation when fluctuations are important and the mean-field approach fails is in the vicinity of the transition itself. At this point fluctuations are critical; for example, a single spin-flip may order all the spins — formally the spin susceptibility goes to infinity.

5.3.1 Hartree-Fock-Bogoliubov-Popov

Our first mean-field theory, Hartree-Fock-Bogoliubov-Popov (HFBP) [123, 124], treats the kinetic and potential terms exactly and replaces some interaction terms with a mean-field. This approach is equally valid in the lattice or harmonic trap, so we start with the generic second-quantized Hamiltonian (Eqn. 2.42) and

make two approximations. First, we assume the field-operator can be written as $\Psi(r, t) = \Phi(r, t) + \tilde{\Psi}(r, t)$, where $\langle \tilde{\Psi}(r, t) \rangle = 0$, and $\Phi(r, t)$ is a complex number. The mean-field term $\Phi(r, t)$ is directly related to the condensate (i.e. a macroscopic occupation of a mode of the single particle density matrix), such that $|\Phi(r, t)|^2$ is the condensate density. Expanding the interaction term of eqn. 2.42 we get,

$$\begin{aligned} \Psi^\dagger \Psi^\dagger \Psi \Psi &= |\Phi|^4 + 2|\Phi|^2 \Phi^* \tilde{\Psi} + (\Phi^*)^2 \tilde{\Psi} \tilde{\Psi} \\ &\quad + 2|\Phi|^2 \Phi \tilde{\Psi}^\dagger + 4|\Phi|^2 \tilde{\Psi}^\dagger \tilde{\Psi} + 2\Phi^* \tilde{\Psi}^\dagger \tilde{\Psi} \tilde{\Psi} \\ &\quad + \Phi^2 \tilde{\Psi}^\dagger \tilde{\Psi}^\dagger + 2\Phi \tilde{\Psi}^\dagger \tilde{\Psi}^\dagger \tilde{\Psi} + \tilde{\Psi}^\dagger \tilde{\Psi}^\dagger \tilde{\Psi} \tilde{\Psi} \end{aligned} \quad (5.28)$$

The mean-field approximation to the interaction is,

$$\begin{aligned} \Psi^\dagger \Psi^\dagger \Psi \Psi &\approx |\Phi|^4 + 2|\Phi|^2 \Phi^* \tilde{\Psi} + (\Phi^*)^2 \tilde{\Psi} \tilde{\Psi} \\ &\quad + 2|\Phi|^2 \Phi \tilde{\Psi}^\dagger + 4|\Phi|^2 \tilde{\Psi}^\dagger \tilde{\Psi} + 4\Phi^* \langle \tilde{\Psi}^\dagger \tilde{\Psi} \rangle \tilde{\Psi} \\ &\quad + \Phi^2 \tilde{\Psi}^\dagger \tilde{\Psi}^\dagger + 4\Phi \langle \tilde{\Psi}^\dagger \tilde{\Psi} \rangle \tilde{\Psi}^\dagger + 4\langle \tilde{\Psi}^\dagger \tilde{\Psi} \rangle \tilde{\Psi}^\dagger \tilde{\Psi} \end{aligned} \quad (5.29)$$

This particular mean-field approximation is the Hartree-Fock-Popov-Bogoliubov approximation [124]. We also make the static Popov approximation, which assumes that the non-condensate density is static ($\langle \tilde{\Psi}^\dagger(r, t) \tilde{\Psi}(r, t) \rangle = \langle \tilde{\Psi}^\dagger(r) \tilde{\Psi}(r) \rangle$). This is the same approximation we made in §2.2.2 to derive the equation of motion for Φ — the generalized time-dependent Gross-Pitaevskii (GP) equation (Eqn. 2.47).

Substituting Eqn. 5.29 back into the second-quantized Hamiltonian (Eqn. 2.42), we get the effective Hamiltonian for $\tilde{\Psi}$ [123]

$$\mathbf{H} = \int dr \tilde{\Psi}^\dagger(r) \left(-\frac{\hbar^2 \nabla^2}{2m} + V(r) - \mu + 2gn(r) \right) \tilde{\Psi}(r) + \frac{g}{2} \int dr \Phi^2(r) \tilde{\Psi}^\dagger(r) \tilde{\Psi}^\dagger(r) + \frac{g}{2} \int dr (\Phi^*)^2(r) \tilde{\Psi}(r) \tilde{\Psi}(r) \quad (5.30)$$

where $n_0(r) = |\Phi(r)|^2$ is condensate density, $n_T(r) = \langle \tilde{\Psi}^\dagger(r) \tilde{\Psi}(r) \rangle$ is the non-condensate density, and $n(r) = n_0(r) + n_T(r)$ is the total density. If we ignore the last two terms of Eqn. 5.30, this reduces to the simple Hartree-Fock mean-field theory discussed in §2.2.2 (Eqn. 2.43). The Hartree-Fock energies are

$$\epsilon_{HF}(r, p) = \frac{p^2}{2m} + V(r) - \mu + 2gn(r). \quad (5.31)$$

If we keep the last two terms in Eqn. 5.30, they are not in the proper form of a single particle Hamiltonian. Therefore, we need to perform a Bogoliubov transformation, and in a semiclassical approximation the solutions are,

$$\epsilon_{BG}(r, p) = (\epsilon_{HF}(r, p)^2 - g^2 n_0(r)^2)^{1/2} \quad (5.32)$$

where the distribution function of non-condensate particles is,

$$n_T(r) = \frac{1}{(2\pi\hbar)^3} \int dp \frac{\epsilon_{HF}(r, p)}{\epsilon(r, p)} \frac{1}{e^{\epsilon(r, p)/k_B T} - 1}. \quad (5.33)$$

Taking into account the Bogoliubov terms means that $N_T = \int dr n_T(r)$ can be non-zero at $T = 0$. This is the depletion of the condensate by interactions.

In thermodynamic equilibrium, the equations, Eqn. 2.47 (GP Equation), Eqn. 5.32, and Eqn. 5.33 form a set of self-consistent equations, given a fixed μ and T . To solve for the condensate density $n_0(r)$ (Eqn. 2.47), we need to know μ and the non-condensate density $n_T(r)$. Yet, we cannot solve for $n_T(r)$ until we know μ , T , $n_0(r)$, and $n_T(r)$ (n_T appears on both the left and right sides of Eqn. 5.33). Therefore, the self-consistency algorithm amounts to solving these different equations iteratively using values from the previous step. For example, in step i , we solve Eqn. 2.47 using the value of $n_{T,i-1}(r)$ from step $i-1$ to get the condensate density $n_{0,i}(r)$. Then we solve for the non-condensate energies (Eqn. 5.32) using $n_{T,i-1}(r)$ and $n_{0,i}(r)$, and calculate $n_{T,i}(r)$ from Eqn. 5.33. We iterate until $n_{T,i}(r) = n_{T,i-1}(r)$ and $n_{0,i}(r) = n_{0,i-1}(r)$ (to within a set tolerance). The iteration loop must be seeded by an initial estimate for $n_{T,-1}$. We furnish this estimate by solving Eqn. 5.33 using the non-interacting energies (i.e., set $g = 0$ in Eqn. 5.32). The semi-ideal model, discussed in §2.2.2, amounts to setting $n_T(r) = 0$ in Eqn. 2.47 and in the non-condensate energies (Eqn. 5.32). In this case we only require two steps. We first solve for the condensate density $n_0(r)$ from Eqn. 2.47, which is only a function of μ , and then we solve for the non-condensate density, which is only a function of μ , T , and $n_0(r)$.

For the lattice we apply the same procedure to the Hubbard model [64]. We get the lattice version of the GP equation for each site i

$$0 = -t \sum_{\langle j \rangle} \phi_j + [\epsilon_i - \mu + UN_{0,i} + 2UN_{T,i}] \phi_i, \quad (5.34)$$

where $\langle j \rangle$ is a sum over all nearest-neighbor sites to i , $N_{0,i}$ is the number of condensate atoms on site i , $N_{T,i}$ is the number of thermal atoms, and ϵ_i is the energy shift of lattice site i (i.e., from a harmonic trap). The semiclassical lattice versions of the Hartree-Fock (Eqn. 5.31) and Bogoliubov (Eqn. 5.32) energies in the LDA are,

$$\epsilon_{HF}(q)_i = -2t \cos(q/q_B \pi) - (\mu - \epsilon_i) + 2gN_i \quad (5.35)$$

$$\epsilon_{BG}(q)_i = (\epsilon_{HF}(q)_i^2 - g^2 N_{0,i}^2)^{1/2}. \quad (5.36)$$

If we also solve the condensate using the Thomas-Fermi approximation, the condensate number (per site) is,

$$N_0 = \text{Max} \left\{ \frac{\mu}{U} - 2N_T, 0 \right\} \quad (5.37)$$

For a uniform system of M sites, the self-consistent value of N_T is

$$N_T = \frac{1}{M} \sum_{q \neq 0} \frac{\epsilon_{HF}(q)}{\epsilon_{BG}(q)} [\coth(\beta \epsilon(q)/2) - 1]. \quad (5.38)$$

In the lattice the condensate depletion, calculated using Eqn. 5.38, can be quite large as U/t increases, which is illustrated in Fig. 5.2. In the HFBP theory, condensate fraction is always non-zero at $T = 0$ for all U/t , and so it does not predict the known Mott-insulator transition. This demonstrates that the interaction mean-field approximation fails for strong interactions [147].




Figure 5.2: HFBP $T = 0$ condensate fraction versus U/t for $\mu = U/2$ in the uniform 3D lattice. The Mott-insulator transition should occur at $U/t \approx 35$, but the HFBP theory fails such that $N_0/N > 0$ for all U/t .

5.3.2 Site-Decoupled

The breakdown of the interaction mean-field in the HFBP theory is best illustrated by considering the exact interaction energy for two states (written in the number basis on a single site), $|\Psi_1\rangle = |1\rangle$ and $|\Psi_2\rangle = \frac{1}{\sqrt{3}}|0\rangle + \frac{1}{\sqrt{3}}|1\rangle + \frac{1}{\sqrt{3}}|2\rangle$. Both these states have the same mean number, $\langle \hat{n} \rangle = \langle \hat{a}^\dagger \hat{a} \rangle = 1$, so the mean-field interaction energy is the same. If we calculate the exact interaction energy, $\langle \mathbf{H}_{int} \rangle = \frac{U}{2} \langle \mathbf{n}(\mathbf{n} - 1) \rangle$ then we get $\langle \Psi_1 | \mathbf{H}_{int} | \Psi_1 \rangle = 0$ and $\langle \Psi_2 | \mathbf{H}_{int} | \Psi_2 \rangle = \frac{U}{3}$. The true interaction energy of $|\Psi_2\rangle$ is an order U higher than state $|\Psi_1\rangle$ because of number fluctuations that are omitted from the HFBP theory. When U is large compared to t , the difference in interaction energy between these two states is critical and demonstrates the requirement to treat interactions exactly.

To construct a mean-field theory where interactions are exact, we start from the atomic limit (§5.2) and add a mean-field tunneling term [68]

$$\mathbf{a}_i^\dagger \mathbf{a}_j \rightarrow \langle \mathbf{a}_i^\dagger \rangle \mathbf{a}_j + \mathbf{a}_i^\dagger \langle \mathbf{a}_j \rangle - \langle \mathbf{a}_i^\dagger \rangle \langle \mathbf{a}_j \rangle, \quad (5.39)$$

where $\langle \mathbf{a} \rangle$ is identified as the condensate. Substituting this mean-field approximation into the single-component Bose-Hubbard model (Eqn. 2.92)

$$\mathbf{H} = \sum_i \left[-t \sum_{\langle j \rangle} \left(\langle \mathbf{a}_j^\dagger \rangle \mathbf{a}_i + \mathbf{a}_i^\dagger \langle \mathbf{a}_j \rangle - \langle \mathbf{a}_i^\dagger \rangle \langle \mathbf{a}_j \rangle \right) + \frac{U}{2} \mathbf{a}_i^\dagger \mathbf{a}_i^\dagger \mathbf{a}_i \mathbf{a}_i + (\epsilon_i - \mu) \mathbf{a}_i^\dagger \mathbf{a}_i \right] \quad (5.40)$$

we get a mean-field Hamiltonian which is diagonal in the number basis on each site. Given Eqn. 5.40, if the system is inhomogeneous we need to solve self-consistently across the whole lattice since the mean-field

terms depend on adjacent sites. However, we typically solve for the uniform system and use LDA in which case the Hamiltonian is

$$\mathbf{H} = \sum_i \left[-zt \left(\mathbf{a}^* \mathbf{a}_i + \mathbf{a}_i^\dagger \mathbf{a} - |\mathbf{a}|^2 \right) + \frac{U}{2} \mathbf{a}_i^\dagger \mathbf{a}_i^\dagger \mathbf{a}_i \mathbf{a}_i - \mu \mathbf{a}_i^\dagger \mathbf{a}_i \right] \quad (5.41)$$

where z is the number of nearest-neighbor sites, and $\mathbf{a} = \langle \Psi_0 | \mathbf{a} | \Psi_0 \rangle$. We must solve Eqn. 5.41 iteratively, because the eigenstates depend on the value of \mathbf{a} , but \mathbf{a} is calculated from these same eigenstates. In step i of the iteration, we solve for the eigenstates of Eqn. 5.41 using \mathbf{a}_{i-1} (the value of \mathbf{a} from step $i-1$), and then use the lowest energy eigenstate to calculate \mathbf{a}_i . We iterate until $\mathbf{a}_i = \mathbf{a}_{i-1}$.




Figure 5.3: Site-Decoupled Mean-Field $T = 0$ condensate fraction versus U/t at $\mu = U/2$ in a uniform 3D lattice. The condensate fraction goes abruptly to zero, indicating the transition to the $n = 1$ Mott-insulator.

If we solve Eqn. 5.41 for $\mu = U/2$ as a function of U/t and plot $|\mathbf{a}|^2/n$ (condensate fraction), we see that the condensate fraction goes to zero at some value of U/t illustrated in Fig. 5.3. This is the superfluid-to-Mott-insulator transition. Once we are in the MI regime and $\mathbf{a} = 0$, Eqn. 5.41 is equivalent to the atomic limit. The full phase diagram calculated using Eqn. 5.41 is illustrated in Fig. 2.11. The dimensionality of the lattice is encoded in the parameter z . For a cubic lattice $z = 6$, for a square lattice $z = 4$, and $z = 2$ in 1D. As expected from a mean-field theory, the discrepancy with QMC increases at lower dimensions where fluctuations start to dominate as discussed in §2.4.1.

To determine finite temperature properties, we include excited states after solving Eqn. 5.41 when calculating \mathbf{a} . We calculate \mathbf{a} including a Boltzmann average over all states, $\mathbf{a} = Z^{-1} \sum_n e^{-E_n/k_B T} \langle \Psi_n | \mathbf{a} | \Psi_n \rangle$ where $Z = \sum_n e^{-E_n/k_B T}$ is the partition function. This theory can also be extended to calculate dynamic properties if t and/or U are time-dependent. We use this time-dependent version (calculated by Stefan Natu at Cornell) to compare to our experimental results in §6.2.

5.4 Exact - Small Systems

An exact solution of the Hubbard model is possible by numerical diagonalization of the Hamiltonian in the Fock basis. However, this solution quickly becomes numerically intractable and is only possible for a small number of sites (< 10). Although the results are exact, the small site Hubbard model does not correspond to the physical system we are trying to model ($N = 10^{23}$) or even our simulator ($N = 10^5$). Still, we gain insight from the solutions and for certain properties they give reasonably close answers to the full solution.

To better understand the size limitations for diagonalization we need to look at how the number of states grows as a function of the number of particles N and lattice sites M . For bosons, the number of states is

$$\frac{(M + N - 1)!}{(M - 1)!N!}, \quad (5.42)$$

and for fermions

$$\frac{M!}{N!(M - N)!} \quad (5.43)$$

If we assume that it is possible to diagonalize a maximum of 5000 states (i.e., 25×10^6 element matrix) then for bosons where $N = M$, this limit corresponds to a maximum of 7 sites (1716 states). For two-component fermions where $N_{\uparrow} = N_{\downarrow} = M/2$ (half-filling), this limit corresponds to a maximum of 8 sites (4900 states). Another option is to go back to a standard approach and solve grand-canonically (i.e., not in a fixed-number basis). Then we need to truncate the number of particles per site. For two-component fermions, the basis is truncated automatically at four states per site, and the basis grows as 4^M . For bosons, if we truncate to 5 particles per site it grows as 5^M . For bosons on 7 sites, this corresponds to 78125 total states, and for fermions on 8 sites this corresponds to 65536 states, so the fixed-number basis is more efficient. The advantage of the grand-canonical basis is that we can compute trap quantities using LDA. The speed at which the basis grows can be seen by considering the number of states required if we double the number of sites. For the bosons this would mean $2 \times 10^7 (6 \times 10^9)$ states for the canonical (grand-canonical) basis and for fermions $2 \times 10^8 (4 \times 10^9)$ states.

One advantage of the exact solution is that we get the complete set of eigenstates and therefore can compute any quantity — for example, higher order correlations. In Fig. 2.12 we show a number of properties for the 5 site, 2D Fermi-Hubbard model. In Fig. 5.4 we show the condensate fraction for the 7 site, 7 particle Bose-Hubbard model with $\mathbf{z} = 6$. The method is also trivially extended to the time-dependent Hubbard model, $i\hbar|\dot{\Psi}\rangle = \mathbf{H}(t)|\Psi\rangle$, by simple matrix multiplication. We use this technique for qualitative understanding of our experimental results in §6.2.

5.5 Comparison

The five main theory tools in our experimental quantum simulation toolkit are non-interacting theory, atomic limit approximation, the HFBP approximation, site-decoupled mean-field theory, and small system exact




Figure 5.4: Condensate fraction versus U/t from exact diagonalization of the 3D 7-Site Bose-Hubbard model with 7 particles at $T = 0$. The model consists of one central site (site #1) connected via tunneling to 6 outer sites. Here condensate fraction is defined as $\langle \sum_{j \neq 1} \mathbf{a}_j^\dagger \mathbf{a}_1 \rangle / \langle \mathbf{a}_1^\dagger \mathbf{a}_1 \rangle$. Although the general downward trend of condensate fraction is observed, there is no sharp transition, which is expected due to the finite system size.




Figure 5.5: Summary of the different theoretical methods for the Bose-Hubbard model. The diagram is similar for the Fermi-Hubbard model with the addition of HFBP-theory and site-decoupled MFT. The regimes of validity and a qualitative level of approximation are shown for each theory. For example, the atomic and non-interacting limits are exact, but the region of validity is small. This diagram illustrates the ‘dead-zone’ in the middle of many theories corresponding to the strongly correlated superfluid and the Mott-insulator transition.

calculations. Because of the specific approximations made in each theory they have certain regimes of validity and describe some properties more accurately than others. In some cases the theory breaks down in a clear way — for example, the lack of a SF-MI transition in the HFBP theory. However, in other cases the regions of validity are not clear. The limitations to these theories is difficult to judge a priori. The only certain way to judge the approximations made in a theory are by direct comparison to experiment. However, we can also learn a lot by comparing the theories against each other. If most theories are in agreement, but there is one that disagrees then we can infer that it is not valid in that regime. Also, there are regimes where we know that certain theories, based on the approximation, should be very close to exact. Therefore, in those regimes we can use those theories as a baseline.

A comparison between the different theories and experiment is done in §7.3 for condensate fraction versus entropy. Also, the different theories, the rough regimes of validity and level of approximation are summarized in Fig. 5.5.

5.6 Beyond our Toolkit

There are a number of theoretical tools that can be applied to the Hubbard models beyond the five previously discussed. However, many of these theories are advanced and must be implemented by specialists. Here we will list a few of the most relevant tools with references to where more details can be found.

In particular, the main advanced theoretical tool that we have access to is quantum Monte Carlo (QMC) calculations. QMC is a classical numerical approach using statistical sampling of the path integral, and, as such, it makes no approximations. The only errors in QMC are statistic. These errors converge to zero for the equilibrium Bose-Hubbard model as the number of sampling points is increased. However, the statistical errors do not converge for a number of fermion problems because of the “sign problem” [153], although new techniques are under development [156]. A number of QMC results and references were discussed in §2.4. Because QMC can solve the equilibrium Bose-Hubbard model with finite classical computational resources it is ideal for validating Bose-Hubbard simulators. A future extension is to use QMC to compare to our experimental work in §7.3.

Another advanced technique extends the mean-field approaches discussed in §5.3, which were all based on a static mean-field parameter. The next-order approximation is to consider a dynamic mean field theory (DMFT) [153]. DMFT is particularly useful for examining the Fermi-Hubbard model, for which the HFBP and site-decoupled MFT do not apply, but can also be applied to the Bose-Hubbard model [144, 145]. Other techniques include strong coupling expansions [146] and high-temperature series expansions [158].

Chapter 6

Simulating the Bose-Hubbard Model Out of Equilibrium

The remaining sections will discuss our experimental work performed on the ^{87}Rb apparatus (§3.1). These experiments fall into two categories: simulating the Bose-Hubbard (BH) model out of equilibrium, which is discussed in this chapter, and the development of simulation probes (§7 and §8).

Out-of-equilibrium experiments are an important subset of the quantum simulation process (ingredient 4). One motivation is that many physical systems are defined by their non-equilibrium properties; for example, the defining characteristic of superfluid ^4He is the ability to transport mass without dissipation. Also, while non-equilibrium physics is often the most challenging problem theoretically, it is intrinsic to our quantum simulator. Therefore, out-of-equilibrium experiments are optimal as a tool to validate new theoretical techniques.

As part of this thesis work, we performed two experiments simulating the BH model out of equilibrium. In the first experiment (§6.1), we measure the decay rate of center-of-mass motion in the lattice. We observe a finite decay rate at zero temperature consistent with quantum phase-slip dissipation. By varying the temperature, we observe the crossover between quantum and thermal phase-slip activation. In the second experiment (§6.2), we measure condensate evolution as we linearly decrease the lattice depth — increase t/U — for varying ramp times. This experiment is a direct probe of many-body dynamics, and in agreement with time-dependent Gutzwiller mean-field calculations, we observe condensate growth with a characteristic timescale on the order of the inverse BH interaction energy U^{-1} .

6.1 Phase Slip Dissipation

In this experiment, published in [1]¹, we use the ^{87}Rb apparatus (§3.1) to simulate the transport (§4.2.3) properties of the BH model as a function of U/t and temperature T (§4.2.2). In particular, to probe dissipation we measure the decay rate (γ) of the center-of-mass velocity (dipole mode). We observe that the decay rate scales as $\log(\gamma) \propto -\sqrt{t/U}$ and that decay persists down to the lowest temperatures we can experimentally achieve. These measurements are consistent with intrinsic dissipation due to phase slips, a well-known dissipation mechanism in macroscopic quantum systems. Phase slips are expected to generate topological

The work presented in §6.1 of this chapter was published in: D. McKay, M. White, M. Pasienski and B. DeMarco. *Phase-slip-induced dissipation in an atomic Bose-Hubbard system*. Nature 453, 76 (2008). As per the Nature publishing group terms, copyright of the article is retained by the authors. Figs. 6.2, 6.3, 6.4, and 6.5 are reproduced from this publication. The work presented in §6.1 of this chapter was published in: Stefan S. Natu, David C. McKay, Brian DeMarco, and Erich J. Mueller. *Evolution of condensate fraction during rapid lattice ramps*. Phys. Rev. A 85, 061601 (R) (2012) ©2012 American Physical Society. Authors are granted the right to include the article in a dissertation without requesting permission from the APS. Fig. 6.9 is reproduced from this publication.

¹Note that in this publication we used J , instead of t , for tunneling.

(i.e., vortices) excitations — we directly measure vortex-like features in 20% of our images. Additionally, we measure the finite-temperature phase-slip decay rate and observe the crossover between quantum and thermal activation of phase slips.

At $T = 0$ the BH model has two phases: the Mott-insulator (MI) phase for $(U/t) \geq (U/t)_c$ and a superfluid (SF) phase for $(U/t) < (U/t)_c$ (§2.4.1). These phases are classified according to their transport properties. In the MI phase, there is no mass transport (i.e., the system is insulating) and in the SF, there is dissipation-free mass transport. Specifically, mass transport is dissipation-free in the SF phase because the flow is metastable to decay. The energy of the moving system is higher than if it were at rest, and so energetically the system should decay, but there is an energy barrier to this process and therefore metastable flow at finite velocity (i.e., superfluidity) is possible. The barrier height is velocity dependent, so this stability has an intrinsic velocity upper bound. As first argued by Landau and later deduced by Bogoliubov for weak-interactions, if the system is traveling faster than $v_c = \min(\epsilon(p)/p)$ ($\epsilon(p)$ is the excitation spectrum) then the system will decay. For weakly interacting gases, the Landau critical velocity is equal to the speed of sound. Superfluidity breaks down when the velocity exceeds the Landau critical velocity.

Although this is a tremendous insight, it does not completely explain the observations in many superfluid and superconducting systems because it neglects the role of vortices, a point first made by Feynman [246]². Despite predicting a sharp breakdown in superflow at the Landau critical velocity, many physical systems display intrinsic dissipation at low velocities and $T = 0$. This dissipation is thought to be driven by phase slips, which were used to describe the intrinsic critical velocity of SF He [247] and resistance in thin superconducting wires [47].

The general idea of a phase slip is illustrated in Fig. 6.1 for a 1D system between two reservoirs of the same phase. In a quantum system, velocity is the gradient of the wavefunction phase (phase is defined modulo 2π), and so velocity is quantized to $v_n = (\hbar/m)(n/L)$, where n is an integer. If a vortex moves through the flow then the velocity can decay by $\Delta n = 1$ [248] — a phase slip. However, there is an energy requirement for a vortex to enter the system, shown as the well depth in Fig. 6.1. A phase slip can be generated by thermal activation, or by quantum tunneling through the barrier. In superconducting wires [48] and superfluid He [249], thermal phase slips are well-understood and can be directly observed. However, the experimental evidence for quantum tunneling of phase slips is inconclusive [49].

The effect of phase slips in the BH model at zero and finite temperature is an open question. While phase slips are predicted to occur in the BH model [46], precise calculations are difficult due to the dynamical nature of the problem which precludes the use of QMC techniques. Therefore, our BH simulator is ideal for investigating phase slips using our transport measurement techniques (§4.2.3). Similar to physical systems (superconductors and superfluid ^4He), we can measure dissipation as a function of temperature (§4.2.2). Uniquely, we can also measure dissipation continuously as a function of the BH parameters. Because our system is “clean”, we are not encumbered by noise and disorder issues, and, in particular, we can probe the quantum-classical crossover. Transport in this regime has been studied by several other groups; see, for

²Another defining characteristic of a SF is that there is a global phase, $\Phi(x)$. The velocity is given by $v = \hbar/m(\nabla\Phi)$, and, thus, $\nabla \times v = 0$. The system is irrotational, unless there is a singularity (vortex), which must be quantized by the single-valued nature of the phase parameter.




Figure 6.1: Schematic of the phase-slip process for a weak link connecting two phase reservoirs. Since the phase at both ends must equal zero (modulo 2π), the quantized velocities are $v = \hbar/m(\nabla\phi) = (h/m) \times (n/L)$, where n is an integer (assuming constant velocity in the weak link). For velocity values intermediate to the quantization condition, a phase kink must exist to satisfy the boundary conditions (see inset). This kink introduces more kinetic energy, and therefore there is a barrier to decay from $n \rightarrow n - 1$. The system can either cross the barrier via thermal activation or by quantum tunneling. Both processes result in a phase slip of 2π as the system relaxes to a lower velocity.

example [44, 51, 52].

6.1.1 Experiment




Figure 6.2: (a) Schematic of the experimental setup for the phase slip experiment. For more information about the apparatus, see §3.1. The atoms are held in a magnetic (harmonic) trap that balances gravity, and a 3D lattice is superimposed. To initiate transport, the harmonic trap frequency is decreased for 7ms, which causes a constant downward force on the atoms. (b) Timing diagram of the experiment. First, the atoms are loaded into the lattice over 100ms. Once in the lattice, an impulse is applied, and then atomic motion evolves for a variable amount of time. Finally, the gas is released from the lattice, and an image is taken to determine the center-of-mass velocity. To image vortices, we apply a magnetic field gradient to counteract gravity during TOF; this enables 80ms of expansion time.

We probe phase slips by first creating Bose-Einstein condensates (BECs) of ^{87}Rb in a harmonic trap³ at varied condensate fraction N_0/N (i.e., varied temperature). Condensate fraction is measured after time-of-flight (TOF) expansion (§4.1.4). Next, we adiabatically load the finite-T gas into the $\lambda = 811.8\text{nm}$, spin-independent lattice as described in §4.2.1. To start center-of-mass (COM) motion (§4.2.3), we apply an impulse to the atoms along gravity. In a harmonically trapped gas, COM motion is not sensitive to 2-body interaction effects because of the generalized Kohn’s theorem [124].

This theorem, however, does not hold with the lattice present. We generate an impulse by lowering the trapping potential for 7ms and then returning it to the original value. Since the atoms are originally at the equilibrium position determined by the trap and gravity, weakening the trap imparts a force on the atoms. The impulse gives the atoms a peak velocity of 0.8–1.8mm/s depending on the lattice depth. Working in this low velocity regime we avoid well-known Landau and dynamical instabilities [250]. After a free evolution time τ up to 200ms, the lattice is bandmapped, and the atoms are imaged after TOF (§4.1.3). We fit the

³This was in winter 2006-2007, so BEC was prepared in a pure magnetic trap; see §3.1 for details.

image to a TF combined with a Gaussian profile (with independent centers), and translate the position after TOF into the velocity in the lattice. The timing sequence is summarized in Fig. 6.2.




Figure 6.3: Representative data from the experiment for two different lattice depths. Each data point represents one experimental cycle with the free evolution time (Fig. 6.2) as the independent variable. The black points are the condensate (TF component of the fit) and the open points the thermal component (Gaussian component of the fit). Velocity is measured using the deviation from the equilibrium position after TOF. The solid line is a fit to the data using Eqn. 6.7, which is used to extract the decay rate γ . (a) $s = 2$, $T/T_C = 0.85$ and 8×10^5 BEC atoms. (b) $s = 6$, $T/T_C = 0.93$ and 2.7×10^5 BEC atoms. For (b), sample images with 0.47 mm field of view are shown as an inset. Regions of high density are red, and regions of low intensity are blue. The narrow BEC component oscillates, while the broad thermal component is stationary (pinned).

In the lattice, the 1D semiclassical equations of motion for an atom at position x with quasimomentum q are [8]

$$\dot{x} = \frac{\partial \epsilon(q)}{\partial q}, \quad (6.1)$$

$$= \bar{v}, \quad (6.2)$$

$$\dot{q} = \frac{\partial V(x)}{\partial x}, \quad (6.3)$$

where ϵ is the band structure, \bar{v} is the mean velocity, and $V(x)$ is a smooth, external potential. These are the lattice-equivalent Newtonian equations of motion. Near the bottom of the band,

$$\frac{\partial \epsilon(q)}{\partial q} = \frac{q^2}{2m^*}, \quad (6.4)$$

where m^* is the effective mass (Eqn. 2.82). For a harmonic potential $V(x) = -kx$,

$$m^* \ddot{x} = -kx. \quad (6.5)$$

Regardless of the source, dissipation enters into Eqn. 6.5 as a force proportional to $-\bar{v}$. For example, ohmic resistance for charge carriers is described by a force $-\rho v$ where ρ is the resistivity. Adding a dissipative force $-m^* \gamma \dot{x}$ into Eqn. 6.5

$$m^* \ddot{x} = -kx - m^* \gamma \dot{x}, \quad (6.6)$$

which is a simple-harmonic oscillator. Therefore, we fit data to the curve,

$$v = v_0 \sin(\omega t + \phi_0) e^{-\gamma t}, \quad (6.7)$$

where v_0 , ω , ϕ_0 , and γ are free-parameters of the fit. The role of phase slips is characterized by the decay rate γ .

Typical oscillation data for the condensate and thermal components, at two different lattice depths and temperatures, are shown in Fig. 6.3. The data fit well to Eqn. 6.7 with non-zero γ . An interesting feature of the data in Fig. 6.3 is that the thermal velocity is smaller and damps more quickly than the condensate velocity. This suggests that the thermal component is pinned to the lattice, which was also seen in [222]. This may be interesting for exploring physics of the Caldeira-Leggett model [251] and for creating second-sound, but we did not investigate this phenomena further.

The oscillation data used to measure γ are checked in two ways for nonlinear decay terms. First, no significant change in the fitted value of γ is measured if the first period of motion is excluded from the fit. Second, fitting the data to a nonlinear damping model with $\gamma = \gamma_0 v^\alpha$, where α is the coefficient of nonlinear damping and γ_0 is the damping rate for $\alpha = 0$, typically increases the reduced χ^2 value of the fit by 10-20%. We find no clear dependence of α on T/T_C or t/U , and the fitted value of α averaged across all of the data is 0.68 ± 0.07 . Ultimately, our sensitivity to detect weak velocity nonlinearities in the decay rate is limited because of finite signal-to-noise ratio in measurements of the COM velocity.

We rule out several technical noise sources as dissipation mechanisms that could explain the damping of COM motion. Anharmonicity in the dipole potential may effectively damp COM motion for large values of s . To check for anharmonic behavior, we measure COM motion when the retro-reflected lattice laser beams are removed (blocked), which eliminates the lattice potential and reduces the depth of the dipole potential by a factor of 4. Using this technique we measure a value of γ consistent with zero for lattice laser intensities corresponding to $s = 4.5$ (if we unblock the retro-reflected beams an $s = 18$ lattice would result)⁴. This eliminates trap anharmonicity as an effective dissipation source for the data in Fig. 6.4. Relative motion between the lattice and the harmonic potential or fluctuations in s (caused by retro-reflecting mirror motion and lattice laser intensity fluctuations, respectively) can lead to the dephasing of dipole mode motion by transferring atoms into states with different values of m^* in excited bands; we do not, however, observe population outside of the lowest energy band. The total spontaneous emission rate per atom is less than 0.3 Hz for $s = 6$, so momentum diffusion caused by scattering light from the optical lattice laser beams is

⁴In [1] we erroneously state that the laser intensities in this measurement corresponds to $s = 9$.

insignificant. The lattice depth varies by less than 3% across the BEC, so spatial variation in the effective mass can play no role in the dissipation timescales measured in our data.

6.1.2 Results




Figure 6.4: Decay rate vs inverse temperature for (a) $s = 2$ and (b) $s = 6$. Each value of γ is obtained from fitting similar data to that shown in Fig. 6.3. The solid line is a fit to thermal activation of phase slips, Eqn. 6.8. Fitted values are given in the main text.

After taking a series of curves similar to Fig. 6.3, we extract γ for different lattice depths (different t/U) and temperatures. In Fig. 6.4 we plot the decay rate γ versus inverse initial temperature in the harmonic trap (in T_C units) for two of these lattice depths⁵. At finite temperature, phase slips are generated by thermal fluctuations; the probability for a thermal fluctuation to have enough energy to get over a barrier of height ΔE is weighted by a Boltzmann factor $e^{-\Delta E/(k_B T)}$. Therefore, we fit to an Arrhenius equation for thermal activation of phase slips with a finite γ as $T \rightarrow 0$ to account for quantum phase slips

$$\gamma(T) = \gamma_0 + (\gamma_\infty - \gamma_0)e^{-\Delta E/(T/T_C)}. \quad (6.8)$$

For $s = 2$, $\Delta = k_B \times (0.5 \pm 0.1)\mu\text{K}$ and $\gamma_0 = 1.4 \pm 0.4\text{Hz}$, and for $s = 6$, $\Delta = k_B \times (0.33 \pm 0.08)\mu\text{K}$ and $\gamma_0 = 13 \pm 2\text{Hz}$. The data shows good agreement to this model and the dissipation rate is non-zero at $T = 0$, which is consistent with the concept of quantum tunneling of phase slips. We omit the highest temperature point taken very close to T_C when we fit to Eqn. 6.8. The considerable increase in γ close to T_C may be

⁵At $s = 2$ there are 14% next-nearest neighbor corrections to the Hubbard model, which should be included in any theoretical treatment. Therefore, this data will not be used to compare to the scaling solution from [46] in Fig. 6.5. Qualitative phase-slip features (e.g., thermal activation) should be unchanged.

due to mutual friction between the pinned thermal component and the BEC [252]. Mutual friction cannot, however, explain the finite γ at $T = 0$, since the thermal component vanishes at low T .




Figure 6.5: Logarithm of the decay rate γ versus $\sqrt{t/U}$ for $T/T_C = 0.76$ and 3×10^5 BEC atoms (filled circles) and $T/T_C = 0.312$ and $N = 1.1 \times 10^6$ BEC atoms (open circles). The equivalent lattice depths are shown at the top of the graph. The data are well fit by a line, indicating $\log \gamma \propto -\sqrt{t/U}$, in agreement with the predicted theoretical scaling from [46]. The offset between the data sets is due to the difference in thermal activation of phase slips. (Inset) Images taken after long expansion at $s = 8$ show evidence of vortex features (highlighted in red) predicted to accompany phase slips. The top row are the raw images and the bottom are the residuals after subtracting a smooth TF fit.

Zero-temperature phase slips are generated by quantum fluctuations, which allow the system to tunnel through the activation barrier. The decay rate grows as t/U decreases (the lattice depth increases); at the MI transition the barrier disappears. For quantum tunneling through a barrier, the tunneling rate is proportional to e^{-S} , where S is the action in imaginary time [253]. A scaling law for this action in the Hubbard model $S \propto \sqrt{t/U}$, was derived in [46], so we expect the logarithm of the phase-slip decay rate to be proportional to $-\sqrt{t/U}$. Therefore, in Fig. 6.5 we plot the logarithm of the decay rate $\log(\gamma_0)$ versus $\sqrt{t/U}$ for two different temperatures. The data fit well to the expected scaling with a slope of $-1.6(0.1)$, consistent between the two temperatures. The higher temperature data is offset upwards in decay rate, consistent with a higher rate of thermal activation of phase slips.

Since a phase slip involves vortices entering into the system, direct imaging of these features is strong evidence of phase slip dissipation. Vortices can be directly imaged in cold atoms system as a density depression in the condensate profile after TOF expansion [254]. However, the vortex size is small, so they can only be resolved at long expansion times (> 50 ms). Additionally, our imaging system is not well-optimized

to image vortices. Nevertheless, we do observe qualitative features consistent with vortices in 20% of our images at $s = 8$, as shown as an inset to Fig. 6.5. These features are not observed if we do not initiate COM motion in the system.

In summary, we simulated transport in the BH model using ultracold atoms to explore phase-slip dissipation in the superfluid regime of the BH phase diagram. We observe evidence for phase slips that are both thermally activated and generated by quantum fluctuations. Compared to similar studies in physical systems (e.g., small-scale superconductors [49]), where the role of quantum phase-slips is obscured by noise, our results provide some of the most convincing evidence for quantum-phase-slip dissipation. There remain unresolved experimental and theoretical questions regarding the exact nature of phase slips in our finite, inhomogeneous system. In the future, a more optimal imaging setup should be able to measure vortex statistics and dynamics [255], which is essential to understanding the phase slip process. New advances in lattice thermometry and cooling will allow us to better clarify the zero-temperature limit of transport. Finally, the recent demonstration of thin channel transport between two cold atom reservoirs [256] could be utilized to study the phase slip problem in a straightforward setup (Fig. 6.1), which would greatly simplify the interpretation of the results.

6.2 Rephasing After Lattice Ramps

In this work, published in [2], we investigate the timescales for condensate fraction growth after the lattice is linearly ramped from the strongly interacting to weakly interacting SF regime. First, we load a pure ^{87}Rb condensate adiabatically into a spin-independent lattice at s_i , and then we linearly ramp down the lattice depth to $s_f = 4$ in a variable amount of time τ_0 . We observe that the condensate fraction increases exponential in τ_0 where the $1/e$ time is approximately 1ms for $s_i = 10 - 14$. In agreement with time-dependent Gutzwiller calculations, we conclude that timescales for condensate fraction growth are comparable to \hbar/U . This precludes the use of bandmapping (§4.1.3) as a precision probe of the lattice condensate fraction.

The ground-state condensate fraction of the lattice gas decreases for smaller t/U (higher lattice depths), eventually vanishing at the MI transition. To experimentally observe these changes, the gas is loaded into the lattice by increasing the lattice depth monotonically over some period of time. Therefore, the condensate fraction must dynamically evolve as the lattice depth increases during the ramp. Understanding the timescales for condensate fraction evolution is important. If we intend to adiabatically load the gas into the lattice (§4.2.1), we must ensure that the ramp time is slow compared to all other dynamical processes in the system. In the opposite limit, bandmapping (§4.1.3), the ramp time must be fast compared to all other processes except the inverse bandgap (\hbar/E_{bg}) — this is the bandmapping condition. Ideally we would prefer to use bandmapping to probe condensate fraction; compared to the momentum distribution, where the gas is spread over several diffraction orders, the quasimomentum distribution has higher signal-to-noise ratio, and is easier to interpret. In order to accurately measure condensate fraction from the bandmapped distribution, the bandmapping ramp time must be sufficiently fast so that quasimomentum redistribution is not possible. For the non-interacting lattice gas, the only timescales in the Hamiltonian are $2\pi/\omega$, the trap oscillation period, and \hbar/t , the tunneling time. These are much slower than the inverse bandgap, and thus the bandmapping condition is easily met. A new timescale \hbar/U enters for the interacting gas, which is much faster than the non-interacting timescales. For example, at $s = 10$, $\hbar/E_{bg} \approx 40\mu\text{s}$, $\hbar/t \approx 15\text{ms}$, and

$\hbar/U \approx 1\text{ms}$. If the condensate evolution is driven by U , then it will be difficult to satisfy the bandmapping condition, and the measured condensate fraction will not be accurate.

From a theoretical perspective, there are a number of open questions related to the dynamics of the BH model [38], and developing non-equilibrium theories is an active area of research. A promising technique is time-dependent Gutzwiller mean-field theory (similar to a time-dependent version of §5.3.2). Comparing this theory to the results of our experimental quantum simulator is an excellent method to validate and test the approximations of this theory. Although our main focus is on local dynamics during fast ramps, in general there are also slower global dynamics such as global mass and entropy redistribution in the trap. These effects are also being explored experimentally [236].

Before we investigate the timescales for condensate fraction evolution in the lattice, it helps to understand the microscopic changes that occur in the ground state for varied t/U . To write down the ground state on each site we use the SDMFT theory (§5.3.2) at different values of t/U . We express the state $|\Psi\rangle$ in the Fock basis (particle number basis), where the set of basis states are written as $\{|n=0\rangle, |1\rangle, |2\rangle, \dots\}$. When the lattice is in the weakly interacting SF regime ($s=4$, $U/t=1.3$), the state with $\langle n \rangle = 1$ is

$$|\Psi_1\rangle = 0.5|0\rangle + 0.64|1\rangle + 0.44|2\rangle + 0.23|3\rangle + 0.1|4\rangle + 0.03|5\rangle. \quad (6.9)$$

For a deep lattice, close to the MI transition ($s=12$, $U/t=27$) the state evolves into

$$|\Psi_2\rangle = 0.25|0\rangle + 0.93|1\rangle + 0.26|2\rangle + 0.03|3\rangle + 0.001|4\rangle. \quad (6.10)$$

In $|\Psi_2\rangle$, stronger interactions reduce number fluctuations as compared to $|\Psi_1\rangle$. These number fluctuations give rise to a non-zero condensate fraction. This is apparent if we consider the generic state $|\Psi\rangle = \sum_{j=0} \alpha_j e^{i\phi_j} |j\rangle$. The condensate fraction for this state is

$$\frac{N_0}{N} = \frac{|\langle \Psi | \mathbf{a} | \Psi \rangle|^2}{\langle \Psi | \mathbf{a}^\dagger \mathbf{a} | \Psi \rangle} \quad (6.11)$$

$$= \frac{\left| \sum_{j=1} \sqrt{j} (\alpha_{j-1} \alpha_j) e^{i(\phi_j - \phi_{j-1})} \right|^2}{\sum_{j=1} j \alpha_j^2}. \quad (6.12)$$

A state with high condensate fraction is in a superposition of many different Fock states (i.e., has large number fluctuations) with fixed relative phases. For $s=4$, the condensate fraction is 0.99, while at $s=12$ it has decreased to 0.36.

6.2.1 Small System Calculation

The change in the local state between $s=12$ to $s=4$, as determined by MFT, indicates that condensate fraction increases as the state evolves into a phase-coherent superposition of Fock states on each site. Although we analyzed the wavefunction at a single site, the Fock-state superposition is equivalent to the atoms becoming delocalized through the lattice. For example, the ground state of the M site non-interacting BH model is $\left(\sum_i^M \hat{a}_i^\dagger \right)^N |0\rangle$, where every atom is maximally delocalized. If we expand this state and examine

terms at a single site, we recover a state similar to $|\Psi_1\rangle$. Atoms delocalize in the BH model by tunneling to neighboring sites; the timescale for this process in the non-interacting BH model is set by \hbar/t . Therefore, one might reason that tunneling sets the timescales for condensate fraction evolution.




Figure 6.6: Small 1D system (5 sites and 5 particles) calculation of the linear lattice ramp from $s = 12$ to $s = 4$ in time τ_0 . (a) Momentum distribution for various τ_0 compared to the $s = 12$ and $s = 4$ ground state momentum distributions. (b) Height of the $p = 0$ peak as a function of τ_0 . The peak starts to increase appreciably around $\tau_0 = 300\mu\text{s}$. To adiabatically follow to the ground state requires $\tau_0 > 2\text{ms}$. (c) Two representative state populations versus the ramp time. The state notation is $|n_1 n_2 n_3 n_4 n_5\rangle = |n_1\rangle \otimes |n_2\rangle \otimes |n_3\rangle \otimes |n_4\rangle \otimes |n_5\rangle$, where $|n_i\rangle$ is a Fock state of particle number n on site i . The $|11111\rangle$ state is highly populated at $s = 12$ because of interactions. To transfer from the $|11111\rangle$ state into the $|11021\rangle$ state requires a single tunneling event from site 3 to site 4. An appreciable fraction of the system has evolved into the state $|11021\rangle$ after a $300\mu\text{s}$ ramp, consistent with the upturn in peak height. This population decreases again for longer ramps as atoms become even more delocalized (e.g., population transfers to the state $|11012\rangle$).

To test this supposition, we calculate the exact time-dependent dynamics for the 1D small system BH model (§5.4) with 5 sites and 5 particles. To start, we first determine the ground state for the initial lattice depth s_i via exact diagonalization. Then we propagate the state forward in time using Euler time-steps

$$|\Psi(\mathbf{t} + \Delta\mathbf{t})\rangle = (\mathbf{I} + \Delta\mathbf{t}/i\hbar\mathbf{H}(\mathbf{t})) |\Psi(\mathbf{t})\rangle. \quad (6.13)$$

The time-dependence of the BH parameters is set to be equivalent to linearly ramping down the lattice depth to s_f in time τ_0 . This approach is valid as long as the time-dependence does not cause band-to-band transitions. In Fig. 6.6 we plot two state populations as a function of τ_0 . Also, we plot the momentum

distribution $n(p)$ for different τ_0 , where

$$n(p) = |w(p)|^2 \langle \Psi(\tau_0) | \sum_{jk} \mathbf{a}_j^\dagger \mathbf{a}_k e^{ipd(j-k)/\hbar} | \Psi(\tau_0) \rangle. \quad (6.14)$$

Even when we ramp down quickly (e.g., $\tau_0 \approx 500\mu\text{s}$), the momentum distribution is close to the ground state momentum distribution. The state populations reveal that only a small admixture of delocalized atoms are required to develop a strongly peaked momentum distribution.

The results suggest that the timescale is set by \hbar/U and not \hbar/t as earlier conjectured. This is a sensible result based on the full energetics of tunneling in this system. Although the matrix element for an atom to tunnel to an adjacent site is t , the energy of the state after the atom has tunneled depends on interactions with all the other atoms in the lattice. For example, tunneling into an occupied site will occur an energy cost of U , the interaction energy (if the atom started on a singly occupied site). A similar argument was made in Ref. [58]. In the two-well model, the problem of tunneling between the wells is formally equivalent to an off-resonance Rabi oscillation (§B.4). For this problem, the oscillation time is set by the detuning and not the coupling; the oscillation is also not full contrast. The same holds for tunneling in the lattice; the detuning is set by U , and so the time for an atom to tunnel is $\approx \hbar/U$, but the atom will not completely transfer to the adjacent site. However, the condensate fraction (Eqn. 6.12) is proportional to the square root of the number of atoms that tunnel. Therefore, a small admixture of higher number states results in a large condensate fraction increase, as borne out by the results of our calculation.

6.2.2 Experiment

To experimentally investigate the time evolution of condensate fraction, we use BECs created in the ^{87}Rb apparatus in the single beam dipole plus magnetic quadrupole trap configuration with $\bar{\nu} = 35.76(8)\text{Hz}$ (see §3.1 for more details). For this measurement, we prepare BEC with the highest possible condensate fraction in our apparatus. A small thermal component is visible in saturated images, and we estimate that the harmonic-trap condensate fraction is greater than 0.81. The atoms are loaded adiabatically into a $\lambda = 812\text{nm}$ spin-independent lattice to an initial lattice depth s_i . After waiting 10ms, we linearly ramp the lattice depth down to $s = 4$ in time τ_0 and then immediately turn off all lattice and trapping potentials. Because we release the atoms directly from the lattice, the images after TOF have diffraction peaks, and so a heuristic fitting procedure is required. The fitting procedure is outlined in §7.3, and to increase the signal-to-noise ratio, we use a combination of high and low OD images as described in §4.1.4. The fitting procedure accurately identifies the number of condensate atoms, but is susceptible to noise when determining the total number of atoms. Therefore, the condensate fraction for a particular data run is determined by taking the condensate number from the fit and dividing by the total number of atoms averaged over all runs. We end the ramp at $s = 4$ so that the system is still well-described by a tight-binding Hamiltonian, which is necessary for comparing to the theory calculations. We observe no significant difference in the measured timescales if the lattice is ramped to $s = 0$, as in standard bandmapping.

We measure condensate fraction versus τ_0 at lattice depths $s_i = \{10, 11, 12, 12.5, 13, 14\}$ with corresponding number, $N = \{104(5), 82(5), 99(5), 71(3), 100(3), 72(2)\} \times 10^3$ averaged over the different τ_0 (typically shot-to-shot fluctuations are 25%). The mean-field prediction for the emergence of the MI phase is $s \approx 13$,




Figure 6.7: (a) Condensate fraction versus ramp down time (τ_0) for $s = 12$. The red line is an exponential fit used to get the exponential time constant for condensate growth τ_{-1} . (b) TOF images from the experiment for different ramp down times labeled on the plot in (a). The top set of images are fully repumped for best signal-to-noise ratio for the non-condensate atoms, and the bottom images are partially repumped so that the OD of the condensate is unsaturated. We use the procedure outlined in §7.3 to measure condensate fraction from the images.

so the data at $s = 13$ and $s = 14$ may have some MI present, but the majority of atoms in the trap are in the SF phase. Results for condensate fraction versus τ_0 are shown for $s = 12$ in Fig. 6.7 along with several of the TOF images. We fit each set of data (condensate fraction versus τ_0) to an exponential function to determine the $1/e$ time τ_{-1} .

6.2.3 Results




Figure 6.8: Experimental results for τ_{-1} vs. lattice depth. The timescales associated with interactions and tunneling are shown, illustrating that τ_{-1} is consistent with h/U and much faster than h/t .

The experimental results for τ_{-1} as a function of the initial lattice depth s_i are illustrated in Fig. 6.8. In the SF regime, τ_{-1} is $\approx 400\mu s$, which is much faster than the tunneling timescale h/t and consistent with interaction timescale h/U . These results are important because the timescales for condensate fraction to evolve is comparable to bandmapping times used in many experiments. For example, sample values from the literature include $200\mu s$ [34], $400\mu s$ [204], $750\mu s$ [3], $1ms$ [139] and $2ms$ [200]. We also observe a slight increase in τ_{-1} for $s \geq 12.5$, which many correspond to a MI phase appearing in the system.

To make quantitative comparison to the data, our theory collaborators Stefan Natu and Eric Mueller at Cornell University calculated the condensate fraction⁶ after the ramp for the full 3D trapped system using time-dependent Gutzwiller theory [257]. In the calculation, the BH model parameters are changed dynamically in a way consistent with a linear ramp of the lattice depth. In addition, the overall trapping potential is adjusted as $\omega = \sqrt{\omega_0^2 + 8s/m\mathbf{w}^2}$, where $\mathbf{w} = 120 \pm 10\mu m$ is the lattice beam waist. The calculation is performed for $N \approx 75,000$ atoms on a $55 \times 55 \times 55$ lattice.

The results are illustrated in Fig. 6.9. The theory was calculated considering two different initial conditions. In case #1 (bottom, dashed line in Fig. 6.9), the calculation starts the ramp in the ground state of the

⁶The condensate fraction is calculated as $\sum_i |\langle \hat{a}_i \rangle|^2 / N$ (summing over all sites).




Figure 6.9: Measured τ_{-1} compared with time-dependent MFT (lines). The dashed bottom line corresponds to case #1 and the top line to case #2 (as discussed in the text). The vertical dashed line indicates when the MI phase is expected to enter into the trap.

experimentally measured lattice depth s_i . There is good agreement with experiment at low lattice depths ($s_i \leq 11$), where the MFT is a good approximation. For deeper lattice depths, the experimental timescales are approximately a factor of two slower than theory. Overall, the same qualitative trend is observed in both theory and experiment. Exact agreement with theory should not be expected — time-dependent Gutzwiller MFT is approximate, and, most importantly, does not consider finite-temperature effects. Although theory and experiment start at the same lattice depth in case #1, the starting condensate fraction in the theory calculation (the ground state condensate fraction at s_i) is larger than in the experiment. To account for this discrepancy, and to put an upper-bound on the theoretical timescales, in case #2 (top, solid line in Fig. 6.9) the theory starts from the same condensate fraction as in the experiment by artificially increasing the initial lattice depth. Indeed, the timescales calculated with these starting conditions are higher. All experimental data points lie between the theoretical curves from the two cases. This demonstrates that the essential features of the condensate fraction evolution are captured by the Gutzwiller theory.

The Gutzwiller theory can provide insight into the mechanisms responsible for setting the condensate fraction timescales. By truncating the Fock basis to 2 particles and diagonalizing for $\mu = U/2$, an analytic expression for the particle/hole gap is found to be

$$\Delta = \frac{U}{4} \left(1 + \sqrt{48z^2 \langle \mathbf{a} \rangle^2 \left(\frac{t}{U} \right)^2 + 1} \right). \quad (6.15)$$

This gap sets the timescale for adiabaticity. If our bandmapping ramp is longer than \hbar/Δ , then the ground state in the deep lattice will adiabatically follow to the ground state of the shallow lattice. Likewise, the condensate fraction will increase as the state changes. Because Δ is a fast timescale, it foils bandmapping as a probe of condensate fraction. It is difficult to be both adiabatic with respect to the band and fast with respect to the particle/hole gap Δ . Therefore, bandmapped distributions must be carefully interpreted when

measuring equilibrium properties. On the other hand, these fast timescales imply that bandmapping is an effective probe of many-body dynamics.

A natural extension to this work is to more closely investigate condensate fraction evolution across the MI-SF transition, which complements quench experiments also carried out in our group [242]. In addition, investigating timescales at finite temperature is an important next step and is related to a larger goal of developing theoretical tools to describe dynamics and thermalization in strongly correlated systems [36]. Finite-temperature issues were already relevant in this work. Even though our experiments were performed with low entropy gases, we could not make accurate, direct comparison to the zero temperature theory. At finite temperatures, a number of new effects become important, including the changing density of states and the suppression of T_C in the strongly interacting regime. Quantum simulation is ideal for studying finite-temperature timescales. Even though these timescales pose a theoretical challenge, they are straightforward to investigate experimentally.

Chapter 7

Developing Direct Thermometry Probes

One of the main ingredients for quantum simulation is the availability of probes to measure the thermodynamic state of the system. Without thermodynamic state information, the simulator cannot fulfill its mandate because direct comparison between theory and the physical systems being modeled (Fig. 1.2) is not possible. Unless we are using an impurity (§8), thermometry of an ultracold gas must be performed by measuring a property of the system with a known dependence on temperature. In the harmonic trap, a number of such probes are available (§4.1), however, these are mainly undeveloped and untested for use in the strongly interacting lattice regime. Direct lattice thermometry is difficult because approximations based on weak interactions break down. A main limitation is that many probes are explicitly based on a theoretical foundation, e.g., that the TOF expansion is given by Eqn. 4.5. This is a problematic starting point given that the goals of quantum simulation are mainly to explore physics in regimes where theory is not established. Additionally, our thermodynamic probes must be decoupled from the physics being simulated so that biases are not introduced into the results. Furthermore, if we are using the simulator to validate several theories, these theories cannot be used to determine thermodynamic information.

To circumvent these issues, a common technique is to probe the thermodynamic state in the harmonic trap and assume the loading process is adiabatic (§4.2.1). The thermodynamic state (T, μ) of the lattice gas is set by the entropy matching condition

$$\begin{aligned} S_{\text{lattice}}(T, \mu, N) &= S_{\text{harmonic}}, \\ N_{\text{lattice}}(T, \mu) &= N. \end{aligned} \tag{7.1}$$

However, there are a number of drawbacks to this approach. For one, it depends on adiabaticity, which is fundamentally unachievable due to light scattering from the lattice. Furthermore, to achieve the low temperatures (and entropies) associated with interesting phases in the lattice (e.g., the Néel state), in-lattice cooling techniques must be developed [36]. Direct thermometry of the lattice gas is vital to assess these techniques. A lack of proper thermometry in the lattice is a pressing issue and has led to controversy regarding the interpretation of results [258].

In this set of experiments, we investigate developing probes for direct thermometry of the lattice gas: bandmapping (§7.1), in-situ RMS width (§7.2), and condensate fraction (§7.3). An altogether different type of probe is to use an impurity gas in thermal constant with the lattice gas (§4.1.7). Although this also

The work presented in §7.1 and §7.2 of this chapter was published in: D. McKay, M. White, and B. DeMarco. *Lattice thermodynamics for ultracold atoms*. Phys. Rev. A 79, 063605 (2009) ©2009 American Physical Society. Authors are granted the right to include the article in a dissertation without requesting permission from the APS. Figs. 7.2, 7.3, 7.4, and 7.6 are reproduced from this publication.

allows for direct thermometry, it will be considered in a dedicated chapter due to its unique challenges and opportunities (§8).

7.1 Bandmapping

In this work, published in [3], we explore thermometry of a thermal Bose gas in a lattice by fitting the bandmapped distribution (§4.1.3) to a quasimomentum distribution with temperature as a free parameter in the fit. We compare the temperature measured from the fit against the temperature expected assuming adiabaticity when loading the atoms into the lattice. We find a large discrepancy between the measured and expected temperatures for $k_B T/t \gtrsim 2$. Using insight from numerical bandmapping calculations, we conclude that the main source of error is the inability of the bandmapping process to faithfully produce the quasimomentum distribution when the temperature of the gas is larger than the bandwidth.

In §4.1.3, we introduced the concept of measuring the thermodynamic state of the lattice gas by fitting the quasimomentum distribution obtained via bandmapping and TOF expansion. Bandmapping is a process whereby the lattice is ramped down adiabatically with respect to the bandgap and fast compared to all other timescales. This maps quasimomentum (q) in the lattice to momentum (p). The momentum distribution is then measured by TOF expansion (§4.1.2). Once we have measured the quasimomentum distribution, we can fit it to Eqn. 4.19 with temperature and chemical potential as free parameters. The crucial step is the mapping of $q \rightarrow p$. Already, in §6.2, we saw that interactions can disrupt this process. In this work, we focus on investigating the limitations at the single-particle level.




Figure 7.1: RMS Quasimomentum width versus temperature calculated from a thermal average of the exact quasimomentum states using Maxwell-Boltzmann statistics for $s = 6$, $d = 400\text{nm}$, and $\nu = 54.5\text{Hz}$ (see Eqn. 2.87 and Eqn. 5.3). At low temperatures, the width is set by the size of the $n = 0$ state (the ground state), and at high temperatures the size saturates to $q_B/\sqrt{3}$ as the band fills.

There are two main potential weaknesses of the bandmapping probe. The first issue is that the quasi-momentum distribution becomes a less sensitive thermodynamic probe as the temperature becomes larger than the tunneling t . To illustrate this point, we plot the RMS width of the quasimomentum distribution versus temperature in Fig. 7.1. The probe becomes very sensitive to non-thermal variations of the width due to noise and imaging distortions since the width saturates in this regime. Furthermore, for $k_B T \gtrsim 2t$ the edge of the distribution becomes critical, and a second issue arises related to the bandmapping process itself. The bandgap decreases (see Fig. 2.5) for states at the edge of the Brillouin zone, and, therefore, it is harder to satisfy the bandmapping adiabaticity condition. Imperfect mapping of $q \rightarrow p$ will smooth out the sharp edge of the Brillouin zone and cause large errors in the measured temperature. It is important to understand the impact of these two issues for bandmapping thermometry.

7.1.1 Numerical Calculation

These issues are all single-particle effects, which are straightforward to explore numerically. To gain insight into the problem, we perform a bandmapping and TOF calculation for the 1D non-interacting harmonically trapped lattice gas by solving the time-dependent Schrödinger equation using a Crank-Nicolson solver (code in §I.4.3). To construct the thermal bandmapped distributions, we use the exact single-particle eigenstates. Analytic solutions for the eigenstates and eigenvalues from Ref. [133] are written out in Eqn. 5.3 and Eqn. 2.87, respectively. The eigenstates given by Eqn. 5.3 are the coefficients f_i^n of the Wannier function on each site. We approximate the Wannier function by a Gaussian so that

$$\Phi_n(x) = \sum_j f_j^n \left([\pi\sigma^2]^{-1/4} e^{-(x-jd)^2/2\sigma^2} \right) \quad (7.2)$$

where $\sigma = (d/\pi)(2s)^{-1/4}$. We evolve each eigenstate for a linear ramp down of the lattice depth in time τ_0 , so the time-dependent potential is

$$V(\mathbf{r}) = \frac{1}{2}m\omega^2 r^2 + s_0 \left(1 - \frac{\mathbf{r}}{\tau_0} \right) \cos^2(kz). \quad (7.3)$$

After bandmapping ($\mathbf{r} = \tau_0$), the state is propagated for 20ms TOF using the free particle propagator (Eqn. 4.5). This procedure results in a set of bandmapped states $\Psi_n(x)$. We show the results for the bandmapped $n = 0$ eigenstate in Fig. 7.2. To construct the thermal bandmapped distribution, we make a thermal average over the propagated states

$$\rho(x) = \sum_n |\Psi_n(x)|^2 \left(e^{[E_n - \mu]/k_B T} - 1 \right)^{-1}. \quad (7.4)$$

For simplicity we assume that $\mu \ll 0$ so that the statistics are Maxwell-Boltzmann. We then fit $\rho(x)$ to the semiclassical 1D Maxwell-Boltzmann quasimomentum distribution

$$\rho(x)_{SC} = A e^{-2 \left(1 - \cos \left[\frac{\pi m}{\hbar q_B} x \right] \right) \left(\frac{t}{k_B T} \right)} \quad (7.5)$$

where \mathbf{r} is the TOF expansion time, and A and $k_B T/t$ are free parameters of the fit.

The distributions for $k_B T = 2.9t$ and $k_B T = 23t$ and corresponding fits, are illustrated in Fig. 7.3. We




Figure 7.2: Example of the bandmapping calculation for the $n = 0$ eigenstate. The in-trap eigenstate is illustrated in the top-left, and the in-trap state after $750\mu\text{s}$ bandmapping is shown in the top-right. The main effect is to remove the density modulations caused by the lattice. The states after 20ms TOF propagation are shown below.




Figure 7.3: Calculated distributions (black line) for (a) $k_B T/t = 2.9$ and (b) $k_B T/t = 23$. The distribution is fit (blue line) using the function given by Eqn. 7.5 to measure $k_B T'/t$. For comparison, the exact quasimomentum distribution is shown (red line) using the exact eigenstates. (c) The results of the measured $k_B T'/t$ from fitting the calculated distributions versus the value $k_B T/t$ used to generate the distribution (using Eqn. 7.4). A significant deviation is observed for $k_B T/t > 5$. The reason for this discrepancy is evident from (a) and (b) — the edge of the calculated distribution is distorted.

also plot the quasimomentum distribution obtained directly from the eigenstates, which should be identical to the bandmapped distribution in the limit that the $q \rightarrow p$ mapping is exact. The difference between these distributions is visually apparent — the bandmapped distribution is smoothed at the edge. To quantify this effect, we plot the temperature of the fit ($k_B T'/t$) as a function of the actual temperature used in Eqn. 7.4 ($k_B T/t$). As $k_B T/t$ increases, $k_B T'/t$ greatly underestimates $k_B T/t$. In this temperature range, the fit is dependent on the shape of the distribution at the edge, and so the smoothing, an artifact of the bandmapping process, dominates and puts an upper bound on the temperature that can be accurately measured. This is even more problematic for the experiment, because other smoothing processes (e.g., finite resolution) are also present.

7.1.2 Experiment and Results

To investigate this problem experimentally we use the ^{87}Rb apparatus (§3.1) in the hybrid single-beam dipole/magnetic quadrupole trap configuration. We prepare ultracold gases in the dipole trap that are not condensed. To access low temperatures so that we can probe the $k_B T \lesssim t$ range in the lattice, we use sufficiently small number of atoms such that $T > T_C$. The temperature in the harmonic trap is measured using the momentum distribution after TOF (§4.1.2). The gas is adiabatically loaded into a $\lambda = 812\text{nm}$, spin-independent 3D lattice to a depth of either $s = 2$ ($U = 0.05E_R$, $t = 0.17E_R$) or $s = 6$ ($U = 0.17E_R$, $t = 0.05E_R$). These depths were selected to be far from the MI transition in order to minimize the role of interactions. While $s = 6$ is well-described by a tight-binding Hamiltonian, for $s = 2$ the next-nearest neighbor tunneling is 14% of t ; this leads to a small deviation from the tight-binding dispersion. The temperature in the lattice is inferred by matching the lattice entropy to the harmonic trap entropy using semiclassical non-interacting thermodynamics (§5.1).




Figure 7.4: Temperature measured by fitting bandmapped distributions (points) versus the temperature expected for adiabatic loading into the lattice (grey, shaded region) at (a) $s = 2$ and (b) $s = 6$. For this data, N varied from 3×10^3 – 1.4×10^5 and fugacity from 0.35–0.75.

Once the atoms are in the lattice, we bandmap in $750\mu\text{s}$ and expand for 8–20ms TOF¹. Images are then fit to Eqn. 4.19 (the 2D semiclassical Bose-Einstein quasimomentum distribution) with $k_B T'/t$ and fugacity as a free parameter. The results are plotted in Fig. 7.4. The shaded region is the predicted temperature

¹Shorter TOFs were used to obtain high enough signal-to-noise ratio when using a small number of particles.

based on the entropy matching calculation. Similar to the calculation, we see that the measured temperature is substantially lower than the predicted temperature when $k_B T \gtrsim 2.5t$. One possibility is that the temperature in the lattice is truly different than we predict, but measurements discussed in §7.2 confirm that this is not the case. Also, the measured temperatures are lower than expected, whereas a failure of adiabaticity during the ramping process would lead to a higher temperature in the lattice. Effects due to interactions can be straightforwardly ruled out as the source of this discrepancy, as the ratio of mean energy to estimated interaction energy per particle is greater than 30.

Therefore, the experimental and calculated data consistently show that bandmapping fails to give accurate quasimomentum distributions at high temperature ($k_B T/t \gtrsim 2t$) when the band starts to be filled. However, bandmapping can still be used for thermometry at low temperatures as in Ref. [66]. An alternative to bandmapping is to turn the lattice off fast compared to all timescales and measure the momentum distribution after expansion. As noted in §4.1.3, the momentum distribution is a periodic function of the quasimomentum distribution. Still, deconvolving the quasimomentum distribution from the momentum distribution can be non-trivial after the imaging process integrates through the distribution along the probe beam. Nevertheless, even if these challenges could be overcome, the quasimomentum distribution is not sensitive to temperature when the band is filled. Therefore, it is prudent to look at other thermometry methods, particularly in-situ (§4.1.5).

7.2 In-situ RMS Width

In this work, also published in [3], we experimentally investigate measuring the temperature from the in-situ RMS width in the temperature regime where bandmapping fails. We verify that this is a reliable technique and is only limited by our ability to image the cloud in-situ.

A complementary approach to measuring temperature using bandmapped distributions is to measure the RMS width of the in-situ distribution (see §4.1.5 for an overview of in-situ probes). This approach addresses the main issues that result in large errors for bandmapping when $k_B T \gtrsim t$. In-situ probes do not rely on any time-dependent mapping operations, such as bandmapping, which ensures that we measure the equilibrium distribution. Also, the harmonic trap energy scale is unbounded. So, unlike quasimomentum which drastically loses sensitivity to temperature when the band starts to fill, the sensitivity of the RMS width is only weakly a function of temperature. Indeed, when the band is full, thermal energy is exclusively added to trap degrees of freedom.

For a single-band non-interacting gas at high temperatures the system can be described by Maxwell-Boltzmann statistics and from equipartition the RMS size is

$$\langle x^2 \rangle = \frac{k_B T}{m\omega^2}. \quad (7.6)$$

Eqn. 7.6 emphasizes the point that the in-situ width grows continuously with temperature. In Fig. 7.5, we plot Eqn. 7.6 compared to the exact width from the non-interacting eigenstates. The agreement is almost exact, except at extremely low temperatures when $k_B T \approx \hbar\omega^*$ ($\omega^* = \omega\sqrt{m/m^*}$), and the width is set by the $n = 0$ eigenstate. The width is a useful probe even when interactions are strong. In Fig. 7.5, we also




Figure 7.5: Calculated in-situ RMS width (black line) versus temperature for the (a) non-interacting gas and (b) in the atomic limit using Maxwell-Boltzmann statistics. The grey, dotted line is the equipartition result $\langle x^2 \rangle = k_B T / m \omega^2$. For (a) the calculation used the exact eigenstates of the non-interacting harmonically trapped lattice gas for $s = 6$, $d = 400\text{nm}$, and $\nu = 54.5\text{Hz}$. There is excellent agreement with the equipartition result until very low temperatures $k_B T \lesssim \hbar \omega^*$ when the size is set by the finite size of the $n = 0$ eigenstate. For (b), the calculation was performed numerically in the atomic limit (§5.2) truncating at 10 particles per site and $U = 0.5 E_R$, $d = 400\text{nm}$, $\nu = 50\text{Hz}$, and $N = 5 \times 10^4$. In the atomic limit, the width is insensitive to temperature until $k_B T \gtrsim 0.2U$, consistent with the MI “melting temperature” defined in Ref. [137]. For $k_B T > U$, the atomic limit converges to the equipartition result.

plot Eqn. 7.6 compared to the width in the atomic limit. Here the width is sensitive to temperature if $k_B T \gtrsim 0.2U$ and equal to the equipartition result (Eqn. 7.6) if $k_B T \gtrsim U$. When $k_B T < 0.2U$ in the strongly interacting gas, temperature information can be extracted from the distribution at the edge of the gas, but this requires high resolution imaging [62]. One assumption of this probe is that the system is in global thermal equilibrium, which is not the case in all experiments [236]. However, for the experiment described in Ref. [236] there is a large MI phase in the system, which is a regime we are not considering here.

Experimentally, we investigate this in-situ temperature probe in the ^{87}Rb apparatus using the same configuration as in the bandmapping experiment (§7.1). The measurements are performed using non-condensed gases loaded into an $s = 6$ lattice. One drawback of in-situ probing is that the fields from the trapping potential can affect the imaging process. This was an issue here since the trap includes a magnetic quadrupole and it takes $\approx 2\text{ms}$ after releasing the atoms before a uniform imaging field is created. Therefore, to measure the in-situ size we turned off the potentials suddenly and let the gas expand for some small time τ . We measure the width using a Gaussian fit for different τ and fit the results to the standard expansion formula $\sqrt{x_0^2 + (v\tau)^2}$ to determine x_0 . Characteristic data are shown in Fig. 7.6. A direct image would be preferred for a more precise implementation of this probe. From x_0 we use Eqn. 7.6 to determine the temperature. Using the same procedure as the bandmapping experiment, we calculate the predicted temperature in the lattice using the harmonic trap entropy and assuming the loading process is adiabatic. The agreement between the measured and predicted temperature is excellent as evident in Fig. 7.6. In this range of temperature at $s = 6$, the quasimomentum probe (Fig. 7.4) did not agree with the predicted temperature by greater than 10 standard deviations.




Figure 7.6: Experimental results using the in-situ RMS width for thermometry. (a) Temperature measured from the RMS width (square points) using $T = m\omega^2\langle x^2 \rangle/k_B$ versus the temperature predicted assuming adiabatic transfer into the lattice (grey, shaded region). There is good agreement between the measured and expected temperatures, especially compared to the level of disagreement for the bandmapping probe in the same regime. (b) Sample expansion data (square points) used to extrapolate the in-situ width (grey curve).

These results, at high-temperatures confirm that the observed inaccuracies in the quasimomentum probe are real, and that the in-situ width is an accurate measure of temperature. Certainly a priority for refining this probe is to reduce uncertainties, which were large mainly due to two issues. The main source of uncertainty is the extrapolation method, which was required because of the magnetic trap. In an all-optical setup, such as the crossed-dipole trap, we can take in-situ images since the magnetic field is already uniform. Second, we could improve our resolution and magnification, which in this experiment was optimized for TOF imaging.

Ultimately, our investigation demonstrates that the in-situ probe is a straightforward method for measuring the temperature despite the complications of the lattice. The main validation of any probe is whether or not it is adapted by other groups. Shortly following the publication of our measurement, this method was successfully compared against the temperature measured using spin-gradient thermometry [59]. Also, probing temperature using the in-situ distribution is now a standard tool in high-resolution experiments [56–58].

7.3 Lattice Condensate Fraction

In this work, we investigate peak fraction, a heuristic measure of condensate fraction, as a thermodynamic probe for the lattice gas in the SF regime of the BH model. We formally define peak fraction based on a procedure for measuring and fitting experimental TOF images. This measure is designed to be as close to the actual condensate fraction as possible and to be non-zero exclusively when the condensate fraction is also non-zero. Therefore, the peak fraction measure is ideal for determining the SF phase boundary. We test this fitting procedure against simulated data. We measure the peak fraction of the gas at $s = 6 - 12$ for a range of entropies $S/Nk_B = 0.25 - 3$ (determined in the harmonic trap). The peak fraction is compared to condensate fraction predicted using the non-interacting (§5.1), HFBP (§5.3.1), and SDMFT (§5.3.2) theories.

There are discrepancies between the measured peak fraction and predicted condensate fraction. In particular, the peak fraction is non-zero when the predicted condensate fraction is zero in certain regimes. This is surprising since most non-idealities (e.g., heating from the lattice beams) result in a lower than expected condensate fraction. This may point to two effects. First, this result could reflect a failure of the peak fraction measure because of sharp peaks in the momentum distribution not associated with the condensate. This possibility was pointed out theoretically [60, 258], but others have argued [259] that this effect is not appreciable in trapped, finite size experiments. Second, this discrepancy could be evidence for long timescales associated with reversible changes to the condensate fraction.

Based on the arguments given at the beginning of this chapter (§7), we are motivated to search for thermodynamic probes that measure a temperature proxy: a unique, easy to measure quantity that is monotonic in and sensitive to temperature, but insensitive to experimental details. This proxy can be later calibrated to temperature in a number of ways. For example, we can make a direct comparison to time-consuming numerical calculations (e.g., QMC calculations of the distribution [67]) at a few points and then interpolate.




Figure 7.7: Various thermodynamic proxies that have been investigated for the TOF distribution of the lattice gas in the SF regime illustrated with experimental data ($s = 8$). (a) Visibility: $\frac{n_1 - n_2}{n_1 + n_2}$ where n_1 (n_2) is the number of atoms in the area indicated by box 1 (box 2). (b) Peak height and width: from the 2D distribution (inset) evaluate a 1D cross section along the dotted line. The peak height (ν) and peak width (σ) are labeled on the plot. These measures may also be applied to the first-order diffraction peaks. (c) Peak fraction: the number of atoms under the peaks (blue) compared to all the atoms (blue plus red). To determine which atoms under the peaks are condensate and which atoms are non-condensate requires careful consideration.

Thermodynamic proxies are well-suited to the TOF distribution of a gas in the SF regime of the BH model for two reasons. First, the strongly interacting SF regime is difficult to treat theoretically because it contains the crossover between theories that treat t as a small parameter and those that treat U as a small parameter (see Fig. 5.5). Only a non-perturbative approach such as QMC can accurately describe the whole parameter space. Therefore, we cannot rely on theory for probes in this regime. Proxies are also ideal for the TOF distribution in the SF regime because there are strong features in the distribution; the defining attribute of the SF is that the momentum distribution is bimodal with sharp, narrow features (‘peaks’) corresponding to atoms in the condensed state. The properties of these peaks are ideal potential proxies, a number of which have been investigated, including visibility [260], peak height, peak width [67], and peak fraction. These methods are illustrated in Fig. 7.7.

In this experiment, we focus on peak fraction, which is an empirical measure that quantifies the fraction of atoms in the narrow peaks of the bimodal distribution. With the appropriate fitting function and procedure, the peak fraction converges to the condensate fraction — the number of atoms in the macroscopically occupied single particle state [53]. In the harmonic trap, a standard peak fraction probe exists that can accurately measure the condensate fraction. As discussed in §4.1.4, this probe is based on a priori theoretical guidance, which is not possible in the lattice. Even though a lattice peak fraction probe is not necessarily a condensate fraction probe, the advantages outlined in §4.1.4 still apply. One of the main advantages is that a peak fraction probe is weakly dependent on the functional form used to identify the narrow condensate peaks in the distribution. Therefore, the peak fraction probe is decoupled from theory and, to some level, local distortions of the distribution during TOF. Since it is a ratio, it is also insensitive to global distortions and certain sources of technical noise. Consequently, peak fraction is an ideal choice as a thermodynamic probe.

7.3.1 Defining the Peak Fraction Fitting Procedure

The premise of the peak fraction probe is to define a procedure that counts the number of condensate atoms in a TOF distribution. Since the distribution is bimodal, identifying regions of the distribution that contain condensate atoms is typically not difficult — the condensate is a sharp, narrow peak. However, these “condensate regions” may also contain non-condensate atoms. For example, in the TOF distribution of the harmonically trapped condensate, discussed in §4.1.4, the condensate atoms are contained in a circle centered on zero momentum with a size defined by the TF radius. However, this region is not exclusively condensate atoms; the distribution of non-condensate atoms is also centered on zero momentum, but continues outward beyond the TF radius. To definitively distinguish condensate from non-condensate atoms one strategy is to completely separate condensate from non-condensate using a time-dependent potential (e.g., Bragg scattering [126]), but this procedure is not always successful and it adds new complications. An added difficulty, unique to the lattice system, is that the condensate is spread over several peaks. In principle, bandmapping can be used to reduce these to a single peak, however, the issues discussed in §6.2 and §7.1 prevent us from pursuing this option². Therefore, we define a peak fraction procedure for TOF distributions after the lattice and all other confining potentials are turned off (‘snapped-off’) as fast as possible (within 10ns). Our goal is to define our measure in such a way that peak fraction can identify the phase boundary and is therefore zero exclusively when the condensate fraction is zero.

Although our goal is to define a procedure that is simple and independent of theory, we can still use theory in limiting cases to motivate our overall strategy. For example, if we neglect interactions during expansion, then the image will be the probe-integrated momentum distribution, which was discussed in §4.1.3. The overall integrated tight-binding momentum distribution is

$$n(p_x, p_y) = \int dp_z |w(p_x)w(p_y)w(p_z)|^2 \sum_{j,k} e^{i\frac{d}{\hbar}\vec{p}\cdot(\vec{j}-\vec{k})} \langle \hat{a}_j^\dagger \hat{a}_k \rangle \quad (7.7)$$

²On a historical note, our original intent for this experiment was to probe peak fraction using bandmapped images. However, during the preliminary phases of this work we noticed some of the issues associated with bandmapping, and our subsequent investigation evolved into the work presented in §6.2.

where $w(p)$ is the Fourier transform of the Wannier function (Eqn. 2.71), which is approximately Gaussian. In the thermal, non-interacting limit, Eqn. 7.7 is given by combining Eqn. 4.15 and Eqn. 4.18 (integrating along a lattice axis)

$$n_{th}(p_x, p_y) \approx A |w(p_x)w(p_y)|^2 \sum_{j,k} \left(\sum_{n=1}^{\infty} \frac{(\mp)^n \zeta^n e^{-2nt\beta\{3-\cos((p_x-P_j)d/\hbar)-\cos((p_y-P_k)d/\hbar)\}} I_0[2nt\beta]}{n^{3/2}} \right), \quad (7.8)$$

where j sums over the reciprocal lattice momenta and quasimomentum distribution is only defined for $|p_x - P_j|, |p_y - P_k| < q_B$. Alternatively, we can make a Gutzwiller approximation $|\Psi\rangle = \prod_j |\phi_j\rangle$, then

$$n(p_x, p_y) = \int dp_z |w(p_x)w(p_y)w(p_z)|^2 \left[N - \sum_j |\langle \mathbf{a}_j \rangle|^2 + \left| \sum_j e^{i\frac{d}{\hbar} \vec{p} \cdot \vec{j}} \langle \mathbf{a}_j \rangle \right|^2 \right], \quad (7.9)$$

where $\langle \hat{a}_j \rangle = \langle \phi_j | \mathbf{a}_j | \phi_j \rangle = a_j$, and $|a_j|^2$ is the number of condensate atoms on site j ; a_j can be calculated from the LDA using SDMFT (§5.3.2). In the limit $t/U \rightarrow 0$, then $a_j \rightarrow 0$ and

$$n(p_x, p_y) \propto |w(p_x)w(p_y)|^2, \quad (7.10)$$

which is just a featureless Gaussian. When $U/t \rightarrow 0, T \rightarrow 0$ (i.e., the limit of a non-interacting, pure condensate), then $\sum_j |\langle \hat{a}_j \rangle|^2 = N$, and

$$n(p_x, p_y) = \int dp_z |w(p_x)w(p_y)w(p_z)|^2 \left| \sum_j e^{i\frac{d}{\hbar} \vec{p} \cdot \vec{j}} \langle \hat{a}_j \rangle \right|^2, \quad (7.11)$$

where the last term is just the square of the Fourier series of the condensate wavefunction. For an infinite, uniform lattice this distribution is a delta function periodic in h/d (the “diffraction peaks”), with a Wannier function envelope. For typical lattice depths, only the first-order peaks are sizable. For a trapped condensate, the width of the condensate peak in TOF images, based on Eqn. 7.11, is set by the inverse size in the trap. These peaks are also broadened by interactions (§4.1.4) and the Fresnel terms relevant for finite expansion times [261]. Combinations of Eqn. 7.8, Eqn. 7.10, and Eqn. 7.11 (with the peaks approximated as a Gaussian or TF profile) have been used by several other groups to determine the peak fraction. Although this is only the zero-order theory, it reveals the basic general features of the TOF distribution of the lattice gas — strong diffraction peaks associated with the condensate, an overall broad Gaussian from incoherent quantum depletion, and a narrow structured distribution for the thermally excited atoms.

When interactions cannot be neglected, the functional form of the expanded condensate peaks and the distribution of non-condensate atoms under these peaks is unresolved. Insight into these issues can be gained from the results of our numerical calculation of the expansion of the harmonically trapped weakly interacting gas (§4.1.4). In the harmonically trapped gas, the condensate expansion is well-described by the GP equation (Eqn. 2.47) and is driven by the condensate mean-field interactions. The condensate distribution remains a TF profile after expansion and is much larger than expected from a Fourier transform of the in-trap density. It seems reasonable to assume that a similar type of dynamics governs the expansion of the lattice condensate. Counter-arguments for neglecting the mean-field interactions during expansion are

given in [261], but we empirically find agreement with the GP expansion (§7.3.2). A complication is that there are multiple condensate peaks, each with a different number of atoms. It is likely that each peak has unique expansion dynamics. Indeed, the peak widths in the experimental data are different; the outer peaks (with a smaller number of atoms) are narrower than the central (zero-order) peak. Another issue associated with multiple peaks is that during early expansion times, when the peaks are spatially overlapped, binary collisions between atoms in different peaks moving relative to each other cause atoms to be scattered out of the peaks [262]. This process will increase the discrepancy between peak fraction and condensate fraction when the condensate number is large. There is little that can be done to account for this except to add a corrective factor using a first-principles calculation.

The main issue is the distribution of non-condensate atoms under the condensate peak. There is no signature in images that can be used to resolve this problem. We could rely on the non-interacting thermal distribution (Eqn. 7.8) to infer the distribution under the peak by fitting to the distribution outside the peak, as is used in the harmonic trap fits (§4.1.4). This conjecture is almost certainly false. QMC calculations of the full interacting non-condensate momentum distribution strongly deviate from the non-interacting case [60]. The best information we have is the density of non-condensate atoms directly adjacent to the peak. Given this information, a reasonable assumption is that the distribution underneath is flat. This is supported by our calculation of non-condensate atoms expanding in the mean-field potential of the condensate when released from the harmonic trap (§4.1.4). These interactions cause the non-condensate distribution under the peak to be more uniform, and a similar effect is likely to occur for the lattice gas.

Based on this discussion, we define our peak fraction fitting procedure in the following way:

1. An area around each condensate peak in the original image (I_1) is masked. This mask is determined ad hoc based on the peak position and typical peak size. Next, the masked image is fit to a Gaussian with an offset and slope. This fit is subtracted from the image; the resulting image is I_2 . The number of atoms from the Gaussian fit is N_{Gauss} . If the condensate fraction is very high (i.e. no broad Gaussian feature is evident in image I_1), then the Gaussian peak height may be constrained to zero during the fit.
2. Each peak in I_2 is fit to a single TF profile (no offset or slope) in a series of separate fits. An ad hoc method is used to minimize the size of the mask around each peak. The number of atoms from all these TF fits is N_{pk} . The TF radii along x and y from the fit to peak i are $\sigma_{x,i}$ and $\sigma_{y,i}$.
3. The average density at the edge of peak i (\bar{n}_i) is measured by summing around the perimeter set by $\sigma_{x,i}$ and $\sigma_{y,i}$. The number of non-condensate atoms under all the peaks is $N_{up} = \sum_i \pi \sigma_{x,i} \sigma_{y,i} \bar{n}_i$.
4. The condensate peaks in I_2 are masked to obtain I_3 . The image I_3 is summed to measure the remaining atoms N_{Σ} .

After finishing these steps, the peak fraction is determined according to

$$\frac{N_{pk} - N_{up}}{N_{pk} + N_{Gauss} + N_{\Sigma}}. \quad (7.12)$$

In practice, we typically use two images to increase the signal-to-noise ratio over a larger dynamic range (2-shot technique discussed in §4.1.4). We take a high OD image to determine N_{Gauss} , N_{Σ} , and N_{up} , and

a low OD image to measure N_{pk} . The fitting procedure is illustrated in Fig. 7.8 step-by-step for typical experimental data.

A slight variation of the above procedure is to replace the denominator of Eqn. 7.12 with the total number measured using a different image. For example, if we ramp the lattice to high s ($t/U \rightarrow 0$), then the TOF distribution will become approximately a single Gaussian, which we can easily fit to determine total N . Another option (if we know N is conserved during the lattice ramp) is to measure N in the harmonic trap where the distribution is well-known. Finding alternative methods for measuring N reduces our biggest source of error: the sum in step #4. Unlike fitting to a smooth function, which acts as a low-pass filter, the sum does not filter noise, such as fringes in our imaging system and photodetector shot noise. The sum is also very sensitive to offsets in the image. Since the imaging area is very large, the atom-number signal can be overwhelmed by this offset. Fortunately, the Gaussian fit in step #1 mostly removes this offset. Residual noise can be reduced by averaging N over several shots if N fluctuations are small.

7.3.2 Simulated Data

To test the above fitting procedure, we apply it to simulated data including realistic technical noise. We use Eqn. 7.10 to represent the quantum depleted non-condensate atoms and Eqn. 7.8 for the thermal non-condensate atoms. For the condensate, we use the momentum-space Wannier function envelope to determine the relative number of atoms in each diffraction peak and then assume that each peak expands as a TF profile according to the scaling discussed in §4.1.4. The exact details of the simulated data are given in §G.4. Sample data is shown in Fig. 7.9 compared to experimental data at the same lattice depth. To replicate noise, we add the simulated data to a region of an experimental image that did not contain any atoms. This effectively captures fringes, photodetection shot noise, and artifacts such as offsets and slopes.

We generate these simulated data samples at $s = 12$ for several different condensate fractions and add them to 3 separate noise samples. The peak fraction is measured using the peak fraction fit procedure. The results are plotted in Fig. 7.9. The simulated data is a simple approximation for the actual TOF data and ignores a number of effects, including all non-condensate interactions (in the lattice and during the expansion), and finite expansion time corrections (i.e., Fresnel terms [261]). Also, the condensate expansion is described heuristically to agree with the experimental data. Therefore, this test is not a proof, but it does benchmark the procedure for many generic features of the experimental data, such as multiple peaks, a broad non-condensate distribution, and noise.

7.3.3 Experiment

To experimentally investigate peak fraction thermometry we use the ^{87}Rb apparatus in the single beam dipole plus magnetic quadrupole trap configuration (see §3.1 for more detail). We use approximately 2×10^5 ^{87}Rb atoms in $|1, -1\rangle$ state at various values of the entropy per particle in the harmonic trap ranging from $S/Nk_B \approx 0.2$ (almost a pure condensate) to $S/Nk_B \approx 3$ (just below T_C). The initial entropy is controlled using our finite temperature techniques (§4.2.2); to achieve lower entropies we evaporate to lower dipole trap depths. Since atoms are lost when we evaporate to lower depths, we decrease the efficiency of earlier evaporation stages to maintain a reasonably constant atom number. After evaporation, we ramp to a con-




Figure 7.8: Applying the fit procedure to experimental data ($s = 10$, 7 shot average). (a) Raw images: high OD (fully repumped) image to acquire high signal-to-noise ratio for the non-condensate, and low OD (partially repumped) image used to obtain an unsaturated image of the condensate peaks. (b) Applying the fit procedure to the high OD data. (i) The condensate peaks are masked and (ii) the remainder of the image is fit to a Gaussian (red area is the fit). The Gaussian fit is subtracted from the image. (iii) Each peak is fit to a TF profile. The result of the fits is shown aggregated onto a single image. We use only the peak widths and positions from the high OD data. (iv) The remaining non-condensate atoms are counted assuming a flat distribution under the peaks. The inferred density directly under the peaks is delineated by the light grey lines. (c) Steps (i)–(iii) are repeated for the low OD data to determine the number of atoms in each peak (blue). The number of non-condensate atoms under each peak from step (iv),(b) is subtracted to measure the condensate number.




Figure 7.9: Simulated lattice data used to test the peak fraction fitting procedure. (a) Samples of simulated data compared to data from the experiment for $s = 6$ and $s = 12$. For simplicity the simulated data is integrated along a lattice axis and rotated to match the orientation of the experimental data. The procedure used to generate the simulated data is summarized in the text and detailed in §G.4. Additionally, the noise added to the simulated data was taken from an empty area of an experimental image. While there are many qualitative similarities between the two sets of images, there are also clear differences. The simulation overestimates the peak width, particularly for the outer diffraction peaks and for all peaks at $s = 12$. The simulated peaks are spherically symmetric, whereas some of the experimental peaks are rotated and elliptical. (b) Measuring the peak fraction for the $s = 12$ simulated data using the procedure outlined in the text. We consider the case for no noise and for three different noise samples. For no noise, the fit is nearly perfect; with noise, errors are at the 10% level (grey line). The magnitude of the error is dependent on the specific noise sample.

stant dipole depth with $\bar{\nu} = 35.76(8)\text{Hz}$ for all measurements. The condensate fraction in the harmonic trap is measured using the 2-shot probe described in §4.1.4. The entropy in the harmonic trap S_0 is calculated using the semi-ideal model (see Eqn. 2.57).

Atoms are loaded into the lattice ($\lambda = 812.0\text{nm}$, $E_R = 2.307 \times 10^{-30}\text{J}$) over 100ms, which is intended to be adiabatic (§4.2.1). We investigate four different depths, $s = 6, 8, 10$, and 12, all corresponding to the SF region of the BH phase diagram (see Fig. 2.11). The lattice increases the harmonic confinement (see §D for more information on trap frequency measurements). The lattice depth is calibrated using Kapitza-Dirac diffraction to within 1% error, with a maximum drift of 6% over a month (see §C.5). The experimental parameters $t U$ and ν are summarized in Table 7.1. Once loaded into the lattice the atoms are held for 10ms, then the potentials are snapped-off (lattice and dipole $< 10\text{ns}$, magnetic quadrupole $< 200\mu\text{s}$), and the atoms are imaged after 20ms TOF expansion. We take images from two subsequent experimental cycles, one image with high OD and the other with low OD using partial repumping (§4.1.1, §4.1.4, see §F.2 for information on the repump calibration). The peak fraction fit procedure described in the previous section (§7.3.1) is applied to the two lattice images. Fig. 7.8 illustrates a step-by-step fit to one set of data. For each lattice depth and initial condensate fraction, we average the results of 7 sets; each set is a high OD/low OD pair in the lattice and harmonic trap corresponding to four images from four separate experiment cycles. The averaged results for each lattice depth and initial entropy are given in Table 7.1.

The averaged high OD lattice images for each lattice depth and initial entropy are shown in Fig. 7.10. A clear trend is evident; the non-condensate increases monotonically with entropy and lattice depth. Sharp diffraction peaks are visible in the majority of images and are only missing in the two highest entropies at $s = 12$ and the highest entropy at $s = 10$. The measured peak fraction versus initial entropy and harmonic trap condensate fraction is shown in Fig. 7.11.

7.3.4 Comparison with Theory and Discussion

In Fig. 7.11, we compare the experimentally measured peak fraction to theoretical predictions for condensate fraction using the non-interacting (§5.1), HFBP (§5.3.1), and SDMFT (§5.3.2) theories. There are significant discrepancies between the theoretical predictions and the experimental data. For all lattice depths at low entropies, the measured peak fraction is lower than predicted and vice-versa at high entropies. The disagreement at high entropies is perhaps the most surprising because most non-ideal behavior (e.g., heating from the lattice beams) is expected to result in lower than predicted peak fraction. The question raised by these data is whether the theory is incorrect, the peak fraction fit procedure is not a good representation of condensate fraction, or that the assumption of adiabaticity is flawed.

The validity of each theory is a difficult question to assess, but an additional advantage of these plots is that we can directly compare different theories against each other. This comparison provides further insight into the regimes of validity for these theories. At the lowest lattice depth ($s = 6$) the theories all agree, so we can be quite confident in the prediction. At $s = 6$ it is not surprising that the non-interacting and HFBP theories are fairly consistent — the gas is far from the MI transition, and so we expect interactions to have a small effect as demonstrated by the small shift between the HFBP and non-interacting curves. It is surprising that the finite temperature extension to the SDMFT is fairly accurate, since this theory is

	$s = 6$	$t = 5.08 \times 10^{-2}$	$U = 0.170$	$\nu = 51.27$	
$(N_0/N)_0$	0.109(5)	0.212(4)	0.536(7)	0.816(8)	0.905(8)
$S (k_B/N)$	2.98(2)	2.56(2)	1.44(2)	0.55(3)	0.28(2)
$N (\times 10^3)$	256(2)	224(2)	186(4)	245(3)	220(3)
$(N_0/N)_{\text{latt}}$	0.061(2)	0.16(1)	0.44(1)	0.68(2)	0.69(2)

	$s = 8$	$t = 3.08 \times 10^{-2}$	$U = 0.227$	$\nu = 55.48$	
$(N_0/N)_0$	0.12(1)	0.21(1)	0.57(8)	0.80(1)	0.90(1)
$S (k_B/N)$	2.95(4)	2.58(3)	1.33(3)	0.60(2)	0.28(4)
$N (\times 10^3)$	259(2)	215(2)	198(3)	239(3)	210(5)
$(N_0/N)_{\text{latt}}$	0.032(2)	0.108(5)	0.38(2)	0.54(2)	0.63(2)

	$s = 10$	$t = 1.92 \times 10^{-2}$	$U = 0.281$	$\nu = 59.4$	
$(N_0/N)_0$	0.109(5)	0.22(1)	0.57(1)	0.81(1)	0.914(4)
$S (k_B/N)$	2.98(2)	2.56(4)	1.34(2)	0.58(3)	0.25(1)
$N (\times 10^3)$	266(4)	207(4)	194(4)	245(7)	215(4)
$(N_0/N)_{\text{latt}}$	0.0023(4)	0.041(4)	0.25(2)	0.44(3)	0.43(3)

	$s = 12$	$t = 1.23 \times 10^{-2}$	$U = 0.332$	$\nu = 63.07$	
$(N_0/N)_0$	0.11(1)	0.22(1)	0.53(1)	0.82(1)	0.91(1)
$S (k_B/N)$	2.96(4)	2.54(4)	1.45(3)	0.55(2)	0.26(3)
$N (\times 10^3)$	258(4)	218(5)	183(4)	253(6)	220(4)
$(N_0/N)_{\text{latt}}$	0	0.0040(4)	0.110(4)	0.27(1)	0.28(2)

Table 7.1: Summary of the experimental lattice parameters, condensate fraction $((N_0/N)_0)$ and entropy in the harmonic trap, number, and measured peak fraction. The experimental parameters s , t , and U are given in E_R units, and ν is in Hz ($\omega = 2\pi \times \nu$) calculated from $\sqrt{f_0^2 + 8sE_R(m\mathbf{w}^2(2\pi)^2)^{-1}}$ where $\mathbf{w} = 120\mu\text{m}$ is the lattice beam waist and $f_0 = 35.76\text{Hz}$ is the bare trap frequency. The lattice depth s is measured experimentally using Kapitza-Dirac diffraction, and t and U are calculated from s .




Figure 7.10: High OD TOF images (averaged from 7 experimental runs), for each of the lattice depths studied versus different initial entropies in the harmonic trap. The top line (and the labeled entropies) correspond to the gas in the harmonic trap before loading into the lattice for the $s = 6$ data, however, the starting conditions were almost identical for all lattice depths. At the lowest entropies, the peaks shrink and the non-condensate grows due to quantum depletion as s increases. At even higher lattice depths than shown here, the transition to the MI will drive the condensate fraction to zero. The MI first appears at $s \approx 13$, but the SF does not disappear across the entire trap until $s \gtrsim 16$. At high entropy, the SF fraction decreases due to thermal fluctuations. As expected, the decrease occurs faster at higher lattice depths because the critical entropy (S_C) is reduced by interactions — $S_C = 0$ at the MI transition.




Figure 7.11: (a) Summary of the peak fraction results versus condensate fraction in the harmonic trap and corresponding entropy. The peak fraction is monotonically increasing with decreasing entropy (except at one point). Also, the peak fraction is always monotonically increasing as lattice depth decreases at all values of the entropy. (b) Comparing the peak fraction results at each lattice depth to theoretical predictions using the non-interacting (§5.1), HFBP (§5.3.1), and SDMFT (§5.3.2) theories. In general, there is poor agreement between the measured peak fraction and the predicted condensate fraction. This is discussed further in the main text.

designed to be exact in the $T \rightarrow 0$, $t/U \rightarrow 0$ and $T \rightarrow 0$, $U/t \rightarrow 0$ limits and it has no band structure. At intermediate depths ($s = 8$ and $s = 10$) the two interacting theories essentially agree, but deviate appreciably from the non-interacting theory. This discrepancy demonstrates the increasing role of interactions at higher s . The agreement between the two interacting theories provides confidence in their accuracy. At $s = 12$ (proximate to the MI transition), all three theories are fairly divergent. It is expected that the HFBP theory gives wrong results in this regime because it does not predict a MI transition. The SDMFT does predict the SF-to-MI transition, but it is unclear whether the entropy predicted by this theory can be trusted. Ultimately, the only calculation that is valid in all regimes is QMC. However, calculating entropy using QMC is challenging, and so calculations have only been performed for the 2D BH model [263]. Indeed, in a similar experiment (Ref. [67]), the lattice-gas entropy, used to calculate temperature by assuming adiabaticity, was also calculated using mean-field theory³ and not QMC. Overall, incorrect theory is not the probable source for the entire range of discrepancies we observe since these issues persist down to $s = 6$ where we have a high degree of confidence in the theoretical assumptions.

To address whether the peak fraction fit procedure can measure condensate fraction, we need to examine separately the low and high entropy regions. At low entropy (i.e., high condensate fraction), there are two main effects not accounted for by the peak fraction fit. The first is that atoms are scattered out of the condensate because of collisions between atoms in different diffraction peaks at the beginning of the expansion. The second is that condensate atoms are misidentified as non-condensate atoms. For example, at high condensate fraction, deviations of the condensate peak from a TF profile (i.e., smoothing at the edge of the condensate peak) is identified incorrectly as non-condensate and results in a measured peak fraction lower than the condensate fraction. Unfortunately, both these effects are hard to incorporate into a fitting procedure and will likely impede accurate thermometry at low entropies. However, they only cause a proportionality error and do not prevent the peak fraction from identifying that a condensate exists. The most significant issue is at high entropies where, if the fit is to blame, it is identifying a condensate when none exists. There are suggestions that such an effect may be due to peaks that exist in the momentum distribution of the normal phase of the BH model [60]. Still, visual analysis of Fig. 7.10 suggests that the peaks present, particularly the first-order diffraction peaks, are truly bimodal features that are distinct from the surrounding distribution and therefore associated with a condensate. To rule out an error such as this, we must apply our peak fraction fit procedure on simulated BH data (i.e., from QMC). For example, in the Ref. [67] experiment, simulated finite-temperature TOF distributions from QMC, taking into account finite expansion time and imaging resolution, were accurately compared to experimental distributions. QMC is not able to calculate the full interaction dynamics during the expansion, but Refs. [67, 261] concluded that TOF interactions are not important. The peak sizes and shapes in our experiment suggest otherwise (Fig. 7.10), however, our central filling is approximately 3–4 atoms per site, while the data in Ref. [67] is unity filling.

The final possibility is that the assumption of adiabaticity is not true. A breakdown of adiabaticity could play a role for the discrepancies we observe. For example, entropy generation from light scattering (§C.4) would result in a lower than expected condensate fraction, consistent with our data at low entropies. This was shown to be an important effect in the Ref. [67] experiment. Adiabaticity may also be violated if the gas is loaded too quickly into the lattice (i.e., fast compared to the timescales for condensate fraction evolution). At zero-temperature, the condensate fraction decreases for increasing lattice depths approaching

³The exact details of this calculation are not described in Ref. [67].




Figure 7.12: Peak fraction at $s = 6$ versus the ramp time into the lattice. Data are taken at three different values for the entropy per particle measured in the harmonic trap. Each point corresponds to a single high OD/low OD image pair. The peak fraction is calculated by measuring the peak number (i.e., the numerator of Eqn. 7.12) at a given ramp time and then dividing by the average number from all images at the same initial entropy (the same procedure used in §6.2). The vertical grey line indicates the ramp time used for the main data in this chapter. At longer ramp times the data continues to decrease monotonically to zero; the $S = 1.2k_B/N$ data goes to zero for a ramp time of 2000ms and the $S = 0.44k_B/N$ data goes to zero for a ramp time of 3000ms.

the MI transition. These changes are caused by coherent interaction effects (i.e. number squeezing), which have fast (1ms) timescales — this was the effect measured in §6.2. There are also slower timescales for global many-body adiabaticity (e.g., mass transport) at $T = 0$. An experimental study ([167]) concluded that ramp times of 100ms are sufficiently adiabatic. At finite-temperatures, condensate fraction also changes because of alterations to the density of states. For example, reversible condensation (in non-lattice experiments) has been observed by changing the density of states using a “dimple trap” [72, 264, 265]. These changes in condensate fraction require collisions out of the condensate [266], which occur at much slower timescales [265].

Timescales are one of the main differences between our experiment and the experiment in Ref. [67]. In our experiment, we load the atoms into the lattice in 100ms (motivated by our own measurements similar to Ref. [167]), but in Ref. [67] the loading time is 300ms. Assessing these loading rates is difficult because we lack both theoretical and experimental studies of finite-temperature dynamics in the lattice. As a preliminary investigation, we measure the peak fraction for $s = 6$ at different loading times for three different entropies (Fig. 7.12). At all entropies there is a general trend towards lower peak fraction (and eventually zero peak fraction) for longer ramp times. At high entropies this improves our agreement with theory, but at low entropies it increases our discrepancy. The complication is that the data in Fig. 7.12 combines two possible effects: heating and thermalization. Separating out these two effects is non-trivial since simple heating assumptions (i.e., light scattering) do not always quantitatively agree with experiment⁴ (e.g., in the Ref. [67] experiment), and the relation between light scattering and heating in a many-body system is not straightforward [37]. Still, if the timescales associated with heating and thermalization are sufficiently different, then theoretical input is not required because two separate timescales will be observed. Unfortunately, there is no strong evidence for a separation of timescales in the data. However, the data in Fig. 7.12 illustrates an important point — if we repeat our main experiment (i.e., the data in Fig. 7.11) with a different lattice loading time, then we will get a different result. Understanding this issue is a major open question from this work.

Therefore, future extensions to this work are twofold. First, our theory collaborators in the group of David Ceperley at the University of Illinois are in the process of calculating entropy curves from QMC and generating simulated data to test our peak fraction fit. The other task is to analyze timescales in the experiment. Part of this task is to experimentally investigate these timescales in more detail than the data presented in Fig. 7.12, but we need theoretical input. Most importantly, this will require developing theoretical techniques to study finite-temperature dynamics of the BH model.

⁴We can calculate the estimated heating rate from the basic light scattering theory (§C.4). For $s = 6$ and $\lambda = 812\text{nm}$, the rate of energy increase per particle is $33\text{nK} \times k_B/s$ (Eqn. C.67), corresponding to a scattering rate of 0.1s^{-1} . During the lattice loading ramp the average lattice intensity is a factor of two smaller than $s = 6$, so divide these numbers by two. If we assume a heat capacity of $3k_B$ (harmonically trapped thermal gas), then the HFPB theory (§5.3.1) predicts an approximate condensate fraction decay during the ramp of $0.1/s$. This is the right order of magnitude, however, we also measured the peak fraction versus load time at $\lambda = 833\text{nm}$ and observed essentially identical results, which supports our assertion that straightforward light scattering mechanisms are insufficient to explain these results. Uncontrolled heating poses a problem because the system does not satisfy ingredient 2 of quantum simulation (i.e., the initial state is not well-defined).

Chapter 8

Developing Impurity Probes

Direct thermometry probes of the lattice gas (discussed in §7) have a distinct disadvantage — they must invoke theory to relate measurable system properties (e.g., the 2D distribution after TOF expansion) to the thermodynamic state. This type of measurement is problematic when we simulate theoretically unsolvable models (i.e., the goal of quantum simulation). Instead, we can alternatively probe using a well-understood auxiliary system (an impurity) in contact with the lattice gas (the primary). This type of probe is advantageous when the impurity properties are related to the lattice gas properties in a straightforward manner. Such an approach is common, for example, in condensed matter physics; a resistor in thermal contact with the sample under study is often used to measure temperature. This is an optimal impurity probe — resistance is simple to measure and usually has a linear temperature dependence.

For quantum simulators using ultracold atoms, an impurity can be a small sample of atoms in another state or from a different species. This type of impurity probe was introduced in §4.1.7. Mutual interactions imprint the properties of the primary system onto the impurity. The specific type of impurity is selected so that its properties are easily measured using our standard set of probes (§4.1).

To create an impurity for our ^{87}Rb Bose-Hubbard simulator (§3.1), we use a sample of atoms in the $|F = 1, m_F = 0\rangle$ state; the technical details of transferring atoms to this state are discussed in §3.1.1. The primary system, which realizes the Bose-Hubbard model, is a degenerate gas of $|1, -1\rangle$ atoms in a lin- \perp -lin spin-dependent lattice at $\lambda = 790\text{nm}$ (§3.1.2, §C.2.3). The depth of the lin- \perp -lin lattice is proportional to $|m_F|$, and so the impurity experiences no lattice potential. The full Hamiltonian for our impurity/primary system is given by Eqn. 2.99. In developing these impurity probes we are exploring the solutions of this model.

In this chapter we focus on developing two specific impurity probes. First, we investigate using the impurity for thermometry (§8.1) by measuring the impurity temperature after it has been in contact with the lattice gas. Ideally, this type of thermometry would provide a theory-free temperature measurement of a Bose-Hubbard simulator. However, we observe a lack of thermalization between the two systems when the primary is loaded into the lattice, and so we mainly investigate thermalization timescales between the primary and impurity. For our second probe (§8.2), we measure the decay of the impurity center-of-mass motion (dipole mode) through the primary. This decay rate may probe the excitation spectrum of the lattice gas, however, further theoretical input is required.

8.1 Impurity Thermometer

In this work, we investigate using a harmonically trapped impurity to perform thermometry on a primary gas simulating the Bose-Hubbard model in a spin-dependent lattice. First, we experimentally confirm that the impurity and primary are overlapped using a combination of low resolution in-situ images and a novel probing technique (§4.1.6). Next, we measure how the impurity temperature changes as the primary system is loaded into the lattice. The impurity temperature is insensitive to the primary lattice depth, and so the two systems are not in thermal contact. To investigate thermalization in more detail, we perform two separate experiments using the lattice to input energy into the primary gas. We then measure the collisional energy exchange with the impurity. The results of these measurements are consistent with thermalization decreasing as the lattice depth increases. We compare these results to a Fermi’s golden rule calculation of the energy exchange rate, which implies that thermalization decreases because of dispersion mismatch between the free impurity and the lattice-bound primary. These results significantly constrain the applicability of an impurity thermometer.

Thermometry is a natural application for an impurity probe since two interacting system will eventually come into thermal equilibrium (irrespective of the interaction details). If we measure the impurity temperature then we know the temperature of the primary without theoretical input, which is desirable for quantum simulation. Of course, this is only beneficial if the impurity has an easy to measure, well-known property that depends on temperature and if the presence of the impurity does not affect the primary system. The canonical example of an impurity thermometer is the standard glass thermometer. In this case, the impurity is an extremely small perturbation to the primary system — the air. A suitable liquid is used, e.g., mercury or alcohol, which has a linear density dependence with temperature.

In this work, we investigate an impurity thermometer based on a small sample of $|1, 0\rangle$ atoms that do not experience a lattice potential. The impurity is co-trapped with a primary system of atoms in the $|1, -1\rangle$ state. The $|1, -1\rangle$ atoms experience a lattice potential and are well-described by the Bose-Hubbard model. The impurity is a weakly interacting Bose gas (§2.2.2), and its temperature can be reliably measured from the momentum distribution (§4.1.2) or condensate fraction (§4.1.4) after TOF. The question of whether the impurity will affect the primary system is difficult to answer. In Ref. [4], we estimate the maximum impurity size from SDMFT heat capacity calculations of the lattice gas. In the deep lattice limit, we estimate an impurity-size upper bound of 500 atoms, which may cause issues with signal-to-noise ratio. Experimentally investigating this effect is left as a future extension to this work.

The more immediate issue for thermometry is whether the system reaches thermal equilibrium during typical experimental timescales. While interactions are necessary for thermalization, they are not sufficient and, in fact, strong interactions could be a hindrance. Thermalization requires interactions which lead to energy exchanging collisions between the impurity and the primary. These collisions may be inhibited at the 2-body level, e.g., if the scattering length is zero, or at the many-body level, e.g., by a gap in the excitation spectrum. The question of thermalization in the impurity/primary system considered here has not been addressed either experimentally or theoretically, so the main body of experimental work presented in this section focuses on this topic.

8.1.1 Overlap

For a system with local interactions, a necessary condition for thermalization is that the impurity and primary must be physically overlapped in the trap. Although the two gases share the same trap minima, they may be prone to phase separation. The gases phase separate when one component increases its potential and/or kinetic energy to minimize the global interaction energy (e.g., by moving to the edge of the trap). In the simplest case (two pure condensates in a harmonic trap) the stability condition is [125]

$$a_{12}^2 < a_{11}a_{22} \quad (8.1)$$

where a_{ij} is the scattering length between states i and j (a_{ii} is the scattering length between atoms in the same state). For the states used in this work, $|1, 0\rangle$ (impurity, state #1) and $|1, -1\rangle$ (primary, state #2), the scattering lengths are $a_{12} = 100.4a_0$, $a_{11} = 100.9a_0$ and $a_{22} = 100.4a_0$ (Table A.3), so Eqn. 8.1 predicts a stable mixture. Therefore, the impurity and primary are overlapped at the start of the experiment before the lattice is applied. However, there will be a new condition when the primary is loaded into the lattice. In general, to determine this condition we need to solve the Hamiltonian given by Eqn. 2.99. A zero-order approximation is to assume the lattice gas is in the atomic limit (§5.2) and the impurity is described by a condensate wavefunction in the TF approximation $\Phi(x)$, where the number of condensate particles in a d^3 volume centered on lattice site i is

$$\tilde{N}_i = d^3 |\Phi(x_i)|^2. \quad (8.2)$$

The approximate form of Eqn. 2.99 is

$$\mathbf{H} = \sum_i \mathbf{H}_i, \quad (8.3)$$

$$\mathbf{H}_i = n_i \left[\frac{U}{2}(n_i - 1) + \frac{1}{2}m\omega_2^2 x_i^2 - \mu_2 \right] + \tilde{N}_i \left[\frac{1}{2}m\omega_1^2 x_i^2 - \mu_1 + \frac{1}{2}g_{11} \frac{\tilde{N}_i}{d^3} \right] + g_{12}n_i \frac{\tilde{N}_i}{d^3}. \quad (8.4)$$

From the general condition for mixture stability given in [125], we can determine that stability in this case requires

$$a_{12}^2 < a_{11} \frac{\pi d}{8} \frac{U}{E_R}. \quad (8.5)$$

Since $U/E_R \sim 0.1$, then $\frac{\pi d}{8} \frac{U}{E_R} \sim 15\text{nm}$ for a 400nm lattice spacing, and the stability condition given by Eqn. 8.5 is easily satisfied. However, it is still important to experimentally test the overlap. Also, in any two-component experiment where the states have different magnetic moments, stray fields could lead to separation.

Experimentally we investigate phase separation in two ways. In the first method, we take an image of the atoms 200 μs after releasing from the x-dipole trap (essentially an in-trap image). We can take an image immediately after turning off the trap because there is a bias field on during the experiment, and so we do not have to wait for the imaging bias field to turn on. If we have a mixture of the states $|1, 0\rangle$ and $|1, -1\rangle$ in the trap, we can selectively image either state by adjusting the repump procedure. To image $|1, -1\rangle$ we use an extremely short off-resonance repump pulse (§F.2). Since the repump has a stronger effect on $|1, -1\rangle$ than $|1, 0\rangle$ and there are more $|1, -1\rangle$, the resulting image is predominantly of the $|1, -1\rangle$ atoms. Alternately, to image just $|1, 0\rangle$, we do a microwave sweep to transfer from $|1, 0\rangle \rightarrow |2, 0\rangle$ and do not pulse




Figure 8.1: In-situ separation between the $|1,0\rangle$ and $|1,-1\rangle$ gas centers (each fit to a Gaussian) in the yz plane. For the mixture (blue circles), the two gases are together in the same trap. The ratio of $|1,0\rangle$ to $|1,-1\rangle$ atoms is approximately 1:3. For the “not mixed” data (red squares), the images are fit in two subsequent experiments with all the atoms in either $|1,0\rangle$ or $|1,-1\rangle$. Each data point is the average of four runs. For reference, the FWHM of the gas is approximately $15\text{--}20\mu\text{m}$.

the repump beam. These images are low resolution and can only determine if there is macroscopic phase separation. We fit each gas to a Gaussian to determine the center, and in Fig. 8.1 we plot the difference between the impurity and primary centers as a function of the lattice depth. There are no obvious signs of phase separation and the separation between the components is less than the gas diameter.

Another possibility is that the impurity and primary do not overlap on microscopic length scales. Since the primary density is concentrated in a narrow region around each lattice site, visualized in Fig. 2.10, one might imagine that the impurity atoms become density modulated to avoid these regions. This is not likely the case since we do not observe diffraction peaks in TOF images of the impurity. To rule this out more definitively, we use the density probe described in §4.1.6. In this probe we use microwaves to transfer the impurity atoms to the $|2,1\rangle$ state, which experiences the same lattice potential as the primary in $|1,-1\rangle$ ¹. Since the Rabi rate depends on the impurity wavefunction overlap with the lattice wavefunction, this probes the impurity at lattice length scales. Results for $s = 12$ are shown in Fig. 8.2, where we transfer $|1,0\rangle$ atoms to $|2,1\rangle$ with and without $|1,-1\rangle$ atoms present. If there is significant phase separation on microscopic length scales we expect to see a strong decrease in the transfer rate when the primary ($|1,-1\rangle$) atoms are also in the trap. Since this is not the case, there is no strong evidence for microscopic phase separation.

8.1.2 Thermometry Test

Since we have confirmed that the impurity and primary are overlapped, the benchmark test of our proposed impurity thermometer is to load the primary gas into the lattice and observe whether the impurity temperature follows changes in the primary temperature. There is some difficulty associated with interpreting the

¹ $|2,1\rangle$ and $|1,-1\rangle$ have the same magnetic moment.




Figure 8.2: Using the microwave density probe (§4.1.6) to measure the impurity density in a $s = 12$ lattice with and without $|1, -1\rangle$ atoms present. In both experiments there are approximately 1.5×10^4 $|1, 0\rangle$ atoms. In the mixture there are $5.3(6) \times 10^4$ $|1, -1\rangle$ atoms. We transfer atoms from $|1, 0\rangle$ to $|2, 1\rangle$ using a $80\mu s$ microwave pulse at various microwave frequencies and plot the ratio transferred. For reference, the free-space transition frequency (no lattice) is 6856.20MHz. We fit both curves to a Lorentzian. For the no mixture data, the peak is 0.32(5), the FWHM is 12(3)kHz, and the center frequency is -47.9(6)kHz (referenced to 6856.18MHz). For the mixture data, the peak is 0.28(4), the FWHM is 11(3)kHz, and the center frequency is -47.2(6)kHz. The data are consistent with no effect from the presence of the $|1, -1\rangle$ atoms.




Figure 8.3: (a) Condensate fraction versus primary lattice depth s for an impurity ($|1, 0\rangle$) and primary ($|1, -1\rangle$) mixture. Each point is an average over several experimental runs. For $s=2-10$ (circle and square points), the lattice was snapped off and condensate fraction was fit using the procedure outlined in §7.3.1. For $s = 12$ and $s = 16$ (triangle and diamond points) the lattice was bandmapped and the condensate fraction was fit using the procedure from §4.1.4. Bandmapping is known to bias condensate fraction to slightly higher values (§6.2), so these points are an upper bound. The lattice turn off does not affect the impurity data, which was fit using the standard harmonic trap procedure in §4.1.4. The lattice gas (primary) condensate fraction displays the characteristic decrease as the lattice depth increases towards the MI transition ($s = 12.9$). However, the condensate fraction of the impurity is remarkably flat, which is a sign that there is no inter-state thermalization. (b) Number versus lattice depth for the impurity and primary. The number in the primary decreases towards larger lattice depths, which may be an indication of spontaneous emission from the lattice beams or a bias in the fitting procedure. (c) Temperature versus lattice depth calculated using the data in (a) and (b). For the impurity, the temperature is calculated using the semi-ideal model (§2.2.2) and agrees well with the temperature obtained by fitting the high momentum tails of the data from (a). For the primary, the temperature is calculated using the SDMFT (§5.3.2). The geometric mean of the x-dipole trap used in this data is $\bar{\nu}_0 = 44.5\text{Hz}$. The primary trap frequency increases as $\omega = (\omega_0^2 + 4sE_R/m\mathfrak{w}^2)^{0.5}$ where $\mathfrak{w} = 120\mu\text{m}$ is the heuristic lattice beam waist (see §D).

results since we lack an adequate independent method to measure the primary temperature. Experimentally, we start with a mixture of $|1, 0\rangle$ impurity atoms and $|1, -1\rangle$ primary atoms. The primary atoms are then loaded into the $\lambda = 790\text{nm}$ spin-dependent lattice over 50ms. This loading time is half our loading time for the $\lambda = 812\text{nm}$ spin-independent lattice (e.g., §7.3) because the heating rate is higher in the spin-dependent lattice (the lattice frequency is detuned closer to resonance). A ramp time of 50ms still allows us to load a condensate into the lattice to the MI regime ($s \sim 16$) and then ramp the lattice off and recover a non-zero condensate fraction. After loading into the lattice, it is turned off suddenly (“snapped off”), releasing the impurity and primary which are spatially separated during expansion using a magnetic field gradient.

The impurity condensate fraction is fit using the well-established technique for a harmonically trapped gas discussed in §4.1.4. There is no standard technique for fitting the primary (lattice) condensate fraction, so we use our procedure developed in this thesis and outlined in §7.3.1. The experimental condensate fractions are shown in (a) of Fig. 8.3. The primary condensate fraction displays the characteristic monotonic decrease with lattice depth and eventually vanishes at $s = 16$. For this lattice, the MI appears at $s = 12.9$ according to a mean-field calculation. On the other hand, the measured impurity condensate fraction is essentially insensitive to the lattice depth with a mean value of 0.47(1).

The more relevant quantity to compare is temperature, although this is difficult because our ability to convert lattice condensate fraction into temperature is inaccurate (§7.3.4). Indeed, our inability to measure temperature from condensate fraction (§7.3) is one of the main motivations for an impurity thermometer. However, the lattice temperature calculated using SDMFT (§5.3.2) should give an estimate for the primary temperature. The impurity temperature is calculated from condensate fraction using the semi-ideal model (§2.2.2) and has a mean value of 44(1)nK. The temperatures are shown in (b) of Fig. 8.3. The primary temperature falls as the lattice depth increases since the tunneling energy scale decreases, however, the impurity temperature clearly does not follow. This is strong evidence for lack of thermalization between these two systems.

When there is slow thermalization the typical solution is to hold the system in contact for a longer duration. However, there is another competing timescale in the problem since the lattice is continually heating both the impurity and primary equally due to spontaneous emission (see §C.4.2). Therefore, after enough time the two gases will come into steady-state with this heat source, and the impurity temperature will not reflect the initial temperature of the primary. To overcome this problem, we require an experiment that can measure thermalization on short timescales, i.e., we need a large temperature gradient. This is not possible by adiabatically loading the primary into the lattice, since we are limited by adiabatic cooling to $\Delta T = T_{\text{Impurity}}$. Also, a measurement technique in which the primary cools the impurity is not optimal because this effect is canceled out by the heating from the lattice — it is difficult to discern whether a null result is due to this cancellation or simply a lack of thermalization. Instead, we need to create a large positive temperature difference between the two systems by selectively heating the primary using the lattice. A schematic of this process is illustrated in Fig. 8.4 and two experimental realizations will be discussed in the following sections.




Figure 8.4: Using selective heating to measure thermalization. (a) Heat is applied to the primary (blue), which starts at a similar temperature to the impurity (red). Both the impurity and primary experience equal heating due to spontaneous emission from the lattice (§C.4.2). This fundamental lattice heat load (Q_{latt}) is smaller than the applied heat load (Q). (b) After some time, the applied heat source is removed and the primary is at a higher temperature T'_p . Heat flows into the impurity, which can be measured. We can measure the temperature increase of the impurity versus time or wait for a constant time after removing the heat and determine the effect of changing Q .

8.1.3 Parametric Heating

To implement the type of heating experiment outlined in Fig. 8.4, we need a method to selectively heat the primary (lattice) gas. Since the impurity is unaffected by the lattice, a time-dependent lattice sequence is an obvious choice for selective heating. In this experiment we use parametric heating, which involves modulating the lattice depth at the frequency corresponding to the energy difference between the ground and second excited band ($\omega_{2,0} = E_{2,0}/\hbar$). Atoms are excited to this band and then decay, thereby releasing energy.

For parametric oscillation, the single-particle Hamiltonian is

$$\mathbf{H} = \frac{\mathbf{p}^2}{2m} + E_R [s + \Delta_s \sin(\omega t)] \cos^2(k\mathbf{x}). \quad (8.6)$$

When $\Delta_s = 0$, this is the lattice Hamiltonian which gives rise to the band structure as described in §2.3. Rewriting Eqn. 8.6 in terms of the Bloch waves of quasimomentum q in 1D

$$\begin{aligned} \mathbf{H} = & \sum_n \int dq \, 2t \left[1 - \cos\left(\pi \frac{q}{q_B}\right) \right] |q_n\rangle \langle q_n| + \\ & E_R \Delta_s \sin(\omega t) \sum_{n,n'} \int \int dq \, dq' \langle q_n | \cos^2(k\mathbf{x}) | q'_{n'} \rangle |q_n\rangle \langle q'_{n'}| \end{aligned} \quad (8.7)$$

where n is the band index. We can rewrite the Bloch waves in terms of plane waves (Eqn. 2.64)

$$|q_n\rangle = e^{iqx/\hbar} \sum_{j=-\infty}^{\infty} b_{j;q,n} e^{i(2kj)x}, \quad (8.8)$$

so the off-diagonal matrix element for the excitation is

$$2\langle q_n | \cos^2(k\mathbf{x}) | q'_{n'} \rangle = \langle q_n | \overset{0}{\cancel{1}} | q'_{n'} \rangle + \langle q_n | \cos(2k\mathbf{x}) | q'_{n'} \rangle, \quad (8.9)$$

$$= \frac{1}{2} \langle q_n | e^{i2k\mathbf{x}} + e^{-i2k\mathbf{x}} | q'_{n'} \rangle, \quad (8.10)$$

$$= \frac{1}{2} \sum_{j,j'} \int dx e^{i(q-q')x/\hbar} b_{j;q,n}^* b_{j';q',n'} \left[e^{i(2k)(j-j'-1)x} + e^{i(2k)(j-j'+1)x} \right], \quad (8.11)$$

$$= \frac{\delta(q-q')}{2} \sum_j \left(b_{j;q,n}^* b_{j+1;q',n'} + b_{j;q,n}^* b_{j-1;q',n'} \right). \quad (8.12)$$

Therefore, the excitation couples states with the same quasimomentum and parity. If we start in a single q state of the ground band and oscillate close to $\omega_{2,0}$ (the bandgap to $n = 2$), then the problem is just a two-level Rabi oscillation (§B.4). In the $s \rightarrow \infty$ limit the resonance frequency is $2\hbar\omega$ (where ω is the well frequency), and this reduces to parametric oscillation of a harmonic oscillator [77]. The single particle picture is not complete because the atoms can decay from the excited band back to the ground band due to interactions and inhomogeneities [79]. When an atom decays from a higher band, the excitation energy is released into the system.

A schematic of the parametric oscillation experiment is shown in Fig. 8.5. First we create an impurity/primary mixture² and load the primary into the lattice over 50ms (as described in §8.1.2). Once the atoms are in the lattice, we modulate one of the beams at the parametric resonance frequency for 2ms. The resonance frequency is determined by a band structure calculation and confirmed at $s = 4$ as shown in Fig. 8.5. After holding for 50ms, we release the atoms and measure the condensate fraction of the impurity as a function of the strength of the parametric oscillation Δ_s , which changes the heat input into the primary (Q in Fig. 8.4). The time in each experiment is fixed, so the heating from light scattering is held constant. In Fig. 8.5 we show the condensate fraction of the impurity as a function of Δ_s at $s = 4$, with and without the primary present. In contrast to the loading data, energy is clearly exchanged from the primary to the impurity. For large Δ_s , the effect saturates, so for quantitative measurements we limit the oscillation strength to the linear region determined by this figure ($\Delta_s \leq 0.5E_R$). We measure for primary lattice depths of $s = 2, 4, 6, 8$, and 11 using oscillation frequencies of 15.5, 18, 20.8, 24.5, and 31kHz respectively.

The full data for the impurity condensate fraction versus modulation depth is shown in (a) of Fig. 8.6 for $s = 4$ and $s = 11$. The decreasing condensate fraction is heuristically fit to a decaying exponential, $Ae^{-\Delta_s/\bar{\Delta}_s}$. The fit parameter $\bar{\Delta}_s$ for each lattice depth is shown in (b) of Fig. 8.6. Higher values of $\bar{\Delta}_s$ indicate a decreasing response of the impurity to the parametric drive. Therefore, the data trends towards less heat transfer to the impurity at higher lattice depths.

These results suggest less thermalization as the lattice depth increases. However, there is considerable uncertainty in the parametric heating process of the $|1, -1\rangle$ primary (lattice) gas. Sample bandmapping images for the primary are shown in (a) of Fig. 8.6. At $s = 4$ the lattice gas is clearly heated by the parametric oscillation as evidenced by the change in condensate fraction between $\Delta_s = 0$ and $\Delta_s = 0.5$. However, heating in the $s = 11$ data is unclear because the band is essentially full even without the parametric oscillation.

²There are $70(5) \times 10^3$ atoms in $|1, 0\rangle$ and $97(7) \times 10^3$ atoms in $|1, -1\rangle$.




Figure 8.5: (a) Lattice depth versus time for the parametric oscillation. The inset is zoomed into the oscillation. After the parametric oscillation, the atoms are held for 50ms to allow energy to transfer from the primary to the impurity. (b) Testing the modulation frequency with only $|1, -1\rangle$. First the lattice is ramped on, then the depth is modulated, then the lattice is ramped off, and the $|1, -1\rangle$ condensate fraction is measured. These data (at $s = 4$) show that the oscillation maximally heats at the predicted resonance of 18kHz. (c) The impurity condensate fraction versus modulation depth at $s = 4$. In the control experiment (red squares) there are only impurity ($|1, 0\rangle$) atoms present. In a 40:60 mixture (blue circles), the modulation heats the primary and this energy is transferred to the impurity, causing the condensate fraction to decrease.




Figure 8.6: (a) Impurity condensate fraction versus parametric oscillation modulation depth Δ_s at $s = 4$ (blue circles) and $s = 11$ (red squares). The data are fit to a heuristic exponential decay $N_0/N(\Delta_s) = Ae^{-\Delta_s/\bar{\Delta}_s}$ (solid curves). To visualize the effect of the parametric oscillation on the primary gas, we show images of the $|1, -1\rangle$ primary atoms (after 0.5ms bandmapping) for $\Delta_s = 0$ and $\Delta_s = 0.5$ (at bottom). For viewability, the two lattice depths have different color scalings. (b) The fit parameter $\bar{\Delta}_s$ versus lattice depth obtained by fitting data similar to (a). This parameter characterizes the effect of the modulation on the impurity condensate fraction.

Therefore, energy can only go into trap degrees of freedom, which may not contribute to thermalization. The excitation process is difficult to address theoretically since we need to understand the band decay rates and how the energy is thermalized within the band; both topics are current areas of study. Ultimately, we need an excitation process which is better understood. We address this issue in the next section.

8.1.4 Dephasing




Figure 8.7: (a) Lattice ramp sequence for the dephasing experiment shown for $s = 4$. To dephase the primary atoms the lattice is ramped up to $20E_R$ and back three times in a row. The ramp is 0.1ms to prevent band transitions. After the dephasing step, we hold the impurity and primary for τ_{hold} . (b) Primary ($|1, -1\rangle$) atoms with and without the dephasing step after bandmapping for different lattice depths $s = 4-12$ ($\tau_{\text{hold}} = 0$). For viewability the two sets of images are scaled differently. Each image is a 10-shot average.

To make a quantitative comparison with theory, we need to understand the temperature gradient between the primary and impurity generated by the excitation process. Unfortunately, no reliable thermometry methods are available for the lattice gas. However, another option is to heat the primary in such a way that effectively $T \rightarrow \infty$. In the lattice gas this is a physical limit because the lattice kinetic energy is bounded. In a uniform single-band lattice, this limit corresponds to a homogeneously filled band; the maximum kinetic energy is $12t$. In the trapped gas, this limit is not physical because the trap degrees of freedom are not bounded, so the gas would become infinitely large and dilute. Nevertheless, we can define a temperature T_q associated with the quasimomentum DOFs, where $T_q \rightarrow \infty$ corresponds to the filled band. In the experimental lattice we also have to contend with higher bands, so in reality the $T_q \rightarrow \infty$ limit is actually $k_B T_q \gg t$, $k_B T_q < E_{\text{band}}$.

Several methods exist to fill the band by dephasing individual lattice sites. Motivated by [241], we ramp our lattice three times from the starting lattice depth s to $20E_R$. The gas at $20E_R$ is strongly out of equilibrium and a combination of interactions and the harmonic trap cause different phase windings at each site. This procedure is shown schematically in (a) of Fig. 8.7. The change to the primary ($|1, -1\rangle$) lattice gas due to the dephasing step is shown using bandmapped images in (b) of Fig. 8.7. Without the dephasing, the gas is condensed at all lattice depths and the cloud is very dense near low quasimomentum (see change

in scale). With the dephasing, the quasimomentum distribution is nearly uniform and similar at each depth.




Figure 8.8: Impurity ($|1,0\rangle$) dephasing data for $s = 4$. (a) 10-shot average of temperature data versus hold time (mean errors are shown). (b) Raw temperature data before averaging. Heating rate is determined by fitting the data to a line (solid curve). (c) Temperature is not fit directly from the raw images. We fit the measured distribution to a TF plus Bose-thermal momentum distribution (Eqn. 4.21) to measure the condensate fraction and number shown in this plot. Temperature is then calculated from these using the semi-ideal model discussed in §2.2.2 using a trap frequency of $\nu = 44.5\text{Hz}$.

Once the dephasing step is performed, we hold the impurity/primary mixture in the lattice for t_{hold} and then release. We measure the impurity condensate fraction and number from TOF images, repeating the experiment 10 times at each hold time. Temperature is calculated from condensate fraction and number. We fit T versus t_{hold} to a line to determine the heating rate. The hold time is short so that the temperature increase is linear and the filled band approximation remains valid. Typical data for $s = 4$ are shown in Fig. 8.8. Two sets of control data are also taken to rule out heating from other sources. In the first control dataset we load the mixture into the lattice, but do not dephase the primary. In the other control we transfer all the atoms to the impurity ($|1,0\rangle$), keeping the impurity number the same as when there is a mixture. Both control tests are consistent with zero heating (see (b),(c) of Fig. 8.9). The null result of the second test is surprising since the heating rate from the lattice itself should be on the order of $0.1\text{--}1\text{nK/ms}$. Likely, atoms that spontaneously scatter off the lattice beams are lost since the impurity atoms are trapped in a shallow trap. This is consistent with the atom loss rate observed. If this is the case, then energy from the

lattice beams is not thermalized. This is a topic that warrants future study.




Figure 8.9: (a) Heating rate of the impurity after dephasing the primary as a function of the primary lattice depth. The heating rate decreases as the lattice depth increases and is essentially zero at $s = 12$. (b),(c) Control data for the heating rate. (b) Impurity heating rate when we load the mixture into the lattice without the dephasing pulses outlined in Fig. 8.7. (c) Impurity heating rate with the lattice on, but without primary atoms. The impurity number is similar between (b) and (c). (d) Impurity ($|1,0\rangle$) and primary ($|1,-1\rangle$) number for the dephasing experiment. Approximately 25% of the atoms are in the impurity. To measure the primary number we fit the dephased images to a Gaussian.

The impurity heating rate versus lattice depth, when the primary is dephased, is shown in (a) of Fig. 8.9. The trend of the data is consistent with the loading (§8.1.2) and parametric heating data (§8.1.3) — all experiments imply that at higher lattice depths energy exchange (thermalization) between the impurity and primary slows down. However, the data still does not constitute hard proof because of uncertainties in the dephasing process. Questions remain about how the lattice gas thermalizes after the dephasing procedure and whether the $T_q \rightarrow \infty$ assumption is valid at higher lattice depths where interactions start to fill the band at $T = 0$. In the next section (§8.1.5), we compare the experimental results to a zero-order thermalization theory. These calculations corroborate the observed trend and identify that mismatch between the dispersion curves can prevent thermalization. The mismatch between the lattice gas dispersion (Eqn. 2.81) and the free condensate Bogoliubov dispersion suppresses energy exchanging collisions at higher

lattice depths. Therefore, using a harmonically trapped impurity for thermometry of the strongly correlated lattice is not feasible. This result also impacts proposed lattice gas cooling schemes that require heat to be transferred to a bath of harmonically trapped atoms (e.g., [69]).

8.1.5 Thermalization Theory

The full calculation for the thermalization between the harmonically trapped impurity and the Bose-Hubbard lattice gas requires a quantum Boltzmann approach (see, e.g., [124]), which is beyond the scope of this work. Instead, we pursue a Fermi's golden rule calculation which models weak collisions between a free impurity condensate with $T_{\text{imp}} = 0$ and a non-interacting lattice gas with $T_{\text{latt}} = \infty$. Although straightforward, this calculation includes an important effect that may be the basis of our experimental result — these two systems have functionally distinct dispersion curves whereas energy exchanging collisions must conserve total momentum and energy. This condition limits thermalization between our impurity and the lattice gas.

To start, we consider the collision between a single lattice atom (in volume V) with quasimomentum \vec{q}_1 and a free condensate of N_{imp} atoms at rest. The collision scatters the lattice particle into quasimomentum \vec{q}_2 and creates a Bogoliubov excitation in the condensate with momentum \vec{p} . From Fermi's golden rule, the rate for this process is

$$\Gamma_{|\vec{q}_1, 0\rangle \rightarrow |\vec{q}_2, \vec{p}\rangle} = \frac{2\pi}{\hbar} |\langle \vec{q}_2, \vec{p} | \mathbf{V}_{\text{int}} | \vec{q}_1, 0 \rangle|^2 \delta(E_{\vec{q}_1} - E_{\vec{q}_2} - \epsilon_{\vec{p}}) \quad (8.13)$$

where $E_{\vec{q}_1}$ is the energy of a lattice particle with quasimomentum \vec{q}_1 (Eqn. 2.81), $\epsilon_{\vec{p}}$ is the energy of a Bogoliubov excitation³ with momentum \vec{p} (see e.g., [125]), and \mathbf{V}_{int} is the contact interaction given by Eqn. 2.16 summing over all condensate particles. The delta function enforces energy conservation in the collision process. The matrix element is⁴

$$|\langle \vec{q}_2, \vec{p} | \mathbf{V}_{\text{int}} | \vec{q}_1, 0 \rangle|^2 = N_{\text{imp}} \left(\frac{4\pi a \hbar^2}{mV} \right)^2 \frac{p^2}{2m\epsilon(\vec{p})} \delta_{\vec{p} + \vec{q}_2, \vec{q}_1} \quad (8.14)$$

where the Kronecker delta imposes momentum conservation.

For thermalization, we are interested in the rate of energy transfer to the impurity due to all possible collisions. Therefore, we need to multiply the collision rate by $\epsilon(\vec{p})$ and integrate over all possible final states \vec{q}_2, \vec{p} , average over initial states \vec{q}_1 , and multiply by the number of lattice atoms N_{latt} . Because $T_{\text{latt}} = \infty$, the initial \vec{q}_1 states are all equally occupied. Therefore, the total energy transfer to the impurity (per unit volume) is

$$\begin{aligned} \frac{\partial \Delta \tilde{E}}{\partial t} &= \frac{2\pi}{\hbar} \left(\frac{4\pi a \hbar^2}{m} \right)^2 \frac{N_{\text{imp}}}{V} \frac{N_{\text{latt}}}{V} \int \frac{d\vec{q}_1}{\int_{BZ} d\vec{q}} \int \frac{d\vec{p}}{(2\pi\hbar)^3} \frac{p^2}{2m\epsilon(\vec{p})} \epsilon(\vec{p}) \\ &\quad \int d\vec{q}_2 \delta_{\vec{p} + \vec{q}_2, \vec{q}_1} \delta(E_{\vec{q}_1} - E_{\vec{q}_2} - \epsilon_{\vec{p}}), \end{aligned} \quad (8.15)$$

$$= 2 \frac{a^2}{m^3} n_{\text{imp}} \frac{\tilde{N}}{d^3} \int \frac{d\vec{q}}{\int_{BZ} d\vec{q}} \int d\vec{p} p^2 \delta(E_{\vec{q}} - E_{\vec{q}-\vec{p}} - \epsilon_{\vec{p}}) \quad (8.16)$$

³At low momentum $\epsilon_{\vec{p}} = c\vec{p}$ where c is the speed of sound and at high momentum $\epsilon_{\vec{p}} = p^2/2m$ (free particle dispersion).

⁴The matrix element is related to the dynamic structure factor, see e.g., [267, 268].

where the quasimomentum index is dropped in the final expression, n_{imp} is the condensate density, and \tilde{N} is the number of lattice atoms per site⁵. For the trapped gas, we calculate the thermalization using LDA and divide by the total number of condensate particles to get the energy transfer rate per particle

$$\frac{\partial \Delta E}{\partial \tau} = 2 \frac{a^2}{m^3} \frac{1}{\int d\vec{r} n_{\text{imp}}(\vec{r})} \int d\vec{r} n_{\text{imp}}(\vec{r}) \frac{\tilde{N}(\vec{r})}{d^3} \int \frac{d\vec{q}}{\int_{BZ} d\vec{q}} \int d\vec{p} p^2 \delta(E_{\vec{q}} - E_{\vec{q}-\vec{p}} - \epsilon_{\vec{p},\vec{r}}). \quad (8.17)$$




Figure 8.10: The normalized energy transfer rate (Eqn. 8.18) compared to the experimentally measured heating rate. A similar trend is observed — thermalization decreases as the lattice gas increases. The black, dashed curve is calculated assuming energy conservation between the Bogoliubov impurity condensate and a non-interacting lattice gas. The red, solid curve is calculated by replacing the energy conserving delta function in Eqn. 8.17 with a Gaussian of width U . This corresponds to assuming energy conservation up to the neglecting interaction energy in the lattice.

We numerically calculate Eqn. 8.17 for the experimental parameters in §8.1.4 and approximate the energy conserving delta function as a Gaussian. In Fig. 8.10 we plot the normalized energy transfer rate (c_0 is the speed of sound at the center of the condensate)

$$\frac{d^3}{2a^2mc_0^3} \frac{\partial \Delta E}{\partial \tau}, \quad (8.18)$$

compared to the experimental data. We do not make a direct quantitative comparison to the data, but instead analyze the trend as the lattice depth increases. Similar to the experimental data, the transfer of energy decreases as the lattice depth decreases — collisions are suppressed as the dispersion curves start to diverge. Two theoretical curves are shown corresponding to different approximations for the width of the

⁵ \tilde{N} is the number of lattice atoms per site averaged over the volume V , i.e., $\tilde{N} = N_{\text{latt}}/(V/d^3)$. In the LDA, $\tilde{N}(r)$ is a continuous variable that smoothly interpolates the discrete variable N_i (the number of lattice atoms in site i).

Gaussian delta function. For the bottom curve, the width is extrapolated to zero (i.e., the true delta function limit). This predicts a faster than observed decrease in thermalization. However, energy conservation should be influenced by interactions in the lattice gas. In the top curve, the Gaussian width is given by the Hubbard interaction energy U . This better follows the experimental trend and may indicate that lattice interactions are modifying the energy spectrum at higher depths.

There are several possible extensions to this calculation that could be considered in the future. One of these extensions is to recalculate for finite temperatures of the impurity condensate. Although we assume that $T_{\text{imp}} = 0$, experimentally the initial impurity condensate fraction is approximately 0.5. We also assume that the density of the impurity and the lattice gas is fixed. However, the density will change as thermalization proceeds. Also, the lattice density is slightly time-dependent due to the dephasing procedure; these changes were neglected in the calculation. Ultimately, it will be important to calculate the dynamics of the dephasing process to confirm that the lattice gas can be approximated as having a homogeneously filled band. Although our original goal of thermometry was not realized, the experiment stimulated understanding thermalization in this system and demonstrates the possibility to study out-of-equilibrium physics in the spin-dependent lattice.

8.2 Impurity Transport Probe

In this work, we measure the decay rate of the center-of-mass (dipole mode) of the harmonically trapped impurity co-trapped with a static Bose-Hubbard lattice gas (primary). To selectively apply an impulse to the impurity we use a combination of magnetic and optical forces, which cancel for the primary $(1, -1)$. After the impulse we wait a variable amount of time in the trap and then image the position after TOF. The decay is measured by fitting the motion to a damped simple harmonic oscillator. We find that the decay rate drops as the lattice depth for the primary atoms increases. This data demonstrates another method type of impurity probe, although theory is required to definitely relate the decay rate to the lattice properties. We conjecture that this probe could be used to measure the excitation spectrum of the lattice gas.

There are many possible applications of an impurity probe beyond thermometry. For example, if impurity/primary interactions are strongly repulsive then impurities can be used to probe defects and vortex cores. Alternatively, if impurity/primary interactions are strongly attractive then impurities can probe the primary flow patterns. If interactions are sufficiently weak, so that the impurity can penetrate into the primary, impurities can be used as a scattering probe. The momentum distribution of the scattered impurities is directly related to the static structure factor. This has been investigated with ultracold atoms in 1D [82]. Here we similarly send impurity atoms through the primary, except we measure the decay rate of this transport (mass current). The advantage of this method is that the signal is unambiguous and the initial conditions are well-controlled compared to the experiments in §8.1. However, relating the decay rate to properties of the primary system requires theoretical input. Nevertheless, the decay rate can probe a number of interesting properties. For example, in a classical gas this decay rate can be used to measure the scattering length between different atomic species, such as ^{41}K and ^{87}Rb [84], and the damping rate can be directly related to the thermalization rate [269]. In addition, the damping rate is a sensitive probe of Fermi statistics [270], integrability in 1D systems [271], spin drag at unitarity [272], superfluidity [273], Gross-Pitaevskii physics [85, 274] and the structure factor [83].




Figure 8.11: (a) Impurity ($|1,0\rangle$) and primary ($|1,-1\rangle$) oscillations along \hat{z} (gravity) in the x-dipole trap at $\nu = 64.0(4)\text{Hz}$ after the spin-dependent impulse. Data are from separate experiments with either all atoms in $|1,0\rangle$ or all atoms in $|1,-1\rangle$. The primary moves less than $0.25\mu\text{m/ms}$, whereas the impurity shows clear oscillations. The data is fit to a damped sinusoid, $v(t) = v_0 e^{-t\gamma} \sin(2\pi\nu t + \phi)$, where $\gamma = 19(2)\text{s}^{-1}$ and $v_0 = 2.0(1)\mu\text{m/ms}$. If the trap were purely harmonic γ should be zero, so the non-zero γ is due to anharmonicities. (b) Residual movement from the spin-dependent kick along \hat{y} . The magnetic gradient (see text) used to cancel the $|1,-1\rangle$ motion along \hat{y} induces some motion in the orthogonal direction.

The main technique developed for the impurity transport probe is the ability to apply a force to the $|1,0\rangle$ impurity atoms along one direction but not to the $|1,-1\rangle$ primary atoms. Applying a force to only the $|1,-1\rangle$ atoms is trivial because $|1,0\rangle$ has no magnetic moment and therefore does not experience a force from magnetic field gradients. Therefore, to only apply a $|1,0\rangle$ force, we need to use a combination of magnetic and optical forces. Specifically, we shift the cross-dipole trap in the direction of gravity (\hat{z} in the coordinates of Fig. 3.1), which imparts an equal force to both the impurity and the primary. Concurrently, we apply a magnetic field gradient using our Stern-Gerlach coil, also along gravity, which only exerts a force on the primary ($|1,-1\rangle$) atoms. The two forces cancel for the primary, so only the impurity experiences a force and therefore starts to move. There is some residual motion for both the impurity and the primary in the orthogonal direction (see Fig. 8.11). The impurity starts to move through the primary system and oscillates due to the harmonic confinement. After a variable amount of time t_{hold} , we image the system after TOF and measure the center-of-mass velocity of the impurity. The velocity curves for the primary and impurity, when they are not mixed together, are shown in Fig. 8.11. There is some decay of the impurity even without the primary present due to anharmonicities in the trapping potential. As described in the caption to Fig. 8.11, we fit the velocity to a damped sinusoid and the baseline damping rate of the control data is $19.0(2)\text{s}^{-1}$.

We perform the impulse procedure when there is mixture of $60(3) \times 10^3$ atoms in the primary and $23(2) \times 10^3$ atoms in the impurity with an initial impurity condensate fraction of approximately 0.5. If we do the impulse without the spin-dependent lattice, this essentially reproduces past experiments investigating out-of-phase condensate oscillations in a harmonic trap [85, 274]. Similar to those experiments, we observe damping of the center-of-mass velocity. Since the impurity starts with condensate fraction less than




Figure 8.12: (a) Impurity velocity versus time for a mixture where the primary is loaded into a $s = 6$ lattice. The data is fit to a damped sinusoid (solid curve). (b) Impurity damping rate versus the lattice depth for the primary. The grey box is the baseline decay rate from the impurity oscillation without the primary atoms present (Fig. 8.11). As the lattice depth increases the damping rate for decreases. Eventually at $s = 14$ the rate is consistent with the background rate.

unity, we fit the images after TOF to a Gaussian plus BEC (TF profile) with separate centers. However, the impurity motion induces large distortions in the cloud, which makes a clean fit difficult. In many cases, the condensate itself is not described by a single peak. Nevertheless, the fits were almost always able to ascertain a broad Gaussian feature and a dominant condensate peak; we fit the condensate velocity to a damped sinusoid to measure the decay rate⁶. Next, we measure the impurity decay rate when the primary is loaded into the lattice. The damping rate as a function of the primary lattice depth and velocity data for $s = 6$ are displayed in Fig. 8.12. When the primary is in the lattice, the impurity decay rate initially grows, and then, as the lattice depth increases, the damping rate decreases (at $s = 14$ the rate consistent with no damping).

The work presented here establishes a new type of impurity probe of the lattice gas. However, interpreting the decay rate data (Fig. 8.12) requires theoretical input — this is left as an extension to this work. Similar to the thermalization data, decay of impurity velocity involves energy exchange between the lattice and impurity. Excitations must be created in the lattice gas in order to damp the impurity motion. Therefore, this technique may be useful as an excitation spectrum probe. This is consistent with the observed trend in the data. At high lattice depths spectral weight is transferred in the lattice gas to a gapped mode (see, e.g., [275]). This gap suppresses the creation of excitations in the lattice and as a result, the impurity motion does not damp. Perhaps this hypothesis can be evaluated in a future theoretical collaboration.

⁶Analyzing the Gaussian motion may be an interesting future extension to this work.

Chapter 9

Conclusions and Outlook

The work presented in this thesis centers on quantum simulation using ultracold atoms trapped in optical lattices. In particular, we focused on developing a toolkit for experimental quantum simulation — tools (§3), probes, and techniques (§4). Most of the experimental work in this thesis was performed on a ^{87}Rb optical lattice apparatus designed to simulate the Bose-Hubbard model (§3.1). This apparatus was essentially complete before this thesis started [41], so my work centered on investigating the non-equilibrium Bose-Hubbard model (§6), developing thermometry probes (§7, §8.1), and realizing a novel 3D spin-dependent lattice in the strongly correlated regime (§8).

While the toolkit is not yet (and may never be) complete, my thesis work has significantly advanced it. We have advanced our understanding of thermometry probes, added new tools (spin-dependent lattice and ^{40}K single-site apparatus), and simulated the BH model out of equilibrium in unexplored regimes. Based on our work, we develop several important conclusions:

The Bose-Hubbard model has intrinsic dissipation: In the transport experiment described in §6, we demonstrated that the clean lattice gas has zero-temperature dissipation for mass transport consistent with the generation of phase slips. This experiment highlights the capabilities of quantum simulation; phase-slip dissipation is a dynamical, many-body effect that is impossible to simulate numerically, but it was straightforward to study using our lattice gas (“quantum simulator”). These results are important for understanding charge transport in superconducting materials thought to be described by the Bose-Hubbard model. There are a number of unexplored subjects that could motivate future experiments. Possible topics include the relationship between dimensionality and decay rate, measuring the decay rate dependence on velocity and density, and understanding why the thermal gas is pinned to the lattice. Furthermore, to better understand the dissipation process, an enhanced setup for direct imaging of phase slips is required. With improved imaging, we could investigate the role of quantum turbulence in finite-temperature dissipation (analogous to superfluid ^4He).

Bandmapping fails as a thermometry probe: Bandmapping is a process whereby the lattice depth is linearly ramped down adiabatically with respect to the bandgap and diabatically (fast) compared to all other timescales to map $q \rightarrow p$ (quasimomentum to momentum). We investigated bandmapping as a probe of the quasimomentum distribution in a single band. In theory, there are many advantages to bandmapping (e.g., to increase signal-to-noise ratio), yet, from the results of two experiments, we conclude that it fails as a probe. In the first experiment, we explored direct thermometry of a thermal lattice gas by fitting bandmapped distributions (§7.1). We found that bandmapping does not reliably produce quasimomentum

distributions because of band transitions at the edge of the Brillouin zone when $k_B T \gtrsim t$. Additionally, the quasimomentum distribution is not sensitive at these temperatures because the kinetic energy bandwidth is finite — additional thermal energy is accommodated as potential and interaction energy. In the second experiment, we measured condensate fraction versus the bandmapping time (§6.2). We found that bandmapping was not accurate because interactions redistribute atoms into the condensate.

In-situ thermometry probes are required for lattice gases: Loading a weakly interacting harmonic-trapped gas into a lattice fundamentally alters the distribution of thermal energy — the lattice potential and interactions break the symmetry between momentum and spatial degrees of freedom which is present in the harmonic potential. In the harmonically trapped Bose-Hubbard model, the kinetic energy is finite, but the trap energy is unbounded. Thermodynamic state information is therefore encoded primarily in spatial properties, e.g., the in-trap size and on-site number fluctuations. We measured temperature using the in-trap size (§7.2), which resulted in an order of magnitude improvement in accuracy compared with bandmapping. In-situ probing is the motivation behind constructing the ^{40}K single-site apparatus (§3.2). The requirement for in-situ probes is evident from the known phases of the two-component Fermi-Hubbard model: the Mott insulator and the antiferromagnet. Both are characterized by signatures in the density distribution: the Mott insulator has reduced number fluctuations on each site, and the antiferromagnet has spin correlations between next-nearest neighbors. In addition, direct in-situ thermometry probes will allow us to assess lattice cooling schemes, which will enable studying low entropy phases (e.g., antiferromagnetism).

Finite-temperature dynamics are an open question: Finite-temperature dynamics are an important element for quantum simulation (§4.2.2). For example, in our BH transport experiment (§6.1), measuring the decay rate versus temperature revealed the crossover to thermal activation of phase slips. However, an incomplete understanding of finite-temperature dynamics is impeding quantum simulation because we are unable to properly assess one of our main assumptions — that a gas can be adiabatically transferred into the lattice. We investigated adiabaticity by measuring the finite-temperature condensate fraction after loading into the lattice (§7.3). We observed significant deviation from theoretical predictions, and, notably, we observe a condensate in the lattice when, according to equilibrium thermodynamics, the temperature should be too high. We also observed that the condensate fraction is strongly dependent on the lattice ramp time — evidence that our adiabaticity assumption is incorrect. Finite-temperature also caused discrepancies between zero-temperature theory and experiment when we investigated fast ramps of the lattice depth (§6.2). Both these results indicate that we require theoretical techniques to study finite-temperature dynamics of the Bose-Hubbard model. These techniques are also needed to assess in-lattice cooling schemes.

3D spin-dependent lattices are a promising new quantum-simulation tool: In this thesis work, we demonstrated the first 3D spin-dependent lattice in the strongly correlated regime (§3.1.2, §8). The depth of this lattice is proportional to $|g_F m_F|$, so atoms with $m_F = 0$ do not experience the lattice. This allows us to simulate the physics of a Hubbard model immersed in a bath of free-particles, for example, to simulate electron-phonon interactions in materials. Furthermore, as explored in this thesis, we can use the free-gas as an impurity probe. The spin-dependent lattice will allow future investigation into proposals for cooling based on emission into the free-particle bath from band decay in the lattice [70].

A harmonically trapped impurity does not thermalize with the lattice gas: Thermometry is an ideal application of the free-gas impurity using the spin-dependent lattice. Indeed, this was proposed as part of this thesis work in Ref. [4]. However, from a series of experiments described in §8.1, we conclude that there is insufficient thermalization between the impurity and lattice gas to enable effective thermometry. Energy-exchanging collisions are suppressed due to a mismatch between the dispersion of the impurity and lattice gas. There are a number of difficulties preventing a more quantitative understanding of the thermalization timescales, namely, fundamental heating due to spontaneous emission from the lattice beams. Heating in the spin-dependent lattice is more severe because the lattice beams must be detuned close to resonance (on the order of the fine-structure splitting). A promising future direction would be to explore impurities in a conservative spin-dependent lattice, for example, based on patterned magnetic lattices [97].

Towards single-site imaging of ^{40}K : A separate aspect of this work was building a new apparatus capable of single-site imaging of a 2D gas of ^{40}K in an optical lattice to simulate the Fermi-Hubbard model (§3.2). We did not achieve this goal, but it was ambitious — no group has demonstrated single-site imaging of ultracold fermions. Consequently, the success of this thesis work was to establish a foundation for the experiment so that we are positioned to, in the near future, become the first group to perform single-site imaging of fermions.

In the decade since the first experimental demonstration of the Mott insulator [33], quantum simulation using ultracold atoms in optical lattices has greatly advanced as a field. However, there is still work to be done for quantum simulation to reach its intended goal of implementing the feedback cycle illustrated in Fig. 1.2 and motivating the search for new materials. The ultimate success of quantum simulation will depend on our ability to attain low entropies, develop more sensitive probes, and to delve into unknown regimes of quantum models.

Appendix A

Properties of ^{87}Rb and ^{40}K

The quantum particles of our simulator are ultracold atoms. Conceptually, the simulator is insensitive to the technical details of the atoms except for their statistics under exchange (i.e., whether they are bosons or fermions). However, from an experimental perspective these details (e.g., mass, magnetic moment, and transition energies) are important parameters required for implementation and analysis of results. In this appendix we summarize relevant properties of the two atoms used in this thesis: ^{87}Rb for our Bose-Hubbard quantum simulator (§3.1) and ^{40}K for our Fermi-Hubbard quantum simulator (§3.2)¹. The physical properties of each are summarized in Table A.1, state and transition properties in §A.1, and collisional properties in §A.2. All physical constants in this section are given by the 2010 CODATA values [5].

	^{87}Rb	^{40}K	Ref.
Atomic Number (Z)	37	19	
Mass	$86.90918053(1)u$ $1.44316065(7) \times 10^{-25} \text{ kg}$	$39.9639985(2)u$ $6.6361769(2) \times 10^{-26} \text{ kg}$	[276]
Abundance	27.83(2)%	0.0117(4)%	[277]
Nuclear Spin (I)	3/2	4	
Nuclear Magnetic Moment (μ_N)	$h \times -1392.824(1) \text{ Hz/Gauss}$	$h \times 247.02(5) \text{ Hz/Gauss}$	[278]
Melting Point	36.31°C	63.38°C	[192, 193]
Vapor Pressure at $T = 300K$	$3.7 \times 10^{-7} \text{ Torr}$	$1.8 \times 10^{-8} \text{ Torr}$	[190, 279]
Vapor Pressure at $T = 400K$	$1.0 \times 10^{-3} \text{ Torr}$	$1.2 \times 10^{-4} \text{ Torr}$	[190, 279]
Half-life	49.2(2) Gy	1.25(1) Gy	[280]

Table A.1: Physical Properties of ^{87}Rb and ^{40}K .

A.1 States

Atoms have a number of discrete quantum states corresponding to different arrangements of the constituent electrons and orientation of the nuclear spin. For the atoms used here, ^{40}K and ^{87}Rb , the states are illustrated in order of increasing energy in Fig. A.1 and Fig. A.2 respectively. For alkali atoms, with a single valence electron, the standard labeling convention for the atomic state is to first specify the state of the outer electron, nL_J , where n is the principle quantum number, L is the orbital angular momentum of the electron ($L = S, P, D(0, 1, 2) \dots$), S is the electron spin angular momentum ($S = 1/2$) and J is the total

¹ ^{87}Rb is also used in the Fermi-Hubbard quantum simulator as a coolant.




Figure A.1: ^{40}K state diagram up to the $5P$ state, including hyperfine details for the states used in this thesis $4S_{1/2}$, $4P_{3/2}$, and $5P_{3/2}$. Transition probabilities are from [192], except for $4P \rightarrow 4S$ (from [281]), and listed as $A = 1/\tau$ in units of μs^{-1} . The total measured lifetime from all decay channels is 134(2)ns for $5P_{3/2}$ [282] and 137.6(1.3)ns for $5P_{1/2}$ [283]. Level energies E referenced from the $4S$ state are from [192] and presented in both frequency units (as E/h) and wavelength units (in italics, as hc/E). The ionization energy, from [192], is shown at the top. Hyperfine splittings ΔE referenced to the level energies are from [278] for $4S_{1/2}$ and $5P_{3/2}$, from [192] for $4P_{3/2}$, and given in frequency units (as $\Delta E/h$). The g-factor for each hyperfine state is calculated using Eqn. A.8. Figure adapted from Ref. [36].




Figure A.2: ^{87}Rb state diagram up to the $5P$ state, including hyperfine details for the states used in this thesis $5S_{1/2}$, and $5P_{3/2}$. Transition probabilities are from [284] and listed as $A = 1/\tau$ in units of μs^{-1} . Level energies E referenced from the $4S$ state are from [193] and presented in both frequency units (as E/h) and wavelength units (in italics, as hc/E). The ionization energy, from [193], is shown at the top. Hyperfine splittings ΔE referenced to the level energies are from [193] and given in frequency units (as $\Delta E/h$). Uncertainties in the $4S$ hyperfine splittings are on the order of 1 mHz. The g -factor for each hyperfine state is calculated using Eqn. A.8.

(orbital plus spin) electron angular momentum². Since $\mathbf{S} = 1/2$, J takes on integer spaced values from $|L + 1/2|$ to $|L - 1/2|$. In the ground state, where $L = 0$, $J = \mathbf{S} = 1/2$. For each atom, the value of n corresponding to the ground state depends on the number of states filled by the electrons in the closed shell, and so $n = 1-6$ for H, Li, Na, K, Rb, and Cs. The electronic ground state of K is therefore $4S_{1/2}$ and Rb is $5S_{1/2}$.

The nucleus also has a total spin quantum number \mathbf{I} (see Table A.1) and so each electronic state is split into a further $(2J + 1)(2\mathbf{I} + 1)$ states depending on the direction of the electronic and nuclear spin angular momenta. The eigenstates in the J manifold are determined by the following Hamiltonian [278]

$$\mathbf{H} = \mathbf{H}_{hf} + \mathbf{H}_{\text{Zeeman}}, \quad (\text{A.1})$$

$$\mathbf{H}_{hf} = A\vec{\mathbf{I}} \cdot \vec{\mathbf{J}} + \mathcal{B} \frac{\frac{3}{2}\vec{\mathbf{I}} \cdot \vec{\mathbf{J}}(2\vec{\mathbf{I}} \cdot \vec{\mathbf{J}} + 1) - \mathbf{I}(\mathbf{I} + 1)J(J + 1)}{2\mathbf{I}(2\mathbf{I} - 1)J(2J - 1)}, \quad (\text{A.2})$$

$$\mathbf{H}_{\text{Zeeman}} = \mu_B g_J \vec{\mathbf{J}} \cdot \vec{\mathbf{B}}. \quad (\text{A.3})$$

\mathbf{H}_{hf} is the hyperfine interaction Hamiltonian, where A characterizes the interaction of the electron with the nuclear magnetic moment and \mathcal{B} with the nuclear electric quadrupole moment ($\mathcal{B} = 0$ when $J = 1/2$). Both A and \mathcal{B} are unique to each isotope and electronic state and are summarized for most alkali atomic states in Ref. [278].

The Hamiltonian $\mathbf{H}_{\text{Zeeman}}$ describes the interaction of the electron spin with the magnetic field³ $\vec{\mathbf{B}}$ (Zeeman effect) where $\mu_B = e\hbar/2m$ is the Bohr magneton ($\approx h \times 1.4 \text{ MHz/Gauss}$) and the gyromagnetic factor is [285]

$$g_J = \frac{3}{2} + \frac{S(S + 1) - L(L + 1)}{2J(J + 1)}. \quad (\text{A.4})$$

When $\vec{\mathbf{B}} = 0$, the Hamiltonian is diagonal in $\vec{\mathbf{F}}^2$ where $\vec{\mathbf{F}} = \vec{\mathbf{I}} + \vec{\mathbf{J}}$ is the total nuclear plus electronic angular momentum. In terms of $K = F(F + 1) - J(J + 1) - \mathbf{I}(\mathbf{I} + 1)$, the energy is

$$E(K) = \frac{A}{2}K + \frac{3}{8}\mathcal{B} \frac{K(K + 1)}{\mathbf{I}(2\mathbf{I} - 1)J(2J - 1)} - \mathcal{B} \frac{\mathbf{I}(\mathbf{I} + 1)J(J + 1)}{2\mathbf{I}(2\mathbf{I} - 1)J(2J - 1)}. \quad (\text{A.5})$$

The splitting between F states is

$$E_F - E_{F-1} = AF + \frac{3}{4}\mathcal{B} \frac{F(1 + 2K - 2F)}{\mathbf{I}(2\mathbf{I} - 1)J(2J - 1)}. \quad (\text{A.6})$$

Each F state has $2F + 1$ degenerate states corresponding to the orientation of F . This degeneracy is lifted when $B \neq 0$ by $\mathbf{H}_{\text{Zeeman}}$. In weak magnetic fields $\mathbf{H}_{\text{Zeeman}}$ can be considered perturbatively, and

$$E(F, m_F) = E(F) + \mu_B g_F m_F |\vec{\mathbf{B}}| \quad (\text{A.7})$$

²All the non-bold numbers ($L, J, \mathbf{S}, \mathbf{I}, F$) are the eigenvalues of the corresponding vector operator squared, e.g., $\vec{\mathbf{L}}^2|\psi\rangle = L(L + 1)|\psi\rangle$. J is the eigenvalue of the operator $\vec{\mathbf{J}}^2$ where $\vec{\mathbf{J}} = \vec{\mathbf{L}} + \vec{\mathbf{S}}$.

³We ignore the nuclear Zeeman effect and assume $g_S = 2$, $g_L = 1$.

where m_F is the eigenvalue of the \mathbf{F}_z operator ($m_F = \{-F, -F+1, \dots, F-1, F\}$) and g_F is [285]

$$g_F = \frac{F(F+1) + J(J+1) - \mathbf{I}(\mathbf{I}+1)}{2F(F+1)} g_J. \quad (\text{A.8})$$

In the large $|\vec{B}|$ limit (Paschen-Back regime), m_J and m_I are eigenstates,

$$E_{m_J, m_I} = A m_J m_I + \mathcal{B} \frac{(3m_J^2 - J(J+1))(3m_I^2 - \mathbf{I}(\mathbf{I}+1))}{4\mathbf{I}(2\mathbf{I}-1)J(2J-1)} + \mu_B g_J m_J |\vec{B}| \quad (\text{A.9})$$

For arbitrary magnetic fields the eigenstates are given by diagonalizing Eqn. A.1. The eigenvalues for the ground states of ^{40}K and ^{87}Rb as a function of the magnetic field are shown in Fig. A.3.

For the alkali atom ground state, when $J = 1/2$ and $L = 0$, there are two hyperfine levels $F = \mathbf{I} \pm 1/2$ split by energy $A(\mathbf{I} + 1/2)$ and with $g_F = \pm(2\mathbf{I} + 1)^{-1}$. There is an analytic formula for the ground state energies at all field strengths, the Breit-Rabi equation [286]

$$E_m = -\frac{A}{4} \pm \frac{A(2\mathbf{I} + 1)}{4} \left(1 + \frac{4m}{2\mathbf{I} + 1} \epsilon + \epsilon^2 \right)^{1/2}, \quad (\text{A.10})$$

$$\epsilon = \frac{2g_J \mu_B |\vec{B}|}{A(2\mathbf{I} + 1)}, \quad (\text{A.11})$$

where m is the magnetization ($m_J + m_I$). Expanding to second order in ϵ ,

$$E_m \approx -\frac{A}{4} \pm \frac{A(2\mathbf{I} + 1)}{4} + g_F m_F \mu_B |\vec{B}| \pm \frac{(g_F \mu_B)^2}{2A(2\mathbf{I} + 1)} ((2\mathbf{I} + 1)^2 - 4m_F^2) B^2. \quad (\text{A.12})$$

Some care has to be taken when applying the Breit-Rabi equation to get the correct sign and to make sure the energy level is physical. The quadratic Zeeman shift from Eqn. A.12 is

$$E_{\text{quad}}(m_F) = \pm \frac{(g_F \mu_B)^2}{AF_{\text{max}}} (F_{\text{max}}^2 - m_F^2) |\vec{B}|^2 \quad (\text{A.13})$$

for the $F = \mathbf{I} \pm 1/2$ manifold where $F_{\text{max}} = \mathbf{I} + 1/2$.

For manipulation and imaging, the ground state is optically coupled to higher energy (excited) states. A review of optical transitions is given in §B.6. For alkali atoms the D1 ($nS_{1/2} \rightarrow nP_{1/2}$) and D2 ($nS_{1/2} \rightarrow nP_{3/2}$) transitions are the strongest (see Figs. A.1 and A.2), and so almost all optical interactions occur along these two transitions. For scattering we exclusively use the D2 transition. In this thesis work, we also explore transitions to the higher $(n+1)P_{3/2}$ level in ^{40}K . The D2 transition properties for ^{87}Rb and ^{40}K , and the $4S_{1/2} \rightarrow 5P_{3/2}$ transition properties for ^{40}K are summarized in Table A.2.

A.2 Collisional Properties

The pairwise interactions between atoms, discussed in §2.1.2, are parameterized by an s-wave scattering length. For pairs of atoms in a particular species and state, the scattering length can be tuned using the magnetic field via a Feshbach resonance. For a broad resonance (see definition in [111]), the scattering length




Figure A.3: (a) $^{87}\text{Rb } 5S_{1/2}$ and (b) $^{40}\text{K } 4S_{1/2}$ hyperfine states versus magnetic field. For ^{87}Rb $A = 3.4173413054522(5)\text{GHz}$ [193] and for ^{40}K $A = -285.7308(24)\text{MHz}$ [278]. The graphs on the right side are zoomed into the low-field regions of each hyperfine manifold. The zero-field energy has been subtracted off: -4.2717GHz for $^{87}\text{Rb } F = 1$, 2.563GHz for $^{87}\text{Rb } F = 2$, -571.46MHz for $^{40}\text{K } F = 9/2$, and 714.33MHz for $^{40}\text{K } F = 7/2$.

	$^{87}\text{Rb}: 5S_{1/2} \rightarrow 5P_{3/2}$	$^{40}\text{K}: 4S_{1/2} \rightarrow 4P_{3/2}$	$^{40}\text{K}: 4S_{1/2} \rightarrow 5P_{3/2}$
Line Width $\Gamma/2\pi$	6.06(2) MHz	6.04(1) MHz	1.19(2) MHz
Branching Ratio	Cycling	Cycling	1/6.4(7)
Cross Section (σ)	0.2907 μm^2	0.2807 μm^2	0.010 μm^2
Saturation Intensity (I_{sat})	1.669(4) mW/cm ²	1.752(3) mW/cm ²	23(2) mW/cm ²
Steady-State Polarization	100%	100%	61(1)%
Depumping Probability	0%	0%	18(2)%
Doppler Temperature	145.5(4) μK	145.0(3) μK	23.5(7) μK

Table A.2: Properties of the D2 transition ($nS_{1/2} \rightarrow nP_{3/2}$) for ^{40}K and ^{87}Rb and the $4S_{1/2} \rightarrow 5P_{3/2}$ transition of ^{40}K . For the D2 transition, the properties are calculated using values from Figs. A.1 and A.2 using Eqn. B.77 for the cross-section, and Eqn. B.73 for I_{sat} . Values for the $4S_{1/2} \rightarrow 5P_{3/2}$ transition are calculated using the procedure described in §E.1. Branching ratio is the ratio of decay events that occur back along the transition versus all decay paths. Steady-state polarization is defined as the fraction of atoms in the stretched ground state when weakly driving the transition with σ^+ light and with infinite repump light (if the transition is not closed). The depumping probability is the probability of decaying from the excited $F' = I + 3/2$ hyperfine level to the dark $F = I - 1/2$ ground hyperfine level (I is the nuclear spin).

tuning is well-described as

$$a(B) = a_{bg} \left(1 - \frac{\Delta}{B - B_0} \right). \quad (\text{A.14})$$

The scattering lengths and Feshbach resonance properties (where applicable) are summarized in Table A.3 for the states used in this thesis.

State 1	State 2	a/a_{bg} (a_0)	B_0 (G)	ΔB	Ref.
$^{87}\text{Rb}: 2, 2\rangle$	$^{87}\text{Rb}: 2, 2\rangle$	98.98(4)	-	-	[287]
$^{87}\text{Rb}: 1, -1\rangle$	$^{87}\text{Rb}: 1, -1\rangle$	100.4(1)	-	-	[287]
$^{87}\text{Rb}: 1, -1\rangle$	$^{87}\text{Rb}: 1, 0\rangle$	100.4(1)	-	-	[287]
$^{87}\text{Rb}: 1, 0\rangle$	$^{87}\text{Rb}: 1, 0\rangle$	100.9(1)	-	-	[287]
$^{40}\text{K}: 9/2, 9/2\rangle$	$^{40}\text{K}: 9/2, 7/2\rangle$	169.7(3)	-	-	[288]
$^{40}\text{K}: 9/2, -9/2\rangle$	$^{40}\text{K}: 9/2, -7/2\rangle$	174	202.10(1)	7.8(6)	[289]
$^{40}\text{K}: 9/2, -9/2\rangle$	$^{40}\text{K}: 9/2, -5/2\rangle$	174	224.21(5)	9.7(6)	[290]
$^{40}\text{K}: 9/2, 9/2\rangle$	$^{87}\text{Rb}: 2, 2\rangle$	-214(4)	-	-	[291]
$^{40}\text{K}: 9/2, -9/2\rangle$	$^{87}\text{Rb}: 1, 1\rangle$	-189	546.6(2)	-3.1	[291]

Table A.3: Zero-field scattering lengths (a) for the ^{40}K and ^{87}Rb states used in this thesis. For states with a Feshbach resonance, we list the parameters necessary to calculate the scattering length near the resonance using Eqn. A.14. The scattering lengths are given in units of the Bohr radius ($a_0=52.917721092(17)\text{pm}$).

The simplified scattering picture presented in §2.1.2 is correct in most situations, however, there are a few relevant deviations to discuss. At high temperatures we must also consider the p-wave component of the interaction, particularly for identical fermions since s-wave collisions are forbidden. The p-wave cross-section (σ_p) for ^{40}K atoms in $|9/2, 9/2\rangle$ was measured in Ref. [202] and peaks to $\sigma_p \approx 2 \times 10^{-11} \text{cm}^2$ at $T = 100 \mu\text{K}$, and then starts to decrease at lower temperatures due to the centrifugal barrier. At low temperatures the cross-section scales as $\sigma_p \propto T^2$. Five collisions are required for p-wave thermalization (three are required for s-wave thermalization).

For s-wave collisions at high energy there are added corrections to the cross-section (Eqn. 2.11) that depend on the details of the interaction potential. One correction is the “effective range” (r_0), and the resulting cross-section is

$$\sigma = \frac{4\pi a^2}{k^2 a^2 + \left(1 - \frac{1}{2} r_0 a k^2\right)}. \quad (\text{A.15})$$

For ^{40}K - ^{87}Rb scattering the effective range was measured to be $381(6)a_0$ [292]. For negative scattering lengths there is also a Ramsauer-Townsend minimum caused by the cancellation of phase shifts at a certain value of k . This minimum occurs at $T \approx 100\text{--}200\mu\text{K}$ for ^{40}K - ^{87}Rb collisions [292].

Although we assume the interactions are diagonal in the state (i.e., Eqn. 2.16), this is only approximately true. The free space eigenstates ($|F_1, m_{F_1}\rangle, |F_2, m_{F_2}\rangle$) are mixed in the interaction region and the eigenstates are the total spin ($\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$) singlet ($\mathbf{S} = 0$) and triplet ($\mathbf{S} = 1$). However, the total hyperfine ($F = F_1 + F_2$) magnitude and projection are eigenstates in all regions [293], so the potential can be written as [98]

$$\mathbf{V}(r_1, r_2) = \delta(r_1 - r_2) \sum_{F=0}^{2f} \frac{4\pi\hbar^2}{M} a_F \mathbf{P}_F \quad (\text{A.16})$$

where \mathbf{P}_F is the projection operator onto the total F subspace. By symmetry, $a_F \neq 0$ for F even (odd) for bosons (fermions). In a given hyperfine manifold the set of a_F need to be calculated from the full scattering potentials (for ^{87}Rb $F = 1$, $a_{F=2} = 100.4(1)a_0$ and $a_{F=0} = 101.8(2)a_0$ [287]). Using the set of a_F values we can calculate the scattering length between two F, m_F states. For example, the collision between two $|1, 0\rangle$ particles is described by the potential

$$\mathbf{V}|1, 0\rangle|1, 0\rangle = \delta(r_1 - r_2) \sum_{F=0}^{2f} \frac{4\pi\hbar^2}{M} a_F \mathbf{P}_F |1, 0\rangle|1, 0\rangle. \quad (\text{A.17})$$

Expanding into the total F subspace

$$|F_1 = 1, m_{F_1} = 0\rangle|1, 0\rangle = \sqrt{\frac{2}{3}}|F = 2, m_F = 0\rangle - \sqrt{\frac{1}{3}}|0, 0\rangle, \quad (\text{A.18})$$

and back

$$|F = 2, m_F = 0\rangle = \sqrt{\frac{2}{3}}|F_1 = 1, m_{F_1} = 0\rangle|1, 0\rangle + \frac{1}{\sqrt{6}}(|1, 1\rangle|1, -1\rangle + |1, -1\rangle|1, 1\rangle), \quad (\text{A.19})$$

$$|F = 0, m_F = 0\rangle = -\frac{1}{\sqrt{3}}|F_1 = 1, m_{F_1} = 0\rangle|1, 0\rangle + \frac{1}{\sqrt{3}}(|1, 1\rangle|1, -1\rangle + |1, -1\rangle|1, 1\rangle), \quad (\text{A.20})$$

then

$$\mathbf{V}|1, 0\rangle|1, 0\rangle = \delta(r_1 - r_2) \frac{4\pi\hbar^2}{M} \left(\frac{2a_{F=2} + a_{F=0}}{3} |1, 0\rangle|1, 0\rangle + \frac{a_{F=2} - a_{F=0}}{3} (|1, 1\rangle|1, -1\rangle + |1, -1\rangle|1, 1\rangle) \right). \quad (\text{A.21})$$

For non-spin changing collisions the effective scattering length of two $|1, 0\rangle$ particles is $(2a_{F=2} + a_{F=0})/3$. However, this full potential illustrates that we can also have spin changing collisions from $|1, 0\rangle|1, 0\rangle \rightarrow |1, 1\rangle|1, -1\rangle$. This process can be tuned out of resonance using a sufficiently strong magnetic field due to the quadratic Zeeman shift (Eqn. A.13). The spin-changing process is minor for ^{87}Rb because all the states

have essentially the same scattering length ($a_{F=2} \approx a_{F=0}$).

Appendix B

Atoms Interacting with Electromagnetic Fields

To interact with the atoms we utilize electromagnetic fields that vary in frequency from DC to 10^{15} Hz. For example, we employ DC magnetic fields for trapping, RF magnetic fields (0-100MHz) to generate transitions between m_F states, microwave magnetic fields (1-10GHz) to transition between F manifolds and light (high frequency electric fields, $\lambda = c/f = 400\text{nm}-10\mu\text{m}$) for trapping, cooling, and imaging.

Atom-field interactions are treated in numerous references, however, all relevant topics are spread throughout the literature in different forms and notations that need to be adapted to our particular conventions and experimental situation. Therefore, in this appendix we summarize key topics regarding the atom-field interaction which are used throughout this thesis and we present formulas in a convenient form for experimental calculation.

B.1 Interaction Hamiltonian

The atom-field interaction Hamiltonian can be greatly simplified by taking the lowest order terms of a multipole expansion (magnetic-dipole and electric-dipole) [294, 295], so that

$$\mathbf{H}_{int} = -\vec{\mathbf{d}} \cdot \vec{E} - \vec{\mu} \cdot \vec{B} \quad (\text{B.1})$$

where $\vec{\mathbf{d}} = -e\vec{\mathbf{r}}$ and $\vec{\mu} = -\mu_B g_J \vec{\mathbf{J}}$ for the alkali atoms considered in this thesis¹. This “dipole approximation” is excellent since the wavelength of the fields we consider are always large compared to the atomic size. We only need to consider higher order terms when a selection rule forbids dipole coupling between a particular set of states (a “dipole-forbidden” transition). However, this situation does not arise for the states used in this thesis.

The atom-field interaction Hamiltonian (Eqn. B.1) operates on the set of atomic states discussed in §A.1, and so we need to add Eqn. B.1 to the atomic interaction Hamiltonian Eqn. A.2. Since Eqn. A.2 is spherically symmetric, the direction in space of the atomic states is ambiguous. An explicit method is to define a constant reference frame, but this becomes unwieldy in the presence of inhomogeneous fields. Instead, we assume that there is a magnetic field present at all times that defines a quantization axis (taken to be along \hat{z}), which is the case in our experiments. Even when the field is inhomogeneous, the z projection is preserved as the atom moves through space (see discussion in §B.2).

¹See §A.1 for the definition of g_J, μ_B and $\vec{\mathbf{J}}$. The operator, $\vec{\mathbf{r}}$, is the position operator of the outer valence electron with respect to the nucleus.

B.2 DC Magnetic Fields

The eigenstates and energies in a uniform DC magnetic field due to the magnetic-dipole (Zeeman) interaction are detailed in §A.1. These energies are plotted versus $|B|$ in Fig. A.3 for ^{40}K and ^{87}Rb . At low fields, the magnetic dipole interaction is a perturbation to the hyperfine interaction (Eqn. A.2) and using the projection theorem [112] we can rewrite it as

$$\mathbf{H}_{\text{Zeeman}} = \mu_B g_F \vec{\mathbf{F}} \cdot \vec{\mathbf{B}} \quad (\text{B.2})$$

where g_F is given by Eqn. A.8. As demonstrated by Fig. A.3, Eqn. B.2 lifts the $(2F+1)$ degeneracy in each F manifold into states $|F, m_F\rangle$ where m_F is the projection onto the field. The energy of each $|F, m_F\rangle$ state is

$$E_{F, m_F} = \mu_B g_F m_F |\vec{\mathbf{B}}|. \quad (\text{B.3})$$

This interaction is important because it results in distinct, robust states for the experiment. Indeed, we do all calculations in the $|F, m_F\rangle$ basis and assume that the state populations are well-defined. Furthermore, since the Zeeman energy shift between the states is almost always the largest term in Eqn. B.1, we can apply all other interactions perturbatively.

When the DC magnetic fields are inhomogeneous, the Zeeman energy shift (Eqn. B.3) is also inhomogeneous

$$E_{F, m_F}(x) = \mu_B g_F m_F |\vec{\mathbf{B}}(x)|. \quad (\text{B.4})$$

Atoms with $g_F m_F > 0$ are trapped in magnetic fields with a local minima in $|B(x)|$ (local maxima are not possible). A few trapping field configurations are discussed in §2.1.1. The form of Eqn. B.4 assumes that the *local* projection is preserved as the atoms move through the inhomogeneous field even as the field changes direction. For cold atoms this is a good assumption under most conditions (see below). Eqn. B.4 illustrates why defining the atomic states with respect to a fixed reference frame is overly complicated; adiabatically the system would be in a spatially-dependent superposition of states. When these states are defined with respect to the local field the system remains in a single state.

To define a bounds on the validity of Eqn. B.4 (i.e., the conditions under which the system can follow the local field direction), consider an atom at $x = 0, \mathbf{t} = 0$ traveling along x with velocity v in the field

$$\vec{\mathbf{B}} = [0, \alpha x, B_0]. \quad (\text{B.5})$$

The time-dependent interaction is

$$\mathbf{H}_{\text{Zeeman}}(\mathbf{t}) = \mu_B g_F (\mathbf{F}_z B_0 + \mathbf{F}_y \alpha v \mathbf{t}). \quad (\text{B.6})$$

The adiabaticity criteria for a general time-dependent Hamiltonian is [296]

$$\left| \frac{\max \left(\frac{\langle j | \partial \mathbf{H} / \partial \mathbf{t} | i \rangle}{\hbar \omega_{ij}} \right)}{\min(\omega_{ij})} \right|^2 \ll 1 \quad (\text{B.7})$$

where i, j are the eigenstates without the time-dependent coupling and $\hbar \omega_{ij}$ is the energy difference between

these states. For the Hamiltonian given by Eqn. B.6, i and j are the states $|F, F\rangle$ and $|F, F-1\rangle$ and $\hbar\omega_{ij} = \mu_B g_F B_0$. These states are coupled by the \mathbf{F}_y operator as

$$\langle F, F-1 | \mathbf{F}_y | F, F \rangle = \sqrt{F/2}. \quad (\text{B.8})$$

Therefore, the adiabaticity criteria (close to $x = 0$) becomes

$$\left| \frac{\alpha}{B_0} \frac{\hbar v}{\mu_B g_F B_0} \sqrt{F/2} \right|^2 \ll 1. \quad (\text{B.9})$$

Since $\alpha/B_0 = \partial\theta/\partial x$ (the change in the field direction), more generally

$$\left| \frac{\partial\theta}{\partial(\vec{x} \cdot \hat{v})} \frac{\hbar|\vec{v}|}{\mu_B g_F B_0} \sqrt{F/2} \right|^2 \ll 1. \quad (\text{B.10})$$

The adiabaticity criterion from Eqn. B.10 indicates that Eqn. B.4 breaks down when the atoms is traveling quickly in regions of small field magnitude. This is a common issue for the quadrupole trapping field, which fully goes to zero at the center of the trap. Fast atoms in the center of this trap do not satisfy Eqn. B.10 and can go from a trapped projection $g_F m_F > 0$, to an untrapped projection $g_F m_F < 0$, also known as a “spin-flip”. When atoms spin-flip they are lost from the trap, a process referred to as Majorana loss. The Majorana loss rate is an important factor when utilizing quadrupole traps and is given by the formula [103]²,

$$\Gamma_{\text{Majorana}} = 1.85 \frac{\hbar}{m} \left(\frac{\mu_B g_F m_F B_{\text{grad}}}{k_B T} \right)^2 \quad (\text{B.11})$$

where B_{grad} is the gradient along the strong axis of the quadrupole field and 1.85 is an empirical factor.

B.3 Oscillating Fields

When the field is oscillating at frequency ω the interaction Hamiltonian gives rise to a number of interesting phenomena. However, solving Eqn. B.1 is more involved because of the time-dependence of the oscillating field. Fortunately, there are several approaches that simplify the problem, each with different advantages and disadvantages. The approach we take here is to quantize the fields and work in the dressed-state picture. Although this approach is conceptually more involved, in practice it simplifies the calculation because the time-dependence is naturally included in the formalism. Furthermore, field-quantization is necessary to describe all possible phenomena (e.g., spontaneous decay).

To quantize the fields we expand into photon operators in the Schrodinger representation [297]

$$\vec{\mathbf{E}}(\vec{\mathbf{X}}) \rightarrow i \sum_k \sum_{\hat{\epsilon}} A_k \left(\hat{\epsilon} \mathbf{a}_{k,\hat{\epsilon}} e^{ik\vec{\mathbf{X}}} - \hat{\epsilon}^* \mathbf{a}_{k,\hat{\epsilon}}^\dagger e^{-ik\vec{\mathbf{X}}} \right), \quad (\text{B.12})$$

$$\vec{\mathbf{B}}(\vec{\mathbf{X}}) \rightarrow \frac{i}{c} \sum_k \sum_{\hat{\epsilon}} A_k \left((\hat{k} \times \hat{\epsilon}) \mathbf{a}_{k,\hat{\epsilon}} e^{ik\vec{\mathbf{X}}} - (\hat{k} \times \hat{\epsilon}^*) \mathbf{a}_{k,\hat{\epsilon}}^\dagger e^{-ik\vec{\mathbf{X}}} \right), \quad (\text{B.13})$$

²This formula was original postulated for atoms in the $F = 1$ manifold, which only has one trapped state, so any violation of Eqn. B.10 leads to loss. Applying this formula to high F states may not be accurate and will not reflect any changes in m_F which still result in a trapped state.

where $A_k = \sqrt{\frac{\hbar\omega_k}{2\epsilon_0 V}}$ is the field mode volume (V is the volume of space) and $\mathbf{a}_{k,\hat{\epsilon}}^\dagger$ creates a “photon” with wavevector k and polarization $\hat{\epsilon}$. In this formalism, we include the field energy in the full Hamiltonian, and so, assuming for simplicity two atomic states ($|g\rangle$ and $|e\rangle$),

$$\begin{aligned} \mathbf{H} = & \frac{\vec{P}^2}{2m} + V(\vec{X}) + \hbar\omega_0 \mathbf{b}^\dagger \mathbf{b} + \sum_k \sum_{\hat{\epsilon}} \hbar\omega_k \mathbf{a}_{k,\hat{\epsilon}}^\dagger \mathbf{a}_{k,\hat{\epsilon}} + \\ & i \sum_k \sum_{\hat{\epsilon}} A_k \left(\hat{\epsilon} \mathbf{a}_{k,\hat{\epsilon}} e^{ik\vec{X}} - \hat{\epsilon}^* \mathbf{a}_{k,\hat{\epsilon}}^\dagger e^{-ik\vec{X}} \right) \left(\langle e|\vec{D}|g\rangle \mathbf{b}^\dagger + \langle g|\vec{D}|e\rangle \mathbf{b} \right) \end{aligned} \quad (\text{B.14})$$

where $\mathbf{b} = |g\rangle\langle e|$ is the lowering operator for the atomic states. For a particular field, only one of the two terms from Eqn. B.1 dominates, so we have written down a general one-term form of the interaction with a generic dipole operator, \vec{D} . If Eqn. B.14 describes an electric-dipole interaction then $\vec{D} = -\vec{d} = e\vec{X}$ and $\mathbf{a}_{k,\hat{\epsilon}}^\dagger$ creates a photon with electric-field polarization $\hat{\epsilon}$. On the other hand, if Eqn. B.14 describes a magnetic-dipole interaction then $\vec{D} = -\vec{\mu} = \frac{\mu_B g_J}{c} \vec{J}$ and $\mathbf{a}_{k,\hat{\epsilon}}^\dagger$ creates a photon with magnetic-field polarization $\hat{\epsilon}$. For completeness we include the atomic center-of-mass operators \vec{P} and \vec{X} , which couple to the field via the phase term. These terms describe the photon imparting momentum to the atom during an absorption or emission process. We often neglect these terms and consider \vec{X} and \vec{P} as classical variables since the timescales for atomic motion are considerably slower than the atomic state timescales.

The fields that we generate in the experiment are predominantly classical fields with peak field amplitude $|E_0|(|B_0|)$ and phase ϕ (i.e., coherent states). Coherent states are written in terms of the photon number states (Fock states) created by operator \mathbf{a}_k^\dagger as [298]

$$|E_0|e^{i\phi} = e^{-\bar{N}/2} \sum_{n=0}^{\infty} \frac{\bar{N}^{n/2} e^{i\phi n}}{\sqrt{n!}} |n_{\omega,k}\rangle \quad (\text{B.15})$$

where \bar{N} is the mean photon number related to $|E_0|(|B_0|)$ as

$$|E_0|^2 = \bar{N} \frac{\hbar\omega_k}{\epsilon_0 V}, \quad (\text{B.16})$$

$$|B_0|^2 = \bar{N} \frac{\hbar\omega_k}{\epsilon_0 c^2 V}. \quad (\text{B.17})$$

A common experimental measure of the field amplitude is the intensity

$$I = \frac{\hbar\omega c \bar{N}}{V}. \quad (\text{B.18})$$

In Eqn. B.14 the interaction couples Fock states $|n_\omega\rangle \rightarrow |n_\omega \pm 1\rangle$, so a full calculation with the coherent state requires keeping an infinite number of terms. For simplicity, we calculate with the Fock state pair $\{|\bar{N}_\omega\rangle, |\bar{N}_\omega - 1\rangle\}$ and assume that the answer is equivalent to the full coherent state result. This is a good approximation when $\bar{N} \gg 1$. However, the full coherent field is important when calculating the expectation value of the dipole operator (see §B.4.1).

The interaction part of Eqn. B.14 expands into four terms

$$\mathbf{a}_{k,\hat{\epsilon}}\mathbf{b}^\dagger, \quad \mathbf{a}_{k,\hat{\epsilon}}^\dagger\mathbf{b}, \quad \mathbf{a}_{k,\hat{\epsilon}}\mathbf{b}, \quad \mathbf{a}_{k,\hat{\epsilon}}^\dagger\mathbf{b}^\dagger. \quad (\text{B.19})$$

The first (second) term couples atomic state from $|g\rangle \rightarrow |e\rangle$ ($|e\rangle \rightarrow |g\rangle$) and removes (adds) a photon to the field. For the last two terms, the opposite process occurs. The energy cost to go from $|g\rangle \rightarrow |e\rangle$ is $\hbar\omega_0$ and the energy to add a photon to the field is $\hbar\omega$. When the detuning $\Delta = \omega - \omega_0$ is small compared to ω_0 ($|\Delta| \ll \omega_0$), the first two terms (co-rotating terms) are near resonance, whereas the last two terms (counter-rotating terms) are far from resonance. Therefore, for small detuning we apply the rotating wave approximation (RWA) and keep only these first two terms. The RWA Hamiltonian for a single field mode is ($\alpha = e^{-ik\vec{\mathbf{X}}} \langle g | \vec{\mathbf{D}} \cdot \hat{\epsilon}^* | e \rangle$)

$$\begin{aligned} \mathbf{H} = & \frac{\vec{\mathbf{P}}^2}{2m} + V(\vec{\mathbf{X}}) + \hbar\omega_0|e\rangle\langle e| + \hbar\omega_k\mathbf{a}_{k,\hat{\epsilon}}^\dagger\mathbf{a}_{k,\hat{\epsilon}} + \\ & iA_k \left(\alpha^*\mathbf{a}_{k,\hat{\epsilon}}\mathbf{b}^\dagger - \alpha\mathbf{a}_{k,\hat{\epsilon}}^\dagger\mathbf{b} \right). \end{aligned} \quad (\text{B.20})$$

B.3.1 Coupling Matrix Element for Alkali Atoms

The atom-field interaction causes transitions between atomic states because of the off-diagonal matrix elements

$$\langle g | \vec{\mathbf{D}} \cdot \hat{\epsilon}^* | e \rangle \quad (\text{B.21})$$

where $\vec{\mathbf{D}}$ is a vector operator and $\hat{\epsilon}$ is the field polarization. By writing $\hat{\epsilon}$ in a specific basis we can determine if the coupling is non-zero based on a set of criteria known as selection rules. To start we write the polarization in the atomic basis

$$\begin{aligned} \hat{\sigma}^+ &= \left(\frac{\hat{x} + i\hat{y}}{\sqrt{2}} \right), \\ \hat{\sigma}^- &= \left(\frac{\hat{x} - i\hat{y}}{\sqrt{2}} \right), \\ \hat{\pi} &= \hat{z} \end{aligned} \quad (\text{B.22})$$

where \hat{x} , \hat{y} , and \hat{z} are defined according to the local quantization axis. The operators obtained by projecting $\vec{\mathbf{D}}$ onto this basis $\hat{\sigma}^+ \cdot \vec{\mathbf{D}} = \mathbf{D}_1^{(1)}$, $\hat{\sigma}^- \cdot \vec{\mathbf{D}} = \mathbf{D}_{-1}^{(1)}$, and $\hat{\sigma}^+ \cdot \vec{\mathbf{D}} = \mathbf{D}_0^{(1)}$ are spherical tensor operators. These are a convenient representation so that we can apply the Wigner-Eckert theorem [112]

$$\langle nFm_F | \mathbf{D}_q^{(1)} | n'F'm_F' \rangle = i(-1)^{F_{max}-m_F} \begin{pmatrix} F & 1 & F' \\ -m_F & q & m_F' \end{pmatrix} \langle nF || \mathbf{D}^{(1)} || n'F' \rangle \quad (\text{B.23})$$

using the alkali notation described in §A.1 and $q = -1, 1, 0$ for σ^+, σ^-, π respectively³. The term in the brackets (...) is the Wigner 3J symbol and is non-zero only if

$$m_F - m_F' = q, \quad (\text{B.24})$$

$$F' = F, F \pm 1. \quad (\text{B.25})$$

³Because of the form of Eqn. B.21, the q in Eqn. B.23 refers to the complex conjugate of the driving polarization.

These are the selection rules of the transition. The final term $\langle nF || \mathbf{D}^{(1)} || n'F' \rangle$ is the reduced matrix element, and it has no angular dependence. Since $F = I + J$ we can further reduce the matrix element to

$$\langle nFJI || \mathbf{D}^{(k)} || n'F'J'I \rangle = (-1)^{J+I+F_{min}+1} \sqrt{(2F+1)(2F'+1)} \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\} \langle nJI || \mathbf{D}^{(k)} || n'J'I \rangle \quad (\text{B.26})$$

where $\{...\}$ is the Wigner 6J symbol.

For magnetic-dipole transitions the reduced operator is just the identity, so the reduced matrix element is $\delta_{n,n'}\delta_{J,J'}$. Therefore, the total matrix element is

$$\langle F, m_F | \mathbf{J}_q^{(1)} | F' m'_F \rangle = i(-1)^{F_{max}-m_F+J+I+F_{min}+1} \sqrt{(2F+1)(2F'+1)} \begin{pmatrix} F & 1 & F' \\ -m_F & q & m'_F \end{pmatrix} \left\{ \begin{matrix} J & F & I \\ F' & J & 1 \end{matrix} \right\}. \quad (\text{B.27})$$

For the electric-dipole transition the matrix element is

$$\begin{aligned} \langle n, J, F, m_F | \mathbf{d}_q^{(1)} | n', J', F' m'_F \rangle &= i(-1)^{F_{max}-m_F+J+I+F_{min}+1} \sqrt{(2F+1)(2F'+1)} \begin{pmatrix} F & 1 & F' \\ -m_F & q & m'_F \end{pmatrix} \dots \\ &\quad \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\} \langle nJ || \mathbf{d} || n'J' \rangle. \end{aligned} \quad (\text{B.28})$$

We are left with a single quantity to calculate — the reduced dipole matrix element. This is given by the decay rate from $|n, J\rangle \rightarrow |n', J'\rangle$ summed over all possible F', m'_F final states (see Eqn. B.101)

$$\Gamma = \frac{\omega_0^3}{3\hbar c^3 \pi \epsilon_0} \frac{|\langle nJ || \mathbf{d} || n'J' \rangle|^2}{2J+1}. \quad (\text{B.29})$$

The factor of $(2J+1)^{-1}$ is conventional and J here always refers to the higher energy state. Using Eqn. B.29 and Eqn. B.28 we can calculate any dipole matrix element using the Wigner symbols and a quantity that can be looked up in the NIST tables (note: $A_{ij} = \Gamma$).

Certain pairs of states have the property that they are a closed decay pair because the excited state can only decay to the lower energy state. Together this pair of states forms a cycling transition. For alkali atoms, a cycling transition occurs between $|i\rangle = |n, 1/2, \tilde{F}, \tilde{F}\rangle$ and $|f\rangle = |n, 3/2, \tilde{F}+1, \tilde{F}+1\rangle$ where n is the ground state principle quantum number and $\tilde{F} = I + 1/2$. The matrix element between these states has the form

$$\begin{aligned} |\langle i | \mathbf{d}_{-1}^{(1)} | f \rangle|^2 &= (2F+1)(2F+3) \left| \begin{pmatrix} F & 1 & F+1 \\ -F & -1 & F+1 \end{pmatrix} \right|^2 \dots \\ &\quad \left| \left\{ \begin{matrix} 1/2 & F & I \\ F+1 & 3/2 & 1 \end{matrix} \right\} \right|^2 ||\mathbf{d}||^2. \end{aligned} \quad (\text{B.30})$$

This can be simplified using the following identities

$$\begin{aligned} \left| \begin{pmatrix} F & 1 & F+1 \\ -F & -1 & F+1 \end{pmatrix} \right|^2 &= \frac{1}{2F+3}, \\ \left| \begin{pmatrix} 1/2 & F & F-1/2 \\ F+1 & 3/2 & 1 \end{pmatrix} \right|^2 &= \frac{1}{4(2F+1)}. \end{aligned} \quad (\text{B.31})$$

Applying these identities to Eqn. B.30 we obtain

$$\begin{aligned} |\langle i | \mathbf{d}_{-1}^{(1)} | f \rangle|^2 &= \frac{(2F+1)(2F+3)}{2F+3} \frac{1}{4(2F+1)} (A\Gamma) \frac{3\hbar c^3 \pi \epsilon_0}{\omega_0^3}, \\ &= \frac{3\hbar c^3 \pi \epsilon_0 \Gamma}{\omega_0^3}. \end{aligned} \quad (\text{B.32})$$

This result is a corollary of the fact that the decay rate from any excited state to all possible ground states is the same.

B.4 Rabi Oscillations and Landau-Zener Transitions

To solve Eqn. B.14 for transitions between atomic states when a near resonance field of frequency ω (with mean photon number N) is applied, we use the RWA Hamiltonian. The RWA Hamiltonian (Eqn. B.20) is block diagonal in the basis $|g, N_{k,\epsilon}\rangle, |e, N_{k,\epsilon}-1\rangle$ (ignoring center-of-mass terms), so we write it out in matrix form and drop the constant $\hbar\omega N$ energy offset

$$\mathbf{H} = \begin{bmatrix} 0 & -i\alpha A_k \sqrt{N} \\ i\alpha^* A_k \sqrt{N} & -\hbar(\omega - \omega_0) \end{bmatrix}. \quad (\text{B.33})$$

We solve for the wavefunction evolution $|\Psi(\mathbf{t})\rangle = \xi_g(\mathbf{t})|g, N\rangle + \xi_e(\mathbf{t})|e, N-1\rangle$ using the Schrödinger equation. Instead of solving for ξ_g and ξ_e directly, a convenient approach is to solve for the density matrix $\boldsymbol{\rho}$, which for pure state evolution is related to $|\Psi\rangle$ by $\boldsymbol{\rho} = |\Psi\rangle\langle\Psi|$. The density matrix is more general than the wavefunction because it can also describe the evolution of statistical ensembles. This will be required when we include non-Hermitian evolution of the state (e.g., due to spontaneous emission) in §B.6. The equations of motion for the atomic density matrix are

$$\begin{aligned} \dot{\boldsymbol{\rho}} &= -\frac{i}{\hbar}[\mathbf{H}, \boldsymbol{\rho}], \\ \dot{\rho}_{ee} &= i\frac{\Omega}{2}\rho_{eg} - i\frac{\Omega^*}{2}\rho_{ge}, \\ \dot{\rho}_{gg} &= -i\frac{\Omega}{2}\rho_{eg} + i\frac{\Omega^*}{2}\rho_{ge}, \\ \dot{\rho}_{eg} = \dot{\rho}_{ge}^* &= i\Delta\rho_{eg} - i\frac{\Omega^*}{2}(\rho_{gg} - \rho_{ee}), \\ \rho_{ee} + \rho_{gg} &= 1, \end{aligned} \quad (\text{B.34})$$

where $\Delta = \omega - \omega_0$ and $\Omega = \frac{2A_k\alpha\sqrt{N}}{\hbar}$. To convert Ω into experimentally measured quantities use Eqn. B.16 or Eqn. B.17.

Starting in the ground state ($\rho_{gg}(0) = 1$) and solving Eqns. B.35 when the field is turned on suddenly at $\mathbf{t} = 0$

$$\rho_{ee}(\mathbf{t}) = \frac{|\Omega|^2}{\Delta^2 + |\Omega|^2} \sin^2 \left(\frac{\sqrt{|\Omega|^2 + \Delta^2}}{2} \mathbf{t} \right). \quad (\text{B.36})$$

Therefore, Ω has a natural interpretation as the on-resonance ($\Delta = 0$) oscillation frequency between atomic states and is referred to as the Rabi frequency. Applying an on-resonance field for $\mathbf{t}_\pi = \pi/|\Omega|$ transfers all the atoms to $|e\rangle$ (“ π -pulse”) and applying for $\mathbf{t}_{\pi/2} = \pi/2|\Omega|$ creates an equal $|g\rangle, |e\rangle$ superposition (“ $\pi/2$ -pulse”). Off-resonance, the effective Rabi frequency increases ($\sqrt{\Delta^2 + |\Omega|^2}$), but the population is only partially transferred to the excited state.

The Rabi frequency and Δ may also be time-dependent. If $\Delta = \gamma\mathbf{t}$, then

$$\mathbf{H} = \begin{bmatrix} 0 & \frac{\hbar\Omega}{2} \\ \frac{\hbar\Omega^*}{2} & -\hbar\gamma\mathbf{t} \end{bmatrix}, \quad (\text{B.37})$$

and as Δ goes to resonance at $\mathbf{t} = 0$ atoms are transferred to the excited state; these are known as Landau-Zener transitions. The fraction of atoms in the ground state at $\mathbf{t} \rightarrow \infty$, if we start with all the atoms in the ground state at $\mathbf{t} = -\infty$, is [299]

$$\rho_{gg} = \exp \left(-\frac{\pi|\Omega|^2}{2\gamma} \right). \quad (\text{B.38})$$

If the detuning changes slowly with respect to the Rabi frequency squared then all the atoms will be transferred from the ground to the excited state. This is equivalent to a π -pulse, but experimentally advantageous for population transfer since it is less sensitive to changes in the resonance frequency during the sweep.

B.4.1 Relating the Density Matrix to the Dipole Operator

The density matrix indicates how the state populations evolve due to the atom-field coupling. Another quantity of interest is the expectation value of the dipole operator $\vec{\mathbf{D}}$. For example, a non-zero electric-dipole expectation value gives rise to a radiated field, which will be discussed further in §B.7.3. The general form of $\vec{\mathbf{D}}$ is

$$\vec{\mathbf{D}} = \sum_q \hat{\sigma}_q \mathbf{D}_q^{(1)}, \quad (\text{B.39})$$

$$\mathbf{D}_q^{(1)} = \langle g | \vec{\mathbf{D}} \cdot \hat{\sigma}_q^* | e \rangle \mathbf{b} + \langle e | \vec{\mathbf{D}} \cdot \hat{\sigma}_q | g \rangle \mathbf{b}^\dagger, \quad (\text{B.40})$$

where $\hat{\sigma}_q$ are the atomic polarizations defined by Eqn. B.22, and \mathbf{b}^\dagger and \mathbf{b} are the raising and lowering operators for the atomic state. The matrix elements $\langle g | \vec{\mathbf{D}} \cdot \hat{\sigma}_q^* | e \rangle$ are detailed in §B.3.1. It is conventional to define the dipole operators as

$$\mathbf{D}_q^- = \langle e | \vec{\mathbf{D}} \cdot \hat{\sigma}_q | g \rangle \mathbf{b}^\dagger, \quad (\text{B.41})$$

$$\mathbf{D}_q^+ = \langle g | \vec{\mathbf{D}} \cdot \hat{\sigma}_q^* | e \rangle \mathbf{b}, \quad (\text{B.42})$$

and

$$\mathbf{D}_q^- \mathbf{D}_q^+ = |\langle e | \vec{\mathbf{D}} \cdot \hat{\sigma}_q | g \rangle|^2 \mathbf{b}^\dagger \mathbf{b}. \quad (\text{B.43})$$

The expectation value of any operator \mathbf{O} is calculated from the density matrix as

$$\langle \mathbf{O} \rangle = \text{Tr}[\boldsymbol{\rho} \mathbf{O}]. \quad (\text{B.44})$$

If we apply Eqn. B.44 using the density matrix from §B.4, which is defined in the basis $|g, N\rangle$, $|e, N-1\rangle$, then

$$\langle \mathbf{O} \rangle = \rho_{eg} \langle g, N | \mathbf{O} | e, N-1 \rangle + \rho_{gg} \langle g, N | \mathbf{O} | g, N \rangle + \rho_{ge} \langle e, N-1 | \mathbf{O} | g, N \rangle + \rho_{ee} \langle e, N-1 | \mathbf{O} | e, N-1 \rangle \quad (\text{B.45})$$

and the off-diagonal terms proportional to ρ_{eg}, ρ_{ge} are zero if the operator does not change the photon number. Since the dipole operator $\vec{\mathbf{D}}$ only operates on the atomic states,

$$\langle \mathbf{D}_q^- \mathbf{D}_q^+ \rangle = |\langle e | \vec{\mathbf{D}} \cdot \hat{\sigma}_q | g \rangle|^2 \rho_{ee}, \quad (\text{B.46})$$

and

$$\langle \mathbf{D}_q^\pm \rangle = 0. \quad (\text{B.47})$$

This is a direct consequence of restricting the field basis to a single Fock state pair. A non-zero dipole expectation value would indicate a classical atomic dipole, which cannot be created by coupling to a non-classical field state (e.g., a single Fock state). Since we are actually driving with a coherent state, we need to include all the coherent state terms when calculating the dipole expectation value. The coherent state (Eqn. B.15) is a superposition of Fock states, so the wavefunction evolution is

$$|\Psi(\mathbf{t})\rangle = e^{-\bar{N}/2} \sum_{n=0}^{\infty} \frac{\bar{N}^{n/2} e^{-i(\omega\mathbf{t}-\phi)n}}{\sqrt{n!}} \left(\frac{\rho_{ge,n}(\mathbf{t})}{\sqrt{\rho_{ee,n}(\mathbf{t})}} |g, n_{\omega,k}\rangle + \sqrt{\rho_{ee,n}(\mathbf{t})} |e, n_{\omega,k-1}\rangle \right) \quad (\text{B.48})$$

where ρ_n are the solutions of Eqn. B.35 with photon number n . Since there are different photon numbers states, we need to keep the phase factor $e^{-i\omega\mathbf{t}n}$ in front of each term (we dropped this term from Eqn. B.33 in §B.4). Using the above form of the wavefunction the expectation value of the dipole operator is

$$\langle \mathbf{D}_q^+ \rangle = \frac{\langle \Psi(\mathbf{t}) | \mathbf{D}_q^+ | \Psi(\mathbf{t}) \rangle}{\langle \Psi(\mathbf{t}) | \Psi(\mathbf{t}) \rangle}, \quad (\text{B.49})$$

$$\simeq \langle e | \vec{\mathbf{D}} \cdot \hat{\sigma}_q | g \rangle e^{-i(\omega\mathbf{t}-\phi)} \rho_{ge,\bar{N}}(\mathbf{t}), \quad (\text{B.50})$$

for $\bar{N} \gg 1$. The \mathbf{D}_q^\pm operator corresponds to a dipole of positive (negative) frequency. The preceding discussion demonstrates that a coherent superposition of Fock states is required to have a non-zero expectation value of the atomic dipole and that the dipole oscillates at the drive frequency. Although we considered all the coherent state terms in this derivation, Eqn. B.50 illustrates that we can determine the dipole operator by solving a single set of equations for the Fock state $|\bar{N}\rangle$. In §B.7.3 we use Eqn. B.50 to calculate the scattered field.

B.5 Dipole Potential (AC Stark Shift)

When there is a non-zero atom-field coupling term the atomic states are not eigenstates of the full Hamiltonian. If the field is turned on slowly, then the atomic states will follow to these new eigenstates. To solve for the energy of these new states, we can diagonalize the RWA Hamiltonian (Eqn. B.33), but this neglects the counter-rotating terms (the last two terms of Eqn. B.19). If we include counter-rotating terms then the Hamiltonian couples all photon Fock states (i.e., the Hamiltonian is not block-diagonal in a finite basis). Nonetheless, for an off-resonant field it is sufficient to calculate the energy shift to second-order in perturbation theory. The second-order ground state ($|g, N\rangle$) energy shift is

$$E_{(g)} = \sum_M \frac{|\langle g, N | \mathbf{H}_{\text{int}} | e, M \rangle|^2}{\hbar(\omega_0 + (M - N)\omega)}. \quad (\text{B.51})$$

The matrix element is only non-zero for $M = N + 1, N - 1$, so

$$E_{(g)} = \frac{|\langle g, N | \mathbf{H}_{\text{int}} | e, N - 1 \rangle|^2}{\hbar(\omega - \omega_0)} - \frac{|\langle g, N | \mathbf{H}_{\text{int}} | e, N + 1 \rangle|^2}{\hbar(\omega_0 + \omega)}. \quad (\text{B.52})$$

Using the notation from the previous section

$$\begin{aligned} E_{(g)} &= \frac{NA_k^2|\alpha|^2}{\hbar(\omega - \omega_0)} - \frac{NA_k^2|\alpha|^2}{\hbar(\omega_0 + \omega)}, \\ &= -\frac{\hbar|\Omega|^2}{4} \left(\frac{1}{\omega_0 - \omega} + \frac{1}{\omega_0 + \omega} \right). \end{aligned} \quad (\text{B.53})$$

To lowest order, the new ground eigenstate is ($\sqrt{N \pm 1} \approx \sqrt{N}$)

$$|g\rangle \rightarrow |g\rangle - \frac{1}{2} \left(\frac{\Omega}{\omega_0 - \omega} + \frac{\Omega^*}{\omega_0 + \omega} \right) |e\rangle. \quad (\text{B.54})$$

The atom-field coupling causes a small mixture with the excited state to shift the ground state energy.

The energy shift (Eqn. B.53) — the “AC Stark shift” — depends on the field strength through the $|\Omega|^2$ prefactor (Eqns. B.16, B.17). If the field strength is spatially varying, then it will form a potential for the atoms. Although both the electric-dipole or magnetic-dipole interaction give rise to this type of potential, typically only an electric field can have a large enough field gradient to impart a significant force (see §C). Therefore, we will only consider electric-dipole potentials in the following discussion.

To calculate the potential for the state $|n, J, F, m_F\rangle$ due to coupling to $|n', J', F' m'_F\rangle$ for an optical field of intensity I , we expand the Rabi rate into the coupling terms

$$\frac{\hbar|\Omega|^2}{4} = \left(\frac{\omega}{2\epsilon_0 V} \right) \sum_q N_q |\langle n, J, F, m_F | \mathbf{d}_q^{(1)} | n', J', F' m'_F \rangle|^2, \quad (\text{B.55})$$

and relate photon number N_q to intensity using Eqn. B.18. The dipole matrix element can be replaced with

Eqns. B.28 and B.29

$$\frac{\hbar|\Omega|^2}{4} = \frac{3c^2\pi\Gamma}{2\omega_0^3}(2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2 (2J'+1) \sum_q I_q \left(\begin{matrix} F & 1 & F' \\ -m_F & q & m'_F \end{matrix} \right)^2 \quad (\text{B.56})$$

where $\hbar\omega_0$ is the energy difference between $|n, J, F, m_F\rangle$ and $|n', J', F' m'_F\rangle$. Then the energy shift of $|n, J, F, m_F\rangle$ is

$$E_{(g)} = -\frac{3c^2\pi\Gamma}{2\omega_0^3}(2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2 (2J'+1) \sum_q I_q \left(\begin{matrix} F & 1 & F' \\ -m_F & q & m'_F \end{matrix} \right)^2 \left(\frac{1}{\omega_0 - \omega} + \frac{1}{\omega_0 + \omega} \right). \quad (\text{B.57})$$

This assumes $|n, J, F, m_F\rangle$ is the lower energy state. The energy shift of $|n', J', F' m'_F\rangle$ is equal and opposite.

B.5.1 Multi-Level

To get the total energy shift of $|n, J, F, m_F\rangle$ we need to sum the energy shifts from all $|n', J', F' m'_F\rangle$ states with non-zero coupling matrix elements. For example, the energy shift of the alkali atom ground state $nS_{1/2}$ is obtained by summing over all P states. If the detuning is larger than the hyperfine splittings inside a particular $n'P_{J'}$ state then we sum Eqn. B.57 over the hyperfine levels F', m'_F after taking the frequency term out of the sum

$$E_{F, m_F} = -\frac{3c^2\pi}{2} \sum_q I_q \sum_{n', J'=\frac{1}{2}, \frac{3}{2}} \frac{\Gamma_{n'P_{J'} \rightarrow nS_{1/2}}}{(\omega_{n'P_{J'}} - \omega_{nS_{1/2}})^3} \left(\frac{1}{\omega_{n'P_{J'}} - \omega_{nS_{1/2}} - \omega} + \frac{1}{\omega_{n'P_{J'}} - \omega_{nS_{1/2}} + \omega} \right) \dots \sum_{F', m'_F} (2J'+1)(2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2 \left(\begin{matrix} F & 1 & F' \\ -m_F & q & m'_F \end{matrix} \right)^2. \quad (\text{B.58})$$

The sum over the Wigner symbols is given by the identity

$$\sum_{F', m'_F} (2J'+1)(2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2 \left(\begin{matrix} F & 1 & F' \\ -m_F & q & m'_F \end{matrix} \right)^2 = \frac{J' + \frac{1}{2} - 2(J'-1)qg_F m_F}{3}, \quad (\text{B.59})$$

and so

$$E_{F, m_F} = -\frac{c^2\pi}{2} \sum_{q, n', J'=\frac{1}{2}, \frac{3}{2}} I_q \frac{\Gamma_{n'P_{J'} \rightarrow nS_{1/2}}}{(\omega_{n'P_{J'}} - \omega_{nS_{1/2}})^3} \left(J' + \frac{1}{2} - 2(J'-1)qg_F m_F \right) \dots \left(\frac{1}{\omega_{n'P_{J'}} - \omega_{nS_{1/2}} - \omega} + \frac{1}{\omega_{n'P_{J'}} - \omega_{nS_{1/2}} + \omega} \right). \quad (\text{B.60})$$

For most laser frequencies, the largest contribution to Eqn. B.60 is from the D1 and D2 transitions ($n' = n$). Keeping only D1 and D2 terms, Eqn. B.60 (for the ground state) reduces to [99],

$$E_{F,m_F} = -\frac{c^2\pi}{2} \sum_q I_q \left[\frac{\Gamma_{3/2}}{\omega_{3/2}^3} (2 - qg_F m_F) \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right) + \frac{\Gamma_{1/2}}{\omega_{1/2}^3} (1 + qg_F m_F) \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) \right]. \quad (\text{B.61})$$

Eqn. B.61 is the general dipole potential we use to describe our optical lattice. For convenience define

$$V_{\text{sc}} = -\frac{c^2\pi}{2} \left[\frac{2\Gamma_{3/2}}{\omega_{3/2}^3} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right) + \frac{\Gamma_{1/2}}{\omega_{1/2}^3} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) \right], \quad (\text{B.62})$$

$$V_{\text{v}} = -\frac{c^2\pi}{2} \left[\frac{\Gamma_{1/2}}{\omega_{1/2}^3} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) - \frac{\Gamma_{3/2}}{\omega_{3/2}^3} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right) \right], \quad (\text{B.63})$$

so that

$$E_{F,m_F} = V_{\text{sc}} I + V_{\text{v}} (I_{\sigma^+} - I_{\sigma^-}) (g_F m_F). \quad (\text{B.64})$$

The first term is the scalar light shift and the second is the vector light shift. The scalar light shift is state independent and proportional to the total intensity, whereas the vector light shift is state dependent and depends on the projection of the intensity onto the atomic basis. In the far-detuned limit $V_{\text{v}} \rightarrow 0$ and $\Delta_{3/2} \simeq \Delta_{1/2}$, so Eqn. B.60 becomes,

$$E_{F,m_F} = -\frac{3c^2\pi I}{2} \sum_{n'} \left[\frac{\Gamma_{n'}}{\omega_{n'}^3} \left(\frac{1}{\Delta_{n'}} + \frac{1}{\omega_{n'} + \omega} \right) \right] \quad (\text{B.65})$$

where

$$\omega_{n'} = \frac{2\omega_{n',3/2} + \omega_{n',1/2}}{3}. \quad (\text{B.66})$$

B.6 Optical Bloch Equations

In the general Hamiltonian (Eqn. B.14) the atom is coupled to all field modes. Although only a specific mode is macroscopically occupied by the applied field, we cannot neglect the empty modes since the excited atomic mode can emit into one of these field modes and return (decay) back to the ground state. To describe the state populations when driving between transitions if there is decay from the excited state we need to add non-Hermitian decay terms to Eqn. B.35 by means of a Liouvillian operator [294]. Here we state these new equations without proof

$$\begin{aligned} \dot{\rho}_{ee} &= i\frac{\Omega}{2}\rho_{eg} - i\frac{\Omega^*}{2}\rho_{ge} - \Gamma\rho_{ee}, \\ \dot{\rho}_{gg} &= -i\frac{\Omega}{2}\rho_{eg} + i\frac{\Omega^*}{2}\rho_{ge} + \Gamma\rho_{ee}, \\ \dot{\rho}_{eg} = \dot{\rho}_{ge}^* &= i\Delta\rho_{eg} - i\frac{\Omega^*}{2}(\rho_{gg} - \rho_{ee}) - \frac{\Gamma}{2}\rho_{eg}. \end{aligned} \quad (\text{B.67})$$

This set of equations describe the average density matrix since individual atoms will have sharp jumps as the excited state decays. The decay term is related to the transition frequency as $\Gamma \propto \omega_0^3$ (see §B.7), so the decay terms are only relevant for states with transition frequencies in the optical range (hence “optical” Bloch equations).

B.6.1 Steady-state Transition Properties

In the steady-state limit (set all derivatives to zero) the excited state fraction is

$$\rho_{ee} = \frac{|\Omega|^2}{\Gamma^2 + 4\Delta^2 + 2|\Omega|^2}, \quad (\text{B.68})$$

and the off-diagonal term is

$$\rho_{eg} = \Omega^* \frac{2\Delta - i\Gamma}{\Gamma^2 + 4\Delta^2 + 2|\Omega|^2}. \quad (\text{B.69})$$

The saturation intensity is defined as

$$I_{sat} = \frac{I\Gamma^2}{2|\Omega|^2}, \quad (\text{B.70})$$

so the excited state fraction is expressed in terms of the dimensionless quantities I/I_{sat} and Δ/Γ

$$\rho_{22} = \frac{1}{2} \frac{\frac{I}{I_{sat}}}{1 + 4\left(\frac{\Delta}{\Gamma}\right)^2 + \frac{I}{I_{sat}}}. \quad (\text{B.71})$$

For $I/I_{sat} \ll 1$ the excited state fraction increases linearly with I , but saturates to $1/2$ as $I/I_{sat} \rightarrow \infty$. We can use Eqn. B.55 to relate the saturation intensity to an atomic matrix element

$$I_{sat} = \frac{\hbar I \Gamma^2}{8} \left(\sum_q \frac{I_q}{2\epsilon_0 c \hbar} |\langle n, J, F, m_F | \mathbf{d}_q^{(1)} | n', J', F' m'_F \rangle|^2 \right)^{-1}. \quad (\text{B.72})$$

For the cycling transition the saturation intensity is (using the identity from Eqn. B.32)

$$I_{sat} = \frac{\hbar \omega_0^3 \Gamma}{12\pi c^2}. \quad (\text{B.73})$$

The decay terms correspond to light scattered out of the driving field. The power scattered out of the field is $P = \rho_{ee} \Gamma \hbar \omega$. This scattering can also be expressed as an absorption

$$\frac{dI}{dz} = -n\sigma I \quad (\text{B.74})$$

where the absorption cross section is ($dI/dz = dI/d\mathbf{t} \cdot d\mathbf{t}/dz$ and $\dot{I} = cP$)

$$\begin{aligned} \sigma &= \frac{P}{nI}, \\ &= \frac{n\rho_{ee}\Gamma\hbar\omega}{nI}, \end{aligned} \quad (\text{B.75})$$

$$= \frac{\hbar\omega\Gamma}{2I_{sat}} \left(1 + 4\left(\frac{\Delta}{\Gamma}\right)^2 + \frac{I}{I_{sat}} \right)^{-1}. \quad (\text{B.76})$$

For $\Delta = 0$ and $I/I_{sat} \ll 1$

$$\sigma = \frac{\hbar\omega\Gamma}{2I_{sat}}, \quad (\text{B.77})$$

which for the cycling transition (substituting Eqn. B.73) is

$$\sigma = \frac{3\lambda^2}{2\pi}. \quad (\text{B.78})$$

B.6.2 Multi-Level OBE




Figure B.1: Several multiple-level decay and driving configurations. The configuration in (c) is difficult to solve using the OBEs, so we split the upper state into two states. This corresponds to neglecting coherent driving effects between the two ground states.

To extend the OBE to multiple levels we add new terms to the density matrix. For the configuration in (a) of Fig. B.1 the density matrix has 9 terms, but we ignore coherences between levels 2 and 3, and 1 and 3, so

$$\begin{aligned} \dot{\rho}_{11} &= i\frac{\Omega}{2}\rho_{12} - i\frac{\Omega^*}{2}\rho_{21} + \Gamma_{13}\rho_{33} + \Gamma\rho_{22}, \\ \dot{\rho}_{22} &= -i\frac{\Omega}{2}\rho_{12} + i\frac{\Omega^*}{2}\rho_{21} - \Gamma_{23}\rho_{22} - \Gamma\rho_{22}, \\ \dot{\rho}_{33} &= \Gamma_{23}\rho_{22} - \Gamma_{13}\rho_{33}, \\ \dot{\rho}_{12} = \dot{\rho}_{21}^* &= -i\Delta\rho_{12} + i\frac{\Omega}{2}\rho_{11} - i\frac{\Omega^*}{2}\rho_{22} - \frac{\Gamma + \Gamma_{23}}{2}\rho_{12}, \\ \rho_{11} + \rho_{22} + \rho_{33} &= 1. \end{aligned} \quad (\text{B.79})$$

In steady state $\Gamma_{13}\rho_{33} = \Gamma_{23}\rho_{22}$, and the equations reduce to two-level OBE equations except the decay rate

is $\Gamma + \Gamma_{23}$. Assuming $\Gamma_{13} \gg \Gamma, \Gamma_{23}$ ($\rho_{11} + \rho_{22} \approx 1$), I_{sat} is

$$\begin{aligned} I_{sat} &= \frac{\hbar\omega_0^3(\Gamma + \Gamma_{23})^2}{12\pi c^2\Gamma}, \\ &= I'_{sat} \left(\frac{\Gamma + \Gamma_{23}}{\Gamma} \right)^2, \end{aligned} \quad (\text{B.80})$$

where I'_{sat} is the two-level value (i.e., $\Gamma_{23} = 0$). The cross-section is

$$\sigma = \frac{3\lambda^2}{2\pi} \frac{\Gamma}{\Gamma + \Gamma_{23}}, \quad (\text{B.81})$$

and the excited state fraction is

$$\rho_{22} = \frac{1}{2} \frac{\frac{I}{I_{sat}}}{1 + 4 \left(\frac{\Delta}{\Gamma + \Gamma_{23}} \right)^2 + \frac{I}{I_{sat}}}. \quad (\text{B.82})$$

If we have multiple driving fields (e.g., (b) and (c) of Fig. B.1), then the equations are significantly more complicated. However, we can obtain approximate solutions for the level populations if we ignore coherences between the driving fields. For example, in (b) the simplified equations are

$$\begin{aligned} \dot{\rho}_{11} &= i\frac{\Omega_1}{2}\rho_{12} - i\frac{\Omega_1^*}{2}\rho_{21} + i\frac{\Omega_2}{2}\rho_{13} - i\frac{\Omega_2^*}{2}\rho_{31} + \Gamma_{13}\rho_{33} + \Gamma\rho_{22}, \\ \dot{\rho}_{22} &= -i\frac{\Omega_1}{2}\rho_{12} + i\frac{\Omega_1^*}{2}\rho_{21} - \Gamma_{23}\rho_{22} - \Gamma\rho_{22}, \\ \dot{\rho}_{33} &= -i\frac{\Omega_2}{2}\rho_{13} + i\frac{\Omega_2^*}{2}\rho_{31} + \Gamma_{23}\rho_{22} - \Gamma_{13}\rho_{33}, \\ \dot{\rho}_{12} = \dot{\rho}_{21}^* &= -i\Delta_1\rho_{12} + i\frac{\Omega_1}{2}\rho_{11} - i\frac{\Omega_1^*}{2}\rho_{22} - \frac{\Gamma + \Gamma_{23}}{2}\rho_{12}, \\ \dot{\rho}_{13} = \dot{\rho}_{31}^* &= -i\Delta_2\rho_{13} + i\frac{\Omega_2}{2}\rho_{11} - i\frac{\Omega_2^*}{2}\rho_{33} - \frac{\Gamma_{13}}{2}\rho_{13}. \end{aligned} \quad (\text{B.83})$$

For the configuration in (c) of Fig. B.1, we create effective two-level systems by splitting the excited state into two separate states. Each new excited state retains the decay paths of the original state. This neglects coherent processes that excite from one ground state and then drive down to the other ground state (e.g., dark states). If the decay rate is larger than the Rabi rate then this is a reasonable description. This approach also allows for easy identification of which excitation path the atom takes to get into the excited

state. The effective OBEs are

$$\begin{aligned}
\dot{\rho}_{11} &= i\frac{\Omega_{13}}{2}\rho_{13b} - i\frac{\Omega_{13}^*}{2}\rho_{3b1} + \Gamma_{13}(\rho_{3a3a} + \rho_{3b3b}), \\
\dot{\rho}_{22} &= i\frac{\Omega_{23}}{2}\rho_{23a} - i\frac{\Omega_{23}^*}{2}\rho_{3a2} + \Gamma_{23}(\rho_{3a3a} + \rho_{3b3b}), \\
\dot{\rho}_{3a3a} &= -i\frac{\Omega_{23}}{2}\rho_{23a} + i\frac{\Omega_{23}^*}{2}\rho_{3a2} - (\Gamma_{13} + \Gamma_{23})\rho_{3a3a}, \\
\dot{\rho}_{3b3b} &= -i\frac{\Omega_{13}}{2}\rho_{13b} + i\frac{\Omega_{13}^*}{2}\rho_{3b1} - (\Gamma_{13} + \Gamma_{23})\rho_{3b3b}, \\
\dot{\rho}_{13b} = \dot{\rho}_{3b1}^* &= -i\Delta_{13}\rho_{13b} + i\frac{\Omega_{13}}{2}\rho_{11} - i\frac{\Omega_{13}^*}{2}\rho_{3b3b} - \frac{\Gamma_{13} + \Gamma_{23}}{2}\rho_{3b3b}, \\
\dot{\rho}_{23a} = \dot{\rho}_{3a2}^* &= -i\Delta_{23}\rho_{23a} + i\frac{\Omega_{23}}{2}\rho_{22} - i\frac{\Omega_{23}^*}{2}\rho_{3a3a} - \frac{\Gamma_{13} + \Gamma_{23}}{2}\rho_{3a3a}, \\
\rho_{11} + \rho_{22} + \rho_{3a3a} + \rho_{3b3b} &= 1.
\end{aligned} \tag{B.84}$$

B.7 Scattering

The decay term introduced into the optical Bloch equations corresponds to the atom scattering a photon from the drive field into the infinite number of empty field modes that the atomic dipole is coupled to in Eqn. B.14. In this section we will calculate this decay rate and also calculate the resulting scattered field. Understanding the direction and coherence of this field is important for imaging (§4.1.1).

B.7.1 Decay Rate

To determine the decay rate of an atom in the excited state we use Fermi's Golden rule starting with the state $|e, 0\rangle$ and decaying to $|g, 1_{\vec{k}, \hat{\epsilon}}\rangle$ (\vec{k} is the scattered photon wavevector and $\hat{\epsilon}$ is the polarization)

$$\Gamma_{|e\rangle \rightarrow |g, 1_{\vec{k}, \hat{\epsilon}}\rangle} = \frac{2\pi}{\hbar} |\langle g, 1_{\vec{k}} | \mathbf{H} | e, 0 \rangle|^2 \delta(\hbar\omega_k - \hbar\omega_0). \tag{B.85}$$

The only non-zero atom-field coupling term in Eqn. B.19 for these states is $\alpha[k] \mathbf{a}^\dagger |g\rangle \langle e|$, and so

$$\Gamma_{|e\rangle \rightarrow |g, 1_{\vec{k}, \hat{\epsilon}}\rangle} = \frac{2\pi}{\hbar} |\langle g, 1_{\vec{k}} | A_k \mathbf{a}_{\omega, k}^\dagger | e, 0 \rangle|^2 \delta(\hbar\omega_k - \hbar\omega_0), \tag{B.86}$$

$$= \frac{2\pi}{\hbar} \left(\frac{\hbar\omega_k}{2\epsilon_0 V} \right) |\langle g | \vec{\mathbf{d}} \cdot \hat{\epsilon} | e \rangle|^2 \delta(\hbar\omega_k - \hbar\omega_0). \tag{B.87}$$

To get the total rate we sum over all \vec{k} and $\hat{\epsilon}$

$$\Gamma_{|e\rangle \rightarrow |g\rangle} = \sum_{\vec{k}, \hat{\epsilon}} \frac{2\pi}{\hbar} \left(\frac{\hbar\omega_k}{2\epsilon_0 V} \right) |\langle g | \vec{\mathbf{d}} \cdot \hat{\epsilon} | e \rangle|^2 \delta(\hbar\omega_k - \hbar\omega_0). \tag{B.88}$$

The \vec{k} sum becomes an integral when $V \rightarrow \infty$ [298]

$$\sum_{\vec{k}} \rightarrow \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \int_0^\infty dk \frac{V k^2}{(2\pi)^3}, \tag{B.89}$$

so ($\omega = ck$)

$$\Gamma_{|e\rangle \rightarrow |g\rangle} = \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \sum_{\hat{\epsilon}} |\langle g | \vec{d} \cdot \hat{\epsilon} | e \rangle|^2 \int_0^\infty dk \frac{V k^2}{(2\pi)^3} \frac{2\pi}{\hbar} \left(\frac{\hbar\omega}{2\epsilon_0 V} \right) \delta(\hbar\omega_k - \hbar\omega_0), \quad (\text{B.90})$$

$$= \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \sum_{\hat{\epsilon}} |\langle g | \vec{d} \cdot \hat{\epsilon} | e \rangle|^2 \int_0^\infty dk \frac{k^2 \omega}{(2\pi)^2} \left(\frac{1}{2\epsilon_0} \right) \delta(\hbar\omega_k - \hbar\omega_0), \quad (\text{B.91})$$

$$= \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \sum_{\hat{\epsilon}} |\langle g | \vec{d} \cdot \hat{\epsilon} | e \rangle|^2 \int_0^\infty d(\hbar\omega) \frac{\omega^3}{\hbar c^3 (2\pi)^2} \left(\frac{1}{2\epsilon_0} \right) \delta(\hbar\omega_k - \hbar\omega_0), \quad (\text{B.92})$$

$$= \frac{\omega_0^3}{2\epsilon_0 \hbar c^3 (2\pi)^2} \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \sum_{\hat{\epsilon}} |\langle g | \vec{d} \cdot \hat{\epsilon} | e \rangle|^2. \quad (\text{B.93})$$

Since the scattered light is TEM polarized, $\hat{\epsilon}$ is a function of ϕ . Writing out Eqn. B.93 in the atomic polarization basis (quantized along z)

$$\Gamma_{|e\rangle \rightarrow |g\rangle} = \frac{\omega_0^3}{2\epsilon_0 \hbar c^3 (2\pi)^2} \left(\sum_q |\langle g | \mathbf{d}_q | e \rangle|^2 \right) \int_0^\pi d\phi \int_0^{2\pi} d\theta \sin(\theta) \alpha_q(\phi) \quad (\text{B.94})$$

where $\alpha_q(\phi)$ is the fraction of q polarized light scattered at angle ϕ . The integral is the total solid angle for q polarized light. Since the total solid angle is 4π and there are 2 TEM polarizations per direction, then the solid angle per q polarized light is $8\pi/3$, and

$$\Gamma_{|e\rangle \rightarrow |g\rangle} = \frac{\omega_0^3}{3\pi\epsilon_0 \hbar c^3} \left(\sum_q |\langle g | \mathbf{d}_q | e \rangle|^2 \right). \quad (\text{B.95})$$

The measured decay rate is the rate for an atom in excited state $|n', J', F', m'_F\rangle$ to decay to any ground state in a certain n, J level. This rate is

$$\Gamma_{|n' J'\rangle \rightarrow |n J\rangle} = \frac{\omega_0^3}{3\pi\epsilon_0 \hbar c^3} \sum_{F, m_F, F', m'_F} \frac{1}{(2J' + 1)(2I + 1)} \left(\sum_q |\langle n, J, F, m_F | \mathbf{d}_q | n', J', F', m'_F \rangle|^2 \right) \quad (\text{B.96})$$

Substituting the result from Eqn. B.28 into the above equation we get

$$\begin{aligned} \Gamma_{|n' J'\rangle \rightarrow |n J\rangle} &= \frac{\omega_0^3}{3\pi\epsilon_0 \hbar c^3} \sum_{F, m_F, F', m'_F, q} \frac{1}{(2J' + 1)(2I + 1)} (2F + 1)(2F' + 1) \left(\begin{array}{ccc} F & 1 & F' \\ -m_F & q & m'_F \end{array} \right)^2 \dots \\ &\quad \left\{ \begin{array}{ccc} J & F & I \\ F' & J' & 1 \end{array} \right\}^2 |\langle n J || \mathbf{d} || n' J' \rangle|^2. \end{aligned} \quad (\text{B.97})$$

Applying the Wigner 3J identity

$$\sum_{m_F, m'_F, q} \left(\begin{array}{ccc} F & 1 & F' \\ -m_F & q & m'_F \end{array} \right)^2 = 1, \quad (\text{B.98})$$

when $F' = F, F \pm 1$, the rate is

$$\Gamma_{|n'J'\rangle \rightarrow |nJ\rangle} = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} \frac{|\langle nJ||\mathbf{d}||n'J'\rangle|^2}{(2J'+1)(2I+1)} \sum_{F,F'} (2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2. \quad (\text{B.99})$$

Applying the Wigner 6J identity

$$\sum_{F,F'} (2F+1)(2F'+1) \left\{ \begin{matrix} J & F & I \\ F' & J' & 1 \end{matrix} \right\}^2 = 2I+1, \quad (\text{B.100})$$

the rate is

$$\Gamma_{|n'J'\rangle \rightarrow |nJ\rangle} = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} \frac{|\langle nJ||\mathbf{d}||n'J'\rangle|^2}{2J'+1}, \quad (\text{B.101})$$

which is the rate to decay from the excited state $|n'J'\rangle$ to the ground state $|nJ\rangle$. For the total decay rate from $|n'J'\rangle$ we must sum over the decay rates to all lower energy states.

B.7.2 Scattering Rate Off-Resonance

In the previous section we derived the decay rate starting in the state $|e\rangle$, however, this is not a typical experimental situation. A common scenario is that $|g\rangle$ and $|e\rangle$ are coupled with an off-resonant field. For example, if we apply a beam for dipole trapping then the initial state is $|i\rangle = \sqrt{\rho_{gg}}|g, N_\omega\rangle + \sqrt{\rho_{ee}}|e, N_\omega - 1\rangle$. The rate to scatter a photon into a new mode \vec{k} and end up in state $|g, N_\omega - 1, 1_{\vec{k}, \epsilon}\rangle$ (using Fermi's golden rule) is

$$\Gamma_{|i\rangle \rightarrow |g, N_\omega - 1, 1_{\vec{k}, \epsilon}\rangle} = \frac{2\pi}{\hbar} |\langle g, N_\omega - 1, 1_{\vec{k}} | \mathbf{H} | (\sqrt{\rho_{gg}}|g, N_\omega\rangle + \sqrt{\rho_{ee}}|e, N_\omega - 1\rangle) |^2 \delta(E_f - E_i), \quad (\text{B.102})$$

$$= \rho_{ee} \left[\frac{2\pi}{\hbar} |\langle g, N_\omega - 1, 1_{\vec{k}} | \mathbf{H} | e, N_\omega - 1 \rangle|^2 \delta(E_f - E_i) \right], \quad (\text{B.103})$$

$$\Gamma_{|i\rangle \rightarrow |g, N_\omega - 1, 1_{\vec{k}, \epsilon}\rangle} = \rho_{ee} \left[\frac{2\pi}{\hbar} |\langle g, N_\omega - 1, 1_{\vec{k}} | \mathbf{H} | e, N_\omega - 1 \rangle|^2 \delta(\hbar\omega_k - \hbar\omega) \right]. \quad (\text{B.104})$$

This rate calculation is similar to before except $\omega_0 \rightarrow \omega$ since the scattered photon is at the frequency of the excitation beam and not the excited state energy. Therefore the off-resonant scattering rate is

$$\Gamma_{|(\sqrt{\rho_{gg}}|g, N_\omega\rangle + \sqrt{\rho_{ee}}|e, N_\omega - 1\rangle) \rightarrow |g, N_\omega - 1, 1_{\vec{k}, \epsilon}\rangle} = \rho_{ee} \Gamma_{|n'J'\rangle \rightarrow |nJ\rangle} \left(\frac{\omega}{\omega_0} \right)^3. \quad (\text{B.105})$$

B.7.3 Properties of the Scattered Light

The physical origin of the scattering process is that the induced atomic dipole acts as a source term in Maxwell's equations. This is a standard problem in any electricity and magnetism textbook [300] and taking the correspondence between the classical solution and the expectation value of the quantum operator⁴

$$\langle \vec{\mathbf{E}}(\vec{x}, \mathbf{t}) \rangle = \frac{k^2}{2\pi\epsilon_0} \frac{e^{ik|\vec{x}|}}{|\vec{x}|} \langle \vec{\mathbf{d}} - (\hat{n} \cdot \vec{\mathbf{d}})\hat{n} \rangle \quad (\text{B.106})$$

⁴As pointed out in [300], the atomic dipole defined here is a factor of 2 larger than the classical dipole defined in [300].

where \hat{n} is a unit vector pointing from the origin. Using the result from §B.4.1 for $\langle \vec{\mathbf{d}} \rangle$, written in terms of the atomic projections q and the density matrix,

$$\langle \vec{\mathbf{E}}_q^+(\vec{x}, \mathbf{t}) \rangle = \frac{k^2}{2\pi\epsilon_0} \frac{e^{ik|\vec{x}|}}{|\vec{x}|} (\hat{\epsilon}_q - (\hat{n} \cdot \hat{\epsilon}_q) \hat{n}) \langle e | \vec{\mathbf{d}} \cdot \hat{\sigma}_q | g \rangle e^{-i(\omega\mathbf{t}-\phi)} \rho_{ge, N}(\mathbf{t}). \quad (\text{B.107})$$

The radiated power per unit solid angle is [300]

$$\frac{dP}{d\Omega} = \langle \vec{\mathbf{E}}^-(\vec{x}, \mathbf{t}) \vec{\mathbf{E}}^+(\vec{x}, \mathbf{t}) \rangle, \quad (\text{B.108})$$

$$= \frac{\omega^4}{4\pi^2 c^3 \epsilon_0} \left\langle \left(\vec{\mathbf{d}}^- \vec{\mathbf{d}}^+ - |\vec{\mathbf{d}}^- \cdot \hat{n}|^2 \right) \right\rangle, \quad (\text{B.109})$$

$$= \frac{\omega^4}{4\pi^2 c^3 \epsilon_0} \sum_q \langle \mathbf{d}_q^- \mathbf{d}_q^+ \rangle (1 - (\hat{n} \cdot \hat{\epsilon}_q)^2). \quad (\text{B.110})$$

Replacing $\langle \mathbf{d}_q^- \mathbf{d}_q^+ \rangle$ with the dipole expectation value calculated in §B.4.1 (Eqn. B.50), the radiated power is

$$\frac{dP}{d\Omega} = \frac{\omega^4}{4\pi^2 c^3 \epsilon_0} \rho_{ee} \sum_q |\langle e | \vec{\mathbf{d}} \cdot \hat{\sigma}_q | g \rangle|^2 (1 - |\hat{n} \cdot \hat{\epsilon}_q|^2). \quad (\text{B.111})$$

If ϕ is the angle with the quantization axis, then $(1 - |\hat{n} \cdot \hat{\epsilon}_0|^2) = \sin^2(\phi)$ and $(1 - |\hat{n} \cdot \hat{\epsilon}_{\pm 1}|^2) = (1 + \cos^2(\phi))/2$. If we are exciting on the cycling transition with σ^+ light, then we can replace the matrix element with Eqn. B.32 and

$$\frac{dP}{d\Omega} = \frac{\omega^4}{4\pi^2 c^3 \epsilon_0} \rho_{ee} \frac{3\hbar c^3 \pi \epsilon_0 \Gamma}{\omega_0^3} \frac{1 + \cos^2(\phi)}{2}, \quad (\text{B.112})$$

$$= \frac{3}{4\pi} \left(\frac{\omega}{\omega_0} \right)^3 (\Gamma \hbar \omega \rho_{ee}) \left(\frac{1 + \cos^2(\phi)}{2} \right). \quad (\text{B.113})$$

The radiated power is twice as strong along the quantization axis versus the perpendicular direction. Integrating over all angles, the total radiated power is

$$P = \left(\frac{\omega}{\omega_0} \right)^3 (\Gamma \hbar \omega \rho_{ee}), \quad (\text{B.114})$$

as expected from the Fermi's golden rule calculation (Eqn. B.105).

For a steady-state driving field close to resonance, we can replace ρ_{eg} in Eqn. B.107 and ρ_{ee} in Eqn. B.114 with the results from §B.6.1, and then the total radiated power is

$$P = \left(\frac{\omega}{\omega_0} \right)^3 (\Gamma \hbar \omega) \frac{|\Omega|^2}{\Gamma^2 + 4\Delta^2 + 2|\Omega|^2}, \quad (\text{B.115})$$

and the electric field is

$$\langle \vec{\mathbf{E}}_q^+(\vec{x}, \mathbf{t}) \rangle = \frac{k^2}{2\pi\epsilon_0} \frac{e^{ik|\vec{x}|}}{|\vec{x}|} (\hat{\epsilon}_q - (\hat{n} \cdot \hat{\epsilon}_q) \hat{n}) \langle e | \vec{\mathbf{d}} \cdot \hat{\sigma}_q | g \rangle e^{-i(\omega\mathbf{t}-\phi)} \left[\Omega^* \frac{2\Delta - i\Gamma}{\Gamma^2 + 4\Delta^2 + 2|\Omega|^2} \right]. \quad (\text{B.116})$$

In the strong drive limit ($|\Omega|^2 \gg \Delta^2, \Gamma^2$) the mean value of the electric field goes to zero, but the radiated

power asymptotes to a finite value — the radiated field is incoherent. The ratio of coherently radiated power to total power is [301]

$$\frac{P_{\text{coherent}}}{P} = \frac{\langle \vec{\mathbf{E}}_q^-(\vec{x}, \mathbf{t}) \rangle \langle \vec{\mathbf{E}}_q^+(\vec{x}, \mathbf{t}) \rangle}{\langle \vec{\mathbf{E}}_q^-(\vec{x}, \mathbf{t}) \vec{\mathbf{E}}_q^+(\vec{x}, \mathbf{t}) \rangle}, \quad (\text{B.117})$$

$$= \frac{|\rho_{eg}|^2}{\rho_{ee}}, \quad (\text{B.118})$$

$$= \frac{4\Delta^2 + \Gamma^2}{\Gamma^2 + 4\Delta^2 + 2|\Omega|^2}, \quad (\text{B.119})$$

$$= 1 - 2\rho_{ee}. \quad (\text{B.120})$$

In the strong-drive limit when $P_{\text{coh}}/P \rightarrow 0$, the radiated field picks up sidebands in the power spectrum at $\pm\Omega$ (the so-called Mollow triplet [301]). The above result is an important consideration for high-resolution fluorescence imaging; resolution is functionally dependent on the coherence of the scattered light.

B.8 Mechanical Effects of Scattering

The full atom-field Hamiltonian (Eqn. B.14) includes the center-of-mass operators. These couple to the field through the phase term in the atom-field interaction, which transfers momentum from the field to the atom and vice-versa. One result of this coupling is that the field can act as a conservative potential to trap or repel the atoms due to the coherent redistribution of momenta among the applied field modes. The form of this potential was detailed in §B.5. Momentum is also transferred to the atom as photons are scattered into previously empty modes. In this section we review how this scattering affects the atomic center-of-mass (mechanical) degrees of freedom.

B.8.1 Heating

The random nature of the scattering direction leads to a kinetic energy increase when the photon exchanges momentum $\hbar k$ with the atom. The atomic momentum after scattering (starting from rest) is $\vec{k} - \vec{k}'$ where \vec{k} is the wavevector of the excitation beam and \vec{k}' is the wavevector of the scattered photon ($|k| = |k'|$). The atomic energy after scattering is

$$\begin{aligned} \Delta E &= \frac{\hbar^2(\vec{k} - \vec{k}')^2}{2m}, \\ &= \frac{\hbar^2(2k^2 - 2\vec{k} \cdot \vec{k}')}{2m}. \end{aligned} \quad (\text{B.121})$$

Integrating over all possible \vec{k}' directions

$$\begin{aligned} \Delta E &= 2 \frac{\hbar^2 k^2}{2m}, \\ &= 2E_R. \end{aligned} \quad (\text{B.122})$$

The rate of energy increase is

$$\dot{E} = 2E_R \rho_{ee} \Gamma \left(\frac{\omega}{\omega_0} \right)^3. \quad (\text{B.123})$$

This calculation ignores heating due to dipole fluctuations if there is a gradient in the electric field. A more detailed formula for the heating rate from a far-detuned field, taking all effects into consideration (assuming ground state $|g\rangle$ and excited states $|e\rangle$), is [168]

$$\begin{aligned} \dot{E} = & \frac{E_R}{\hbar^2} \sum_e \Gamma_{|e\rangle \rightarrow |g\rangle} \left(\frac{\omega}{\omega_e} \right)^3 \left(\frac{1}{\omega_e - \omega} + \frac{1}{\omega_e + \omega} \right)^2 \times \\ & |\langle g | \vec{d} | e \rangle \cdot \vec{E}|^2 \left(1 + \frac{1}{k^2} \left| \frac{\nabla(\langle g | \vec{d} | e \rangle \cdot \vec{E})}{\langle g | \vec{d} | e \rangle \cdot \vec{E}} \right|^2 \right). \end{aligned} \quad (\text{B.124})$$

The matrix elements in the heating rate are similar to the matrix elements for the dipole potential. Summing over the D1 and D2 excited states, assuming the detuning is large compared to the hyperfine splittings, then (see §B.5)

$$\begin{aligned} \dot{E} = & E_R \sum_q \frac{c^2 \pi}{2\hbar} I_q \left(1 + \frac{1}{k^2} \left| \frac{\nabla E_q}{E_q} \right|^2 \right) \left[\left(\frac{\omega}{\omega_{3/2}} \right)^3 \frac{\Gamma_{3/2}^2}{\omega_{3/2}^3} (2 - qg_F m_F) \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right)^2 + \right. \\ & \left. \left(\frac{\omega}{\omega_{1/2}} \right)^3 \frac{\Gamma_{1/2}^2}{\omega_{1/2}^3} (1 + qg_F m_F) \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right)^2 \right]. \end{aligned} \quad (\text{B.125})$$

For convenience define the following

$$H_{\text{sc}} = \frac{c^2 \pi \omega^3}{2\hbar} \left[\frac{2\Gamma_{3/2}}{\omega_{3/2}^6} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right)^2 + \left(\frac{\Gamma_{1/2}}{\Gamma_{3/2}} \right) \frac{\Gamma_{1/2}}{\omega_{1/2}^6} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right)^2 \right], \quad (\text{B.126})$$

$$H_{\text{v}} = \frac{c^2 \pi \omega^3}{2\hbar} \left[\left(\frac{\Gamma_{1/2}}{\Gamma_{3/2}} \right) \frac{\Gamma_{1/2}}{\omega_{1/2}^6} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right)^2 - \frac{\Gamma_{3/2}}{\omega_{3/2}^6} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right)^2 \right], \quad (\text{B.127})$$

so that

$$\dot{E} = E_R \Gamma_{3/2} \left[H_{\text{sc}} I \left(1 + \frac{1}{k^2} \left| \frac{\nabla E}{E} \right|^2 \right) + H_{\text{v}} (g_F m_F) \sum_q q I_q \left(1 + \frac{1}{k^2} \left| \frac{\nabla E_q}{E_q} \right|^2 \right) \right]. \quad (\text{B.128})$$

For a scalar plane wave $E = E_0 e^{ikz}$ ($|\nabla E|^2/k^2|E|^2 = 1$)

$$\dot{E} = 2E_R \Gamma_{3/2} H_{\text{sc}} I_0. \quad (\text{B.129})$$

From Eqn. B.54 (the eigenstate of the atom coupled to a far-detuned beam), we can write Eqn. B.129 in terms of the excited state fractions in the D1 and D2 states

$$\dot{E} = 2E_R \left[\rho_{D2} \Gamma_{3/2} \left(\frac{\omega}{\omega_{3/2}} \right)^3 + \rho_{D1} \Gamma_{1/2} \left(\frac{\omega}{\omega_{1/2}} \right)^3 \right]. \quad (\text{B.130})$$

Therefore, the heating from a scalar plane wave is equal to the recoil heating rate Eqn. B.123. However, in lattice geometries Eqn. B.125 does not equal Eqn. B.123; these geometries are discussed in §C.4.




Figure B.2: Schematic of a MOT (along one direction) for a simplified atom with $F = 0$ in the ground state and $F' = 1$ in the excited state. The image is described in the main text.

B.8.2 Scattering Force - Laser Cooling and MOT

Since each scattering event changes the momentum of the atom, this results in an average force

$$\begin{aligned}
 \vec{F} &= \langle \dot{\vec{P}} \rangle, \\
 &= \hbar \langle (\vec{k} - \vec{k}') \rangle \rho_{ee} \Gamma \left(\frac{\omega}{\omega_0} \right)^3, \\
 &= \hbar \vec{k} \rho_{ee} \Gamma \left(\frac{\omega}{\omega_0} \right)^3.
 \end{aligned} \tag{B.131}$$

The scattering force is negligible unless the driving field is close to resonance ($\omega \approx \omega_0$). Then ρ_{ee} is given by Eqn. B.71 and the force is

$$\vec{F} = \frac{\hbar \vec{k} \Gamma}{2 I_{sat}} \frac{I}{1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}}}. \tag{B.132}$$

If Eqn. B.132 depends on position then the scattering force can trap the atoms. There are several elements of Eqn. B.132 that could depend on position I , I_{sat} , and Δ . Since I_{sat} depends on the polarization of the light (due to selection rules), it can change if the direction of the quantizing field is spatially dependent. The detuning Δ is spatially dependent (due to the Zeeman shift) if there is an inhomogeneous magnetic field. A magneto-optical trap (MOT) uses these two spatially-dependent terms in the scattering force to trap the atoms as shown schematically in Fig. B.2. The MOT requires a linearly changing field $B_z = B'z$. Since our convention is that m_F is defined with respect to the local field direction, the states flip at $z = 0$. As the field magnitude grows the energy levels split further apart due to the Zeeman shift. The atoms can scatter from two counter propagating beams with opposite σ polarizations (in the atomic basis). The beams are detuned in frequency $-\Delta$ from the $B = 0$ transition. For $z > 0$, the σ^- beam, which couples $m_F = 0 \rightarrow m_{F'} = -1$

and has $\hat{k} = -\hat{z}$, is shifted closer to resonance by the magnetic field. The scattering force from this beam pushes the atoms back to $z = 0$. For $z < 0$, the process is flipped. Quantitatively, the force is

$$\vec{F} = -\frac{\hbar k \Gamma I}{2I_{sat}} \left(\frac{1}{1 + 4 \left(\frac{\Delta - B' g_F m_F z}{\Gamma} \right)^2 + \frac{I}{I_{sat}}} - \frac{1}{1 + 4 \left(\frac{\Delta + B' g_F m_F z}{\Gamma} \right)^2 + \frac{I}{I_{sat}}} \right), \quad (\text{B.133})$$

$$= -8\hbar k \frac{I}{I_{sat}} \frac{\Delta}{\Gamma} B' g_F m_F \frac{z}{\left(1 + 4 \left(\frac{\Delta - B' g_F m_F z}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right) \left(1 + 4 \left(\frac{\Delta + B' g_F m_F z}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right)}, \quad (\text{B.134})$$

$$\approx -8\hbar k \frac{I}{I_{sat}} \frac{\Delta}{\Gamma} B' g_F m_F \frac{z}{\left(1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right)^2}. \quad (\text{B.135})$$

The MOT therefore forms a harmonic trap.

The force is also velocity-dependent. For an atom moving at velocity \vec{v} , the laser frequency, in the atomic frame of reference, is $\omega' = \omega + \vec{k} \cdot \vec{v}$ due to the Doppler shift. Ignoring polarizations and field detunings, for two counter-propagating beams the force is

$$\vec{F} = -\frac{\hbar k \Gamma I}{2I_{sat}} \left(\frac{1}{1 + 4 \left(\frac{\Delta - kv}{\Gamma} \right)^2 + \frac{I}{I_{sat}}} - \frac{1}{1 + 4 \left(\frac{\Delta + kv}{\Gamma} \right)^2 + \frac{I}{I_{sat}}} \right), \quad (\text{B.136})$$

$$= -8\hbar k^2 \frac{I}{I_{sat}} \frac{\Delta}{\Gamma} \frac{v}{\left(1 + 4 \left(\frac{\Delta - kv}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right) \left(1 + 4 \left(\frac{\Delta + kv}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right)}, \quad (\text{B.137})$$

$$\approx -8\hbar k^2 \frac{I}{I_{sat}} \frac{\Delta}{\Gamma} \frac{v}{\left(1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right)^2}. \quad (\text{B.138})$$

If $\Delta > 0$ (red-detuned), then this force will damp the atomic motion and lead to cooling. The atom will not damp to zero velocity because of recoil heating (Eqn. B.123). Combining the two effects, the energy change is

$$\dot{E} = \dot{E}_{heat} + F \cdot v, \quad (\text{B.139})$$

$$= 2E_R \frac{I/I_{sat}}{1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}}} \Gamma - 8\hbar k^2 \frac{I}{I_{sat}} \frac{\Delta}{\Gamma} \frac{v^2}{\left(1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right)^2}. \quad (\text{B.140})$$

In steady-state $\dot{E} = 0$ and

$$mv^2 = \left(1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + \frac{I}{I_{sat}} \right) \frac{\hbar \Gamma}{8 \Delta}. \quad (\text{B.141})$$

The minimum occurs when $\Delta = \Gamma/2$ and

$$mv^2 = \frac{\hbar \Gamma}{2}. \quad (\text{B.142})$$

From equipartition $k_B T = mv^2$, so the effective temperature is

$$T = \frac{\hbar\Gamma}{2k_B}, \tag{B.143}$$

which is the Doppler temperature for laser cooling of a two-level atom. For an atom with multiple ground states (e.g., an alkali atom), sub-Doppler laser cooling is possible [\[302\]](#).

Appendix C

Optical Lattice and Optical-Dipole Trap

The atom-field coupling shifts the atomic energy levels (§B.5). This shift can be employed to produce conservative potentials (dipole potentials) for manipulating the atoms. To generate an appreciable force on the atoms, the field must vary over short distances. Therefore these potentials are typically formed using light, which can have feature sizes on the order of $\lambda \simeq 1\mu\text{m}$. An optical-dipole potential can result from a light field with a spatially-dependent intensity and/or polarization (Eqn. B.61). In this appendix we consider both cases.

Optical potentials are a key component to our quantum simulation apparatus. Simple traps are formed by focusing a Gaussian laser beam, which creates a harmonic potential near the center of the beam. We review the properties of Gaussian beams in §C.3. More interesting geometries are obtained by interfering two or more beams, leading to periodic structures with periodicity on the order of λ . Atoms in these “optical lattices” facilitate simulating interesting models in condensed matter physics. In this appendix we summarize the common lattice geometries used in this thesis. We perform a detailed derivation for lin- θ -lin spin-dependent lattices, which has not been reproduced elsewhere. We also discuss some technical details for implementing optical-dipole potentials such as heating in the potential, and calibration and alignment.

C.1 Scalar Optical Lattice

The general dipole potential (Eqn. B.61) depends on the intensity and the polarization (in the atomic basis $\hat{\sigma}^\pm, \hat{\pi}$) of the optical fields. However, if all the fields are $\hat{\pi}$ polarized or if the detunings are large, then we only need to keep the first term of Eqn. B.64 — the scalar light shift. The scalar potential is proportional to the total intensity

$$V(\vec{x}) = V_{\text{sc}} I(\vec{x}) \quad (\text{C.1})$$

where V_{sc} is defined by Eqn. B.62. For optical-lattice potentials, the scalar potential is more widely applicable since it applies to all linearly polarized counter-propagating fields. The scalar optical-dipole potential (Eqn. C.1) is a straightforward starting point to explain most of the interesting features of optical lattices and is valid for many of our experimental lattice configurations. For simplicity, we will consider scalar plane waves $E(\vec{x}) = e^{i\vec{k}\cdot\vec{x} - i\omega t}$, although in practice our lattices are formed from focused Gaussian beams. Gaussian beam lattices are discussed in §C.3.2.




Figure C.1: Scalar lattice setups: (a) two beams intersecting at an angle, (b) retro-reflected beams, and (c) four beams at 90 degrees with different frequencies. (d) and (e) illustrate the time-averaged intensity pattern for two cases of setup (c). In (d), $\Delta_1 = \Delta_2$, $\Delta_3 = \Delta_4$, but $|\Delta_1 - \Delta_3|$ is large. The resulting intensity is the sum of two 1D retro-reflected intensities with a lattice spacing of $\lambda/2$. In (e), $\Delta_1 = \Delta_2 = \Delta_3 = \Delta_4$, so all terms interfere. The lattice spacing is λ along x and y and $\lambda/\sqrt{2}$ along $x = y$.

C.1.1 Two Fields Intersecting an Angle

The canonical optical lattice configuration is two identical fields intersecting at angle θ_L (illustrated in (a) of Fig. C.1). The two dimensional electric field in this configuration is

$$E(x, y) = E_0 e^{i(\cos(\theta_L)y + \omega t)} \left(e^{ik \sin(\theta_L)x} + e^{-ik \sin(\theta_L)x + i\Delta\phi} \right) \quad (\text{C.2})$$

where $\Delta\phi$ is the phase difference between the fields. The potential is proportional to the time-averaged intensity

$$\begin{aligned} \frac{I(x, y)}{c\epsilon_0} &= |E(x, y)|^2, \\ &= |E_0|^2 \left(e^{ik \sin(\theta_L)x} + e^{-ik \sin(\theta_L)x + i\Delta\phi} \right) \left(e^{-ik \sin(\theta_L)x} + e^{ik \sin(\theta_L)x - i\Delta\phi} \right), \\ &= |E_0|^2 \left(2 + e^{2i(k \sin(\theta_L)x + \Delta\phi)} + e^{-2i(k \sin(\theta_L)x + \Delta\phi)} \right), \\ &= 4|E_0|^2 \cos^2(k \sin(\theta_L)x + \Delta\phi), \end{aligned} \quad (\text{C.3})$$

$$= 4|E_0|^2 \cos^2 \left(\frac{\pi x}{d} + \Delta\phi \right), \quad (\text{C.4})$$

where $d = \frac{\lambda}{2 \sin(\theta_L)}$, and the absolute position of the lattice minima is set by $\Delta\phi$. This intensity highlights the basic principles of the optical lattice — when two fields intersect at some angle, the interference creates a sinusoidal intensity pattern with a periodicity of d . The periodicity is controlled by changing either λ or θ_L ; the lower bound is $\lambda/2$ consistent with diffraction theory. When $\theta_L = \pi/2$, the configuration can be obtained by retro-reflecting the beam (illustrated in (b) of Fig. C.1). Since the beams are from the same source there is intrinsic phase stability; any phase shift to the beam before the lattice region does not change the lattice intensity. Only phase shifts in the retro beam path (labeled in Fig. C.1) can affect the lattice position.

C.1.2 Four fields at 90 Degrees with Offset Frequencies

To create a lattice along more than one direction we need additional beams. In the retro-reflected configuration we need four beams to create lattices along two directions. The general arrangement, illustrated in (c) of Fig. C.1, is that all these four beams have different frequencies¹. Assuming $|\Delta| \ll \omega$, so that the wavevectors are all approximately equal, the overall field is

$$E(x, y) = E_0 e^{i\omega t} \left(e^{-iky + i\Delta_1 t} + e^{iky + i\Delta_2 t} + e^{ikx + i\Delta_3 t} + e^{-ikx + i\Delta_4 t} \right). \quad (\text{C.5})$$

¹Although the retro-reflected beam is from the same laser source as the forward beam, an AOM in the retro path can add a frequency shift to the retro beam.

The intensity, time-averaging over oscillating terms of order ω , but leaving time-dependent terms of order Δ , is

$$\begin{aligned}
\frac{I}{c\epsilon_0} &= |E_0|^2 \left(e^{-iky+i\Delta_1 t} + e^{iky+i\Delta_2 t} + e^{ikx+i\Delta_3 t} + e^{-ikx+i\Delta_4 t} \right) \\
&\quad \left(e^{iky-i\Delta_1 t} + e^{-iky-i\Delta_2 t} + e^{-ikx-i\Delta_3 t} + e^{ikx-i\Delta_4 t} \right), \\
&= |E_0|^2 \{ 4 + 2 \cos[2ky + (\Delta_2 - \Delta_1)t] + 2 \cos[k(x+y) + (\Delta_3 - \Delta_1)t] + \\
&\quad \cos[k(x-y) + (\Delta_1 - \Delta_4)t] + 2 \cos[k(x-y) + (\Delta_3 - \Delta_2)t] + \\
&\quad 2 \cos[k(x+y) + (\Delta_2 - \Delta_4)t] + 2 \cos[2kx + (\Delta_3 - \Delta_4)t] \}.
\end{aligned} \tag{C.6}$$

The resulting intensity is a time-dependent lattice. We will discuss a few common choices for Δ .

If all the frequencies are the same ($\Delta = 0$), then the intensity is

$$\frac{I}{c\epsilon_0} = 2|E_0|^2 [2 + \cos(2ky) + 2 \cos(k(x+y)) + 2 \cos(k(x-y)) + \cos(2kx)]. \tag{C.7}$$

All the beams are mutually interfering and this intensity is plotted in (e) of Fig. C.1. The intensity along x at $y = 0$ is

$$\begin{aligned}
\frac{I}{c\epsilon_0} &= 2|E_0|^2 [3 + 4 \cos(kx) + \cos(2kx)], \\
&\quad 4|E_0|^2 [1 + 2 \cos(kx) + \cos^2(kx)], \\
&\quad 4|E_0|^2 [1 + \cos(kx)]^2,
\end{aligned} \tag{C.8}$$

which has a periodicity of λ . The intensity along $x = y$ is

$$\begin{aligned}
\frac{I}{c\epsilon_0} &= 8|E_0|^2 [1 + \cos(2kx)], \\
&\quad 16|\tilde{E}_0|^2 \cos^2(kx)
\end{aligned} \tag{C.9}$$

which has a periodicity of $\lambda/\sqrt{2}$.

If the frequencies between the x and y counter-propagating beam pairs are sufficiently different then the cross terms will time average to zero, and the intensity is

$$\begin{aligned}
\frac{I}{c\epsilon_0} &= |E_0|^2 \{ 4 + 2 \cos[2ky + (\Delta_2 - \Delta_1)t] + 2 \cos[2kx + (\Delta_3 - \Delta_4)t] \} \\
&= 4|E_0|^2 \{ \cos^2[ky + (\Delta_2 - \Delta_1)t/2] + \cos^2[kx + (\Delta_3 - \Delta_4)t/2] \},
\end{aligned} \tag{C.10}$$

which is the sum of the intensities from two retro-reflected beams. Therefore, to superimpose lattice potentials, we can add frequency offsets between any fields that we do not want to interfere. If there is a frequency difference between the reflected beams ($\Delta_2 - \Delta_1 \neq 0$) then the lattice potential is moving with a constant velocity. This effect can be used to study transport of atoms in the lattice. If this frequency detuning is too large, this potential will also time average to zero. Typically, the forward and retro beams are at the same

frequency, so

$$\frac{I}{c\epsilon_0} = 4|\tilde{E}_0|^2 [\cos^2(ky) + \cos^2(kx)], \quad (\text{C.11})$$

which is a standard 2D square lattice with $d = \lambda/2$. To create a 3D lattice we add a third retro-reflected field along the direction in and out of the page. This is how the spin-independent lattice potential is created for experiments in this thesis.

C.2 The Lin- θ -Lin Retro-Reflected Lattice Potential




Figure C.2: Illustrating the Lin- θ -Lin lattice beam setup. Figure reproduced from Ref. [4] (© 2010 IOP Science).

In the general case, all terms in Eqn. B.61 contribute and the full potential is (Eqn. B.64),

$$V = V_{\text{sc}}I + V_v(I_{\sigma+} - I_{\sigma-})(g_F m_F). \quad (\text{C.12})$$

A lattice potential can be produced if the polarization of the lattice beams (in the atomic basis) is spatially dependent. This type of potential occurs for the lin- θ -lin lattice configuration. In this configuration, the lattice beams are retro-reflected and both beams are linearly polarized, but the retro beam polarization is rotated by an angle (θ_P) with respect to the forward lattice beam polarization. The lin- θ -lin lattice setup is illustrated in Fig. C.2. When $\theta_P = 0$, the lin- θ -lin lattice is equivalent to the scalar potential discussed in §C.1.

The total intensity is straightforward to calculate. The intensity of the forward lattice beam with the retro beam blocked is

$$I_0 = c\epsilon_0|E_0|^2, \quad (\text{C.13})$$

and the lattice intensity is

$$\frac{I}{I_0} = |e^{ikz} + \cos(\theta_P)e^{-ikz}|^2 + \sin^2(\theta_P), \quad (\text{C.14})$$

$$= 1 + \cos^2(\theta_P) + 2\cos(\theta_P)\cos(2kz) + \sin^2(\theta_P), \quad (\text{C.15})$$

$$= 2[1 + \cos(\theta_P)\cos(2kz)]. \quad (\text{C.16})$$

When $\theta_P = \pi/2$, there is no scalar lattice potential because the intensity is featureless.

To calculate the vector potential (second term of Eqn. C.12), we solve for the intensity of the field

projected onto $\hat{\sigma}^\pm$. To start, we define \vec{k} to be along z in some coordinate system

$$\vec{k} = [0, 0, 1], \quad (\text{C.17})$$

and define a polarization basis in the plane orthogonal to \vec{k} (the field is TEM)

$$\vec{p}_1 = [1, 0, 0], \quad (\text{C.18})$$

$$\vec{p}_2 = [0, 1, 0]. \quad (\text{C.19})$$

We set the coordinate system by defining the forward beam polarizations as

$$\vec{p}_k = \vec{p}_1. \quad (\text{C.20})$$

Since the retro-polarization is the forward polarization rotated by θ_P

$$\vec{p}_{-k} = \cos(\theta_P)\vec{p}_k + \sin(\theta_P)\vec{p}_{k\perp} \quad (\text{C.21})$$

where $\vec{p}_{k\perp} = \vec{p}_2$. We define the magnetic field direction as rotated from our current coordinate system by angles θ_B and ϕ_B . Given this definition of \vec{B} , the atomic polarizations are

$$\hat{\sigma}^\pm = \frac{1}{\sqrt{2}} [\cos(\phi_B) \cos(\theta_B) \mp i \sin(\phi_B), \sin(\phi_B) \cos(\theta_B) \pm i \cos(\phi_B), -\sin(\theta_B)]. \quad (\text{C.22})$$

Define the projections of the polarizations onto the atomic basis

$$\begin{aligned} \gamma &= \hat{\sigma}^+ \cdot \vec{p}_k, \\ &= |\gamma| e^{i\theta_\gamma}, \\ &= \frac{1}{\sqrt{2}} (\cos(\phi_B) \cos(\theta_B) - i \sin(\phi_B)), \\ \eta &= \hat{\sigma}^+ \cdot \vec{p}_{k\perp}, \\ &= |\eta| e^{i\theta_\eta}, \\ &= \frac{1}{\sqrt{2}} (\sin(\phi_B) \cos(\theta_B) + i \cos(\phi_B)). \end{aligned} \quad (\text{C.23})$$

Since $\hat{\sigma}^- = (\hat{\sigma}^+)^*$, the projections onto $\hat{\sigma}^-$ are γ^* and η^* . From the above definitions, the σ^+ intensity of the lattice field is

$$\begin{aligned} \frac{I_{\sigma^+}}{I_0} &= |\gamma e^{ikz} + (\cos(\theta_P)\gamma + \sin(\theta_P)\eta) e^{-ikz}|^2, \\ &= |\gamma|^2 (1 + \cos^2(\theta_P)) + |\eta|^2 \sin^2(\theta_P) + 2|\gamma||\eta| \sin(\theta_P) \cos(\theta_P) \cos(\theta_\gamma - \theta_\eta) + \\ &\quad 2|\gamma|^2 \cos(\theta_P) \cos(2kz) + 2|\gamma||\eta| \sin(\theta_P) \cos(2kz + \theta_\gamma - \theta_\eta), \end{aligned} \quad (\text{C.24})$$

and the σ^- intensity I_{σ^-} follows by replacing $\theta_\gamma \rightarrow -\theta_\gamma$ and $\theta_\eta \rightarrow -\theta_\eta$ in I_{σ^+}

$$\begin{aligned} \frac{I_{\sigma^-}}{I_0} &= |\gamma|^2(1 + \cos^2(\theta_P)) + |\eta|^2 \sin^2(\theta_P) + 2|\gamma||\eta| \sin(\theta_P) \cos(\theta_P) \cos(\theta_\gamma - \theta_\eta) + \\ &\quad 2|\gamma|^2 \cos(\theta_P) \cos(2kz) + 2|\gamma||\eta| \sin(\theta_P) \cos(2kz - \theta_\gamma + \theta_\eta). \end{aligned} \quad (C.25)$$

The vector potential is proportional to the difference of the $\hat{\sigma}$ intensities (the vector intensity)

$$\frac{I_{\sigma^-} - I_{\sigma^+}}{I_0} = 2|\gamma||\eta| \sin(\theta_P) (\cos(2kz + (\theta_\eta - \theta_\gamma)) - \cos(2kz - (\theta_\eta - \theta_\gamma))), \quad (C.26)$$

$$= 4|\gamma||\eta| \sin(\theta_P) \sin(2kz) \sin(\theta_\gamma - \theta_\eta). \quad (C.27)$$

Because of the projection factor $|\gamma||\eta| \sin(\theta_\gamma - \theta_\eta)$ this appears to contain a complicated dependence on the field direction. However, this factor can be calculated in terms of the angles using Eqn. C.23

$$\begin{aligned} |\gamma||\eta| \sin(\theta_\gamma - \theta_\eta) &= |\gamma||\eta| (\sin(\theta_\gamma) \cos(\theta_\eta) - \cos(\theta_\gamma) \sin(\theta_\eta)), \\ &= \text{Im}(\gamma) \text{Re}(\eta) - \text{Re}(\gamma) \text{Im}(\eta), \\ &= \frac{1}{2} (-\sin(\phi_B) \sin(\phi)_B \cos(\theta_B) - \cos(\phi_B) \cos(\theta_B) \cos(\phi_B)), \\ &= -\frac{1}{2} \cos(\theta_B) (\sin^2(\phi_B) + \cos^2(\phi_B)), \\ &= -\frac{1}{2} \cos(\theta_B). \end{aligned} \quad (C.28)$$

Since we defined $\hat{k} = [0, 0, 1]$ this is $-\frac{1}{2} \hat{k} \cdot \hat{B}$, which is frame independent. Therefore, the vector intensity is

$$\frac{I_{\sigma^-} - I_{\sigma^+}}{I_0} = -2 \sin(\theta_P) \sin(2kz) (\hat{k} \cdot \hat{B}). \quad (C.29)$$

Substituting the scalar and vector intensities into Eqn. C.12, the potential is

$$V = 2V_{\text{sc}} I_0 (1 + \cos(\theta) \cos(2kz)) + 2(g_F m_F) V_v I_0 \sin(\theta_P) \sin(2kz) (\hat{k} \cdot \hat{B}). \quad (C.30)$$

We can rewrite the potential as a single lattice term

$$V = 2V_{\text{sc}} I_0 + \frac{sE_R}{2} \cos \left[2kz + \tan^{-1} \left(\frac{g_F m_F V_v (\hat{k} \cdot \hat{B})}{V_{\text{sc}}} \tan(\theta_P) \right) \right] \quad (C.31)$$

where the lattice depth is

$$sE_R = 4I_0 \sqrt{V_{\text{sc}}^2 \cos^2(\theta_P) + (g_F m_F)^2 V_v^2 (\hat{k} \cdot \hat{B})^2 \sin^2(\theta_P)}. \quad (C.32)$$

In the following we will discuss the potential for specific values of θ_P and λ .

C.2.1 Lin-Lin

When $\theta_P = 0$ the lattice is

$$V = sE_R \cos^2(kz) \quad (\text{C.33})$$

where the lattice depth is

$$sE_R = 4I_0 V_{\text{sc}}, \quad (\text{C.34})$$

$$= -2c^2 \pi I_0 \left[\frac{2\Gamma_{3/2}}{\omega_{3/2}^3} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right) + \frac{\Gamma_{1/2}}{\omega_{1/2}^3} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) \right]. \quad (\text{C.35})$$

The lattice depth is proportional to the total intensity, which is the scalar lattice limit. This applies for any linear polarization, even if the linear polarization is not $\hat{\pi}$. This is a general feature of counter-propagating beams; the vector potentials terms (if any) cancel out between the forward and retro beams.

C.2.2 θ Independent Lattice Depth

When the following condition is met

$$V_{\text{sc}}^2 = (g_F m_F)^2 V_v^2 (\hat{k} \cdot \hat{B})^2, \quad (\text{C.36})$$

then

$$V = 2I_0 V_{\text{sc}} + 2I_0 V_{\text{sc}} \sqrt{\cos^2(\theta_P) + \sin^2(\theta_P) \cos(2kz + \tan^{-1}(\pm \tan(\theta_P)))}, \quad (\text{C.37})$$

$$= 2I_0 V_{\text{sc}} [1 + \cos(2kz \pm \theta)] \quad (\text{C.38})$$

where \pm depends on the sign of the ratio of terms in Eqn. C.31. Changing the polarization angle θ_P continuously shifts the location of the lattice minima. This is the principle behind a number of spin-dependent transport experiments (e.g., [73, 74]). For ^{87}Rb in the $|1 \pm 1\rangle$ state and $\hat{k} \cdot \hat{B} = 1$, the condition given by Eqn. C.36 is satisfied when $\lambda \approx 785\text{nm}$.

C.2.3 Lin- \perp -Lin

When $\theta_P = \pi/2$, the potential is referred to as the lin- \perp -lin (lin-perpendicular-lin) lattice. The lin- \perp -lin lattice potential is

$$\begin{aligned} V &= 2V_{\text{sc}} I_0 + (g_F m_F) 2V_v I_0 (\hat{k} \cdot \hat{B}) \sin(2kz), \\ \frac{V}{c^2 \pi I_0} &= \left[\frac{2\Gamma_{3/2}}{\omega_{3/2}^3} \left(\frac{1}{\Delta_{3/2}} + \frac{1}{\omega_{3/2} + \omega} \right) + \frac{\Gamma_{1/2}}{\omega_{1/2}^3} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) \right] + \\ &\quad (g_F m_F) (\hat{k} \cdot \hat{B}) \left[\frac{\Gamma_{3/2}}{\omega_{3/2}^3} \left(\frac{1}{\Delta_{3/2}} - \frac{1}{\omega_{3/2} + \omega} \right) + \frac{\Gamma_{1/2}}{\omega_{1/2}^3} \left(\frac{1}{\Delta_{1/2}} + \frac{1}{\omega_{1/2} + \omega} \right) \right] \sin(2kz), \end{aligned} \quad (\text{C.39})$$

and the lattice depth and location are both spin dependent. In general, there is a constant offset given by the scalar potential. Depending on the wavelength, this offset can be quite large. For example, at large detunings the vector shift is proportional Δ^{-2} whereas the scalar shift is proportional to Δ^{-1} . When the scalar shift

is large compared to the vector lattice depth, the lin- \perp -lin configuration requires pure polarization. If θ_P deviates slightly from $\pi/2$, a large scalar lattice may be formed. This effect can be mitigated by tuning the wavelength so that $V_{sc} = 0$ ($\lambda \approx 790\text{nm}$ for ^{87}Rb).

C.3 Gaussian Beams




Figure C.3: Schematic of a Gaussian beam propagating along z . In the radial direction the intensity has a Gaussian shape defined by a e^{-2} width of w (the beam waist). The beam focuses to a new waist after a lens (the new waist is given by Eqn. C.45) and always propagates with a Gaussian cross-section. Because of diffraction, the beam expands away from the focus; the distance away from the focus after which waist has expanded by $\sqrt{2}$ is the Rayleigh length z_0 .

In practice, the lattice is not constructed from pure plane waves since these require infinite power for a finite intensity. Instead, real propagating optical fields are beams (i.e. a field where the power is concentrated close to the propagation axis). Beams will eventually diverge due to diffraction. Large beams, which diffract slowly, are used for propagation over long distances, and then are focused at the atoms to create sufficiently high intensities required to generate lattices of appreciable depth. Beams propagating along z are described by the electric field $\vec{E} = \vec{f}(\vec{x})e^{ikz}$ where \vec{f} obeys the paraxial wave equation [100]

$$\frac{\partial^2 \vec{f}}{\partial x^2} + \frac{\partial^2 \vec{f}}{\partial y^2} - 2ik \frac{\partial \vec{f}}{\partial z} = 0 \quad (\text{C.40})$$

The paraxial wave equation is an approximation of the full wave equation when the power is close to the propagation axis.

The lowest-order solution of Eqn. C.40 is a beam with a Gaussian cross-section [100]

$$E(z, r) = \sqrt{\frac{2P}{\epsilon_0 c \pi w(z)^2}} \exp \left[i \left(kz - \tan^{-1}(z/z_0) + k \frac{r^2}{2R(z)} \right) - \frac{r^2}{w(z)^2} \right] \quad (\text{C.41})$$

where $z = 0$ is defined as the focus of the beam and P is the beam power. The function $w(z)$ is the waist and defines the e^{-1} size of the electric field

$$w(z) = w_0 \sqrt{1 + \frac{z^2}{z_0^2}}. \quad (\text{C.42})$$

$R(z)$ is the radius of curvature,

$$R(z) = z + \frac{z_0^2}{z}, \quad (\text{C.43})$$

and z_0 is the Rayleigh length

$$z_0 = \frac{\pi \mathbf{w}_0^2}{\lambda}. \quad (\text{C.44})$$

The only free parameter is \mathbf{w}_0 , which is the beam waist at the focus. Consistent with diffraction theory, the beam expands away from the focus and the waist is $\sqrt{2}$ times larger z_0 away from the focus. The effect of optical elements (e.g., a lens) on the beam are described by the ABCD matrix theory (see e.g., [100]). If we propagate a collimated beam of waist \mathbf{w}_i through a lens of focal length f , the new waist is

$$\mathbf{w}_f = \frac{\lambda f}{\pi \mathbf{w}_i} \frac{1}{\sqrt{1 + \frac{f^2}{z_i^2}}}. \quad (\text{C.45})$$

C.3.1 Dipole Trap

The scalar potential (Eqn. B.62) from a focused Gaussian beam is

$$V(z, r) = V_{\text{sc}} \frac{2P z_0^2}{\pi \mathbf{w}_0^2 (z_0^2 + z^2)} \exp \left[-2 \frac{r^2 z_0^2}{\mathbf{w}_0^2 (z^2 + z_0^2)} \right]. \quad (\text{C.46})$$

If we expand Eqn. C.46 around $z, \rho = 0$ the potential is

$$\frac{V(z, r)}{V_{\text{sc}}} \approx \frac{2P}{\pi \mathbf{w}_0^2} - \frac{4P}{\pi \mathbf{w}_0^4} r^2 - \frac{2P}{\pi \mathbf{w}_0^2 z_0^2} z^2, \quad (\text{C.47})$$

which is a harmonic potential with trap frequencies

$$\omega_r = \frac{2}{\mathbf{w}_0^2} \sqrt{\frac{2V_{\text{sc}} P}{m\pi}}, \quad (\text{C.48})$$

$$\omega_z = \frac{2\lambda}{\pi \mathbf{w}_0^3} \sqrt{\frac{V_{\text{sc}} P}{\pi m}}, \quad (\text{C.49})$$

$$\frac{\omega_r}{\omega_z} = \frac{\mathbf{w}_0}{\lambda} \sqrt{2\pi}. \quad (\text{C.50})$$

Typically $\mathbf{w}_0 \gg \lambda$, so the radial confinement is much stronger than the axial confinement. Characteristic experimental values for the dipole trap depth and trap frequencies are given in §2.1.1. This is an anharmonic potential, so the confinement will change away from the center. The radial frequency decreases as

$$\omega'_r = \omega_r e^{-2r^2/\mathbf{w}_0^2}. \quad (\text{C.51})$$

The frequency at $r = 0.1\mathbf{w}_0$ is 90% of the frequency at $r = 0$.

C.3.2 Gaussian Beam Lattices

To create optical lattices in the experiment we interfere focused Gaussian beams. These Gaussian beam lattices differ in certain ways from the plane waves lattices considered in §C.1. For one, there are new phase terms, however, these are not important since the lattice potential for the atoms is formed close to the focus ($z \ll z_0$). The main Gaussian beam effect is that the intensity decreases radially. If we consider a retro-reflected scalar lattice with a forward beam power P and beam waist \mathbf{w}_1 , and a retro beam power αP

($\alpha \leq 1$) and beam waist \mathbf{w}_2 , then the electric field and intensity are

$$E(z, r) = \sqrt{\frac{2P}{\epsilon_0 c \pi \mathbf{w}_1^2}} e^{-r^2/\mathbf{w}_1^2} e^{ikz} \left[1 + \sqrt{\alpha} \frac{\mathbf{w}_1}{\mathbf{w}_2} e^{-2ikz} e^{-r^2\left(\frac{1}{\mathbf{w}_2^2} - \frac{1}{\mathbf{w}_1^2}\right)} \right], \quad (\text{C.52})$$

$$I(z, r) = \frac{2P}{\pi \mathbf{w}_1^2} e^{-2r^2/\mathbf{w}_1^2} \left[1 + \alpha \frac{\mathbf{w}_1^2}{\mathbf{w}_2^2} e^{-2r^2\left(\frac{1}{\mathbf{w}_2^2} - \frac{1}{\mathbf{w}_1^2}\right)} + 2\sqrt{\alpha} \frac{\mathbf{w}_1}{\mathbf{w}_2} e^{-r^2\left(\frac{1}{\mathbf{w}_2^2} - \frac{1}{\mathbf{w}_1^2}\right)} \cos(2kz) \right]. \quad (\text{C.53})$$

We rewrite the potential in terms of the lattice depth ($s < 0$ for red-detuned, $s > 0$ for blue-detuned)

$$sE_R = V_{\text{sc}} \frac{8P\sqrt{\alpha}}{\pi \mathbf{w}_1 \mathbf{w}_2}, \quad (\text{C.54})$$

$$\frac{V}{E_R} = \frac{s}{4\sqrt{\alpha}} \frac{\mathbf{w}_2}{\mathbf{w}_1} e^{-2r^2/\mathbf{w}_1^2} + \frac{s\sqrt{\alpha}}{4} \frac{\mathbf{w}_1}{\mathbf{w}_2} e^{-2r^2/\mathbf{w}_2^2} + \frac{s}{2} e^{-r^2\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2}\right)} \cos(2kz), \quad (\text{C.55})$$

$$= -\frac{s}{2} e^{-r^2\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2}\right)} + \frac{s}{4\sqrt{\alpha}} \frac{\mathbf{w}_2}{\mathbf{w}_1} e^{-2r^2/\mathbf{w}_1^2} + \frac{s\sqrt{\alpha}}{4} \frac{\mathbf{w}_1}{\mathbf{w}_2} e^{-2r^2/\mathbf{w}_2^2} + s e^{-r^2\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2}\right)} \cos^2(kz). \quad (\text{C.56})$$

The first three terms are due to imperfect mode matching between the forward and retro beams. We can expand these terms to order r^2

$$\frac{V}{E_R} = \frac{s}{2} \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right) r^2 - \frac{s}{2\sqrt{\alpha}} \frac{\mathbf{w}_2}{\mathbf{w}_1^3} r^2 - \frac{s\sqrt{\alpha}}{2} \frac{\mathbf{w}_1}{\mathbf{w}_2^3} r^2 + s e^{-r^2\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2}\right)} \cos^2(kz). \quad (\text{C.57})$$

The last term, the lattice potential, is coupled to the Gaussian envelope of the beam. One effect of the Gaussian envelope is that the lattice depth decreases radially. For example, $20\mu\text{m}$ off-center ($\mathbf{w}_1 = \mathbf{w}_2 = 120\mu\text{m}$), the lattice depth is 5% less than in the middle. In terms of Hubbard parameters, a 5% decrease in lattice depth corresponds to an $\approx 10\%$ increase in t and a 2% decrease in U . The Gaussian envelope also forms a harmonic potential for atoms trapped in the lattice minima. Along the z direction, in the tight-binding limit (§2.3.2), the atomic wavefunction is a Wannier function (Eqn. 2.71). The radial energy shift of the Wannier function is

$$\frac{\epsilon}{E_R} = s e^{-r^2\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2}\right)} \int dz |w(z)|^2 \cos^2(kz), \quad (\text{C.58})$$

For red-detuned lattices, this energy shift is equivalent to a confining harmonic potential with trap frequency

$$\omega^2 = \frac{2sE_R}{m} \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right) \left(1 - \frac{1}{2\sqrt{s}} \right), \quad (\text{C.59})$$

and for blue-detuned lattices the harmonic potential is deconfining with frequency

$$\omega^2 = -\frac{\sqrt{s}E_R}{m} \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right). \quad (\text{C.60})$$

The blue-detuned lattice deconfinement is weaker than the red-detuned confinement because in the blue-detuned lattice the Wannier functions are peaked in the minima of the beam intensity. The overall red-

detuned lattice trap frequency is

$$\frac{m\omega^2}{sE_R} = 2 \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right) \left(1 - \frac{1}{2\sqrt{s}} \right) + \frac{1}{\sqrt{\alpha}} \frac{\mathbf{w}_2}{\mathbf{w}_1^3} + \sqrt{\alpha} \frac{\mathbf{w}_1}{\mathbf{w}_2^3} - \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right). \quad (\text{C.61})$$

The first term is valid for all lin-lin lattices, however, the last three terms are only valid for a $\hat{\pi} - \hat{\pi}$ lattice or large detunings. These terms represent the residual dipole trap formed because of mode mismatch between the forward and retro beams. In general, they will have some vector shift contribution. When the beams are perfectly matched ($\alpha = 1, \mathbf{w}_1 = \mathbf{w}_2$), the last three terms cancel.

For the spin-dependent lin- \perp -lin lattice, the confinement has a slightly different form because the lattice is a combination of red and blue detuning. For ^{87}Rb atoms and $\lambda = 790\text{nm}$, the lattice is exactly half attractive and half repulsive, so the confinement from the lattice term is

$$\omega^2 = \frac{sE_R}{m} \left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right) \left(1 - \frac{1}{\sqrt{s}} \right). \quad (\text{C.62})$$

C.4 Heating

Energy diffusion due to spontaneous scattering (§B.8) is a fundamental source of heating in optical potentials. The energy transferred to the atomic mechanical degrees of freedom from the trapping beam per unit time is given generally by Eqn. B.125. In this section, we will compute this rate for several lattice geometries. For the dipole trap this rate is essentially given by the plane wave heating rate Eqn. B.129. There is also technical heating, e.g., due to intensity and vibration noise [77]. In principle, technical noise can be eliminated, whereas scattering is an unavoidable consequence of the optical potential.

C.4.1 Scalar Lattice Heating

For the case of two lattice beams intersecting at an angle (§C.1.1), the electric field is given by Eqn. C.2 and the intensity by Eqn. C.4. The gradient of the electric field is

$$\begin{aligned} \vec{\nabla} E(x, y) &= E_0 i k e^{ik(\cos(\alpha)y + \sin(\alpha)x)} \left[\sin(\alpha) \left(1 - e^{-2ik \sin(\alpha)x} \right) \hat{x} + \dots \right. \\ &\quad \left. \cos(\alpha) \left(1 + e^{-2ik \sin(\alpha)x} \right) \hat{y} \right], \end{aligned} \quad (\text{C.63})$$

$$|\vec{\nabla} E(x, y)|^2 = 4|E_0|^2 k^2 \left[\sin^2(\alpha) \sin^2(k \sin(\alpha)x) + \cos^2(\alpha) \cos^2(k \sin(\alpha)x) \right]. \quad (\text{C.64})$$

The heating rate is proportional to

$$\left[|E|^2 + \frac{|\nabla E|^2}{k^2} \right] = 4|E_0|^2 \left[\sin^2(\alpha) \sin^2(k \sin(\alpha)x) + (1 + \cos^2(\alpha)) \cos^2(k \sin(\alpha)x) \right], \quad (\text{C.65})$$

$$= 4|E_0|^2 \left[\sin^2(\alpha) + 2 \cos^2(\alpha) \cos^2(k \sin(\alpha)x) \right]. \quad (\text{C.66})$$

The full heating rate is

$$\dot{E} = 4E_R \Gamma_{3/2} I_0 H_{\text{sc}} \left[\sin^2(\alpha) + 2 \cos^2(\alpha) \cos^2(k \sin(\alpha)x) \right] \quad (\text{C.67})$$

where H_{sc} is defined by Eqn. B.126. The ratio of the heating rate to lattice depth in terms of $d = \lambda/2 \sin(\alpha)$, λ , and the lattice depth V_{latt} is

$$\frac{\dot{E}}{V_{\text{latt}}} = E_{R_d} \left[1 + 2 \left(\left(\frac{2d}{\lambda} \right)^2 - 1 \right) \cos^2(k \sin(\alpha)x) \right] \frac{\Gamma_{3/2} H_{\text{sc}}}{V_{\text{sc}}} \quad (\text{C.68})$$

where E_{R_d} is the recoil energy for a photon of $\lambda = 2d$.

Although Eqn. C.68 depends on position, the atoms are essentially localized to particular values of x by the lattice potential. For red-detuned lattices, the atoms are localized to $x = n\pi[k \sin(\alpha)]^{-1}$ (n is an integer), and for blue-detuned lattices $x = (1+2n)\pi[2k \sin(\alpha)]^{-1}$. For ^{87}Rb , Fig. C.4 illustrates the wavelength which minimizes Eqn. C.68 for a lattice of spacing d . When the retro wavelength is close to resonance, the heating rate is lower if we use an angled blue-detuned lattice. However, outside this range it is optimal to always use a retro configuration. The absolute lowest ratio of heating rate to depth is achieved for large red-detunings.

For large-detunings and a retro-reflected lattice Eqn. C.68 simplifies to

$$\frac{\dot{E}}{V_{\text{latt}}} = E_{R_d} \frac{\Gamma H_{\text{sc}}}{V_{\text{sc}}}, \quad (\text{C.69})$$

$$\simeq E_{R_d} \Gamma \left(\frac{\lambda_0}{\lambda} \right)^3 \left| \frac{1}{\Delta} + \frac{1}{\omega_0 + \omega} \right|. \quad (\text{C.70})$$

As $\lambda \rightarrow \infty$, the heating rate to depth can be made arbitrarily small.

C.4.2 Lin- θ -Lin Lattice Heating

For the general lin- θ -lin retro-reflected lattice we need to use the full heating rate (Eqn. B.128)

$$\dot{E} = E_R \Gamma_{3/2} \left[H_{\text{sc}} I \left(1 + \frac{1}{k^2} \left| \frac{\nabla E}{E} \right|^2 \right) + H_v(g_F m_F) \sum_q q I_q \left(1 + \frac{1}{k^2} \left| \frac{\nabla E_q}{E_q} \right|^2 \right) \right]. \quad (\text{C.71})$$

To calculate the vector term use our general notation from §C.2. The square of the gradient for the $\hat{\sigma}^+$ projection of the field is

$$\frac{\nabla E_{\sigma^+}}{E_0} = ik\gamma e^{ikz} - ik(\cos(\theta_P)\gamma + \sin(\theta_P)\eta)e^{-ikz}, \quad (\text{C.72})$$

$$\begin{aligned} \frac{|\nabla E_{\sigma^+}|^2}{|E_0|^2 k^2} &= |\gamma|^2(1 + \cos^2(\theta_P)) + \sin^2(\theta_P)|\eta|^2 - 2|\gamma|^2 \cos(\theta_P) \cos(2kz) - \\ &\quad 2\sin(\theta_P)|\gamma||\eta| \cos(2kz + \theta_\gamma - \theta_\eta) + 2\cos(\theta_P) \sin(\theta_P)|\eta||\gamma| \cos(\theta_\gamma - \theta_\eta). \end{aligned} \quad (\text{C.73})$$




Figure C.4: (a) Wavelength which minimizes Eqn. C.68, the ratio of heating to lattice depth, for a ^{87}Rb scalar lattice of spacing d . (b) Heating rate to depth in the retro-reflected lattice versus the optimal lattice. For these plots we only consider the $D1$ and $D2$ transitions.

Using the result for $I_{\sigma+}$ given by Eqn. C.24

$$\begin{aligned} \frac{1}{|E_0|^2} \left(|E_{\sigma+}|^2 + \frac{|\nabla E_{\sigma+}|^2}{k^2} \right) &= |\gamma|^2(1 + \cos^2(\theta_P)) + |\eta|^2 \sin^2(\theta_P) + 2|\gamma||\eta| \sin(\theta_P) \cos(\theta_P) \cos(\theta_\gamma - \theta_\eta) + \\ &\quad \frac{2|\gamma|^2 \cos(\theta_P) \cos(2kz) + 2|\gamma||\eta| \sin(\theta_P) \cos(2kz + \theta_\gamma - \theta_\eta) +}{|\gamma|^2(1 + \cos^2(\theta_P)) + \sin^2(\theta_P)|\eta|^2 - 2|\gamma|^2 \cos(\theta_P) \cos(2kz) -} \\ &\quad \frac{2 \sin(\theta_P) |\gamma||\eta| \cos(2kz + \theta_\gamma - \theta_\eta) +}{2 \cos(\theta_P) \sin(\theta_P) |\eta||\gamma| \cos(\theta_\gamma - \theta_\eta)}, \end{aligned} \quad (\text{C.74})$$

$$\begin{aligned} &= 2|\gamma|^2(1 + \cos^2(\theta_P)) + 2|\eta|^2 \sin^2(\theta_P) + \\ &\quad 4 \cos(\theta_P) \sin(\theta_P) |\eta||\gamma| \cos(\theta_\gamma - \theta_\eta). \end{aligned} \quad (\text{C.75})$$

Since the same term for $\hat{\sigma}^-$ is given by making the substitution $\theta_\gamma \rightarrow -\theta_\gamma$ and $\theta_\eta \rightarrow -\theta_\eta$

$$\left(|E_{\sigma+}|^2 + \frac{|\nabla E_{\sigma+}|^2}{k^2} \right) = \left(|E_{\sigma-}|^2 + \frac{|\nabla E_{\sigma-}|^2}{k^2} \right) \quad (\text{C.76})$$

and the vector term cancels to zero because of the sum over q .

To calculate the heating for the scalar term start from the electric field,

$$E = E_0(e^{ikz} + \cos(\theta_P)e^{-ikz})\hat{x} + E_0 \sin^2(\theta_P)e^{-ikz}\hat{y} \quad (\text{C.77})$$

where \hat{x} is defined along the forward polarization. The gradient squared is

$$\frac{|\nabla E|^2}{|E_0|^2 k^2} = 1 + \cos^2(\theta_P) - 2 \cos(\theta_P) \cos(2kz) + \sin^2(\theta_P), \quad (\text{C.78})$$

$$= 2(1 - \cos(\theta_P) \cos(2kz)), \quad (\text{C.79})$$

and so

$$\left(|E|^2 + \frac{|\nabla E|^2}{k^2} \right) = 2|E_0|^2(1 + \cos(\theta_P) \cos(2kz)) + 2|E_0|^2(1 - \cos(\theta_P) \cos(2kz)) \quad (\text{C.80})$$

$$= 4|E_0|^2. \quad (\text{C.81})$$

Substituting into the full heat equation

$$\dot{E} = 4E_R \Gamma_{3/2} I_0 H_{sc} \quad (\text{C.82})$$

which is the exact same expression as the heating in the scalar retro-reflected lattice, Eqn. C.67. Heating in the lin- θ -lin lattice is independent of $m_F g_F$, θ_P or $\hat{k} \cdot \hat{B}$. For large detunings, the heating rate scales as Δ^{-2} . However, unlike the scalar lattice depth, the vector lattice depth scales as Δ^{-2} and not as Δ^{-1} . Therefore, the heating rate to lattice depth does not decrease at large detunings.

C.4.3 Impact of Heating on Hubbard Experiments

To understand the relevance of the heating rate derived above (Eqn. C.82) it must be compared against the energy and timescales in our experiments. For simulating the Hubbard model these energy and timescales are associated with combinations of the Hubbard parameters t and U . A full understanding of the timescales is difficult, but an estimated figure of merit is presented in Ref. [36]. It is also possible that the heating rate is not the appropriate benchmark and instead we want to look at the related rate of entropy generation. Calculations to this effect are presented in Ref. [37].

C.5 Calibrating the Lattice with Diffraction




Figure C.5: Calibrating lattice depth using diffraction. (a) Measured ratio of atoms in the $\pm 2\hbar k$ diffraction peaks (inset) to the number of atoms in the central (undiffracted) peak as a function of the pulse time at fixed power. The data is fit to Eqn. C.87 to determine the bandgap, which is related to the lattice depth using a numerical bandstructure calculation. (b) Lattice depth measured using the procedure in (a) for two different pulse powers (in units of photodiode voltage). The slope of the fit is the calibration.

Diffraction is used to calibrate the depth of the lattice to the experimentally measured beam power. After preparing a BEC in the harmonic trap (§3.1) we pulse the lattice along one direction for a short time ($< 100\mu\text{s}$). The atoms are then released and imaged after TOF. Clear diffraction peaks are visible at $\pm 2\hbar k$ (see Fig. C.5). The number of atoms in the diffraction peaks as a function of the pulse time τ is used to measure the lattice depth for a fixed pulse power.

The lattice eigenstates form a bandstructure (§2.3.1) and when we turn on the lattice quickly with respect to the band gap, the atoms are projected into the various bands. Because of parity, a BEC with zero momentum ($p = 0$) can only project into even bands with zero quasimomentum ($q = 0$). We assume that the initial state projects predominantly into the ground ($n = 0$) and second-excited ($n = 2$) bands and therefore we can treat the time evolution as an effective two-level system. The evolution of the state after pulse time τ is

$$|\Psi(\tau)\rangle = \gamma_0|q_0, n_0\rangle + e^{i\omega_{02}\tau}\gamma_2|q_0, n_2\rangle \quad (\text{C.83})$$

where $\hbar\omega_{02}$ is the bandgap between $n = 0$ and $n = 2$ and γ_0, γ_2 are the BEC projections onto the band states $|q_0, n_0\rangle$ and $|q_0, n_2\rangle$. The lattice states can be decomposed into a series of momentum states at $2j\hbar k$ (j is an integer)

$$|n\rangle = \sum_{j=-\infty}^{\infty} \alpha_{n,j} |p = 2j\hbar k\rangle. \quad (\text{C.84})$$

After the pulse, the probability to be in $p = 0$ is

$$\begin{aligned} |\langle p = 0 | \Psi(\tau) \rangle|^2 &= |\gamma_0 \alpha_{0,n_0} + e^{i\omega_{02}t} \gamma_2 \alpha_{0,n_2}|^2, \\ &= |\gamma_0 \alpha_{0,n_0}|^2 + |\gamma_2 \alpha_{0,n_2}|^2 + 2\gamma_0 \gamma_2 \alpha_{0,n_0} \alpha_{0,n_2} \cos(\omega_{02}t), \end{aligned} \quad (\text{C.85})$$

$$= A + B \cos(\omega_{02}t). \quad (\text{C.86})$$

To normalize to the total number we look at the diffraction ratio, the number of atoms in both diffraction peaks at $\pm 2\hbar k$ over the number of atoms in the original peak at $0\hbar k$. We assume that the main momentum components are $0, \pm 2$ so that

$$\frac{N_{2\hbar k} + N_{-2\hbar k}}{N_{0\hbar k}} = \frac{1 - [A + B \cos(\omega t)]}{A + B \cos(\omega t)}. \quad (\text{C.87})$$

Fig. C.5 shows diffraction data fit to Eqn. C.87. A band structure calculation then converts the fit parameter ω into the lattice depth.

The pulse power must be constant from shot-to-shot during the calibration, but the diffraction pulses are too fast for active feedback. For accurate calibration we post-select diffraction runs where the pulse power was within an acceptable range (typically $\pm 5\%$ of the average value). To get the overall calibration we take 2–3 diffraction data sets at different overall pulse powers. Then we fit the lattice depth versus pulse power data to a line with the intercept fixed to zero, as shown in Fig. C.5 (b).

C.6 Alignment and Setup of the Spin-Dependent Lattice




Figure C.6: Optimizing the (a) wavelength and (b) waveplate angle for the spin-dependent lattice

To align the spin-dependent lattice (§3.1.2) we follow the same basic procedure for the spin-independent lattice outlined in [41] except for a few added steps. To start we do a coarse alignment of the forward and retro beams using fluorescence. Next, we do fine alignment of the forward beam by blocking the retro and loading the atoms into the combined dipole plus lattice potential. The forward beam is aligned when the position of the atoms does not change as the lattice beam is turned on. However, for the spin-dependent lattice we cannot test this alignment at $\lambda_{latt} \approx 790\text{nm}$ because there is no dipole force. Instead, we first align at $\lambda = 782\text{nm}$ where the dipole force is strong and attractive, and then align at $\lambda = 785\text{nm}$ to be as close as possible to λ_{latt} . Empirically, the change in alignment over the 5nm change in wavelength is negligible.

Next, we need to determine the exact value of λ_{latt} where the scalar light shift disappears. The most sensitive measurement of this shift is diffraction using the lin-lin lattice. First we align the lin-lin retro beam at $\lambda = 785\text{nm}$ by maximizing diffraction. Then we look at the diffraction ratio versus λ . If the lattice is perfectly lin-lin, it is state-independent. However, since the polarization gets slightly rotated as it goes through the glass vacuum cell, the lattice is not perfectly lin-lin and is therefore slightly state-dependent. Fortunately, the $m_F = 0$ state is only sensitive to the scalar light shift. Therefore, we use the state $|1, 0\rangle$ to probe diffraction versus λ ; the data is shown in Fig. C.6. Fitting the data to a parabola to get the minima, the wavelength is $\lambda = 790.01(1)\text{nm}$. The theoretical value is obtained by setting Eqn. B.62 to zero, ignoring counter-rotating terms,

$$\begin{aligned}\lambda &= \frac{2\lambda_{1/2}\Gamma_{1/2}\lambda_{3/2}^4 + \lambda_{3/2}\Gamma_{1/2}\lambda_{1/2}^4}{2\Gamma_{3/2}\lambda_{3/2}^4 + \Gamma_{1/2}\lambda_{1/2}^4}, \\ &= 789.998(7)\text{nm},\end{aligned}\tag{C.88}$$

using values from [193]. Including the counter rotating terms we get $\lambda = 790.028(7)\text{nm}$. Therefore, our measured value is in good agreement with the theoretical values. Our experimental value also has an unknown systematic error from our wavemeter, which has not been calibrated in several years.

Once we determine the correct wavelength, we put in the quarter-waveplate (see Fig. 3.1) and maximize the diffraction versus waveplate angle. At the optimal angle, we realign the retro to maximize diffraction. These last two steps are iterated in case there is some small cross-talk between the waveplate angle and retro alignment. In practice, the biggest jump in alignment is when the waveplate is first put into the retro path, and further alignment is fairly insensitive to the angle. Typical data for the waveplate optimization is shown in Fig. C.6.

C.7 X-Dipole Trap Alignment and Heating

In the ^{87}Rb apparatus, part of the trapping potential is a crossed dipole trap (“x-dipole”) as described in §3.1. In addition the atoms are also trapped by a magnetic quadrupole potential. The setup is shown in Fig. C.7. The alignment procedure for the forward beam is given in [41], and so here we will assume that the forward beam is aligned and we will summarize the additional steps required to align the crossing beam. Coarse alignment of the crossing beam is the same as the forward beam; resonant light is overlapped along the beam path and the beam is imaged in the cell via fluorescence.




Figure C.7: Schematic of the x-dipole trap setup. The forward beam is sent through the side of the science cell at 45° and focused to the atoms. After the cell, the forward beam is collimated and recycled around before being focused to the atoms at as close to a 90° angle to the forward beam as possible given the physical constraints (crossing beam). Note that the crossing and forward beams are not in the same plane as drawn here. The crossing beam is coming up out of the page and enters the cell from the bottom (see Fig. 3.1). The crossing beam is rotated by a half-waveplate to minimize heating (see main text). The entire setup is only intensity servoed on the forward beam, so the ratio of crossing to forward beam power is fixed.

Fine alignment proceeds in two stages. In stage 1 we load atoms into the x-dipole plus magnetic quadrupole trap and then significantly relax the quadrupole gradient and increase the dipole power (to compensate for gravity). If the beams are not crossing, then the atoms will significantly expand along the axis of the forward beam, which provides almost no confinement. Once the crossing beam starts to become aligned the size of the cloud shrinks and localizes to the crossing region. The trap is free to form anywhere along the forward beam axis, so we shift the crossing beam until the atoms remain near the original position. In the final stage, the ultra-fine alignment, we keep the quadrupole trap on to balance gravity, and measure the dipole trap depth at constant beam power. First, with the crossing beam blocked, we ramp down the dipole trap linearly in $\simeq 3$ s to a depth just below the point where there are no remaining atoms. Then, we unblock the crossing beam and optimize the atom number with alignment. For best results we iterate the final depth (i.e., once the atom number is optimized, repeat the optimization with a lower dipole depth).

To minimize heating, we rotate the polarization of the crossing beam so that there is no residual interference between the beams. Because of the angle that the beam enters the cell, there is significant coupling between crossing beam power at the atoms and the waveplate angle. We cannot simply optimize number at a certain point in the evaporation cycle versus waveplate angle because more atoms may be transferred if the crossing beam has more power. Instead, we looked at cloud width and atom number as a function of the hold time immediately after the transfer into the x-dipole trap for different waveplate angles.

Appendix D

Characterizing Trap Frequencies for the Lattice System

In our experiments, the atoms are trapped in a harmonic potential

$$V(\vec{x}) = \frac{1}{2}M \left[\sum_i \omega_i^2 (\vec{x} \cdot \hat{\omega}_i)^2 + \sum_i \omega_{k_i}^2 [s_i] \left(\vec{x} - (\hat{k}_i \cdot \vec{x}) \hat{k}_i \right)^2 \right], \quad (\text{D.1})$$

which is a combination of a fixed trap (magnetic and/or optical-dipole) and confinement from the lattice beams (as described in §C.3). The fixed trap has principle axes $\hat{\omega}_i$ and trap frequencies ω_i . Each lattice beam provides harmonic confinement in the direction orthogonal to the wavevector \hat{k}_i with a trap frequency $\omega_{k_i}[s_i]$, which depends on the lattice depth (s_i). Calculating the trap frequencies and principle axes of the full harmonic potential (Eqn. D.1), as a function of the lattice depth, is critical for comparing experiment and theory. To do this calculation, we must diagonalize Eqn. D.1, which requires full knowledge of $\hat{\omega}_i, \omega_i$ and $\hat{k}_i, \omega_{k_i}[s_i]$. However, our ability to directly measure these quantities is limited since we can only observe trap oscillations projected onto our 2D imaging plane(s). Additionally, the bare trap frequencies entering into Eqn. D.1 are obscured when the atoms are loaded into the lattice by the effective mass, which is a function of lattice depth, and interaction effects (e.g., §6.1). This appendix will summarize the series of trap frequency measurements that can be used to fully characterize Eqn. D.1. We then compare to our heuristic calculation for the geometric mean trap frequency as a function of lattice depth,

$$\bar{\omega} = \sqrt{(\omega_1 \omega_2 \omega_3)^{2/3} + \frac{8sE_R}{m\mathbf{w}^2}} \quad (\text{D.2})$$

where \mathbf{w} is an empirical beam waist determined via this calculation ($\mathbf{w} = 120\mu\text{m}$).

D.1 Experimentally Measuring the Trap Frequency

To measure the trap frequency we create a BEC, give it a “kick”, and let the BEC oscillate in the trap for time τ . Then we measure at the velocity by imaging after TOF (§4.1.2). A sample oscillation is shown in Fig. D.1. In the following we detail several of the methods for kicking the atoms in the ^{87}Rb apparatus. We use the coordinate system defined in Fig. 3.1.

D.1.1 ^{87}Rb Apparatus Kicks

To start motion of the BEC in the trap we need to provide an impulse to the atoms by applying a uniform force for a short period of time. One method for applying this impulse is to shift the trap position. For a harmonic trap this is equivalent to a uniform force $\vec{F} = m\omega\vec{x}_0$ where \vec{x}_0 is the displacement vector. We have




Figure D.1: Example of a trap oscillation measurement along \hat{z} in the ^{87}Rb apparatus using a Stern-Gerlach coil kick. The curve is a fit to a single frequency sinusoid with amplitude, offset, frequency and phase as free parameters.

several tools at our disposal to create this displacement depending on the particular type of trap.

The Stern-Gerlach coil¹, a single many-turn coil used to create a field gradient in the direction of gravity (\hat{z}), can be pulsed on for a short amount of time. The gradient along \hat{z} from this coil applies a uniform force on the atoms in that direction, proportional to $g_F m_F$. The coil also creates a bias field, which can indirectly provide an impulse to the atoms. If the atoms are initially trapped in a quadrupole field, then the bias field from the AG coil shifts the quadrupole zero. This generates a force on the atoms because they are off-center in the quadrupole field. A typical kick size for the AG Coil is a few amps for a few ms. An example of an AG coil kick can be seen in Fig. D.1. The AG coil is best used for measuring the trap frequency along \hat{z} .

The bias coil, a coil for producing uniform fields along \hat{y} , can be used to induce a kick by also shifting the zero position of the quadrupole field. If the quadrupole field is off then they can both be pulsed on to provide a kick in the \hat{y} direction. If the bias coil is off before the pulse, then the time and current specified in the control program will not be the actual time and current due to windup lag in the servo. Therefore, it is important to directly monitor the coil current during a kick. Specifying “10ms” and “8.5A” in the program empirically works well when atoms are confined in the single-beam plus magnetic quadrupole trap. The imaging coil, a coil for generating a uniform field for imaging along \hat{x} , can be used to generate a similar kick along \hat{x} .

Changing the frequency of the dipole beam servo AOM changes the angle of the dipole beam before it hits the focusing lens which in turn shifts the position of the dipole beam at the atoms. By quickly shifting the position of the beam this imparts a force on the atoms along \hat{z} . An advantage of the dipole trap kick is that unlike all the previous methods which utilize magnetic fields, this kick is state independent. If the

¹This coil is also referred to as the Anti-Gravity (AG) coil.

dipole trap is supporting against gravity (e.g., in the x-dipole trap with the magnetic quadrupole off), then a kick can be generated along \hat{z} by quickly increasing or decreasing the beam power.

D.1.2 Measuring Along a Non-Principle Axis




Figure D.2: Two frequency oscillation after giving a kick along \hat{y} and projecting the velocity along the \hat{x} direction. The curve is a two-frequency fit.

Imaging the atoms after TOF (§4.1.1) projects the 3D motion in the trap onto a 2D imaging plane. In the ^{87}Rb apparatus we can measure projections onto the xy and xz planes. As an example for the type of data we expect, consider the following potential,

$$V = \frac{1}{2}m\omega_z^2 z^2 + \frac{1}{2}m\omega_1^2 \left(\frac{x+y}{\sqrt{2}} \right)^2 + \frac{1}{2}m\omega_2^2 \left(\frac{x-y}{\sqrt{2}} \right)^2. \quad (\text{D.3})$$

If we apply force along \hat{y} then the atoms will start with a non-zero velocity along \hat{y} . The initial conditions are $\vec{x}(\mathbf{t}=0) = 0$ and $\vec{v}(\mathbf{t}=0) = v_0\hat{y}$. If we solve the equations of motion, then

$$z(\mathbf{t}) = 0, \quad (\text{D.4})$$

$$\frac{x(\mathbf{t}) + y(\mathbf{t})}{\sqrt{2}} = \frac{v_0}{\sqrt{2}\omega_1} \sin(\omega_1 \mathbf{t}), \quad (\text{D.5})$$

$$\frac{x(\mathbf{t}) - y(\mathbf{t})}{\sqrt{2}} = -\frac{v_0}{\sqrt{2}\omega_2} \sin(\omega_2 \mathbf{t}), \quad (\text{D.6})$$

so if we project onto yz by imaging we get

$$y(\mathbf{t}) = v_0 \left(\frac{\sin(\omega_1 t)}{\omega_1} + \frac{\sin(\omega_2 t)}{\omega_2} \right). \quad (\text{D.7})$$

$$z(\mathbf{t}) = 0 \quad (\text{D.8})$$

Therefore in a non-principle direction we will see a multi-frequency curve and fitting we can measure two

frequencies. An example of a multi-frequency curve measured experimentally is presented in Fig. D.2. To measure the three principle frequencies of the full trap we take curves similar to the data in Fig. D.2 using different kicks until we identify three unique frequencies.

D.2 Parameterizing the Fixed Trap

The fixed trap consists of a combination of separate magnetic and/or optical potentials. We know the trap axes of these individual traps based on their symmetries and the geometry of the physical setup. To calculate the axes (required for adding in the lattice confinement) and frequencies of the full trap, we need to know the frequencies of these individual traps. For certain traps we can calculate the trap frequencies theoretically, however empirical measurements are preferred. The complication is that the trap frequencies that we measure are a convolution of all the individual traps. To separate out the individual contributions, we need to vary the relative strengths of these traps and model how the full trap frequencies should change.

D.2.1 Example: ^{87}Rb Apparatus Hybrid-Dipole Trap




Figure D.3: Parameterizing the hybrid-dipole trap by fitting the measured principle frequencies to Eqns. D.12, D.13 and D.14 as a function of dipole trap power. For the fit to ω_1 (top-left) we get $A = 1.54(2)$. For the fit to ω_2 (top-right) we get $\omega_A = 16(4)$, $\omega_B = 49(9)$, $A = 1.4(5)$. For the fit to ω_3 (bottom) we get $\omega_B = 39.7(1.4)$, $A = 1.44(6)$. This fit was not sensitive to ω_A so it was set to the value from the ω_2 fit.

The hybrid-dipole trap consists of a magnetic quadrupole field (strong axis along \hat{y}) and a single horizontally propagating dipole beam ($\lambda = 1064\text{nm}$). The quadrupole field is set so that the gradient along \hat{z} just

balances gravity for atoms in the $F = 1, m_F = -1$ state. Therefore, all the confinement along \hat{z} is provided by the dipole trap. The dipole beam sets the position of the trap along \hat{z} , which is away from the quadrupole zero so that there is a finite magnetic field and no Majorana losses. Because we are away from the center of the quadrupole field, the quadrupole field provides harmonic confinement along \hat{x} and \hat{y} . The angle of the dipole beam with the \hat{x} and \hat{y} axes is almost perfectly 45° since the beam is physically constrained by input and output apertures. Therefore, we can model the trap as ($m = 1$)

$$V(\vec{x}) = \frac{1}{2}\omega_A x^2 + \frac{1}{2}\omega_B y^2 + \frac{1}{2}APz^2 + \frac{1}{4}AP(x+y)^2 \quad (\text{D.9})$$

where ω_A, ω_B (the quadrupole trap frequencies) and A are free parameters. A is a function of the dipole beam shape, wavelength and power loss from the point where the power (P) is measured.

We can rewrite Eqn. D.9 into a diagonalized form,

$$U = \frac{1}{2}\omega_1^2 z^2 + \frac{1}{2}\omega_2^2 (\cos \phi x + \sin \phi y)^2 + \frac{1}{2}\omega_3^2 (\sin \phi x - \cos \phi y)^2 \quad (\text{D.10})$$

The physical parameters in Eqn. D.9 are related to $\omega_1, \omega_2, \omega_3$ and ϕ as,

$$\tan(2\phi) = \frac{AP}{\omega_A^2 - \omega_B^2} \quad (\text{D.11})$$

$$\omega_1 = \sqrt{AP} \quad (\text{D.12})$$

$$\omega_2 = \frac{1}{2}\sqrt{\omega_A^2 + \omega_B^2 + AP - \sqrt{(AP)^2 + (\omega_B^2 - \omega_A^2)^2}} \quad (\text{D.13})$$

$$\omega_3 = \frac{1}{2}\sqrt{\omega_A^2 + \omega_B^2 + AP + \sqrt{(AP)^2 + (\omega_B^2 - \omega_A^2)^2}} \quad (\text{D.14})$$

From these equations we see that, as expected, if $P \rightarrow \infty$, the dipole trap sets the principle axes ($\phi = \pi/4$).

In Fig. D.3 we fit the frequencies measured in the experiment, as we vary the dipole beam power P , to Eqns. D.12, D.13, and D.14 with A, ω_A , and ω_B as free parameters. From the fit to the vertical trap frequency we get that $A = 1.54(2)$, and from the two horizontal frequencies we get $A = 1.4(5), 1.44(6)$. This agreement is reasonable and also very good among the two horizontal curves. For the values of ω_A and ω_B there is only one reliable value for ω_A (16(4)Hz) and for ω_B we measure 49(9)Hz from one fit and 39.7(1.4)Hz from the other. Although the error in these parameters is high, we are only measuring these parameters to determine the principle axes (defined by the angle ϕ). Once we have determined the principle axes, the frequencies along those axes are from the empirical trap frequency measurements, which typically have $< 1\%$ error. Ultimately, the final frequencies of the full trap (fixed plus lattice) will be only weakly dependent on ϕ .

D.3 Parameterizing the Lattice Confinement

The lattice is created by retro-reflecting a focused laser beam as discussed in §C.3.2. Therefore, the potential is characterized by the forward beam and retro beam power P and αP ($\alpha \leq 1$), the beam waists w_1 and w_2 ,

the polarizations, and the detuning. For a 1D spin-independent, red-detuned lattice with wavevector along \hat{k}_i , the harmonic confinement is given by Eqn. C.61 (dropping $1/\sqrt{s}$ terms),

$$V_{\text{latt},i}(\vec{x}) = \frac{sE_R}{2} \left[\left(\frac{1}{\mathbf{w}_1^2} + \frac{1}{\mathbf{w}_2^2} \right) + \frac{\mathbf{w}_2}{\sqrt{\alpha}\mathbf{w}_1^3} + \sqrt{\alpha}\frac{\mathbf{w}_1}{\mathbf{w}_2^3} \right] \left(\vec{x} - (\hat{k}_i \cdot \vec{x})\hat{k}_i \right)^2. \quad (\text{D.15})$$

We know \hat{k}_i well from the geometry of the physical setup. The lattice depth s is independently measured using diffraction. Therefore, to fully characterize the lattice confinement we need to measure α , \mathbf{w}_1 , and \mathbf{w}_2 . These are fixed quantities that only need to be measured once.

To measure α we need the ratio of the forward and retro beam powers, however, there are two complications. First, P refers to the power when the beam is in vacuum, so it cannot be directly measured. Second, the retro beam cannot be measured without blocking the forward beam. To overcome to the first problem we measure the power directly before (P_B) and after (P_A) the glass vacuum cell and assume the main loss is due to Fresnel reflections at the four surfaces. Since the Fresnel losses are polarization dependent, they change slightly after each surface as the polarization rotates. However, this effect is small and the forward power is well-approximated as $P_1 = P_B \times \sqrt{P_A/P_B}$. The retro beam travels through several optical elements after the vacuum cell (mirrors and lenses) before being reflected. Losses in these are characterized by measuring the power just before the retro-mirror, P_R . Then the retro power at the atoms is approximately $P_2 = P_R \times P_R/P_A \times \sqrt{P_A/P_B}$. The ratio of retro to forward beam power at the atoms is $\alpha = P_R/P_B \times P_R/P_A$.

To measure \mathbf{w}_1 and \mathbf{w}_2 we can use two more pieces of information. First, we can measure the lattice power required for a certain lattice depth. The lattice depth to power ratio ($\xi = P/sE_R$) is a fixed constant. Then, the product of beam waists is (Eqn. C.54)

$$\mathbf{w}_1\mathbf{w}_2 = \frac{V_{\text{sc}}8\sqrt{\alpha}}{\pi}\xi. \quad (\text{D.16})$$

Next, we can measure the trap frequency with the retro blocked and the power set so that if we unblock the retro we measure a lattice depth s_0 . The trap frequency from the lattice beam with the retro blocked is

$$\omega_{bl}^2 = s_0E_R \frac{\mathbf{w}_1\mathbf{w}_2}{\sqrt{\alpha}\mathbf{w}_1^4}. \quad (\text{D.17})$$

We cannot directly measure ω_{bl} because we need to add the forward lattice beam dipole trap to the fixed trap given by Eqn. D.10. After measuring the principle axes of the full trap we can use the methods of §D.4 to solve for ω_{bl} . Substituting Eqn. D.17 into Eqn. D.15

$$V_{\text{latt},i}(\vec{x}) = \frac{1}{2} \frac{s}{s_0} \omega_{bl}^2 \left[\sqrt{\alpha} \frac{\mathbf{w}_1}{\mathbf{w}_2} + 1 + \sqrt{\alpha} \left(\frac{\mathbf{w}_1}{\mathbf{w}_2} \right)^3 + \alpha \left(\frac{\mathbf{w}_1}{\mathbf{w}_2} \right)^4 \right] \left(\vec{x} - (\hat{k}_i \cdot \vec{x})\hat{k}_i \right)^2, \quad (\text{D.18})$$

so we just need to scale the blocked frequency measurement by s/s_0 and a fixed (with respect to lattice depth) proportionality constant. The ratio of waists is

$$\left(\frac{\mathbf{w}_1}{\mathbf{w}_2} \right)^2 = \left(\frac{\pi}{8V_{\text{sc}}\alpha\xi} \right) \left(\frac{s_0E_R}{\omega_{bl}^2} \right) \quad (\text{D.19})$$

If we have perfect mode matching ($\alpha = 1, \mathbf{w}_1 = \mathbf{w}_2$), then

$$V_{\text{latt},i}(\vec{x}) = \frac{1}{2} \frac{s}{s_0} (2\omega_{bl})^2 \left(\vec{x} - (\hat{k}_i \cdot \vec{x}) \hat{k}_i \right)^2, \quad (\text{D.20})$$

and the confinement can be determined from a single blocked retro measurement.

D.4 Trap Plus Lattice Confinement

Using the fixed trap axes and frequencies ($\{\omega_i, \hat{\omega}_i\}$) given by Eqn. D.10 and the lattice wavevectors and trap frequencies ($\{\hat{k}_i, \omega_{k_i}(s_i)\}$) we can determine the new trap axes and frequencies as a function of the lattice depth using Eqn. D.1. A coordinate system defined along the lattice wavevectors $\tilde{x} = (\hat{k}_1 \cdot \hat{x})$, $\tilde{y} = (\hat{k}_2 \cdot \hat{x})$, and $\tilde{z} = (\hat{k}_3 \cdot \hat{x})$ is natural to use because these are the principle harmonic axes in the $s \rightarrow \infty$ limit. We can write the full harmonic potential in matrix form as

$$V = \frac{1}{2} \begin{pmatrix} \tilde{x} & \tilde{y} & \tilde{z} \end{pmatrix} \begin{pmatrix} \frac{\partial V_{\text{trap}}^2}{\partial \tilde{x}^2} + (\omega_{k2}^2 + \omega_{k3}^2) & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{x} \partial \tilde{y}} & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{x} \partial \tilde{z}} \\ \frac{\partial V_{\text{trap}}^2}{\partial \tilde{y} \partial \tilde{x}} & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{y}^2} + (\omega_{k1}^2 + \omega_{k3}^2) & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{y} \partial \tilde{z}} \\ \frac{\partial V_{\text{trap}}^2}{\partial \tilde{z} \partial \tilde{x}} & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{z} \partial \tilde{y}} & \frac{\partial V_{\text{trap}}^2}{\partial \tilde{z}^2} + (\omega_{k1}^2 + \omega_{k2}^2) \end{pmatrix} \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix} \quad (\text{D.21})$$

where V_{trap} is the fixed trap. The eigenvalues of this matrix are the principle trap frequencies squared.

D.4.1 Example: ^{87}Rb Apparatus Hybrid-Dipole and $\lambda = 812\text{nm}$ Lattice

In the ^{87}Rb apparatus the lattice beam wavevectors are given by Eqn. 3.4 and the fixed trap potential is given by Eqn. D.10. Using the wavevectors as the coordinate system the full potential is

$$\begin{aligned} U = & \frac{1}{2} \omega_1^2 \left(\frac{1}{2}(\tilde{x} + \tilde{y}) + \frac{1}{\sqrt{2}}\tilde{z} \right)^2 + \frac{1}{2} \omega_2^2 \left(\frac{1}{\sqrt{2}} \cos \phi (\tilde{y} - \tilde{x}) + \sin \phi \left(\frac{1}{2}(\tilde{x} + \tilde{y}) - \frac{1}{\sqrt{2}}\tilde{z} \right) \right)^2 + \dots \\ & \frac{1}{2} \omega_3^2 \left(\frac{1}{\sqrt{2}} \sin \phi (\tilde{y} - \tilde{x}) - \cos \phi \left(\frac{1}{2}(\tilde{x} + \tilde{y}) - \frac{1}{\sqrt{2}}\tilde{z} \right) \right)^2 + \dots \\ & \frac{1}{2} (\omega_{k1}^2 + \omega_{k2}^2) \tilde{z}^2 + \frac{1}{2} (\omega_{k1}^2 + \omega_{k3}^2) \tilde{y}^2 + \frac{1}{2} (\omega_{k2}^2 + \omega_{k3}^2) \tilde{x}^2. \end{aligned} \quad (\text{D.22})$$

For a dipole power of $P = 500\text{mW}$ the measured fixed trap frequencies are $\omega_1 = 2\pi \times 34.8\text{Hz}$, $\omega_2 = 2\pi \times 27\text{Hz}$, and $\omega_3 = 2\pi \times 48.65\text{Hz}$. From the fits to the model potential (Eqn. D.9) as a function of dipole power (Fig. D.3), we infer that $\phi = -0.329$.

To determine the lattice parameters we use the procedure outlined in §D.3. To measure the frequency of the individual retro-blocked lattice beams we invert Eqn. D.22. The results are summarized in Table D.1.

Using the fixed trap and lattice parameters we calculate $\bar{\omega}$ as a function of the lattice depth in Fig. D.4. We compare to the heuristic geometric mean frequency (Eqn. D.2) with $\mathbf{w} = 120\mu\text{m}$. The heuristic calculation is good to within 2% for $s \leq 15$ and 3% for $s \leq 30$. Therefore, the full calculation validates the heuristic approach.

	P/s (mW/ E_R)	α	ν_{meas} ($s = 20$)	ν_{latt}	ν_{calc} ($s = 20$)	$w_1(\mu m)$	$w_2(\mu m)$
Beam #1	1.57	0.46	{36.1, 52.1, 59.5}	40.62	{36.9, 51.5, 59.9}	108	170
Beam #2	1.15	0.61	{27.9, 41.4, 51.7}	21.72	{28.6, 41.2, 50.1}	136	114
Beam #3	1.88	0.88	{36.9, 40, 52.7}	26.9	{36.1, 39.9, 53.5}	138	219

Table D.1: Lattice results for $\lambda = 812.03nm$. P is the forward beam power at the atoms and α is P_R/P where P_R is the retro beam power at the atoms. ω_{meas} are the trap frequencies measured with the lattice on at $s = 20$ and the retro blocked. ω_{latt} is the lattice beam frequency with the retro blocked at $s = 20$ as calculated from Eqn. D.22 to get the total trap frequencies ω_{meas} . For comparison, ω_{calc} are the calculated trap frequencies given ω_{latt} . From the powers and trap frequency we calculate the forward and retro beam waists w_1 and w_2 .




Figure D.4: (a) Calculation of the principle trap frequencies versus lattice depth using Eqn. D.22 and the experimentally measured trap and lattice parameters summarized in Table D.1. These frequencies are compared to the naive calculation, Eqn. D.2 using $w = 120\mu m$. (b) The error in the naive calculation versus lattice depth. By definition the error must be zero when there is no lattice. For the range of lattice depths considering in this thesis, the error is less than 2%.

Appendix E

K $4S \rightarrow 5P$ Laser System and Magnetic-Optical Trap

To achieve single-site imaging (§3.2.10, §4.1.5) for our ^{40}K Fermi-Hubbard simulator (§3.2), we plan to image on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition of potassium at 404.5nm (see state diagram, Fig. A.1). By imaging on a transition with a smaller wavelength, versus the standard D2 transition at 767nm, we can improve the limit of our resolution by almost a factor of two. However, to image at 405nm requires developing a new set of locked lasers at this novel wavelength. Fortunately, this task is simplified by the availability of inexpensive GaN diode lasers in the 395-410nm range. This appendix describes our laser system capable of generating 30mW of single-mode, single-frequency light at 405nm.

Resolution is not the only consideration for single-site imaging. We also need to have high signal-to-noise, which implies collecting as many scattered photons as possible in comparison to spurious counts from e.g., detector dark counts and background scattering. Even absent these sources, there is fundamental photon shot noise. To collect a large number of photons we need high scattering rates and long imaging times. High scattering rates are obtained by using transitions with large linewidths. The D2 transition has the largest linewidth $\Gamma_{4S \rightarrow 4P} = 2\pi \times 6.04\text{MHz}$, which is more than thirty times larger than the linewidth of the $4S \rightarrow 5P$ transition at 405nm, $\Gamma_{4S \rightarrow 5P} = 2\pi \times 185\text{kHz}$ (the total linewidth of the $5P$ state from all decay channels is $\Gamma = 2\pi \times 1.19\text{MHz}$). On the other hand, narrow linewidths are advantageous for long imaging times since we need to simultaneously cool with the excitation light. For the simplest two-level cooling theory (§B.8.2) the minimum temperature T_D (the Doppler temperature), is $\frac{\hbar\Gamma}{2k_B}$ (Eqn. B.143). This suggests a significant temperature improvement for 405nm cooling. Yet, small Γ is not always beneficial because in many alkali atoms, such as ^{87}Rb , the D2 cooling is not limited to T_D due to sub-Doppler processes. However, in ^{40}K sub-Doppler is weak [172, 303] or difficult to observe [24]. Since the lattice depth during imaging needs to be at least several times the laser cooling temperature, 405nm imaging on the $4S \rightarrow 5P$ transition should greatly reduce the power demands of the lattice and allow for longer imaging times.

To test our premise for imaging and cooling on the $4S \rightarrow 5P$ transition we explored 405nm laser cooling and magneto-optical trapping of ^{40}K in free space with our vapor cell MOT (§3.2.6). In addition, the predicted cooling benefits may be useful as an enhanced cooling step after the collection MOT stage to increase the initial phase-space density before evaporative cooling. In fact, this is a standard technique for earth alkaline atoms, which have narrow, cycling transitions ($\Gamma < \text{kHz}$), due to intercombination lines and electric-dipole forbidden transitions. However, unlike narrow line cooling in earth alkaline atoms,

The work presented in this appendix was published in: D. C. McKay, D. Jervis, D. J. Fine, J. W. Simpson-Porco, G. J. A. Edge, and J. H. Thywissen. *Low-temperature high-density magneto-optical trapping of potassium using the open $4S \rightarrow 5P$ transition at 405 nm*. Phys. Rev. A 84, 063420 (2011). Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Figs. E.2, E.13, E.14, E.15, E.17, and E.18 are reproduced from this publication.

$4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F = 11/2$ in ^{40}K is an open and cascading transition. The $5P_{3/2}$ state can decay via a three-photon cascade to the $F = 7/2$ ground state, which cannot be excited by the same laser-cooling light. Still cooling has been demonstrated on the nS ground state to the $(n+1)P$ excited state in two systems, He^* [89,90] and ^6Li [91]. In the case of He^* , a MOT on the $2S \rightarrow 3P$ transition at 389 nm, which is cascading but closed, was shown to have a lower two-body loss rate β , lower re-absorption rates, and a larger cooling force per recoil, resulting in increased density. No reduction in temperature was observed since the He^* 2P and 3P excited states have the same lifetime. In the case of ^6Li , a MOT on the open $2S \rightarrow 3P$ transition at 323 nm did have a reduced temperature, but not an increased density. Unlike lithium and metastable helium, potassium has a D-state decay channel that could perturb laser cooling more significantly (see Fig. A.1).

In §E.3 we will describe our cooling and MOT experiment and results. To summarize we start by accumulating atoms using a standard D2 MOT ($4S_{1/2} \rightarrow 4P_{3/2}$) which has a high capture velocity and volume. Then we transfer to a blue MOT for $\approx 30\text{ms}$ of cooling and compression. We observe temperatures as low as $63(6)\mu\text{K}$ and a 10-fold increase in density. For optimal parameters this translates to an almost 20 times increase in phase-space density.

E.1 $4S \rightarrow 5P$ Calculation

To understand cooling and imaging on the $4S \rightarrow 5P$ transition, in particular to estimate the power required for our laser setup, we performed a multi-level optical Bloch equation (OBE) calculation; an overview is discussed in §B.6.2. In the calculation we include all the electronic levels indicated in Fig. A.1 and their corresponding $(2J+1)(2I+1)$ hyperfine levels — a total of 234 individual $|n, J, F, m_F\rangle$ states. The OBE's form a set of linear equations, $\dot{\mathbf{y}} = \mathbf{A}\mathbf{y}$. To solve, we diagonalize the matrix \mathbf{A} so that $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$ where \mathbf{D} is diagonal, and then $\mathbf{y}(\mathbf{t}) = \mathbf{P}e^{\mathbf{D}\mathbf{t}}\mathbf{P}^{-1}\mathbf{y}_0$. To get the steady-state populations we let $\mathbf{t} \rightarrow \infty$.

To calculate the transition properties (Table A.2), we include a single excitation beam into the calculation from $4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F' = 11/2$. Since the transition is not closed we also need a repump beam which connects $4S_{1/2}, F = 7/2 \rightarrow 4P_{3/2}, F' = 9/2$. We define the steady-state ρ_{ee} as the total population in all $5P_{3/2}$ states and then we calculate transition properties assuming the two-level formulas. For example, we define I_{sat} as the intensity required to get 35% of the population into the excited state, which is equivalent to the two-level definition (Eqn. B.71). To calculate the cross-section we use Eqn. B.75 with the calculated ρ_{ee} where Γ is the total scattering rate $2\pi \times 1.19\text{MHz}$.

For cooling and trapping calculations we include two counter-propagating 405nm cooling beams and two counter-propagating 767nm repump beams. As discussed in §B.6.2, we split the excited states so that the two counter-propagating beams couple to different states. This also allows us to separate excitations caused by different beams, which is important for determining the scattering force¹. The state schematic for the calculation is shown in Fig. E.1. The dashed excitation beams are coming from the left and the solid excitation beams from the right.

¹We ignore the dipole force and coherences that can form between ground states which together can lead to sub-Doppler cooling mechanisms. This approximation is valid when calculating the force curves at large velocities and displacements. It may also be valid in all regions because sub-Doppler cooling may be interrupted by the depolarizing effects of the three-photon cascade. An exact calculation of these effects is beyond the treatment presented in this thesis.




Figure E.1: State scheme for the $4S \rightarrow 5P$ cooling and trapping calculation. The solid blue (red) and dashed blue (red) lines indicate the cooling (repump) excitation beams. The solid grey lines are the possible decay paths. For the purposes of the calculation the $5P_{3/2}$ and $4P_{3/2}$ states have been divided into several quasi-states as described in §B.6.2. See Fig. A.1 for the transition rates and state energies.

Using the state labeling from Fig. E.1, the total force on the atom is

$$\vec{F} = \hbar ((n_a - n_b)k_{405}\Gamma_{5P_{3/2}} + (n_{aa} - n_{bb})k_{767}\Gamma_{4P_{3/2}}) \quad (\text{E.1})$$

where k is the wavevector $2\pi/\lambda$, and the total recoil heating is

$$\dot{E} = \dot{E}_{abs} + \dot{E}_{decay}, \quad (\text{E.2})$$

$$\dot{E}_{abs} = \frac{\hbar^2}{2M} ((n_a + n_b)k_{405}^2\Gamma_{5P_{3/2}} + (n_{aa} + n_{bb})k_{767}^2\Gamma_{4P_{3/2}}), \quad (\text{E.3})$$

$$\dot{E}_{decay} = \frac{1}{2Mc^2} \sum_{i,j < i} n_i \Gamma_{i \rightarrow j} (E_i - E_j)^2, \quad (\text{E.4})$$

where i is a sum over all states and j is a sum over all states to which i can decay. To calculate the force as a function of the atomic velocity v , we change the effective frequency of the dashed excitation beams in Fig. E.1 to be $\omega' = \omega - kv$ and the solid excitation beams to be $\omega' = \omega + kv$. To calculate the spatially dependent force due to a magnetic field gradient, we solve for the force at a fixed field by changing the energy of each m_F level due to the Zeeman shift (§A.1). Then the force is $\vec{F}(x) = \vec{F}(B = \alpha x)$. We can do a similar calculation for the 767nm MOT by only including states up to $4P$ and having four effective $4P_{3/2}$ states (2 cooling, 2 repump). The force curves for both MOTs are shown in Fig. E.2.




Figure E.2: Calculated force curves (in units of Mg) versus (a) position at $v = 0$ and (b) velocity at $x = 0$ for excitation on the $4S \rightarrow 5P$ (blue) and $4S \rightarrow 4P$ (red) transitions. For both curves the gradient is 10G/cm, the intensity per beam is $I_{sat}/5$ (see Table A.2) and the detuning is $-\Gamma/2$. For the 405nm data there is a repump beam at 767nm with 10% of the 404.5nm intensity. Near zero velocity, the damping rate is $\gamma = 1.8 \times 10^4 \text{ s}^{-1}$ for the red and $\gamma = 6.9 \times 10^4 \text{ s}^{-1}$ for the blue. Near zero displacement, the spring constant corresponds to an undamped trap frequency $\omega/(2\pi) = 730 \text{ Hz}$ for the red and $\omega/(2\pi) = 650 \text{ Hz}$ for the blue.




Figure E.3: Temperature versus detuning for cooling on the $4S \rightarrow 5P$ transition. The black curve is the full OBE calculation for $I = I_{sat}/50$ (the repump intensity is 10% of the cooling beam intensity). The red curve and blue curves are from the two-level Doppler theory, Eqn. B.141 with $\Gamma = 2\pi \times 1.19\text{MHz}$. The red curve is for $I \ll I_{sat}$ and the blue curve is for $I = I_{sat}$. Grey lines highlight T_D and the Doppler detuning.

Steady-state laser cooling occurs when the heating and cooling rates are equal,

$$\dot{E} = -F(v)v. \quad (\text{E.5})$$

Near $v = 0$, $F(v) = -\gamma v$, so

$$v^2 = \frac{\dot{E}}{\gamma}. \quad (\text{E.6})$$

From equipartition $T = Mv^2/k_B$, and so

$$T = \frac{M}{k_B} \frac{\dot{E}}{\gamma}. \quad (\text{E.7})$$

In general, T is a function of the detuning and intensity of the cooling and repump beams. At low intensities, we plot the temperature as a function of detuning in Fig. E.3 and compare it to the two-level theory, Eqn. B.141. We see that the calculated temperature is slightly lower than the two-level theory because recoil heating from the three-photon cascade is less than the heating from a single photon kick at the same total energy. From the calculation the lowest temperature is $23.5\mu\text{K}$ at $\Delta = -\Gamma/2$. This detuning coincides with the detuning at the lowest temperature in the two-level theory (the Doppler detuning). To demonstrate the effect of high intensity on temperature, we also plot temperature versus detuning at $I = I_{sat}$ (for the two-level theory). The minimum temperature is approximately twice as high and occurs at a larger value of the detuning.

E.1.1 Ionization




Figure E.4: For each K state, up to $5P_{3/2}$, we list the wavelength corresponding to a photon with the minimum ionization energy.

Fig. E.4 lists the wavelength of a photon with the minimum photoionize energy for the various states of K. In a D2 MOT, where atoms are in the $4S$ and $4P$ levels and there is only 767nm light, there is no photoionization; 767nm photons are not sufficient to ionize the $4P$ state. However, in the $4S \rightarrow 5P$ MOT, the 405nm cooling light can ionize all excited state levels and the 767nm repump light can ionize atoms in

the $5P$ level.

The photoionization rate from state i due to a beam of intensity I and frequency ω is [304]

$$R = \rho_{ii}\sigma_i \left(\frac{I}{\hbar\omega} \right) \quad (\text{E.8})$$

where σ_i is the photoionization cross-section for state i and ρ_{ii} is the state population. There are no measured cross-sections for the ionization of the $5P$ states of K, just for the equivalent state, $6P$, of Rb [305–307]. However, the K $5P$ and Rb $6P$ values should be comparable since the ionization cross-section for the K $4P$ state [308,309] is within a factor of 2–3 of the Rb $5P$ cross-section [304]. Using $\sigma = 2 \times 10^{-18} \text{cm}^2$, measured for Rb $6P$ at $\lambda = 350 \text{nm}$ [307], we estimate the 404.5nm ionization rate of $5P_{3/2}$ to be 0.5s^{-1} if all the atoms are in the $5P$ state ($I_{404.5} = 100 \text{mW/cm}^2$). Using $\sigma = 1.5 \times 10^{-17} \text{cm}^2$, measured for Rb $6P$ at $\lambda = 690 \text{nm}$ [307], and we estimate the 767nm ionization rate of $5P_{3/2}$ to be 0.33s^{-1} ($I_{766} = 5 \text{mW/cm}^2$). Since we cool for $< 100 \text{ms}$ these ionization rates should not be an issue and in §E.3 this conclusion is experimentally confirmed.

E.2 Laser System and Setup

For imaging, cooling and trapping we have developed a 405nm laser system based on GaN laser diodes capable of delivering up to 30mW of single-mode, single-frequency light to the experiment. GaN laser diodes have been previously used before for K spectroscopy [283,310–314], but never for K cooling and trapping. The overall system is illustrated in Fig. E.5.

E.2.1 Master Laser and Saturation Spectroscopy

The source of single-frequency light is a Toptica DL 100 external cavity diode laser² (ECDL) with $P=10 \text{mW}$ output. The spectrum of the master laser measured with a scanning cavity (§E.2.3) shows a single peak with a linewidth below the cavity resolution. After a few years of operation, small peaks appeared in the spectrum which were removed by re-aligning the feedback.

The schematic of the laser setup is shown in Fig. E.5. Part of the master laser light is used to lock the laser to a ^{39}K saturation spectroscopy feature (see Fig. E.6). For saturation spectroscopy we use a natural abundance K vapor cell that is 15cm long and heated to 100°C inside a thermally insulated box. We heat the cell to increase the K density, which is proportional to the vapor pressure. According to [190] the vapor pressure increases a factor of $\approx 10^{3.5}$ per 100°C at $T = 50^\circ\text{C}$. Higher temperatures are required for saturation spectroscopy on $4S \rightarrow 5P$ than $4S \rightarrow 4P$ ($T \approx 50 - 60^\circ\text{C}$) because the $4S \rightarrow 5P$ cross-section in the Doppler broadened gas is nearly 300 times smaller. To obtain an error signal we frequency modulate the pump beam at f_{mod} using an acousto-optical modulator (AOM) and then demodulating the probe signal at f_{mod} using the mixer circuit shown in §H.3 (modulation transfer spectroscopy). The error signal height is temperature optimized, as shown in Fig. E.6. If the cell is overheated, the optical density becomes too high.

⁴⁰K features are not observed in the error signal because of the low natural abundance of ⁴⁰K. Therefore, we lock to the closest possible feature, the $4S_{1/2}, F = 1 \rightarrow 5P_{3/2}$ transition of ³⁹K. The ³⁹K excited state

²Diffraction grating feedback in Littrow configuration.




Figure E.5: Schematic of the 405nm laser system for generating 30mW of single-mode, locked laser light to the experiment. The single-frequency source is a Toptica ECDL (master laser). The light out of the master laser is split two ways. Part of the light goes to a modulation transfer spectroscopy setup for locking to the $4S_{1/2}, F = 1 \rightarrow 5P_{3/2}$ transition of ^{39}K . The frequency is shifted by a series of AOMs so that the maser laser is near-resonant with the $4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F' = 11/2$ transition of ^{40}K . The remaining master laser light is used to injection lock a high power diode ($P = 100\text{mW}$). To monitor the injection we use a homebuilt scanning cavity (§E.2.3). The injected diode light is sent through a single-mode polarization maintaining fiber to the MOT setup.




Figure E.6: (a) $4S_{1/2} \rightarrow 5P_{3/2}$ saturation spectroscopy in a heated, natural abundance potassium cell. The Doppler envelope has been subtracted off and the signal filtered to illustrate the saturation spectroscopy features. The three main peaks are from ^{39}K and the smaller peaks are from ^{41}K . We lock to the ^{39}K $F = 1$ peak, which is the closest ^{39}K feature to the ^{40}K transitions, which are not visible. (b) Optimizing the cell temperature using the peak-peak value of the error signal.

hyperfine manifold is not resolved. To lock the ECDL master laser, we feedback to the current and PZT of the Toptica controller using homebuilt servo electronics (see §H.1 for the schematic). When the laser is locked to the $F = 1$ ^{39}K peak, it is $\approx 500\text{MHz}$ red-detuned from the ^{40}K $4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F = 11/2$ transition. To compensate for this difference we shift the frequency of the light with AOMs. These are standard TeO_2 AOMs with AR coatings that are only designed to go down to 450nm. Because of this, there is some power loss from the coatings and the damage threshold at 404.5nm is quite low. Therefore, we avoid using these AOMs after the high power diode to prevent damage. Empirically, a few mW of the master laser light focused with an $f = 100\text{mm}$ lens has not resulted in any damage.

E.2.2 Testing the Master Laser — Probing $4S \rightarrow 5P$ in the 767nm MOT

To make a ^{40}K MOT on the $4S \rightarrow 5P$ transition our cooling laser has to be narrow, stable and at the correct absolute frequency. To achieve stability we lock to the $F = 1$ saturation spectroscopy feature of ^{39}K and then shift to the ^{40}K $4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F' = 11/2$ transition frequency using AOMs. However, the isotope shift³ between these transitions is not a priori known. For the $4S_{1/2} \rightarrow 5P_{1/2}$ transition, the isotope shift has been measured to be 235(2)MHz [313] and 207(13)MHz [315]. We can estimate the $4S_{1/2} \rightarrow 5P_{3/2}$ isotope shift by taking the $4S_{1/2} \rightarrow 4P_{3/2}$ shift 125.6MHz [316], and then multiply by the ratio of the transition wavelengths 767/404.5, to get 238MHz. To get the exact shift in our setup we need to probe ^{40}K atoms. Probing atoms can also test our linewidth, which has to be on the order of the $5P_{3/2}$ linewidth ($2\pi \times 1.19\text{MHz}$) for cooling to be effective. An important test of the master laser linewidth and frequency shift is whether we can probe ^{40}K atoms and resolve the $5P_{3/2}$ hyperfine structure (see Fig. A.1 and Fig. E.8).

To test the master laser and determine the isotope shift we probe a ^{40}K $4S \rightarrow 4P$ MOT (§3.2.6) with the 405nm master laser. The setup is shown in Fig. E.7. We send the laser through the MOT and collect all fluorescence, 767nm and any possible 405nm, on a photodiode. Using the same AOM that shifts the 405nm frequency, we amplitude modulate the beam at 1kHz. Then we use a lock-in amplifier to detect any

³The isotope shift is defined as the shift of the bare transition frequency (no hyperfine splittings) between different isotopes.




Figure E.7: Laser setup for probing the $4S \rightarrow 5P$ transition with the 767nm MOT. Master laser light is sent through a double-pass AOM setup, which increases the frequency by approximately 400MHz and amplitude modulates the beam at 1kHz . Then the light is sent via a single-mode fiber through a ^{40}K MOT operating on the $4S \rightarrow 4P$ transition at 767nm . Fluorescence from the MOT is collected on a photodiode and a lock-in detector is used to measure the 1kHz component of the photodiode signal. Data is obtained for various frequencies of the 405nm light by scanning F1 (see data in Fig. E.8).

photodiode signal that appears at 1kHz.




Figure E.8: Results of probing the $^{40}\text{K } 4S_{1/2}, F = 9/2 \rightarrow 5P_{3/2}, F'$ transition. (a) The state structure being probed. For reference we show the ^{39}K locking line (light blue). Frequencies are given with respect to the bare $^{40}\text{K } 4S_{1/2} \rightarrow 5P_{3/2}$ transition assuming an isotope shift of 240MHz. (b) The lock-in signal (inverted) as a function of the frequency of the probe beam with respect to the locking transition. The data is given by the solid blue line and the dashed black line is a three Gaussian fit.

The results of this probe experiment are shown in Fig. E.8. The signal indicates an isotope shift of $\approx 240\text{MHz}$, similar to our estimate. To reach the transition frequency required for a MOT, $F = 9/2 \rightarrow F' = 11/2$, we need to shift the laser $\approx 510\text{MHz}$ from the $^{39}\text{K}, F = 1$ reference line. The measured probe signal is not due to 405nm fluorescence since the peaks vanish if we repeat the experiment with a blue filter over the photodiode. Instead, the signal appears to be a decrease (quench) in the 767nm fluorescence as the probe beam comes on resonance. One explanation is that atoms are pumped into the dark $F = 7/2$ ground state by the 405nm light and therefore these atoms cannot fluoresce until pumped back into $F = 9/2$ by the repump light. The plot also puts an upper bound of $\approx 3\text{MHz}$ on the laser linewidth. The actual linewidth is likely much less as the signal is significantly broadened from the magnetic field gradients required for the MOT.

E.2.3 Scanning Cavity

To characterize the fine ($<1\text{GHz}$) features of the laser spectra we use a scanning cavity. This consists of two spherical mirrors in confocal configuration⁴, with a piezo element to scan the cavity length. As we scan the cavity length, a single-frequency laser will produce a sharp peak in the cavity signal (a photodiode after the cavity) at a specific cavity length. If we change the laser frequency by the free-spectral range $c/4L$, the peak will appear at the same value of the cavity length. Similarly, a single-frequency will reappear as a new peak after the cavity length changes by $\lambda/4$. The full width at half maximum (FWHM) of the peak is set by the cavity finesse F where $\text{FWHM} = c(2LF)^{-1}$ and $F = \pi\sqrt{R}(1 - R)^{-1}$ where R is the mirror reflectivity.

⁴In confocal configuration, the cavity length is equal to the mirror radius of curvature R_{curv} , and all cavity modes are resonant at the same cavity length.




Figure E.9: (a) Custom made 405nm scanning cavity. Mirrors (#1) are Layertec 100599, $R_{curv} = 50\text{mm}$, $D = 0.5\text{''}$ spherical mirrors which are HR coated at 370-410 and 700-825nm ($R > 99.8\%$ at 400nm and 800nm). The cavity is scanned by applying voltage to a ring piezo (#2, PiezoMechanik HPSt 500/15-8/5). The signal is measured by a 0.5'' tube mounted SMA photodiode (#3, Thorlabs SM05PD1B). The Thorlabs SA-201 cavity controller is used to drive the piezo and amplify the photodiode signal. All mounting parts are from Thorlabs; the cavity bars are 6mm diameter stainless steel rods (Thorlabs ER series). (b) Cavity trace of the Toptica master laser scanning the piezo over a 30V range. The two peaks correspond to the free-spectral range (FSR) of the cavity, 1.5GHz. Zooming in on a peaks, we see that there is some asymmetry, likely due to imperfect alignment and/or the cavity length slightly off from the confocal configuration. Fitting the peak to a Gaussian we measure a FWHM of 11MHz giving a finesse of approximately 270.

When we started this project, a low cost scanning cavity was not available at 405nm; subsequently at least one has entered the market (Thorlabs SA200-3B). Therefore, we constructed our own using off the shelf components as illustrated and discussed in Fig. E.9. Although the cavity body consists of four thin rods, the cavity worked as intended, except that it is very sensitive to frequency drifts due to temperature changes and air currents. Enclosing the cavity in a box significantly improved this issue. A sample cavity trace is shown in Fig. E.9.

E.2.4 High Power Injection Locked Diode

To obtain the required 405nm power for cooling and imaging we injection lock a high power (up to 100mW) GaN diode [317,318]. Without injection the spectra of these diodes is extremely poor and, for the diodes we tested, the center wavelength is approximately 407–408nm (see Fig. E.10). The diode wavelength tunes as a function of current (temperature) as $+0.02\text{nm}/\text{mA}$ ($+0.05\text{nm}/^\circ\text{C}$). This temperature tuning is an order of magnitude worse than our InGaAs diodes at 767nm. To get to 404.5nm we would need to cool the diode from room temperature to $\approx -30^\circ\text{C}$.




Figure E.10: High power diode spectra measured using an Ocean Optics spectrometer with 0.1nm resolution. (a) Spectra as a function of temperature for 108mW output power. (b) Spectra as a function of current at $T = -5^\circ\text{C}$. In each plot the spectra after injection is shown for comparison; with sufficient injection power we can pull the diode wavelength up to 2nm. (c) Required injection power as a function of the diode temperature for an output power of 95mW.




Figure E.11: High power diode mount and enclosure (top lid not shown). Parts are described in the text. (Bottom inset) Protection circuit for the diode, based on a similar circuit in [173].

To reach temperatures below the dew point (approximately 9°C in the lab), the diode must be thermally insulated from the environment and free of moisture to prevent ice formation. To accomplish this, the diode is enclosed in a hermetically sealed box on a mount that is only in thermal contact with the air in the box. The box has desiccant to minimize moisture and is also designed to be vacuum pumped to improve thermal isolation and eliminate moisture entirely. The laser enclosure and diode mount is illustrated in Fig. E.11. The diode is heat sunk to a custom aluminum (Al) diode mount using an Al faceplate to maintain good thermal contact and to keep the diode in place. No thermal grease is used directly on the diode. A 10k thermistor (Thorlabs TH10K) is glued into the top of the diode mount for temperature monitoring and stabilization. Four Thorlabs ER1 rods are threaded into the diode mount to hold the aspheric collimating lens (Thorlabs C671TME-405). The mount is put on top of two thermo-electric coolers (TECs), wired in series, which are sunk straight to the Al baseplate. Boron nitride thermal grease is applied to the diode mount and TECs. The mount is screwed to the baseplate using nylon hardware to maintain thermal isolation. Copper pipes were soldered to the baseplate for water cooling. Since Al/Cu soldering is difficult, in the future, it would be sufficient to use thermal epoxy as the heat load is small. Another option would be to construct the baseplate from copper or brass. Heat is removed from the baseplate using a computer CPU chiller (Coolance Exos LT), which runs liquid coolant through the baseplate and dissipates the heat at a heat exchanger using fans. There is no active cooling and therefore this chiller is limited to keeping the baseplate at room temperature.

The baseplate plus diode mount is covered by an Al box (Digikey HM845-ND). A hole in the bottom of the box goes around the diode mount and the box sits on a Buna-N O-ring (Shore A Durometer 40) coated in Dow Corning high vacuum grease. The box screws directly into the baseplate to complete the seal. The box lid also has an O-ring that is coated in vacuum grease. The light goes out through a window on the front of the box sealed with Torr Seal glue. Electrical connections for the diode current, thermistor and TECs are

made using magnet wire which is individually passed through holes in the box and sealed with Torr Seal. We avoid standard BNC feedthroughs since they are not hermetically sealed [319]. A copper tube with a Swage-Loc ball valve and CF20 connector goes through the side of the box and is sealed with Torr Seal. The plan was to connect the box to a roughing pump, seal the system with the ball valve, and then remove the pump. However, the vacuum seal did not hold. The box can be pumped down, but a few seconds after the pump is turned off, the box returns to room pressure. Nevertheless, the seal is effective at keeping moisture out and a homemade mesh pouch of desiccant ($\approx 10\text{cm}^3$) is sufficient to keep ice off the diode.

The diode is cooled using TECs, which are characterized by three parameters, A , B , and C . These parameters relate the heat pumped Q , the voltage V , the current I , and the temperatures at the hot and cold side of the TEC T_H and T_C ($\Delta T = T_H - T_C$) as

$$Q = AT_C I - BI - C\Delta T, \quad (\text{E.9})$$

$$V = 2BI + A\Delta T. \quad (\text{E.10})$$

Typically a TEC datasheet does not give A , B , and C , but rather Q_{max} (Q when $\Delta T = 0$ at a given T_{C1}) and ΔT_{max} (ΔT when $Q = 0$ at a given I_{max} , V_{max} and T_{C2}). These values can be used to get A , B , and C using the relations

$$A = \frac{-\Delta T_{max} + \sqrt{(\Delta T_{max})^2 + \frac{2T_{C1}T_{C2}V_{max}I_{max}}{Q_{max}}}}{\frac{T_{C1}T_{C2}I_{max}}{Q_{max}}}, \quad (\text{E.11})$$

$$B = \frac{V_{max} - A\Delta T_{max}}{2I_{max}}, \quad (\text{E.12})$$

$$C = \frac{AT_{C2}I_{max} - BI_{max}^2}{\Delta T_{max}}. \quad (\text{E.13})$$

Given A , B and C it is straightforward to use Eqns. E.9 and E.10 to calculate the required I and V to run the TEC for a given heat load and temperature gradient. For most TECs, $\Delta T_{max} \approx 60^\circ\text{C}$ when $T_{H2} = 25^\circ\text{C}$, so for a larger temperature gradient the TECs must be stacked, which is the configuration we use for cooling the diode mount. Once we know the parameters of the top and bottom device we need to solve Eqns. E.9 and E.10 setting the hot side temperature of the top device to the cold side temperature for the bottom device. The bottom TEC needs to pull all the heat from the system being cooled (the diode mount) plus the heat generated by the top TEC ($Q_{load} + I_{top}V_{top}$). Therefore the optimal configuration is for the top TEC to have low resistance and the bottom TEC to be large so that it is capable to pulling large heat loads. The heat generated by the bottom TEC is less important since it is in thermal contact with a large heat sink.

We modeled different pairs of TECs that are available by solving Eqns. E.9 and E.10 where each TEC has the same current since we put the TECs electrically in series. One of the main inputs into the calculation is the expected heat load from the diode mount. When running, heat is generated by the diode itself, however, the main contribution is convective heat transfer from the air in the box. We estimate the convective load to be approximately $4W$ when the diode mount is -30°C . Therefore, for the calculation we looked at four sets of parameters; $\{T_C = -20^\circ\text{C}, T_H = 20^\circ\text{C}, Q = 3W\}$, $\{T_C = -20^\circ\text{C}, T_H = 20^\circ\text{C}, Q = 5W\}$, $\{T_C = -30^\circ\text{C}, T_H = 20^\circ\text{C}, Q = 3W\}$, and $\{T_C = -40^\circ\text{C}, T_H = 20^\circ\text{C}, Q = 3W\}$. The best combination of TECs is Digikey 102-1680-ND and 102-1681-ND which can meet all the above parameters within the current

and voltage constraints of the Thorlabs ITC510 TEC controller (4A, 8V). An alternative design, which may be more flexible, is to run the bottom TEC from a high current supply and only feedback to the current of the top TEC [319].

Ultimately the diode can be cooled to $< -20^\circ\text{C}$ without pumping, but it was important to add a layer of fiberglass insulation to the mount and collimating lens assembly. To inject the high power diode we send some of the master light back through an isolator into the diode. The diode spectra before after injection is shown in Fig. E.10. Although we do not cool the diode all the way to 404.5nm, the injection is able to pull the diode the remaining distance. In Fig. E.10 we show the minimum required injection power as a function of temperature. This plot justifies cooling the diode as the injection power continues to drop until the diode is cooled below -10°C .

E.3 $4S \rightarrow 5P$ MOT and Laser Cooling

In this section we summarize our results for trapping and cooling ^{40}K on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition at 404.5nm. These results are published in Ref. [88]. To realize a MOT on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition, we first load ^{40}K from an enriched K vapor using our high capture velocity (high power, large beams) 767nm MOT described previously in §3.2.6. After 10s we turn on the 405nm, $4S \rightarrow 5P$ cooling beams with up to 30mW of total power. The blue beams are mixed with the 767nm beams using dichroic mirrors, as illustrated in Fig. E.13. The two MOTs overlap for 3ms before we turn off the 767nm cooling light; we leave on the 767nm repump light at a reduced power. The state structure for the $4S \rightarrow 5P$ MOT is shown in Fig. E.12 and a schematic of the physical setup is shown and described in Fig. E.13.




Figure E.12: (a) State diagram for the $4S \rightarrow 5P$ MOT. The blue line indicates the 405nm cooling beam and the red line indicates the 767nm repump beam. The cooling beam is locked to the $4S_{1/2}$, $F = 1 \rightarrow 5P_{3/2}$ transition of ^{39}K (see Fig. E.8) and the repump beam is locked to the $4S_{1/2} \rightarrow 4P_{3/2}$ crossover transition of ^{39}K (see Fig. 3.11). (b) Possible decay paths for atoms excited by one of the MOT beams. Atoms excited to the $5P_{3/2}$ state can decay straight back to $4S_{1/2}$ or via a three-photon cascade by decaying to either $3D$ or $5S$. The repump beam excites atoms to $4P_{3/2}$, which can only decay straight back to $4S_{1/2}$. For transition rates see Fig. A.1.




Figure E.13: $4S \rightarrow 5P$ MOT setup and timing diagram. The 405nm light is combined with the 767nm MOT light using a 2 inch dichroic mirror (DM). Each MOT beam is retro-reflected through two sets of dichroic wave plates (DWP). The DWP is a quarter wave retardation for one wavelength and no retardation for the other. There are three orthogonal beam paths in total; for clarity only one is shown. The quadrupole field strong axis is aligned parallel to the direction of gravity.




Figure E.14: Characterizing the $4S \rightarrow 5P$ MOT. Absorption images taken on the $4S \rightarrow 4P$ cycling transition of the (a) $4S \rightarrow 4P$ MOT and the (b) $4S \rightarrow 5P$ MOT illustrate that the gas is significantly compressed in the 405nm MOT. (c) We plot the RMS Gaussian width of the vertical direction (open, red circles) and horizontal direction (black squares) versus the time after the transfer to the 405nm MOT. The vertical direction is aligned with the strong axis of the quadrupole field. The widths exponentially relax to lower values in 1.8ms (4ms) for the vertical (horizontal). (d) To determine the required 6-beam 405nm intensity we measure the sag of the 405nm MOT and the transfer efficiency from the 767nm MOT. The data is fit to an exponential and the e^{-1} intensity of the fit is comparable to I_{sat} . For sufficient intensity the transfer efficiency is nearly 100%. (e) Similar data looking at the 767nm repump intensity. Again, $I > I_{sat}$ is sufficient. As this is a retro-reflected MOT, we estimate the 6-beam intensities by measuring the power before and after the cell. However, given uncertainties in losses and the atom position in the beam, quoted values of the intensity should be taken as good to within a factor of two.

After loading in the 405nm MOT, the cloud size shrinks exponentially to a new steady state value (Fig. E.14). The exponential timescales ($\tau = 2 - 4\text{ms}$) are consistent with the damping rates predicted by the MOT calculations (§E.1). To characterize the required power for the 405nm cooling beams and the 767nm repump beams, we look at the fraction of atoms transferred to the 405nm MOT and the trap sag under gravity versus the total 6-beam 405nm and 767nm intensity. These results are also displayed in Fig. E.14 and indicate that the required intensity is approximately I_{sat} (see Table A.2). If $I > I_{sat}$, there is nearly perfect transfer efficiency into the 405nm MOT. Since our beams are small and the MOT is retro-reflected this intensity corresponds to around 2mW of total 405nm power, which is easily produced by our laser system.

E.3.1 One and Two-Body Loss Rates




Figure E.15: Atom number versus hold time in the $4S \rightarrow 5P$ MOT. The main data is fit to two exponential timescales as described in the text. (Top inset) Peak density for short hold times is approximately constant. (Bottom inset) For low atom number, an exponential fit to the data gives the one-body lifetime. The lifetime decreases as we increase the 6-beam 405nm intensity. We fit the data to a model including photoionization losses (Eqn. E.19).

If we measure number versus the time in the 405nm MOT (Fig. E.15), there are no signs of loading even after a hundred-fold reduction in number. This is expected since the capture rate scales as $\Gamma^4 d^4$ for small beam diameters d . Since there is no loading, measuring N versus t provides a clean method for measuring loss processes since we can ignore loading terms in the rate equation [320],

$$\frac{dN}{dt} = -\frac{N}{\tau} - \beta \int n(\vec{x})^2 d^3x \quad (\text{E.14})$$

where τ is the one-body lifetime, β is the two-body loss coefficient and n is density. The data is Fig. E.15 fits well to an exponential with a different slope at high and low densities. At low densities (long hold time)

an exponential fit to N vs. \mathbf{t} returns τ . In the high-density regime, the peak density appears roughly fixed (see inset to Fig. E.15) and $n \approx n_0 e^{-r^2/2\sigma(\mathbf{t})^2}$, so

$$\frac{dN}{d\mathbf{t}} = -\frac{N}{\tau} - \beta n_0^2 \int e^{-r^2/\sigma(\mathbf{t})^2} d^3x, \quad (\text{E.15})$$

$$= -\frac{N}{\tau} - \beta n_0^2 \pi^{3/2} \sigma(\mathbf{t})^3, \quad (\text{E.16})$$

$$= -\frac{N}{\tau} - \beta \frac{n_0}{\sqrt{8}} N, \quad (\text{E.17})$$

$$= -\left(\frac{1}{\tau} + \beta \frac{n_0}{\sqrt{8}}\right) N. \quad (\text{E.18})$$

In this limit, N versus \mathbf{t} also fits to an exponential and we can get β from the exponential time constant if we know τ and n_0 . To measure τ we fit the low density regime and to measure n_0 we take an absorption images of the MOT. At $\Delta = -2\text{MHz}$, we find $\beta = 2.0(1) \times 10^{-10} \text{cm}^3/\text{s}$ for $I \approx 75 \text{mW}/\text{cm}^2$, and $\beta = 1.4(1) \times 10^{-10} \text{cm}^3/\text{s}$ for $I \approx 20 \text{mW}/\text{cm}^2$. These are relatively high two-body loss rates compared to ^{40}K MOTs on the D2 transition, which are typically in the range of 10^{-12} to $10^{-10} \text{cm}^3/\text{s}$ [188, 189, 321]. Therefore, density-dependent loss rates in the potassium $nS \rightarrow (n+1)P$ MOT, unlike metastable helium [89, 90], are not smaller than in the $nS \rightarrow nP$ MOT. However, for the $< 100\text{ms}$ timescales required for blue MOT cooling this loss rate is negligible.




Figure E.16: (a) Fluorescence per atom from the 405nm MOT at $\Delta = -2\text{MHz}$ and $I_{405} = 76 \text{mW}/\text{cm}^2$ versus repump power. The fluorescence clearly asymptotes to a fixed value once the system is not repump limited. The data is fit to an exponential (black line). In the data shown here, the asymptotic value is 0.048. (b) Fluorescence per atom for infinite 767nm repump (asymptotic value the curve in (a)) versus 405nm intensity I_{405} . This curve is fit to $f_{\text{max}} I_{405} / (I' + I_{405})$ (black line) to get the maximum fluorescence value f_{max} ; for this data $f_{\text{max}} = 0.186$. From f_{max} the excited state fraction is determined as $\frac{1}{2} \frac{f}{f_{\text{max}}}$ where f is the measured fluorescence. For example, in part (a) we measure a fluorescence value of $f = 0.035$ when the repump power is 20mW and so the excited state fraction is approximately 10%.

The decrease of the one-body lifetime, τ , versus total intensity (inset to Fig. E.15) indicates possible photoionization. To convert the loss rate into a photoionization cross-section using Eqn. E.8 we need to know the excited state fraction, ρ_{ee} . This can be estimated by looking at MOT fluorescence since the fluorescence per atom f is proportional to ρ_{ee} . If we can determine the maximum possible fluorescence f_{max} , which occurs when $\rho_{ee} = 1/2$, then $\rho_{ee} = \frac{1}{2} \frac{f}{f_{\text{max}}}$. To get f_{max} we have to look at the fluorescence versus

405nm intensity for infinite repump power. This process is explained further in Fig. E.16. Once we know f_{\max} , we can get the photoionization cross-section, σ , by fitting the lifetime versus 405nm intensity I to the formula

$$\tau = \left(\frac{1}{\tau_{bk}} + \sigma \frac{I}{hc/\lambda} \frac{f(I)}{2f_{\max}} \right)^{-1} \quad (\text{E.19})$$

where τ_{bk} (the vacuum lifetime) is a free parameter. Fitting the data from the inset to Fig. E.15 we get that $\sigma = 8 \times 10^{-18} \text{cm}^2$ for the $5P_{3/2}$ ionization at 404.5nm. We quote this as an upper bound because we cannot rule out other loss processes and there is considerable uncertainty in our ability to measure the intensity. This cross-section is comparable to the Rb value we used as an estimate in §E.1.1. Overall, the photoionization losses are on the seconds timescale and have no effect on cooling. Ionization will place an upper bound on the number of photons we can collect for imaging. For an objective with NA=0.6, fluorescence imaging at $I_{\text{sat}}/4$ could collect 10^4 blue photons with a 1% upper bound on the probability of ionization.

E.3.2 Temperature and Density




Figure E.17: (a) $4S \rightarrow 5P$ MOT temperature and (b) density for several magnetic quadrupole gradients: 5G/cm (black squares), 7.5G/cm (red circles), and 10G/cm (blue triangles). Temperature is shown along a weak axis of the quadrupole (horizontal direction); temperatures along the strong axis (along gravity) were 1.82(15) times higher.

The temperature and density performance of the MOT is summarized in Fig. E.17. A minimum temperature of 63(6) μK is seen for a 5G/cm gradient. As the gradient increases, so does the temperature and density. At all gradients the minimum temperature and maximum density occurs at $\Delta \approx -\Gamma$. One unconventional aspect of the measurement is that the vertical temperature (along gravity) is 1.82(15) times higher than the horizontal temperature. This may be due to differences in the gradient along these directions; the strong axis of the quadrupole is vertical, so there is twice the gradient in that direction. Nevertheless, this issue warrants further study.

The optimal phase space density, which scales as $n_0 T^{-3/2}$, occurs when $T = 75\mu\text{K}$ and $n_0 = 1.2 \times 10^{10} \text{cm}^{-3}$ (7.5G/cm, $\Delta = -1.5\text{MHz}$). In comparison, the 767nm MOT has a typical temperature and density of $T = 180\mu\text{K}$ and $n_0 = 2 \times 10^9 \text{cm}^{-3}$. Therefore, the 405nm MOT realizes a twenty-fold enhancement in phase space density. The biggest contribution to this improvement is the nearly ten-fold increase in density.

We attribute the increase in density to three factors. First, the temperature in the MOT is lower. Second, expulsive forces due to re-absorption are reduced, because the same spring constant is achieved with a scattering rate that is six times smaller. The open transition may further reduce re-absorption by pumping to the $F = 7/2$ ground states. Third, the optical density seen by both the incident and the scattered blue photons is reduced by a factor of nearly thirty.




Figure E.18: $4S \rightarrow 5P$ molasses temperature along the horizontal direction versus detuning. Temperature along the strong axis of the quadrupole field (along gravity) are 1.96(11) times higher. The black curve is a theoretical calculation at low intensities from Fig. E.3.

In Fig. E.17 the minimum temperature appears to decrease monotonically with decreasing gradient, which suggests that temperature may asymptote to the Doppler temperature $T_D = 23.5\mu\text{K}$ in the zero-field (molasses) limit. However, the minimum experimental molasses temperature (Fig. E.18) is only $85\mu\text{K}$. This minimum temperature occurs around the expected Doppler detuning $\Delta/\Gamma = -0.6(2)$. This temperature is more than three times the Doppler temperature calculate in §E.1, which may be due to a low cooling rate as compared to the short molasses time available. Indeed we observe that the temperature asymmetry between the vertical and horizontal axes in the MOT persists during the molasses phase, when no gradient asymmetry remains. Even the asymptotic temperature might not be expected to reach the Doppler limit: while the calculation in §E.1 includes multi-level effects, it neglects three-dimensional effects, reabsorption, and heating due the intensity fluctuations. The latter effect has been shown in [322] to limit one-dimensional Doppler cooling to several times T_D in Sr, when Sisyphus and other sub-Doppler mechanisms are absent [322]. In the case of $4S \rightarrow 5P$ cooling in ^{40}K , polarization gradient cooling may be interrupted by the optical pumping and depolarization effects of the three-photon cascade.

In summary, we have demonstrated the feasibility of cooling on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition of ^{40}K at $\lambda = 404.5\text{nm}$. This is a promising first step towards our ultimate goal of site resolved imaging in the lattice using the 405nm scattered light.

Appendix F

Imaging

In this appendix we review several imaging topics.

F.1 Number Calibration of Absorption Images

For resonant absorption imaging on the cycling transition (§4.1.1), which is the predominant imaging technique used in this thesis, the relation between optical depth (Eqn. 4.4) and density is given by the cross-section on the cycling transition (Eqn. B.78). The experimental cross-section is smaller than the cycling transition cross-section due to deviations from ideal imaging (e.g., imperfect polarization). To empirically calibrate absorption images we use an experimental observable that depends on particle number. Of the several possible observables, we use T_C to calibrate N , which is a common choice in many experiments.




Figure F.1: T_C data used to calibrate number for the ^{87}Rb apparatus. The intercept measured from the data is 3.21(7), so $\gamma^{-1} = 7.9(5)$. This data has only been partially repumped, so we also need to account for the repump factor of 5.4(5) and therefore the on-resonance number calibration is $\gamma_0^{-1} = 1.46(17)$.

Our procedure in the ^{87}Rb apparatus is to evaporate until the system is just below T_C and take alternating high and low OD images (see §F.2). We fit the high momentum tails of the thermal distribution in the high

OD image to measure temperature. We fit the low OD image to get N_0/N (the condensate fraction) and \tilde{N} (the number assuming the cycling transition cross-section). \tilde{N} is related to the actual number (N) as $\tilde{N} = \gamma N$ where γ is the calibration factor. Since the BEC number drifts as we run the experiment, this is a natural method to extrapolate back to $N_0/N = 0$. As seen in Fig. F.1, we plot $T(\gamma N)^{-1/3}$ versus γN_0 and fit to a line to get the intercept (A) at $\gamma N_0 = 0$. In our trap ($\bar{\nu} = 35.76(8)\text{Hz}$), $T_C/N^{-1/3} = 1.613$ and so $\gamma = (1.613/A)^3$.

F.2 Partial Repumping

With absorption imaging there is an upper bound on the optical depth that can be imaged of $OD \simeq 3$ (see §4.1.1). To circumvent this issue we can control the number of atoms that are imaged by partial repumping. For ^{87}Rb , the cycling imaging transition is $F = 2, m_F = \pm 2 \rightarrow F' = 3, m_{F'} = \pm 3$ as shown in Fig. F.2. Atoms in $F = 1$ are dark to this imaging light and will not be imaged. Since we perform our experiments with atoms in $F = 1$, we have to actively transfer atoms to $F = 2$ for imaging using repump light on the $F = 1 \rightarrow F' = 2$ transition. When the repump light is on-resonance we transfer all the atoms to $F = 2$ and the resulting absorption image has high OD. However, if the repump is detuned from resonance we can vary the repump pulse length ($4\mu\text{s}$ – $600\mu\text{s}$) to control the fraction of atoms that are imaged where

$$\frac{N_{low}}{N_{high}} = 1 - e^{-\tau/\tau_0}. \quad (\text{F.1})$$

We refer to this as a low OD image. High and low OD images are shown in Fig. F.2.

Often there is complementary information in both the high and low OD images, so we take both types of images in sequential experimental runs. For example, the high OD image has good signal-to-noise ratio in the high momentum region, which is ideal for measuring temperature, but the condensate peak is saturated. The low OD image is therefore required to measure the condensate number.

To calibrate the partial repumping we prepare a thermal cloud and take alternating high OD and low OD images to get the fraction of atoms repumped. One technique is to set the atom number so that the high OD image is just below the point of being saturated and then take both images at the same expansion time. The problem is that at very low repump times there can be very little signal in the low OD image. To get higher signal, we can take the low OD image at shorter expansion times, but one has to be careful because the repump intensity may vary as a function of position. A sample repump calibration curve is shown in Fig. F.3.

F.3 Resolution

For time-of-flight imaging, resolution is not a major issue since the feature sizes after expansion are much larger than the wavelength of the imaging light. However, this is not the case when imaging in-situ and particularly if we want to achieve single-site resolution of a gas in an optical lattice (§4.1.5). This section will briefly overview several high resolution topics as a reference for the main text and to motivate design considerations for the ^{40}K apparatus where the goal is to image atoms with single-site resolution (§3.2.10).




Figure F.2: (a) Imaging and repump transition for ^{87}Rb . Atoms start in $F = 1$ during the experiment and are repumped to $F = 2$ for imaging. We can transfer a controllable fraction to $F = 2$. (b) Cross-section through an image with all the atoms repumped (top) and only a fraction repumped (bottom). The dense middle part of the condensate is completely saturated in the top image, but the signal-to-noise ratio at the edge is higher.




Figure F.3: Sample calibration curve for partial repumping. The fraction repumped, taken from alternating high OD/low OD shots, is fit to Eqn. F.1 with τ_0 as a free parameter.

The resolution of an imaging system is limited by two effects, diffraction and aberrations. Aberrations are technical deviations from ideal imaging and may be random, e.g., due to phase fluctuations in poor quality glass, or smooth, e.g., spherical aberrations due to lens shape. In theory, aberrations may be eliminated by judicious choice of optics and/or corrective optics at which point the resolution becomes diffraction limited.

At the diffraction limit (DL), the only parameter defining the imaging system is the wavelength (λ) and numerical aperture (NA), $NA = n \sin(\theta)$, where n is the index of refraction between the object and the lens and θ is the acceptance half-angle of the imaging lens (see Fig. 4.1). The NA specifies the size of the point spread function (PSF), which is the light intensity distribution in the imaging plane of a point source in the object plane. If we define a normalized variable

$$\xi = (2\pi NA) \frac{r}{\lambda}, \quad (\text{F.2})$$

at the imaging plane (on focus), the PSF is [205]

$$PSF(\xi) = \frac{1}{\pi} \left(\frac{J_1(\xi)}{\xi} \right)^2 \quad (\text{F.3})$$

where J_1 is the Bessel function of the first kind. In the form of Eqn. F.3, the PSF is normalized to unity. In order to discuss resolution, we need to image two objects. If we consider two incoherent point sources separated by distance ξ_0 , then the image is given by adding the two point spread functions (Eqn. F.3). As shown in Fig. F.4, when the objects are closer together the PSFs overlap and it becomes more difficult to resolve whether there are two separate sources. There is no absolute definition of resolution — it is a subtle concept that depends on any a priori knowledge of the object and noise. For example in Fig. F.4 we also show the two point sources at a fixed distance with increasing noise. It is clear that as the noise increases, distinguishing the two objects becomes more difficult. Therefore, determining a metric to define when they are unresolvable requires careful consideration. Nevertheless, to establish a universal imaging language, there are several standard definitions of resolution. The most common is the Rayleigh criterion, which defines resolution as the distance from the center of the peak to the first zero in the PSF,

$$\sigma = \frac{0.6098\lambda}{NA} \quad (\text{F.4})$$

When two objects are separated by σ the contrast is 15.3% (without noise).

From the PSF we can derive the image $f(\xi)$ from an arbitrary incoherent source distribution $H(\xi)$

$$f(\xi) = \int_{-\infty}^{\infty} d\xi' PSF(\xi - \xi') H(\xi'). \quad (\text{F.5})$$

The following are some common examples. The edge spread function (ESF) is the image formed from a source that has an infinitely sharp edge

$$LSF(\xi) = \int_{-\infty}^0 d\xi' PSF(\xi - \xi'). \quad (\text{F.6})$$

The resolution is the distance it takes the ESF to drop from 92.7% to 7.4% of the value at $-\infty$. The




Figure F.4: (a) Image of two incoherent point sources as a function of the distance between the sources in ξ units (Eqn. F.2). The Rayleigh criterion of resolution is defined as a separation of $\xi_0 = 3.83$. For point sources separated by less than the Rayleigh resolution some contrast is still evident. (b) Pixelized image with photon shot noise for two point sources separated by $\xi_0 = 3.4$ varying the max photon number in a bin. Noise makes it increasingly more difficult to detect two individual objects.

modulation transfer function (MTF) is the image formed from a sinusoidally varying source with spatial period ξ_0

$$MTF(\xi, \xi_0) = \int_{-\infty}^{\infty} d\xi' PSF(\xi - \xi') \sin^2(\pi\xi'/\xi_0). \quad (F.7)$$

When the period of the MTF is equal to σ the contrast is 7%. A similar function is the contrast transfer function (CTF), which is the image formed from a source with a square wave intensity. If the spatial period is ξ_0 then

$$CTF(\xi, \xi_0) = \sum_{i=-\infty}^{\infty} \int_{\xi_0 i}^{\xi_0(i+1/2)} d\xi' PSF(\xi - \xi'). \quad (F.8)$$

When the period of the CTF is equal to σ the contrast is 8.8%. For a diffraction limited imaging system the ESF, MTF, and CTF do not provide additional information. However, they are each unique experimental measures that can be used to characterize a non-ideal imaging system. Deviations of these measures from the diffraction limit indicates the presence of aberrations. If an imaging system resolution is limited by aberrations, these measures are useful because they can still define a reasonably accurate resolution by connecting back to the diffraction-limited result. For example, we can define the aberration-limited resolution based on the CTF frequency with a contrast of 8.8%.

An important consideration when imaging a lattice system, with many planes of atoms, is how the PSF changes away from the imaging plane (i.e., out of focus). If we define the normalized distance away from the focus (z) as

$$\gamma = (2\pi NA^2) \frac{z}{\lambda} \quad (F.9)$$

then the PSF as a function of ξ and γ is [205]

$$PSF(\xi, \gamma) = \left[\sum_{i=0}^{\infty} (-1)^i \left(\frac{\gamma}{\xi} \right)^{2i} \frac{J_{1+2i}(\gamma)}{\gamma} \right]^2 + \left[\sum_{i=0}^{\infty} (-1)^i \left(\frac{\gamma}{\xi} \right)^{1+2i} \frac{J_{2+2i}(\gamma)}{\gamma} \right]^2. \quad (F.10)$$

This function is plotted in Fig. F.5. The depth of focus is defined roughly as $\gamma = \pi$ (i.e., the depth at which the peak height is approximately 80% of its value at $\gamma = 0$). An important note is the scaling of resolution versus depth of focus. Resolution scales as NA^{-1} whereas depth of focus as NA^{-2} .




Figure F.5: (a) Image of the PSF versus the radial coordinate ξ at different focus depths γ . (b) The peak height versus depth normalized to the height when $\gamma = 0$.

Appendix G

Supporting Calculations

The purpose of this appendix is to provide supporting calculations and more detail for topics discussed in the main thesis text. These details may be useful to readers attempting similar calculations. The sections of this appendix are independent of each other.

G.1 K $4S$ - $5P$ Magic Wavelength for Dipole Trapping

Our plan for the ^{40}K apparatus (§3.2) is to implement single-site imaging of atoms in the lattice on the $4S_{1/2} \rightarrow 5P_{3/2}$ transition at 404.5nm (§4.1.5). In order to facilitate long imaging times, the excitation light must double as laser cooling light. Free-space laser cooling on this transition was demonstrated as part of this thesis work; see §E for details. Laser cooling in the lattice is more complicated because the lattice potential is generated by coupling the ground and excited states with the lattice light (§B.5). The potential depends strongly on the whether the atom is predominantly in the ground or excited state and can change as we drive on the imaging transition.

In the simplest case, the two-level atom, the atomic ground state follows to a new dressed ground state with an energy shift proportional to the lattice light intensity. The atomic excited state also follows to a dressed data; the excited state energy shift is equal and opposite to the atomic ground state energy shift. For fluorescence imaging we must transition between the ground and excited states to scatter photons. If the excited state experiences the opposite potential then the atom, which starts out in the lowest energy region of the light, will be in the highest energy region when it transitions to the excited state. The atom experiences an effective potential

$$V_{eff} = \rho_{gg}V_g + \rho_{ee}V_e, \quad (\text{G.1})$$

and if $V_e = -V_g$ then

$$V_{eff} = V_g(\rho_{gg} - \rho_{ee}). \quad (\text{G.2})$$

For strong driving $\rho_{ee} \approx \rho_{gg}$, so $V_{eff} \rightarrow 0$ and the atom will be unconfined.

Fortunately, alkali atoms are not two-level systems. For two coupled states (1 and 2), it is not necessarily true that $V_1 = -V_2$ because state 2 may also couple to higher energy states. The shift of the two states will be equal at specific wavelengths. These wavelengths are referred to as a “magic wavelengths”. At a magic wavelength the effective potential is state-independent. The purpose of this calculation is to determine these magic wavelengths for the pair of imaging states in the ^{40}K experiment: $4S_{1/2}$ and $5P_{3/2}$. There are technical and fundamental constraints on the possible wavelengths we can use. To minimize heating

during the simulation the wavelength must be far and red-detuned from $4S_{1/2} \rightarrow 4P$. (i.e., $\lambda \gg 800\text{nm}$). To prevent ionization directly from the lattice light in the $5P_{3/2}$ state $\lambda > 972\text{nm}$. The final constraint is that a high-power ($> 20\text{W}$) laser needs to be available. To satisfy these three conditions, we focus in on wavelengths around $\lambda = 1064\text{nm}$ ($\pm 10\text{nm}$), where ytterbium doped fiber lasers are available.

For any two pairs of states $|nJFm_F\rangle$ and $|n'J'F'm'_F\rangle$, the energy shift of the lower energy state is given by Eqn. B.57 where Γ is the decay rate from the higher to lower J manifold. The energy shift of the higher energy state is the negative of Eqn. B.57. To get the full shift of state $|nJFm_F\rangle$, we sum over the energy shift due to the coupling with all other states. For our two states considered here, $4S_{1/2}$ and $5P_{3/2}$, Fig. G.1 illustrates these couplings. $4S_{1/2}$ can couple to any P states, however, given our choice of wavelength, the energy shift is predominantly due to coupling to the $4P$ and $5P$ states. The $5P_{3/2}$ state can couple to any S and D states. Because the $4S_{1/2}$ state is far from resonance when $\lambda \simeq 1\mu\text{m}$, the magic wavelength will be due to a resonant shift in the $5P_{3/2}$ energy. From Fig. G.1, we can identify several candidate $5P_{3/2}$ transitions: $13S$, $14S$, $11D$, and $12D$. Unfortunately, information about the transition rates to these states is unknown, however, we can extrapolate from the known rates at lower n , as described in the caption of Fig. G.1. It is preferable to couple to a D state because there is less polarization dependence. The $5P_{3/2}$ shift is not spin-independent at large detunings when coupling to an S state. For example, if we are in the $5P_{3/2}$, $F = 11/2$, $m_F = 11/2$ state then only a σ^- beam can couple to an S state. Ultimately, we decided to use the $12D$ resonance.




Figure G.2: (a) $4S_{1/2}$ and $5P_{3/2}$ energy shift for an unpolarized beam with $1\mu\text{m} < \lambda < 2\mu\text{m}$. (b) Energy shift around the $12D$ magic wavelength, $\lambda = 1053.582\text{nm}$.

In Fig. G.2 we plot the energy shift for the $4S_{1/2}$ and $5P_{3/2}$ states versus the wavelength of an unpolarized beam in the range $1\mu\text{m} < \lambda < 2\mu\text{m}$. Throughout this regime the $4S_{1/2}$ shift is entirely negative and smoothly varies from $-0.075\mu\text{K}\text{mm}^2\text{W}^{-1}$ at $\lambda = 2\mu\text{m}$ to $-0.156\mu\text{K}\text{mm}^2\text{W}^{-1}$ at $\lambda = 1\mu\text{m}$. The $5P_{3/2}$ shift is mostly positive, however, there are numerous resonances. Zooming into the region around the $12D$ peak, we see that the $4S_{1/2}$ and $5P_{3/2}$ shifts are equal when $\lambda = 1053.582\text{nm}$. Because the $5P_{3/2} \rightarrow 12D$ transition is so weak, we need to tune λ within 25GHz of resonance to get a large enough negative shift. Fig. G.2 is a good

indication of our tuning constraints. Since the transition is so narrow, we need to have frequency stability of $< 1\text{GHz}$.




Figure G.3: Energy shift of all intermediate states between $5P_{3/2}$ and $4S_{1/2}$ for an unpolarized beam at the magic wavelength, $\lambda = 1053.582\text{nm}$.

The $5P_{3/2}$ state can also decay via a 3-photon cascade either along the path $5P \rightarrow 3D \rightarrow 4P \rightarrow 4S$ or $5P \rightarrow 5S \rightarrow 4P \rightarrow 4S$ (see Fig. A.1). Therefore, we need to consider the energy shift of these intermediate states. If the lattice wavelength is at the magic wavelength ($\lambda = 1053.582\text{nm}$) then the differential shift for each state along the cascade is shown in Fig. G.3. In all the other states the energy shift is positive, however, the atom spends little time in these intermediate states.

G.2 Quadrupole to Dipole Transfer

In a magnetic quadrupole trap at low temperatures, atoms undergo Majorana spin-flips at the trap center where the magnetic field goes to zero and are lost from the trap (§B.2). Unless the zero is “plugged”, the atoms must be transferred to a dipole trap away from the field zero before Majorana losses become severe. This is the procedure used in the ^{87}Rb apparatus (§3.1). However, when moving to a dipole trap there is a compromise between transferring at low or high temperatures. At low temperature, the fraction of atoms captured by the dipole trap is large, but Majorana losses lead to a decrease in number. At high temperatures, Majorana losses are not a concern, but a significant fraction of atoms are not transferred to the dipole trap, which has a small depth compared to the quadrupole trap. To better understand this process, these calculations model the relationship between transfer efficiency, initial temperature and the dipole trap parameters. One purpose is to assess the feasibility of using this procedure for cooling ^{40}K in the ^{40}K apparatus (§3.2).

To model the transfer we use Maxwell-Boltzmann statistics and assume adiabaticity. These assumptions are justified since our system is far from degeneracy and the transfer is slow compared to trap and collisional




Figure G.4: Stages of the quadrupole-dipole transfer: (1) a pure quadrupole trap, (2) the dipole trap is turned on, (3) the quadrupole gradient is reduced and the dipole power is increased, and (4) the quadrupole trap is balancing gravity.

timescales. The different stages of the transfer are shown in Fig. G.4. In general, the full time-dependent potential is

$$V(\mathbf{r}) = V_{quad}(\mathbf{r}) + V_{dipole}(\mathbf{r}) - mgz, \quad (\text{G.3})$$

$$V_{quad}(\mathbf{r}) = A(\mathbf{r}) \sqrt{\left(\frac{x}{2}\right)^2 + \left(\frac{y}{2}\right)^2 + z^2}, \quad (\text{G.4})$$

$$V_{dipole}(\mathbf{r}) = \frac{2P(\mathbf{r})V_{sc}}{\pi \mathbf{w}^2} e^{-2\left(\left(\frac{x}{\mathbf{w}_x}\right)^2 + \left(\frac{y}{\mathbf{w}_y}\right)^2 + \left(\frac{z-z_0}{\mathbf{w}_z}\right)^2\right)}, \quad (\text{G.5})$$

where we assume the strong direction of the quadrupole field is along gravity. The dipole trap considered here has Gaussian confinement along all three directions. In the calculation, we set $\mathbf{w}_y \gg \mathbf{w}_x, \mathbf{w}_z$, which models a single beam trap with finite volume for numerical tractability. At $\mathbf{t} = 0$ the trap is a pure quadrupole ($A(0) = A_0$ and $P(0) = 0$). After the transfer at $\mathbf{t} \rightarrow \infty$ the quadrupole is just balancing gravity and the dipole trap is at maximum depth ($A(\infty) = mg$ and $P(\infty) = P_0$). Since the dipole trap has a limited depth, some of the atoms will escape. In the calculation we let $A(\infty) = 1.02mg$ so that the partition function integral will converge. Since we assume the transfer is adiabatic, we only need to consider equilibrium thermodynamics. The n-particle partition function is

$$Z_N = \frac{1}{N!} Z_1^N, \quad (\text{G.6})$$

$$(\text{G.7})$$

where

$$Z_1 = \Lambda^{-3} \int d^3x e^{-\beta V(\vec{x})}, \quad (\text{G.8})$$

and Λ is the de Broglie wavelength

$$\Lambda = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}. \quad (\text{G.9})$$

The entropy is

$$\frac{S}{Nk_B} = \ln\left(\frac{Z_1}{N}\right) + 1 + \frac{T}{Z_1} \frac{\partial Z_1}{\partial T}. \quad (\text{G.10})$$

The general procedure is to equate the entropies of the different potentials to calculate the change in temperature. At $\tau = 0$ the single-particle partition function is

$$Z_1 = \frac{32\pi(k_B T)^3}{A_0^3 \Lambda^3} \frac{1}{\left(1 - \frac{mg}{A_0}\right)^2 \left(1 + \frac{mg}{A_0}\right)^2}, \quad (\text{G.11})$$

and entropy is

$$\frac{S}{Nk_B} = \ln\left(\frac{Z_1}{N}\right) + \frac{11}{2}. \quad (\text{G.12})$$

At $\tau \rightarrow \infty$ we need to calculate the partition function numerically using Eqn. G.8.




Figure G.5: Final temperature after transfer to the dipole trap versus initial temperature in a 300G/cm quadrupole trap. The $\lambda = 1064\text{nm}$ dipole trap has a beam waist of $80\mu\text{m}$ and is traveling along y . All calculations are for ^{40}K .

The most efficient method for determining the final temperature is to construct entropy versus temperature tables and then interpolate. Using this method, the final temperature versus initial temperature is shown in Fig. G.5 for different dipole trap powers P_0 . There are three regimes depending on the initial temperature. For high initial temperatures the atoms are predominantly still in the weak quadrupole potential after the transfer. Therefore, the final temperature is independent of the dipole trap. The temperature decreases due to adiabatic decompression of the quadrupole trap. In the intermediate regime atoms start to load into the dipole trap. The final temperature becomes less sensitive to the initial temperature because the temperature is set by the depth of the dipole trap (empirically $k_B T_f$ is $\approx 10\%$ of the dipole trap depth). Finally, when the initial gas is very cold, there is 100% transfer to the dipole trap. The atoms are located

near the bottom of the dipole trap, which is well-described as a harmonic potential. The ratio of initial and final temperatures can be calculated by assuming transfer from a quadrupole to a harmonic trap. The harmonic trap frequencies are larger for higher dipole powers.




Figure G.6: (a) Quadrupole to dipole transfer efficiency versus initial quadrupole temperature (300 G/cm) and Majorana loss rate (top axis). The grey dashed line indicates where the Majorana loss rate is 1s^{-1} , which is around the point that these losses impede evaporation. We see that for $P_0 \geq 5\text{W}$, we can transfer approximately half the atoms at this point. (b) Final phase-space density (PSD) at the bottom of the dipole trap versus initial temperature. For cold initial temperatures there is 100% transfer so the PSD is just given by atoms in a harmonic trap. Since the phase-space density is not a function of the trap frequency, all dipole powers give the same PSD (in the low temperature regime). For higher temperatures, the dipole trap has a dimple effect and can lead to a significant PSD increase.

An important experimental benchmark is the transfer fraction. We define atoms as transferred to the dipole trap if their energy is less than the dipole trap depth (the bottom of the dipole trap is defined to be $V = 0$). The transfer fraction is then

$$\frac{1}{Z} \frac{1}{h^3} \int e^{-V(x)/k_B T} dx \int_0^{\sqrt{2m(E-V(x))}} dp p^2 e^{-\frac{p^2}{2mk_B T}} \quad (\text{G.13})$$

where E is the cutoff and the integrand is zero if $U(x) > E$. This can be expanded as

$$\frac{\Lambda^{-3}}{Z\sqrt{\pi}} dx \int e^{-V(x)/k_B T} \left(-2\sqrt{\frac{E-V(x)}{k_B T}} e^{-\frac{E-V(x)}{k_B T}} + \sqrt{\pi} \text{erf} \left(\sqrt{\frac{E-V(x)}{k_B T}} \right) \right). \quad (\text{G.14})$$

Using the above definition, the transfer efficiency versus initial temperature for different dipole power is shown in Fig. G.6. The Majorana loss rate (Eqn. B.11) is shown on the top axis of this plot. This rate formula agrees well with experiments in the $F = 1$ state of ^{87}Rb , but likely underestimates the actual rate for ^{40}K . From Fig. G.6 we see that imperfect transfer is not necessarily negative because there is an evaporation type effect which leads to a phase space density improvement.

The calculation agrees well with the observation from the ^{87}Rb apparatus that we can transfer approx-

imately 50% of the atoms into a 7W dipole trap ($w = 100\mu\text{m}$) when the initial cloud is $T \approx 20\mu\text{K}$. For fermions, such as ^{40}K , an interesting question is whether this method would work with a polarized sample where only p-wave collisions are allowed. This would be ideal because a mixture is difficult to maintain and create in a magnetic trap. However, p-wave collisions in the polarized sample decrease rapidly when $T < 40\mu\text{K}$ [173]. This may cause the adiabaticity assumption to break down during the transfer.

G.3 Expansion in Condensate Mean-Field

In §4.1.4, we discussed using condensate fraction after time-of-flight (TOF) as a thermodynamic probe. Because of interactions during TOF an analytical fit to determine the condensate fraction is not possible, so several heuristic methods have been developed. To evaluate these methods and to gain insight into the distribution after TOF, we numerically calculate the expanded distribution keeping the main interaction terms.

In general, the full interacting expansion requires a time-dependent many-body approach. However, since interactions are weak and dominated by the dense condensate, we can make several reasonable approximations that greatly simplify the calculation yet retain the most important features of the expansion. First, we assume that the condensate expansion is described by the GP equation (Eqn. 2.47), ignoring interactions with the non-condensate atoms. Although several experiments have observed that non-condensate atoms exert a pressure on the condensate (e.g., [126]), this effect is small and does not alter the main features of the full distribution. Using this assumption, the condensate expansion is only a function of condensate number N_0 , and not temperature T or total number N .

Next, we assume that as the thermal atoms expand, they experience the condensate mean-field potential, but we neglect any interactions between thermal atoms (semi-ideal model, §2.2.2). The overall time-dependent potential is

$$V(r, \mathbf{t}) = \frac{1}{2}m\omega^2(\mathbf{t})r^2 + 2U_0N_0|\Phi(r, \mathbf{t})|^2 \quad (\text{G.15})$$

where $U_0 = 4\pi\hbar^2a/m$, $|\Phi(r, \mathbf{t})|^2$ is the time-dependent condensate, and $\omega(\tau) = \omega_0$ for $\mathbf{t} < 0$ and $\omega(\mathbf{t}) = 0$ for $\mathbf{t} \geq 0$. In the trap the non-condensate distribution is

$$n_T(r) = \sum_n \left(e^{E_n/k_B T} - 1 \right)^{-1} |\psi_n(r)|^2 \quad (\text{G.16})$$

where ψ_n are the eigenstates of the semi-ideal model. Based on our assumption, each of these eigenstates expands independently so that

$$n_T(r, \mathbf{t}) = \sum_n \left(e^{E_n/k_B T} - 1 \right)^{-1} |\psi_n(r, \mathbf{t})|^2. \quad (\text{G.17})$$

Therefore, once we solve the time-dependent single particle Schrödinger equation for each eigenstate, we can construct the distribution for any arbitrary T . A similar calculation was performed in Ref. [223] except they treated the thermal atoms as classical particles. They showed a distinct change in the thermal atom distribution near the edge of the condensate which agreed with experiment. The final distribution that we

calculate is

$$n(r, \mathbf{t}) = N_0 |\Phi(r, \mathbf{t})|^2 + \sum_n \left(e^{E_n/k_B T} - 1 \right)^{-1} |\psi_n(r, \mathbf{t})|^2. \quad (\text{G.18})$$

In this calculation the total condensate number is 10^5 and we use $a = 5.5\text{nm}$, $m = m_{Rb}$, and $\nu_0 = 50\text{Hz}$.

G.3.1 Spherical Coordinates

In a spherically symmetric trap, the problem reduces to 1D, which greatly decreases the computational time and resources required. To write the equations in the 1D radial coordinate, we start with the full 3D Schrödinger equation

$$i\hbar \frac{\partial \Psi(\vec{x}, \mathbf{t})}{\partial \mathbf{t}} = \left(-\frac{\hbar^2}{2m} \vec{\nabla}^2 + V_{ext}(r, \mathbf{t}) \right) \Psi(\vec{x}, \mathbf{t}). \quad (\text{G.19})$$

In spherical coordinates the Laplacian is

$$\nabla^2 = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial}{\partial r} \right) + \frac{1}{r^2} \left[\frac{1}{\sin^2(\theta)} \frac{\partial^2}{\partial \phi^2} + \frac{1}{\sin(\theta)} \frac{\partial}{\partial \theta} \left(\sin(\theta) \frac{\partial}{\partial \theta} \right) \right]. \quad (\text{G.20})$$

The terms in the square bracket comprise the total angular momentum operator ($-\mathbf{L}^2/\hbar^2$). The spherical harmonics are the eigenstates of this operator ($\mathbf{L}^2 Y_l^m(\theta, \phi) = l(l+1)$). Rewriting the wavefunction as $\Psi(\vec{x}, \mathbf{t}) = \psi(r, \mathbf{t}) Y_l^m(\theta, \phi)$ we obtain an effective 1D equation for the radial wavefunction ($\psi(r, \mathbf{t})$)

$$i\hbar \frac{\partial \psi(r, \mathbf{t})}{\partial \mathbf{t}} = \left(-\frac{\hbar^2}{2m} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial}{\partial r} \right) + \frac{\hbar^2 l(l+1)}{2mr^2} + V_{ext}(r, \tau) \right) \psi(r, \mathbf{t}). \quad (\text{G.21})$$

G.3.2 Condensate Expansion

The condensate is described by the spherical GP equation

$$i\hbar \frac{\partial \Phi(r, \mathbf{t})}{\partial \mathbf{t}} = \left(-\frac{\hbar^2}{2m} \left(\frac{\partial^2}{\partial r^2} + 2r \frac{\partial}{\partial r} \right) + V_{ext}(r, \mathbf{t}) + \frac{4\pi\hbar^2 a N}{m} |\Phi(r, \mathbf{t})|^2 \right) \Phi(r, \mathbf{t}) \quad (\text{G.22})$$

where the condensate is normalized such that $4\pi \int_0^\infty r^2 |\Phi(r, t)|^2 = 1$. To get the in-trap condensate, we solve for the ground state of Eqn. G.22 with $V_{ext} = \frac{1}{2} m \omega_0^2 r^2$ by performing imaginary time propagation using the code in §I.4.2. Then we do forward time propagation of $\Phi(r, \mathbf{t})$ with $V_{ext} = 0$.

We compare the numerical solution of Eqn. G.22 to the standard approximate solution

$$|\Phi(r)|^2 = \frac{m}{4\pi\hbar^2 a} \left(\mu - \frac{1}{2} m \omega_0^2 r^2 \right) \quad (\text{G.23})$$

which we get by setting the kinetic energy term in Eqn. G.22 to zero (Thomas-Fermi approximation). As the condensate expands during TOF, the in-trap profile in Eqn. G.23 (Thomas-Fermi profile) is just rescaled as [221]

$$|\Phi(r, \mathbf{t})|^2 = \frac{1}{\lambda^3} |\Phi(r/\lambda)|^2, \quad (\text{G.24})$$

$$\ddot{\lambda} = \frac{\omega_0}{\lambda^4}. \quad (\text{G.25})$$




Figure G.7: (a) TF radius of the condensate versus expansion time (black curve). The TF radius is obtained by fitting the distribution that was numerically propagated using the GP equation to a TF profile. The data fits well to a heuristic expansion function given by Eqn. G.26 (red, dashed curve). At long times the condensate expands ballistically (gray, dashed curve). (b) Condensate distribution in-trap (black curve) and after 10ms expansion (red curve).

The numerical solutions are shown in Fig. G.7. If we fit the numerical data at each expansion time to the profile in Eqn. G.23 with the width as a free parameter, then the width versus expansion time fits well to the function

$$r_{TF}(\tau_0) = r_a + (r_b^2 + \alpha_a \tau_0 + (\alpha_b \tau_0)^2)^{0.5} \quad (\text{G.26})$$

where $r_a = 4.53$, $r_b = 4.13$, $\alpha_a = -3.44$ and $\alpha_b = 2.15$. If we fit the scaling solution to the same formula the agreement with the numerical calculation is within 0.5%.

G.3.3 Non-Condensate Expansion

To get the initial eigenstates in the trap, we diagonalize Eqn. G.21 with the potential given by Eqn. G.15 at $\mathbf{t} = 0$, so

$$\begin{aligned} E_{n,l} \psi_{n,l}(r) &= -\frac{\hbar^2}{2m dr^2} (\psi_{n,l}(r-dr) + \psi_{n,l}(r+dr) - 2\psi_{n,l}(r)) - \dots \\ &\quad -\frac{\hbar^2}{m(2dr)} \frac{1}{r} (\psi_{n,l}(r+dr) - \psi_{n,l}(r-dr)) + \dots \\ &\quad \left(\frac{1}{2} M \omega_0^2 r^2 + \frac{\hbar^2}{2mr^2} l(l+1) + \frac{8\hbar^2 \pi a N}{m} |\Phi(r)|^2 \right) \psi_{n,l}(r), \end{aligned} \quad (\text{G.27})$$

which can be solved with straightforward matrix math in the discretized r basis. Next we forward propagate each $\psi_{n,l}$ using the discretized time-dependent version of Eqn. G.21

$$\begin{aligned} i\hbar \frac{\psi_{n,l}(r, \mathbf{t} + d\mathbf{t}) - \psi_{n,l}(r, \mathbf{t} - d\mathbf{t})}{2d\mathbf{t}} &= -\frac{\hbar^2}{2m dr^2} (\psi_{n,l}(r-dr, \mathbf{t}) + \psi_{n,l}(r+dr, \mathbf{t}) - 2\psi_{n,l}(r, \mathbf{t})) \\ &\quad -\frac{\hbar^2}{m(2dr)} \frac{1}{r} (\psi_{n,l}(r+dr, \mathbf{t}) - \psi_{n,l}(r-dr, \mathbf{t})) + \\ &\quad \left(\frac{\hbar^2}{2mr^2} l(l+1) + \frac{8\hbar^2 \pi a N}{m} |\Phi(r, \mathbf{t})|^2 \right) \psi_{n,l}(r, \mathbf{t}). \end{aligned} \quad (\text{G.28})$$

In the calculation, $\Phi(r, \mathbf{r})$ is given by the Thomas-Fermi distribution with the widths given by Eqn. G.26. If $a = 0$ in Eqn. G.28, then this equation solves for the momentum distribution of the eigenstates (i.e., the Fourier transform). Solving for the non-condensate expansion is computationally intensive because we need to include a large number of states (the occupation of each state is less than one). Fortunately, for states with higher l values there is a $2l + 1$ degeneracy, so not all states need to be calculated.

For the actual calculation we propagated 17134 states comprising all states with energy less than $2\langle E \rangle$ where $\langle E \rangle$ is the mean thermal energy for a non-interacting gas with $T = 105\text{nK}$ ($N_0/N \approx 0.5$ at $T = 105\text{nK}$ for $N_0 = 10^5$). The maximum l state was $l = 260$ and including the $2l + 1$ degeneracy this corresponds to almost 3 million total states. For the in-trap eigenstates we used 1500 points with $dr = 33.56\text{nm}$. We propagated each state for $\mathbf{t} = 10\text{ms}$ with $d\mathbf{t} = 300\text{ns}$. Sample in-trap and propagated eigenstates are shown in Fig. G.8.




Figure G.8: (a) Sample non-condensate states in the trap. (b) Sample non-condensate states after 10ms expansion. The black curve is the lowest energy non-condensate state. It is pushed away from the center of the trap because of the condensate mean-field.

G.3.4 Results

Once we have numerically calculated the propagated eigenstates we construct the thermal distribution using Eqn. G.17. Fig. G.9 displays these distributions for $T = 70\text{nK}$ both in-trap and after 10ms expansion. Although the high momentum tails of the distribution are well-described by a non-interacting fit, near the condensate edge there is noticeable deviation. Under the condensate there is a prominent decrease in the non-condensate density (a “hole”) due to the condensate mean-field repulsion. This hole appears in-trap and after propagation. In the experiment, we measure the integrated 2D atomic distribution, which is also shown in Fig. G.9. The integration smooths the hole, but the distribution still deviates from the non-interacting distribution, which is peaked.




Figure G.9: Thermal distributions after numerically solving the mean-field expansion. (a) Condensate and non-condensate thermal distribution in-trap at $T = 70\text{nK}$. The thermal distribution is obtained using Eqn. G.16 after diagonalizing Eqn. G.27 to solve for the eigenstates. (b) Condensate and non-condensate distribution after expanding for 10ms. The non-condensate distribution was obtained by propagating the eigenstates using Eqn. G.28. For reference, the semiclassical non-interacting momentum distribution is shown. The non-interacting and mean-field distribution converge at the edge of the trap where the influence of the condensate is small. (c) The expanded distribution integrated along the imaging probe direction. The integration minimizes the indent in the non-condensate distribution caused by the repulsive condensate interactions.




Figure G.10: (a) Cross-section of simulated condensate fraction data at $s = 6$ for $\nu = 50\text{Hz}$ and $N = 1.5 \times 10^5$ atoms after 20ms expansion. The condensate distribution is determined using the procedure outlined in the main text. (b) Comparable experimental data, averaged over 7 shots. The experimental peaks are higher and narrower, indicated that our procedure over-estimates the contribution of the mean-field to the expansion. However, interactions must play a role in the experimental data as the first-order diffraction peaks are narrower than the central peak.




Figure G.11: Noise samples used in Fig. 7.9 taken from experimental images without atoms.

G.4 Simulated Bose-Hubbard TOF Distributions

In §7.3.2 we use simulated data to assess our lattice fitting procedure. In this section we will describe how that data is generated. The form of the simulated distribution is motivated by two limits. In the finite T non-interacting limit, the bimodal distribution is composed of a thermal non-condensate distribution given by Eqn. 7.8 and condensate peaks at the diffraction orders $(\pm 2n\hbar k)$. The entire distribution is multiplied by the Fourier transform of the Wannier function (Eqn. 7.10, essentially a Gaussian). As $T \rightarrow 0$ the thermal distribution disappears and only the condensate peaks remain. However, as $T \rightarrow 0$ in the interacting limit, described by the SDMFT (§5.3.2), the distribution remains bimodal. In addition to the condensate peaks, there is a quantum depleted non-condensate given by the Fourier transform of the Wannier function. Therefore, combining these two limits, our simulated distribution consists of a thermal non-condensate (Eqn. 7.8, N_T), a Gaussian quantum depleted non-condensate (Eqn. 7.10, N_{NC}), and the condensate (N_0). We use the SDMFT to determine the ratio of these three quantities for a fixed total number $N = N_T + N_{NC} + N_0$. The $T = 0$ condensate fraction gives the ratio N_{NC}/N_0 , which we assume is temperature independent, and the finite T condensate fraction gives the ratio N_0/N .

The above procedure gives the total number of condensate atoms. We use the Wannier envelope to determine how these atoms are divided into diffraction peaks, but we still need to decide on the form of the condensate peak after expansion. If we ignore interactions during expansion, then the condensate after expansion is the free propagated in-trap wavefunction given by SDMFT. If we drop the Fresnel terms, then this is just the Fourier transform. The peak height given by this procedure is far greater than the peak height observed experimentally. In Ref. [261] it was pointed out that these Fresnel terms should not be dropped for finite expansion times. If we include Fresnel terms, the peak height drops by an order of magnitude at 20ms (our experimental expansion time). However, the peaks are still higher than we measure experimentally. Also, we observe that the central peak is wider than the first order diffraction peaks. Since the central peak has more atoms, this suggests interaction effects need to be considered.

To take interactions into consideration during the condensate peak expansion, we have devised a heuristic method that qualitatively captures relevant features. First we use the Wannier envelope to determine the number of atoms in peak j , $N_{0,j}$. Then we use the lattice TF approximation to determine the condensate size if the number of atoms in the condensate were $N_{0,j}$ [323]

$$r_{TF}(0) = \sqrt{\frac{2\mu}{m\omega^2}}, \quad (\text{G.29})$$

$$\mu = \left(\frac{15}{16} \frac{d^3 m^{3/2} N_{0,j} U \omega^3}{\sqrt{2\pi}} \right)^{2/5}. \quad (\text{G.30})$$

Finally, the TF radius after expansion time τ is given by the scaling equation (Eqn. 4.20). For a spherically symmetric trap of $\nu = 50\text{Hz}$, the numerical solution of the scaling equation is.

$$\frac{r_{TF}(\mathbf{t})}{r_{TF}(0)} = 0.4 + \sqrt{0.605^2 - 0.01625\mathbf{t} + (0.25829\mathbf{t})^2}. \quad (\text{G.31})$$

In Fig. G.10 we compare the generated data to experimental data at $s = 6$. The quantitative agreement is not exact, but this gives a good approximation for testing our fit procedure. As expected, the widths are

overestimated for the first order diffraction peaks where the mean-field interaction effects are small.

To simulate experimental noise we take an image without atoms from the experiment and add it to the simulated data. This effectively captures the fringes and photon shot noise from the experiment. In Fig. G.11 we show the three noise samples used in §7.3.2. Final simulated distributions with noise are shown in Fig. 7.9.

G.5 Deriving the Probe Integrated Quasimomentum Distribution

In this section we derive Eqn. 4.19, the 2D quasimomentum distribution after integrating along the probe direction in the ^{87}Rb apparatus. To keep the derivation free of extraneous symbols, we drop all prefactors in front of the distributions. These are irrelevant since we will ultimately leave the prefactor as a free parameter in the fit. The starting point of the derivation is the 3D quasimomentum distribution for the lattice plus harmonically trapped gas (Eqn. 4.17)

$$n(q_1, q_2, q_3) \propto \text{Li}_{3/2} \left[\mp \zeta e^{-2t\beta \sum_i \{1 - \cos(q_i \pi / q_B)\}} \right] \quad (\text{G.32})$$

where q_i are the quasimomentum along the lattice axes, $\beta = (k_B T)^{-1}$, and ζ is the fugacity. In Eqn. 3.4 these lattice axes are written in terms of the fixed coordinate system defined by Fig. 3.1. We integrate the 3D quasimomentum distribution along \hat{x} , the probe direction, and image in the yz plane. The boundary condition for the integral is set by the condition that $|q_i| < q_B$ (i.e., the quasimomentum is only defined in the first Brillouin zone). In the following we assume that the physical coordinates are written in terms of the dimensionless wavevector coordinates after expansion time, τ , i.e., $x \rightarrow x(md/\hbar\tau)$. Therefore, the boundary conditions on the spatial coordinates are

$$|y - z| < \sqrt{2}\pi, \quad (\text{G.33})$$

$$\left| x + \frac{1}{\sqrt{2}}(y + z) \right| < \sqrt{2}\pi, \quad (\text{G.34})$$

$$\left| -x + \frac{1}{\sqrt{2}}(y + z) \right| < \sqrt{2}\pi. \quad (\text{G.35})$$

Since the integral is along \hat{x} we can arrange these conditions into the form

$$|y - z| < \sqrt{2}\pi, \quad (\text{G.36})$$

$$|y + z| < 2\pi, \quad (\text{G.37})$$

$$|x| < \sqrt{2} \left(\pi - \frac{1}{2}|y + z| \right). \quad (\text{G.38})$$

We define new variables in the yz plane

$$\alpha_1 = \frac{y+z}{2}, \quad (\text{G.39})$$

$$\alpha_2 = \frac{y-z}{\sqrt{2}}, \quad (\text{G.40})$$

$$x = x/\sqrt{2}. \quad (\text{G.41})$$

The 2D quasimomentum distribution in terms of these variables is

$$n(\alpha_1, \alpha_2) \propto \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx \operatorname{Li}_{3/2} \left[\zeta e^{-2t\beta\{3-\cos(\alpha_2)-\cos(x+\alpha_1)-\cos(-x+\alpha_1)\}} \right], \quad (\text{G.42})$$

$$\propto \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx \sum_{n=1}^{\infty} \frac{\zeta^n e^{-2tn\beta\{3-\cos(\alpha_2)-\cos(x+\alpha_1)-\cos(-x+\alpha_1)\}}}{n^{3/2}}, \quad (\text{G.43})$$

$$\propto \sum_{n=1}^{\infty} \frac{\zeta^n e^{-2tn\beta\{3-\cos(\alpha_2)\}}}{n^{3/2}} \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx e^{2tn\beta\{\cos(x+\alpha_1)+\cos(-x+\alpha_1)\}}. \quad (\text{G.44})$$

Expanding the cosines in the integrand

$$\int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx e^{2tn\beta\{\cos(x+\alpha_1)+\cos(-x+\alpha_1)\}} = \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx e^{4tn\beta \cos(x) \cos(\alpha_1)}. \quad (\text{G.45})$$

To solve the above we substitute in a series expression for the exponential (Jacobi-Anger expansion)

$$e^{\xi \cos(x)} = I_0(\xi) + \sum_{n=-\infty, n \neq 0}^{\infty} I_n(\xi) e^{inx}, \quad (\text{G.46})$$

and so

$$\begin{aligned} \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx e^{4tn\beta \cos(x) \cos(\alpha_1)} &= \int_{-(\pi-|\alpha_1|)}^{\pi-|\alpha_1|} dx \left(I_0(4tn\beta \cos(\alpha_1)) + \sum_{n=-\infty, n \neq 0}^{\infty} I_n(4tn\beta \cos(\alpha_1)) e^{inx} \right), \\ &= 2(\pi - |\alpha_1|) I_0(4tn\beta \cos(\alpha_1)) + \\ &\quad \sum_{n=1}^{\infty} I_n(4tn\beta \cos(\alpha_1)) \frac{4 \sin(n(\pi - |\alpha_1|))}{n}. \end{aligned} \quad (\text{G.47})$$

The full distribution is

$$\begin{aligned} n(\alpha_1, \alpha_2) \propto \sum_{n=1}^{\infty} \frac{\zeta^n e^{-2tn\beta\{3-\cos(\alpha_2)\}}}{n^{3/2}} \\ \left[2(\pi - |\alpha_1|) I_0(4tn\beta \cos(\alpha_1)) + \sum_{n=1}^{\infty} I_n(4tn\beta \cos(\alpha_1)) \frac{4 \sin(n(\pi - |\alpha_1|))}{n} \right], \end{aligned} \quad (\text{G.48})$$

which is a doubly infinite sum. An example of the distribution given by Eqn. G.48 is shown in Fig. 4.3.

Appendix H

Circuits

This appendix overviews three specific circuits developed for this thesis work. Other circuits developed for the ^{87}Rb apparatus may be found in [41] and for the ^{40}K apparatus at <http://www.physics.utoronto.ca/~astummer/>. For further information about the circuits presented in this appendix (e.g., board layouts), please contact the author.

H.1 Laser Lock

The laser lock circuit is a generic laser servo designed to provide piezo and current feedback for stabilizing the laser frequencies in the ^{40}K apparatus. The input to the circuit is an error signal (e.g., from a saturated-absorption spectroscopy setup) and the error signal zero crossing (lock-point) is set using a potentiometer on the front panel of the circuit. Switches on the front panel enable/disable feedback to the laser; the circuit generates proportional plus integrator (PI) error output for piezo feedback and a high pass proportional output for current feedback.

The piezo output has an front panel adjustable bias, which centers the laser on the error feature. To scan across this feature there is a front panel adjustable ramp (supplied internally or externally) that can be switched on/off. One distinct feature of this circuit is an integrator fault detect. If the integrator is engaged and the absolute value of the output goes beyond a user-set voltage, a fault condition is tripped and the integrator is disabled until it is manually switched off and reset. A front panel LED alerts the user when the integrator fault trips. This feature helps to prevent large changes to the piezo voltage which can cause the laser to mode hop.

This lock circuit has two optional settings: a double integrator on the piezo and a fast (up to 10MHz) proportional feedback. The double integrator may be useful when absolute laser frequency stability is required. The fast error is for direct feedback to the diode for narrowing the linewidth of lasers with fast noise processes. In practice, we have not used these features, but they are available for future consideration.

The front and back panels of an assembled laser lock are illustrated in Fig. H.1 with locking instructions. The circuit schematic is shown in Figs. H.2, H.3, and H.4.




Figure H.1: Photograph of the assembled laser lock circuit front and back panel. To lock, first supply an error signal input to the circuit and connect the outputs to the laser controller. Next, disable the piezo and current servo integrators using the switches on the front, and turn on the ramp. Use the ramp adjust knob on the front panel to increase/decrease the ramp range. Use the piezo bias to center the error feature. When the error feature is centered, turn off the ramp and turn on the piezo integrator. If the lock does not “capture” then the integrator will rail and the fault light will turn on. Turn the piezo integrator off, press the reset button, and repeat the previous steps. If the piezo locks successfully then turn on the current error. DIP switches on the board invert the piezo and current error outputs if the sign of the error signal is incorrect.




Figure H.2: Pg. 1 of the laser lock circuit schematic.

Figure H.3: Pg. 2 of the laser lock circuit schematic.




Figure H.4: Pg. 3 of the laser lock circuit schematic.

H.2 Experiment Monitor and Interlock

This circuit is a dual purpose monitor and interlock. It can monitor the MOT fluorescence, a water flowmeter voltage, and twelve thermistors. Using a rotary switch on the front panel (see Fig. H.5), one can select any of these signals and view them on a LCD voltmeter display (top-left of Fig. H.5). The water flow and thermistors are compared against individual setpoints. If these setpoints are exceeded then an interlock is tripped.

The first part of this circuit, the MOT photodiode monitor, subtracts and then amplifies the signals from two unamplified photodiodes. The first of these photodiode is optimized to collect the MOT fluorescence and the other photodiode collects stray light from the MOT beams. The stray light photodiode input is subtracted from the fluorescence photodiode input to obtain a pure fluorescence signal (i.e., a signal that is zero if the MOT beams are on, but there are no atoms). This signal is then amplified to produce a clean, high contrast monitor of the MOT size. The schematic for this circuit is shown in Fig. H.6.

On the interlock board, the water flow and twelve thermistor inputs are compared to individual setpoints. For the thermistors this is accomplished by incorporating each thermistor into a Wheatstone bridge. When the setpoint is exceeded, the interlock trips and remains tripped until the input goes back below the setpoint and the interlock is manually reset via a push button on the front panel. For historic reasons the thermistors are in three separate groups. The circuit can be setup so that each groups trips independent of the other groups (e.g., if the thermistors correspond to coils run by separate supplies). However, we run the circuit in the mode where any thermistor trips all the interlocks. The water flow is always set to trip all interlocks. The interlock circuit schematics are shown in Figs. H.7, H.8, and H.9.

When an interlock is tripped, an output from the board goes to TTL high. We use this TTL signal to drive three relays, shown in Fig. H.10. These relays are connected to shutoff switches on our power supplies.

H.3 Mixer

To create an error signal for laser locking, one method is to frequency modulate the pump or probe beam of a saturated-absorption spectroscopy setup at f_0 and then demodulate the probe photodiode signal at f_0 . This type of modulation/demodulation spectroscopy is an effective method to generate a signal with a true zero crossings that is immune to first-order temperature and pressure drifts in the atomic vapor cell. This circuit mixes the probe photodiode signal with the modulation signal and then low pass filters the output to perform demodulation. The schematic is shown in Fig. H.11




Figure H.5: MOT photodiode monitor and interlock front panel.




Figure H.6: MOT photodiode Monitor Schematic.




Figure H.8: Pg. 2 of the interlock.




Figure H.9: Pg. 3 of the interlock.




Figure H.10: Relay board for the interlock. The output of the interlock board are a set of TTL signals. This board, separate from the interlock PCB, uses the TTL signals to drive a set of relays, shown as resistors in the schematics. These relays are hooked up to the shutoff inputs of the various power supplies in the lab.




Figure H.11: Mixer circuit schematic. The signal is mixed using part SRA-6+ from Mini-Circuits. Out #1 is sent to the laser servo for locking and Out #2 is a filtered, amplified monitor output.

Appendix I

Code

This appendix contains MATLAB code developed for this thesis.

I.1 SDMFT Code

This code implements the theory discussed in §5.3.2. The following function, `sdmft_solve`, computes the properties at a single site given temperature `T` and chemical potential `mu`.

```
%=====
%-----
%Inputs:
%mu: Chemical Potential
%T: Temperature
%U: Bose-Hubbard interaction parameter
%J: Bose-Hubbard tunneling parameter
%Outputs:
%y(1): N
%y(2): H-mu*N
%y(3): S
%y(4): |<a>|^2 (Condensate number, N0)
%y(5): <a>
%y(6): <N^2>
%y(7): (<N^2>-<N>^2)^(0.5)
%y(8): Heat Capacity
%-----

function y = sdmft_solve(mu,T,U,J,flag)

%Number of fock states for the calculation
nvals = 0:15;
nlength = length(nvals);

if T ==0
    %scale everything by kT for numerical reasons
    beta = 1/(1.38E-23*T);
    mu = mu*beta;
    U = U*beta;
    J = J*beta;
else
    %zero temp, scale everything by U
    U0 = U;
```

```

    mu = mu/U0;
    U = 1;
    J = J/U0;
end

%guess mean field parameter
a = 2.7;
%coordination number
z = 6;

%Do maximum 100 iterations
for i = 1:100

    %Fock basis Hamiltonian
    H0 = zeros((nlength+2),(nlength+2));

    %Add terms to the Hamiltonian
    for j = 2:(nlength+1)
        H0(j,j) = U/2*nvals(j-1)*(nvals(j-1)-1)-nvals(j-1)*mu+z*a*conj(a)*J;
        H0(j+1,j) = -z*J*a*sqrt(nvals(j-1)+1);
        H0(j-1,j) = -z*J*conj(a)*sqrt(nvals(j-1));
    end

    %Take off padding
    H = H0(2:(nlength+1),2:(nlength+1));

    %Diagonalize
    [estates evals] = eig(H);

    %Calculate mean field parameter
    aold = a;
    a = 0;

    if T = 0

        %Finite temperature so calculate by summing states with a Boltzmann factor
        for j = 1:length(evals(1,:))
            for k = 2:nlength
                a = a + exp(-(evals(j,j)))*sqrt(nvals(k))*conj(estates(k-1,j))*estates(k,j);
            end
        end

        a = a/sum(exp(-(diag(evals)))));

    else
        %Zero temperature so calculate using the ground state
        for j = 1:1
            for k = 2:nlength
                a = a + sqrt(nvals(k))*conj(estates(k-1,j))*estates(k,j);
            end
        end
    end
end

```

```

    %Has the mean-field parameter converged?
    if abs(aold-a)/abs(aold) < 1E-6
        break
    end

end

if T ==0
    %make a matrix with boltzmann factors and the squares of the eigenstates
    part_func = sum(exp(-(diag(evals))));
    boltz_states = (conj(estates).*estates)*diag(exp(-(diag(evals))))/part_func;
    energ_states = boltz_states*evals;
    energ_states2 = boltz_states*(evals.^2);
    n_states = diag(nvals)*boltz_states;
    n2_states = diag(nvals.^2)*boltz_states;
else
    n_states = nvals'.*abs(estates(:,1)).^2;
    n2_states = (nvals.^2)'.*abs(estates(:,1)).^2;
    energ_states = evals(1,1)*U0;
    energ_states2 = evals(1,1)^2*U0^2;
end

%calculate N
y(1) = sum(sum(n_states));
%calculate H-uN
y(2) = sum(sum(energ_states));

%calculate S (kb*log(Z)+U/T)
if T ==0
    y(3) = log(part_func)+y(2);
else
    y(3) = 0;
end

%calculate condensate number
y(4) = a*conj(a);
%<a>
y(5) = a;
%<n^2>
y(6) = sum(sum(n2_states));
%variance
y(7) = sqrt(y(6)-y(1)^2);
%this is the heat capacity
y(8) = sum(sum(energ_states2))-sum(sum(energ_states)).^2;

end
%=====

```

The following function takes the above code and computes the properties of the trapped gas using LDA.

```

%=====
%-----
%Inputs:

```

```

%mu: Chemical Potential
%T: Temperature
%U: Bose-Hubbard interaction parameter
%J: Bose-Hubbard tunneling parameter
%m: Mass
%w: Trap frequency
%d: Lattice spacing
%Outputs:
%y(1): N
%y(2): S
%y(3): N0 (Condensate Number)
%y(4): N0/N
%y(5): Energy
%y(6): Heat Capacity
%-----

function y = sdmft_solve_trap(mu,T,U,J,m,w,d)

%Number of sites radially to sum over
nsites = 60;
%Check our radial estimate against an exact sum
exact_sum = 0;
%Trap Energy
trapenerg = 0.5*m*w^2*d^2;

if exact_sum

    %Do a radial approximation to the trap

    r = 0:nsites;

    %calculate the chemical potential in the LDA
    mulda = mu-trapenerg*(r.^2);

    ntest = zeros(1,length(mulda));
    stest = zeros(1,length(mulda));
    n0test = zeros(1,length(mulda));
    etest = zeros(1,length(mulda));
    ctest = zeros(1,length(mulda));
    n0test2 = zeros(1,length(mulda));

    %calculate n,s,n0 for the different chemical potentials across the
    %lattice
    for i = 1:length(mulda)
        y2 = sdmft_solve(mulda(i),t,U,J);
        ntest(i) = y2(1);
        stest(i) = y2(3);
        n0test(i) = y2(4);
        etest(i) = y2(2);
        ctest(i) = y2(8);
    end
end

```



```

    %calculate the totals across the lattice by weighting with the
    %appropriate area factors (4*pi*r^2)
    y(1) = sum(4*pi*ntest.*r.^2);
    y(2) = sum(4*pi*stest.*r.^2);
    y(3) = sum(4*pi*n0test.*r.^2);
    y(5) = sum(4*pi*etest.*r.^2)*t+mu/1.38E-23*y(1);
    y(6) = sum(4*pi*ctest.*r.^2)/y(1);
else

    %Do the exact 3D sum
    y(1) = 0;
    y(2) = 0;
    y(3) = 0;

    for i = 1:nsites
        for j = i:nsites
            for k = j:nsites
                mulda = mu-trapenerg*(i^2+j^2+k^2);
                y2 = sdmft_solve(mulda,t,U,J);
                y(1) = y(1) + 6*y2(1);
                y(2) = y(2) + 6*y2(3);
                y(3) = y(3) + 6*y2(4);
            end
        end
        disp(i);
    end
    y(1:3) = 8*y(1:3);
end
y(4) = y(3)/y(1);
end
%=====

```

I.2 HFBP

This code implements the HFBP theory discussed in §5.3.1. To run the code we first need to create an array with the non-interacting energies of a uniform lattice in k-space (quasimomentum space).

```

%=====
%-----
%Inputs:
%nsites:  Number of sites in the lattice array
%Outputs:
%y:  Energies (per unit 4t) in one octant of a lattice in k space
%-----
function y = make_lattice_array(nsites)

j = 0:nsites;
[jx, jy, jz] = meshgrid(j,j,j);
y = ((1-cos(pi*jx/nsites))*(1-cos(pi*jy/nsites))*(1-cos(pi*jz/nsites)));

end

```

```
%=====
```

Next, we use the array created from the above function to compute the HFBP energies in a uniform lattice. Once these have been computed, we divide by the number of sites to get the quantities per site in the uniform lattice. The following function `hfbp_solve`, also contains a subfunction `nT_solve`.

```
%=====
%-----
%Inputs:
%mu: Chemical Potential
%T: Temperature
%U: Bose-Hubbard interaction parameter
%J: Bose-Hubbard tunneling parameter
%latt_array: Lattice k-space energies from "make_lattice_array.m"
%Outputs:
%y(1): N per site
%y(2): N0 per site
%y(3): S per site
%-----

function y = hfbp_solve(mu,T,U,J,latt_array)

beta = 1/1.38E-23/T;
e0 = 0;

%Nsites here refers to the number of sites on either
%side of zero
nsites = (length(latt_array(1,1,:))-1)*2;
nsites_mid = 1;

%Use latt_array to compute the non-interacting energy
ej = 2*J*latt_array+e0;

%The absolute ground state cannot contribute to these sums because it is the condensate.
%Set it to a really high energy so that it is essentially removed from the
%calculation.
ej(nsites_mid,nsites_mid,nsites_mid) = 1;

%For the HFBP solution we have to solve self-consistently for the value of nT (thermal number)
%from mu and T. If we have nT,mu and T we can calculate n (total number) and n0 (condensate number).
%Seed with self-consistency loop with the non-interacting value of nT
nT_solve(0,0);

%maximum number of iterations
n_max_iters = 30;

dx = 1;

%Solve by interpolation. Make the region bigger or smaller depending on if
%the results are getting better or worse
for i = 1:n_max_iters
```

```

%This is our best guess for the thermal number
nT_seed = nT;

%create an 10 element array of thermal number guesses around our best guess in a
%region of size 1-dx to 1+dx
nT_guess = linspace(max([(1-dx),0])*nT_seed,(1+dx)*nT_seed,10);
nT2 = zeros(1,10);

%solve for the thermal number generated using our guesses
for j = 1:10
    nT_solve(nT_guess(j),U);
    nT2(j) = nT;
end

%interpolate to the thermal number where the guess and the result are
%the same
nT_new = interp1(spline(nT_guess,nT_guess-nT2,linspace(nT_guess(1),...
nT_guess(end),1000)),linspace(nT_guess(1),nT_guess(end),1000),0);

%nT cannot be negative
if nT_new<0
    nT_new=0;
end

%Need to make the guess interval bigger if the interpolation fails
if isnan(nT_new)
    dx = dx*3;
    nT = nT_seed;
    continue;
end

%solve for thermal number using the interpolated guess value
nT_solve(nT_new,U);

%check self-consistency to 0.1%
if (sqrt((nT-nT_new)^2/nT_new^2) < 10^(-3) || (nT==0 && nT_new==0))
    break;
end

%If we were able to interpolate, but with not enough precision then
%We need to calculate in a smaller region for higher precision
nT = nT_new;
dx = dx/2;

if i==n_max_iters
    disp(['Did not converge for mu ' num2str(mu)])
end

end

%calculate S (units of kB)
s = (8*sum(sum(sum(beta*Ej.*(exp(beta*Ej)-1).^(-1)-log(1-exp(-beta*Ej)))))-...
4*sum(sum(sum(beta*Ej(1, :, :).*(exp(beta*Ej(1, :, :))-1).^(-1)-log(1-exp(-beta*Ej(1, :, :))))))-...

```

```

4*sum(sum(sum(beta*Ej(:,1,:).*(exp(beta*Ej(:,1,:))-1).^(-1)-log(1-exp(-beta*Ej(:,1,:)))))-...
4*sum(sum(sum(beta*Ej(:,,1).*(exp(beta*Ej(:,,1))-1).^(-1)-log(1-exp(-beta*Ej(:,,1)))))+...
2*sum(sum(sum(beta*Ej(1,1,:).*(exp(beta*Ej(1,1,:))-1).^(-1)-log(1-exp(-beta*Ej(1,1,:)))))+...
2*sum(sum(sum(beta*Ej(1,:,1).*(exp(beta*Ej(1,:,1))-1).^(-1)-log(1-exp(-beta*Ej(1,:,1)))))+...
2*sum(sum(sum(beta*Ej(:,1,1).*(exp(beta*Ej(:,1,1))-1).^(-1)-log(1-exp(-beta*Ej(:,1,1))))))...
/(nsites^3-1);

if (nT==0&& n0==0)
    s=0;
end

y(1) = nT+n0;
y(2) = n0;
y(3) = s;

%This is a subfunction which calculates the thermal number from a guess value for the thermal
%number nT_0 for a given interaction strength U0
function nT.solve(nT_0,U0)

    mu0 = mu;

    %From the thermal number guess and chemical potential get n0
    if U0==0
        %If non-interacting then set condensate and chemical potential to zero
        %For the purposes of this calculation
        n0 = 0;
        mu0 = 0;
    else
        n0 = mu0/U-2*nT_0;
    end

    %n0 cannot be negative
    if n0<0
        n0=0;
    end

    %HF energy
    HFj = (ej-e0)-mu0+2*U*(n0+nT_0);

    %HFBP energy
    Ej = -(U*n0)^2+HFj.^2).^^(0.5);

    %solve for the thermal number
    %odd sum because latt_array was one octant of the full k-space
    nT = (8*sum(0.5*(HFj(:)./Ej(:).*coth(beta*Ej(:)/2)-1))-...
        4*sum(sum(0.5*(HFj(1,,:)./Ej(1,,:).*coth(beta*Ej(1,,:)/2)-1)))-...
        4*sum(sum(0.5*(HFj(:,1,:)./Ej(:,1,:).*coth(beta*Ej(:,1,:)/2)-1)))-...
        4*sum(sum(0.5*(HFj(:,,1)./Ej(:,,1).*coth(beta*Ej(:,,1)/2)-1)))+...
        2*sum(sum(0.5*(HFj(1,1,:)./Ej(1,1,:).*coth(beta*Ej(1,1,:)/2)-1)))+...
        2*sum(sum(0.5*(HFj(1,:,1)./Ej(1,:,1).*coth(beta*Ej(1,:,1)/2)-1)))+...
        2*sum(sum(0.5*(HFj(:,1,1)./Ej(:,1,1).*coth(beta*Ej(:,1,1)/2)-1))))/(nsites^3-1);

```

```

    %nT cannot be negative
    if nT<0
        nT=0;
    end

end

end

%=====

Solve for the trapped gas using LDA.

%=====
%-----
%Inputs:
%mu: Chemical Potential
%T: Temperature
%U: Bose-Hubbard interaction parameter
%J: Bose-Hubbard tunneling parameter
%m: Mass
%w: Trap frequency
%d: Lattice Spacing
%lattice_array: Lattice k-space energies from "make_lattice_array.m"
%Outputs:
%y(1): N per site
%y(2): NO per site
%y(3): S per site
%-----

function y = hfbp_solve_trap(mu,T,U,J,m,w,d,lattice_array)

%Outputs
%y(1): total N
%y(2): total S
%y(3): total NO
%y(4): cond frac

%Number of lattice sites to sum over
nsites = 200;

%get the trap energy
trapenerg = 0.5*m*w^2*d^2;

r = 0:nsites;

%calculate the chemical potential in the LDA
mulda = mu-trapenerg*(r.^2);

%calculate n,s,n0 for the different chemical potentials across the
%lattice
ntest = zeros(1,length(mulda));
stest = zeros(1,length(mulda));

```

```

n0test = zeros(1,length(mulda));

%calculate n,s,n0 for the different chemical potentials across the
%lattice
for i = 1:length(mulda)

    y2 = hfbp_solve(mulda(i),T,U,J,latt_array);
    ntest(i) = y2(1);
    stest(i) = y2(3);
    n0test(i) = y2(2);

    %no more sites needed
    if i>10 && (4*pi*ntest(i)*r(i)^2) < 1
        %disp(i)
        break;
    end

end

%calculate the totals across the lattice by weighting with the
%appropriate area factors (4*pi*r^2)
y(1) = sum(4*pi*ntest.*r.^2);
y(2) = sum(4*pi*stest.*r.^2)*1.38E-23;
y(3) = sum(4*pi*n0test.*r.^2);
y(4) = y(3)/y(1);

end
%=====

```

I.3 Small System Exact Hubbard Solver

The following code performs exact diagonalization to solve small system Hubbard models.

I.3.1 1D BH Model

To solve the exact system, we first need a function to generate a matrix of all possible states in the Fock basis. The following function generates this matrix recursively. The state matrix is independent of the dimensionality of the problem and only depends on the total number of particles and the number of sites.

```

%=====
%-----
%Inputs:
%n:  Number of sites
%m:  Number of particles
%max_m:  Maximum number of particle per site
%Outputs:
%y:  n x (number of states) matrix where the columns represent each state
%and the rows are the number of particle in site i for that state
%-----

```

```

function y = create_state_matrix_bosons(n,m,max_m)

if nargin < 3
    max_m = m;
end

y = [];
%no sites
if n==0
    return;
end

%trivial case
if n==1
    if m>max_m
        return;
    end
    y = m;
    return;
end

%no atoms
if m==0
    y = zeros(n,1);
    return;
end

count = 1;
%create matrix
psi = zeros(n,factorial(n+m-1)/factorial(n-1)/factorial(m));

if m>max_m
    mm = max_m;
else
    mm = m;
end

for j = 0:mm
    %recursively get all the possible states if we put j particles in
    %site i
    psi2 = create_state_matrix(n-1,m-j,max_m);
    psi(1,count:(count+size(psi2,2)-1))=j;
    psi(2:end,count:(count+size(psi2,2)-1)) = psi2;

    %readjust counter
    count = count+size(psi2,2);
end
y = psi(:,1:(count-1));

return;
%=====

```

Next, we need to construct the BH Hamiltonian for 1D given the state matrix generated in the previous function. The diagonal terms of the matrix are proportional to U and the off-diagonal are proportional to t (specified as J in the code). This matrix is sparse, i.e., most elements of the matrix are zero. Therefore, this function exports the coordinates of the non-zero off-diagonal elements, so that we can easily replace these terms if t changes. Once we have the Hamiltonian we can diagonalize to get the eigenstates and/or use matrix multiplication to time-propagate an initial state.

```
%=====
%-----
%Inputs:
%J: Bose-Hubbard tunneling parameter
%U: Bose-Hubbard interaction parameter
%state_matrix: State matrix from function "create_state_matrix_bosons"
%nsites: number of sites (needs to be the same as in state_matrix)
%Outputs:
%H: Hamiltonian
%Jrep: an nx2 matrix given the matrix elements in H which need to be
%replaced if J changes
%-----

function [H Jrep] = construct_BH_matrix(J,U,state_matrix,nsites)

%construct an n-site BH matrix in a 1D chain with a fixed particle number
%The first row of state_matrix is the one end of the chain and the last row
%the end of the chain
H = zeros(size(state_matrix,2));
Jrep = zeros(size(state_matrix,2),2);
jrep_count = 1;

for i = 1:size(state_matrix,2)
    for j = i:size(state_matrix,2)
        %tunneling terms
        for k = 1:nsites

            test_vector = state_matrix(:,i);

            %tunnel to the "right"
            if k == nsites
                test_vector(k) = test_vector(k)-1;
                test_vector(k+1) = test_vector(k+1)+1;

                if (state_matrix(:,j)-test_vector)

                    if H(i,j)==0
                        Jrep(jrep_count,:) = [i j];
                        jrep_count = jrep_count+1;
                    end

                    H(i,j) = H(i,j) - J*sqrt(state_matrix(k+1,j)*state_matrix(k,i));
                end
            end
        end
    end
end
```



```

test_vector = state_matrix(:,i);

%tunnel to the "left"
if k =1
    test_vector(k) = test_vector(k)-1;
    test_vector(k-1) = test_vector(k-1)+1;

    if (state_matrix(:,j)-test_vector)

        if H(i,j)==0
            Jrep(jrep_count,:) = [i j];
            jrep_count = jrep_count+1;
        end
        H(i,j) = H(i,j) - J*sqrt(state_matrix(k-1,j)*state_matrix(k,i));

    end
end
end
%interaction terms
if (i==j)
    H(i,j) = U/2*sum(state_matrix(:,i).*(state_matrix(:,i)-1));
end
end
end

for i = 1:size(state_matrix,2)
    for j = (i+1):size(state_matrix,2)
        H(j,i) = H(i,j);
    end
end
Jrep = Jrep(1:(jrep_count-1),:);
end
%=====

```

The following function constructs the matrix corresponding to the momentum field operator. With this matrix we can determine the density at momentum p using matrix multiplication.

```

%=====
%-----
%Inputs:
%p: Momentum value to construct the matrix for
%state_matrix: State matrix from function "create_state_matrix.bosons"
%nsites: number of sites (needs to be the same as in state_matrix)
%Outputs:
%y: Momentum field operator matrix
%-----

function y = create_p_matrix(p,state_matrix,nsites)

%create the field operator matrix for momentum
y = zeros(size(state_matrix,2));

%also symmetric so only due top triangle

```

```

for i = 1:size(state_matrix,2);
    for j = i:size(state_matrix,2)

        %tunneling terms
        for k = 1:nsites
            for kk = 1:nsites

                test_vector = state_matrix(:,i);
                test_vector(k) = test_vector(k)-1;
                test_vector(kk) = test_vector(kk)+1;

                if (state_matrix(:,j)-test_vector)
                    y(i,j) = y(i,j)+exp(sqrt(-1)*p*(k-kk))*sqrt(state_matrix(k,i)*state_matrix(kk,j));
                end
            end
        end
    end
end

for i = 1:size(state_matrix,2);
    for j = (i+1):size(state_matrix,2)
        y(j,i) = conj(y(i,j));
    end
end
end
end
%=====

```

I.3.2 5 Site 2D FH Model

There are new considerations for a small system FH solver as compared to the BH solver. First, we need to include two states, otherwise the solutions are trivial. Also, the number of Fock states per site is naturally truncated by Pauli blocking, so creating the state matrix is more straightforward. However, one has to be extremely careful about the sign of the creation/annihilation operators acting on the states when constructing the Hamiltonian and/or any operator matrix. The following function creates the state matrix and FH Hamiltonian for 5 sites in 2D (one site in the center, connected to five outside sites). The state matrix is organized so that each row corresponds to a given site. The first row is the center site and all other rows correspond to the outside sites. Two state matrices are generated, one for state 1 (“spin up”) and the other for state 2 (“spin down”).

```

%=====
%-----
%Inputs:
%J: Fermi-Hubbard tunneling parameter
%U: Fermi-Hubbard interaction parameter between states 1 and 2
%mu1: Chemical Potential for state 1
%mu2: Chemical Potential for state 2
%Outputs:
%H: Hamiltonian
%Psi1: state matrix for state 1
%Psi2: state matrix for state 2

```

```

%-----

function [H Psi1 Psi2] = construct_FH_matrix(J,U,mu1,mu2)

%construct a 5 site FH matrix with the basis in each site (2D)
%|00>,|10>,|01>,|11>

n_sites = 3;

%create state matrix
%for each spin...can only be 1 or 0
state_matrix1 = zeros(n_sites,4^n_sites);
state_matrix2 = zeros(n_sites,4^n_sites);

state_indices = 0:(4^n_sites-1);

for i = 1:n_sites
    state_matrix1(i,:) = floor(mod(state_indices/4^(n_sites-i),2));
    state_matrix2(i,:) = floor(mod(state_indices/4^(n_sites-i),4)/2);
end

H = zeros(4^n_sites);

for i = 1:(4^n_sites)
    for j = (i+1):(4^n_sites)
        for k = 1:1
            %tunneling terms
            if k==1
                kk2 = 2:n_sites;
            elseif k==2
                kk2 = 4;
            else
                kk2 = 5;
            end

            for kk = kk2
                for l = 1:2
                    if l==1 %up spin

                        test_vector1 = state_matrix1(:,i);
                        test_vector2 = state_matrix1(:,j);
                        test_vector3 = state_matrix2(:,i);

                        %other spin state needs to be the same
                        if sum((state_matrix2(:,i)==state_matrix2(:,j))) =n_sites
                            continue;
                        end

                    else %down spin

                        test_vector1 = state_matrix2(:,i);
                        test_vector2 = state_matrix2(:,j);
                        test_vector3 = state_matrix1(:,i);
                    end
                end
            end
        end
    end
end

```

```

        %other spin state needs to be the same
        if sum((state_matrix1(:,i)==state_matrix1(:,j))) ==n_sites
            continue;
        end

    end

    %subtract particle from center site and add to site k
    if (test_vector1(k)==1 && test_vector1(kk)==0)

        %this is to take into account the phase associated
        %with the anti-commuting fermion operators...I'm
        %not sure if I did this exactly right...sort of
        %trial and error
        prefactor = sum(test_vector1(2:(kk-1)))+sum(test_vector3(1:(kk-1)));

        if l==2
            prefactor = prefactor + test_vector3(1)+test_vector3(kk);
        end

        test_vector1(k) = 0;
        test_vector1(kk) = 1;

        if sum((test_vector1==test_vector2))==n_sites
            H(i,j) = H(i,j) - (-1)^(prefactor)*J;
        end

        continue;

    end

    %add particle to center site from site k
    if (test_vector1(kk)==1 && test_vector1(k)==0)

        prefactor = sum(test_vector1(2:(kk-1)))+sum(test_vector3(1:(kk-1)));

        if l==2
            prefactor = prefactor + test_vector3(1)+test_vector3(kk);
        end

        test_vector1(kk) = 0;
        test_vector1(k) = 1;

        if sum((test_vector1==test_vector2))==n_sites
            H(i,j) = H(i,j) - (-1)^(prefactor)*J;
        end

    end

end
end
end
end
end

```

```

%interaction and chemical potential terms
H(i,i) = U*sum(state_matrix1(:,i).*state_matrix2(:,i))-...
    mu1*sum(state_matrix1(:,i))-mu2*sum(state_matrix2(:,i));

end

%make Hermitian
for i = 1:(4^n_sites)
    for j = (i+1):(4^n_sites)
        H(j,i) = conj(H(i,j));
    end
end

Psi1 = state_matrix1;
Psi2 = state_matrix2;
end
%=====

```

This functions constructs a matrix corresponding to the spin correlation operator $\sum_{j=2}^4 \vec{S}_j \cdot \vec{S}_1$.

```

%=====
%-----
%Inputs:
%Psi1: state matrix for state 1
%Psi2: state matrix for state 2
%Outputs:
%S: Spin correlation matrix
%-----
function [S] = construct_spin_corr_matrix(Psi1,Psi2)

%construct the spin correlation matrix
n_sites = 3;
S = zeros(4^n_sites);

for i = 1:(4^n_sites)
    for j = i:(4^n_sites)
        %Sx_1*Sx_j+Sy_1*Sy_j
        for kk = 2:n_sites
            for ii = 0:1
                for jj = 0:1
                    %spin up
                    test_vector1 = Psi1(:,i);
                    %spin down
                    test_vector2 = Psi2(:,i);
                    %check we have one spin at site 1 and one spin at
                    %site kk
                    if (Psi1(1,i)==ii && Psi2(1,i)==(~ii) && Psi1(kk,i)==jj && Psi2(kk,i)==(~jj))

                        %flip the spins at site 1 and site kk
                        test_vector1(1) = ~test_vector1(1);
                        test_vector1(kk) = ~test_vector1(kk);
                        test_vector2(1) = ~test_vector2(1);

```

```

        test_vector2(kk) = ~test_vector2(kk);

        if (sum((test_vector1==Psi1(:,j)))==n_sites && ...
            sum((test_vector2==Psi2(:,j)))==n_sites)

            %sigma_x just adds a 1
            S(i,j) = S(i,j) + 1;
            %sigma_y has a sign
            if ii==jj
                S(i,j) = S(i,j) - 1;
            else
                S(i,j) = S(i,j) + 1;
            end
        end
    end
end
end
end

end

for kk = 2:n_sites
    %Sz_1*Sz_j
    S(i,i) = S(i,i) + (Psi1(1,i)-Psi2(1,i))*(Psi1(kk,i)-Psi2(kk,i));
end
end

%make Hermitian
for i = 1:(4^n_sites)
    for j = (i+1):(4^n_sites)
        S(j,i) = S(i,j);
    end
end

end

end
%=====

```

I.4 Wavefunction Solvers

This section contains different sets of code to solve for the wavefunction. The GP Solver (§I.4.2) is optimized to solve the GP equation (Eqn. 2.47), which is essentially the Schrödinger equation with a non-linear term. However, the GP Solver can also be used to solve the single particle problem if the interaction term is set to zero. However, the Crank-Nicolson solver (§I.4.3) is optimized for the non-interacting problem.

I.4.1 Schrodinger Equation Eigenstates

This function computes the eigenstates of the 1D or radial Schrodinger equation given an arbitrary single particle potential sent in as a function handle.

```

%=====
%-----
%Inputs:
%xb: Upper bound on x for the solution
%pts: Number of points to use (i.e., dx = xb/pts)
%m: mass
%Ufunc: handle to a function that gives the potential. Potential needs to
%have the form @(x,t,flag) where x is the position, t is time. If flag==1
%then the function needs to return the characteristic length scale of the
%potential. If flag==2 then it needs to return whether or not it is
%time-dependent.
%L: If L<0 then this is a 1D problem. If L>=0, then this is a spherically
%symmetric problem and we are solving in the radial variable. L is the
%total angular momentum
%nstates: Number of eigenstates to solve for (if nstates=1 then ground
%state is returned)
%Outputs:
%x: Position vector
%E: Eigenvalues
%Psi: Eigenstates
%-----
function [E Psi] = solve_estates(xb,pts,m,Ufunc,L,nstates)

%Solve for the eigenstates of a potential Ufunc on the bound +- xb with
%number of points ('pts'). This is the finite different method

%-----Define Constants-----
hbar=1.055e-34;
momentum_prefactor = hbar^2/2/m;

%set up grid
dx=2*xb/(pts-1); %step size
v = [0:pts-1]';
x = v*dx - xb;

%spherically symmetric so x is dx to 2*xb
if L>=0
    x = x(1:end)+xb+dx;
end

if Ufunc(0,0,2)
    warning('Potential is not time independent. t=0 will be used.');
```

```

end

%Make Hamiltonian matrix
H = diag(Ufunc(x,0,0)/momentum_prefactor+2/dx^2+(L*(L+1)./x.^2).*(L>0));

%Put in off-diagonal momentum terms using finite different
for i = 1:size(H,1)

    if i = 1
        H(i,i-1) = -1/dx^2 + (L>=0)/dx/x(i);
    end
end

```

```

        if i == size(H,1)
            H(i,i+1) = -1/dx^2 - (L>=0)/dx/x(i);
        end
    end

    %diagonalize
    [Psi E] = eigs(H,nstates,'sm');
    E = diag(E)*momentum_prefactor;
    %Sort by energy
    [E IX] = sort(E);

    Psi = Psi(:,IX);
end
%=====

```

I.4.2 GP Solver

The following is code used to solve the GP Equation (Eqn. 2.47), but this code can also be used to solve the single particle Schrödinger equation if interactions are set to zero. Earlier versions of this code were provided by Professor Brian P. Anderson at the University of Arizona. The code is divided into code optimized for the 3D solution and the 1D radial solution. However, if we just want to solve for the condensate TF expansion, then we can use the scaling equations, Eqn. 4.20. The following is a small function that can be used to solve these equations numerically.

```

%=====
%-----
%Inputs:
%omega: Trap frequency vector
%t: Time vector
%Outputs:
%t3: Output time vector (same as t)
%lambda: Scaling solutions at each time t. Is a six vector where elements
%1:3 are the lambda derivatives and 4:6 are the lambdas.
%-----

function [t3 lambda] = solve_tf_scaling_expansion(omega,t)

%Solve the TF expansion using the equations of Castin and Dum (PRL 77,
%5315)
[t3 lambda] = ode23(@(t,y)(tf_eqns(t,y,omega)),t,[0 0 0 1 1 1]);
end

%Sub function sent to ode23
function yprime = tf_eqns(t2,y,omega)
    %6 equations
    yprime(4:6) = y(1:3);
    yprime(1:3) = omega(1:3).^2./(y(4)*y(5)*y(6))./y(4:6)';
    yprime = yprime';
end

```



```
%=====
```

Radial GP

The following code solves the radial GP equation (if $L \geq 0$) or 1D GP equation (if $L < 0$). For stability the time propagation is $\Psi(t+dt) = \Psi(t-dt) + 2dti\mathbf{H}/\hbar\Psi(t)$ where the kinetic energy terms of \mathbf{H} are computed using a finite difference method. If the propagation time is negative this code will perform imaginary time propagation which will converge to the ground state. A good guess value is the non-interacting ground state.

```
%=====
%-----
%Inputs:
%Psi0: Initial wavefunction (guess for imaginary time propagation)
%dx: step size
%N: Number of particles
%m: mass
%a: scattering length
%Ufunc: handle to a function that gives the potential. Potential needs to
%have the form @(x,t,flag) where x is the position, t is time. If flag==1
%then the function needs to return the characteristic length scale of the
%potential. If flag==2 then it needs to return whether or not it is
%time-dependent.
%dt: Time step
%tmax: Maximum propagation time. If tmax < 0 then do imaginary time propagation and propagate
%until convergence or until count*dt > -tmax. If tmax > 0 then do regular
%propagation.
%L: If L<0 then this is a 1D problem. If L>=0, then this is a spherically
%symmetric problem and we are solving in the radial variable. L is the
%total angular momentum
%flag: If 1 then save the wavefunction every 1000 time propagation steps.
%Otherwise only the end result is output.
%Outputs:
%x: Position vector
%y: Wavefunction (see flag)
%t: time vector
%-----

% Variant of Visscher scheme for solving GPE
% See PRE 62, 7438 (2000) and PRE 62, 1382 (2000)

function [x y t] = GP(Psi0,dx,N,m,a,Ufunc,dt,tmax,L,flag)

%constants
%experimentalists use physical units!
hbar=1.055e-34;
num = N;

%trap length scale
ax=Ufunc(0,0,1);
```

```

xb=dx*(size(Psi0,2)-1)/2; %space extends to +/- xb
pts=size(Psi0,2); %number of grid points
npsi = size(Psi0,1);
if npsi==1
    sumdir = 1;
else
    sumdir = 2;
end

x=ones(npsi,1)*linspace(-xb,xb,pts); %grid
x = reshape(x,1,1,size(x,1),size(x,2));
if L>=0
    x = x+xb+dx; %solve for x=[dx 2xb+dx]
end
dx2=dx*dx;

%Is the potential time-independent? If so then we don't have to update the
%Hamiltonian matrix every time step
timeind = Ufunc(0,0,2);

U=Ufunc(x,0,0); %potential

%interaction parameter...note psi is normalized to unity in this
%calculation
g=4*pi*hbar^2*a*num/m;
if L<0
    %renormalization parameter...see Jamie William's thesis
    alpha=0.269/num^0.4/(a^4*ax^16)^0.1;
else
    alpha = 1;
end

if a==0
    alpha = 0;
end

if L<0
    normalizer = 1;
else
    normalizer = x.^2*4*pi;
end

%kinetic energy
ke=-hbar*hbar/2/m/dx2;

%kinetic energy prefactor
ke2 = hbar*hbar/2/m;

%timestep prefactor
tprop = 2*dt/hbar*[1 -1];

if tmax<0
    %imaginary time propagation

```

```

        tprop = -tprop;
        tprop(2) = 0;
        prop_index = [1 2];
        main_index = 2;
    else
        prop_index = [2 1];
        main_index = 1;
    end

    if L<0 %use periodic boundary conditions
        leftoffset = [pts 1:(pts-1)];
        rightoffset = [2:pts 1];
    else
        leftoffset = [1 1:(pts-1)];
        rightoffset = [2:pts pts];
    end

    %Solve
    %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
    Psi = zeros(3,2,size(Psi0,1),size(Psi0,2));
    Psi(1,1, :, :)=real(Psi0);
    Psi(1,2, :, :)=imag(Psi0);

    %setup and normalize wavefunction
    AA = (sqrt(sum(squeeze((abs(Psi(1,1, :, :)).^2+0*abs(Psi(1,2, :, :)).^2).*...
        normalizer),sumdir)*dx))*ones(1,pts);
    AA = reshape(AA,1,1,size(AA,1),size(AA,2));
    Psi(1,1, :, :)=Psi(1,1, :, :)./AA;
    Psi(1,2, :, :)=Psi(1,2, :, :)./AA;

    count=-1;
    stepsize=1000; %make plots every this many steps (if flag=1 then save this many steps)
    t = zeros(1,ceil(tmax/dt/stepsize));

    if flag
        y = zeros(ceil(tmax/dt/stepsize),length(x));
        y(1, :) = squeeze(Psi(1,1, :, :)+sqrt(-1)*Psi(1,2, :, :));
        t(1) = 0;
        count2 = 1;
    end

    olderr=1e-6;
    newerr=1e6;

    Psi(2,1, :, :)=Psi(1,1, :, :);
    Psi(2,2, :, :)=Psi(1,2, :, :);

    %-----Procedure-----
    %There are 6 wavefunctions we keep track of t=t0-dt,t=t0 and t=t0+dt (real
    %and imaginary)
    %Psi(t0+dt) = Psi(t0-dt) + iH/hbar Psi(t0)
    total.potential = (U+ke2*(L>0).*L*(L+1)./x.^2);

```

```

if alpha==0
    V=0;
end

%this finds where the wavefunction starts so that we can replace the
%centrifugal barrier with an infinite potential (and force the wavefunction
%to zero)
if L>=0

    minx = pts;
    for i = 1:npsi
        %find where the wavefunction starts to set the centrifugal barrier
        x.indices = 1:length(x(1,:));
        x2 = x.indices((abs(Psi(1,1,i,:)).^2>(max(abs(Psi(1,1,i,:)).^2)/100)));
        minx = min(minx,floor(x2(1)/2));
    end
end

while (1)
    count=count+1;
    %enforce boundary conditions
    if L<0 %solved on -x to x

        for i = 1:2
            for j = 1:2
                Psi(i,j,:,1)=0;
                Psi(i,j,:,pts)=0;
                Psi(i,j,:,2)=0;
                Psi(i,j,:,pts-1)=0;
            end
        end

    else %solved on x=0 to 2xb

        for i = 1:2
            for j = 1:2
                Psi(i,j,:,1)=Psi(i,j,:,2); %derivative at the origin must be zero
                Psi(i,j,:,pts-1:pts)=0; %must be zero at infinity
                %Psi(i,j,:,pts) = Psi(i,j,:,pts-1);
            end
        end

        %for high L need to set the wavefunction near the origin
        %to zero because of the centrifugal barrier
        if minx ==0
            Psi(:, :, :, 1:minx) = 0;
        end

    end

    if alpha ==0
        den=Psi(2,1, :, :).^2+Psi(2,2, :, :).^2; %density
        V=g*alpha*den; %interaction energy
    end
end

```

```

end

%time dependent potential
if timeind
    U=Ufunc(x,dt*count,0); %potential
end

%note this non-staggered time evolution seems necessary for stability
%(Tzu-Chieh found the same thing)
if alpha ==0 || timeind
    total_potential = (U+V+ke2*(L>0).*L*(L+1)./x.^2);
end

%time propagate
for i = 1:2

    prop_psi = Psi(2,prop_index(i),:,:);

    Psi(3,i,:,:) = Psi(main_index,i,:,:) + tprop(i)*(ke*(prop_psi(1,1,:,leftoffset)+...
        prop_psi(1,1,:,rightoffset)-2*prop_psi(L>=0)./x.*...
        (prop_psi(1,1,:,rightoffset)-prop_psi(1,1,:,leftoffset))*dx)+...
        total_potential.*prop_psi);

end

%renormalize (only if we are solving for a case with interactions)
if alpha ==0
    AA = (sqrt(sum(squeeze((abs(Psi(3,1,:,:)).^2+...
        abs(Psi(3,2,:,:)).^2).*normalizer),sumdir)*dx))*ones(1,pts);
    AA = reshape(AA,1,1,size(AA,1),size(AA,2));
    Psi(3,1,:,:) = Psi(3,1,:,:) ./ AA;
    Psi(3,2,:,:) = Psi(3,2,:,:) ./ AA;
end

if rem(count,stepsize)
    figure(1);
    plot(squeeze(x(1,1,1,:))/1e-6,squeeze(abs(Psi(3,1,1,:)).^2+abs(Psi(3,2,1,:)).^2));

    if flag
        y(count2,:) = squeeze(Psi(3,1,1,:) + sqrt(-1)*Psi(3,2,1,:));
        t(count2) = count*dt;
        count2 = count2 + 1;
    end
end

%check break conditions
olderr=newerr;
newerr=sum((Psi(3,1,:,:) - Psi(main_index,1,:,:)).^2.*normalizer*dx,4);

%break conditions
if tmax<0

    %converging

```

```

        if max(newerr)<(dt^2)
            disp(count)
            break;
        end

        if ((dt*count) > (-tmax))
            %did not converge
            disp('Warning: Did Not Converge')
            break;
        end

        elseif (dt*count > tmax)
            break;
        end

        %go forward in time dt
        for i = 1:2
            for j = 1:2
                Psi(i,j,,:) = Psi(i+1,j,,:);
            end
        end

    end

end

if flag
    y(count2,:) = squeeze(Psi(3,1,1,:)) + sqrt(-1)*Psi(3,2,1,:));
else
    y = squeeze(Psi(3,1,,:)) + sqrt(-1)*Psi(3,2,,:));
end
x = squeeze(x);

end

%=====

```

The above code was used to time-propagate the non-condensate wavefunctions expanding in the condensate mean-field potential given by the following,

```

%=====
%-----
%Inputs:
%x: Position
%t: time
%omega: trap frequency (before release)
%m: Mass
%NBEC: Number of atoms in BEC
%a: Scattering length
%flag: If 1 output the characteristic length scale of the potential. If 2
%output whether time-independent. If 3 then return the TF profile, but not
%the potential
%Outputs:
%Psi0: Guess for initial wavefunction
%dx: step size
%-----

```

```

function y = TF.potential(x,t,omega,m,NBEC,a,flag)

hbar=1.055e-34;

%harmonic oscillator length
aosc = sqrt(hbar/m/omega);

%compute the BEC chem potential in trap (t<=0)
mu = 15^(2/5)/2*(NBEC*a/aosc)^(2/5)*hbar*omega;

%interaction
U0 = 4*pi*hbar^2*a/m;

%TF radius in trap (t<=0)
rTF = sqrt(2*mu/m/omega^2);

%time dependent solution (from fit to N=1E5 data)
rTF2 = rTF*(4.126+(4.53^2-3.439*t/1E-3+(2.146*t/1E-3).^2).^(0.5))/(4.126+4.53);

%return the characteristic length scale
if flag==1
    y=rTF2; %Thomas Fermi radius
    return;
elseif flag==2
    y=1; %time dependent
    return;
end

%TF profile
y = mu/U0*(1-(x/rTF2).^2);

%set zero if outside the Thomas-Fermi radius
y = (y>0).*y;
y = y*(rTF/rTF2)^3;

if flag==3 %just want the TF profile...not the potential
    return;
end

%interaction potential
y = y*2*U0;
%=====

```

3D GP

For solving the 3D GP equation we use a split step method. First compute the propagation due to the spatially on-diagonal terms (potential and interaction energy) of the hamiltonian for time step $dt/2$. Then compute the kinetic energy using an FFT. Finally, propagate the other half of the on diagonal terms for $dt/2$.

```

%=====
%-----

```

```

%Inputs:
%Psi0: Initial wavefunction (guess for imaginary time propagation)
%dx: step size
%N: Number of particles
%m: mass
%a: scattering length
%Ufunc: handle to a function that gives the potential. Potential needs to
%have the form @(x,t,flag) where x is the position, t is time. If flag==1
%then the function needs to return the characteristic length scale of the
%potential. If flag==2 then it needs to return whether or not it is
%time-dependent.
%prop: Set to i for normal time propagation and 1 for imaginary time
%propagation
%dt: time step
%tmax: Maximum propagation time. If tmax < 0 then do imaginary time propagation and propagate
%until convergence or until count*dt > -tmax. If tmax > 0 then do regular
%propagation.
%Outputs:
%[x y z]: Position vectors
%f: Wavefunction after tmax
%-----

function [x y z f] = GPsplitstep_3D(Psi0,dx,N,m,a,Ufunc,prop,dt,tmax)

%-----Define Constants-----
hbar=1.055e-34;
num=N; %atom number
xb=dx*(length(Psi0)-1)/2; %space extends to +/- xb
pts=length(Psi0); %number of grid points

%Set up the grid
nmid = floor(pts/2);
v = [0:pts-1]';
x = v*dx - xb;

%setup 3D grid
[x y z] = meshgrid(x,x,x);

%setup momentum grid
kmax = 2*pi/dx;
dk = kmax/pts;
p = find(v > nmid);
v(p) = pts - v(p);
kx = v*dk;

%setup 3D momentum grid
[kx ky kz] = meshgrid(kx,kx,kx);
K = exp(-1 * prop * dt * (kx.^2+ky.^2+kz.^2) * hbar/2/m); %propagator for Kinetic energy

timeind = Ufunc(0,0,0,0,2);
if timeind
    U=Ufunc(x,y,z,0,0); %potential
end

```



```

g=4*pi*hbar^2*a*num/m; %interaction parameter...note psi is normalized to unity in this calculation
%Solve
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
psi = Psi0;

%psi normalized to one
psi=psi/sqrt(sum(abs(psi(:)).^2)*dx^3);

count=-1;
stepsize=10; %make plots every this many steps
olderr=1e-6;
newerr=1e6;

while (1)

    count=count+1;
    psiold=psi;
    den=abs(psi).^2; %density

    if a==0
        V = 0*den;
    else
        V = g*den; %get density for nonlinear part of f equation
    end

    %time dependent potential (need to evaluate every time step)
    if timeind
        U=Ufunc(x,y,z,dt*count,0); %potential
    end

    %note this non-staggered time evolution seems necessary for stability
    %(Tzu-Chieh found the same thing)
    psi = psi.*exp(-1* prop * (U+V)*dt/2/hbar);

    %propagate kinetic energy
    psi = fftn(ifftn(psi).*K); %straight forward for 1D case

    den = abs(psi).^2;

    if a==0
        V = 0*den;
    else
        V=g*den; %get density for nonlinear part of f equation
    end

    psi = psi.*exp(-1*prop*(U+V)*dt/2/hbar); %propagate OTHER HALF of potential and interaction energy

    %renormalize
    psi=psi/sqrt(sum(abs(psi(:)).^2)*dx^3);
    if rem(count,stepsize)
        figure(1);

```

```

        plot(squeeze(x(nmid,:),nmid))/1e-6,abs(squeeze(psi(nmid,:),nmid)).^2);
        disp(['Count' num2str(count)]);
    end

    olderr=newerr;
    newerr=sum((abs(psi(:)-psiold(:)).^2)*dx^3);

    %break conditions
    if tmax<0
        if abs(newerr)<(dt^2)
            disp(count)
            break;
        end

        if rem(count,stepsize)
            disp(dt*count);
            disp(newerr)
        end

        if ((dt*count) > (-tmax))
            %did not converge
            disp('Warning: Did Not Converge')
            break;
        end

    elseif (dt*count > tmax)
        break;
    end
end
f = psi;
%=====

```

I.4.3 Crank-Nicolson

For solving the single-particle Schrödinger equation the Crank-Nicolson method ($\Psi(t+dt) = \frac{I-idt/2\hbar H}{I+idt/2\hbar H}\Psi(t)$) is very stable and conserves unitary evolution. It can be computationally intensity because we need to invert the Hamiltonian. However, for time-independent problems this only needs to be done once. The following is code to implement the Crank-Nicolson time evolution.

```

%=====
%-----
%Inputs:
%psi_init: Initial wavefunction
%x: Spatial points in 1D
%t: time vector
%vfunc: handle to a function that gives the potential. Potential needs to
%have the form @(x,t) where x is the position, t is time.
%time_indep: 1 if the potential is time-independent
%Outputs:
%y: Wavefunction at each time step

```

```

%-----
function y = cn.solve(psi_init,x,t,vfunc,time_indep)

%psi_init:  initial wavefunction (column vector)
%x:  row vector of x values
%t:  row vector of times
%vfunc:  potential energy function of the form vfunc(numpts,deltax,t)
%time_indep:  is the potential time independent (1 or 0)

%NOTE: column and time vectors must be spaced evenly and must be increasing

hbar = 1.05E-34;
mRb = 87*1.67E-27;

dt = t(2)-t(1);
dx = x(2)-x(1);

psi_cur = psi_init;
df = i*dt*hbar/dx^2/2/mRb;

y = zeros(length(psi_cur),length(t));

A1_blank = eye(length(psi_cur)+2)*(1-2*df);
A2_blank = eye(length(psi_cur)+2)*(1+2*df);

if time_indep

    %if the potential is time-independent then only invert the matrix once
    A1 = A1_blank;
    A2 = A2_blank;

    %returns with a 0 on each end, must return a col vector
    v_vec = vfunc(x,0)*i*dt/2;

    %just need to add off-diagonal terms
    for ii = 2:(length(psi_cur)+1)
        A1(ii,ii-1) = df;
        A1(ii,ii+1) = df;
        A2(ii,ii-1) = -df;
        A2(ii,ii+1) = -df;
        A1(ii,ii) = A1(ii,ii)-v_vec(ii);
        A2(ii,ii) = A2(ii,ii)+v_vec(ii);
    end

    A1 = A1(2:(length(psi_cur)+1),2:(length(psi_cur)+1));
    A2 = A2(2:(length(psi_cur)+1),2:(length(psi_cur)+1));
    A3 = A1/A2;

    for j = 1:length(t)
        psi_cur = A3*psi_cur;
        y(:,j) = psi_cur;
    end
end

```

```

else
    %need one extra time step
    t = [t (t(end)+dt)];

    for j = 1:(length(t)-1)

        %returns with a 0 on each end, must return a col vector
        v_vec = vfunc(x,t(j))*i*dt/2;
        v_vec2 = vfunc(x,t(j+1))*i*dt/2;

        A1 = A1_blank;
        A2 = A2_blank;

        %just need to add off-diagonal terms

        for ii = 2:(length(psi_cur)+1)
            A1(ii,ii-1) = df;
            A1(ii,ii+1) = df;
            A2(ii,ii-1) = -df;
            A2(ii,ii+1) = -df;
            A1(ii,ii) = A1(ii,ii)-v_vec(ii);
            A2(ii,ii) = A2(ii,ii)+v_vec2(ii);
        end

        A1 = A1(2:(length(psi_cur)+1),2:(length(psi_cur)+1));
        A2 = A2(2:(length(psi_cur)+1),2:(length(psi_cur)+1));

        psi_cur = A1/A2*psi_cur;

        y(:,j) = psi_cur;

    end

end

%=====

```

I.5 Band Structure

The following is code to calculate the bandstructure and Bloch wavefunctions of a cubic lattice.

```

%=====
%-----
%Inputs:
%s: Lattice depth (in Er)
%q: Quasimomentum (vector with M elements)
%nband: Bands to solve for (vector with L elements)..must be integers
%Outputs:
%Eq: eigenenergies at q (n x L matrix)
%Psiq: Bloch wavefunctions in plane wave basis (n x L x M) where n is the truncation

```

```

%-----
function [Eq Psiq] = solve_lattice_bloch(s, q, nband)

n_trunc = 7;
Eq = zeros(length(q),length(nband));
Psiq = zeros(2*n_trunc+1,length(q),length(nband));

plane_waves = -n_trunc:n_trunc;

%go through for each value of q
for i = 1:length(q)

    %construct matrix
    H = zeros(length(plane_waves)+2);

    for j = 2:(length(plane_waves)+1)
        H(j,j) = (2*plane_waves(j-1))^2+(4*plane_waves(j-1)*q(i))+q(i)^2+s/2;
        H(j,j+1) = -s/4;
        H(j+1,j) = -s/4;
    end

    H = H(2:(length(plane_waves)+1),2:(length(plane_waves)+1));

    %diagonalize
    [V D] = eig(H);

    %output the results
    for j = 1:length(nband)
        Eq(i,j) = D(nband(j)+1,nband(j)+1);
        Psiq(:,i,j) = V(:,nband(j)+1);
        %make sure overall wavefunction phase is consistent
        %this works for even "n"
        if sum(V(:,1))<0
            Psiq(:,i,j) = -1*Psiq(:,i,j);
        end
    end
end
end
%=====

```

The solutions of the bandstructure code can be used to construct Wannier function using the following function.

```

%=====
%-----
%Inputs:
%Psiq: matrix of q eigenstates (in a single band)
%q: vector of q values
%x: Position row vector (in units of the lattice spacing)
%R: Position of Wannier (integer lattice units)
%Outputs:
%Wannier: Wannier wavefunction at x
%-----
function Wannier = construct_wannier(Psiq,q,x,R)

```

```

%get the plane wave states
k_states = 2*pi*(-((size(Psiq,1)-1)/2):((size(Psiq,1)-1)/2))';
plane_waves = exp(sqrt(-1)*k_states*x);
%the quasimomentum plane waves
plane_q_waves = exp(sqrt(-1)*pi*q'*x);
%construct the bloch wavefunctions
bloch_waves = ((Psiq')*plane_waves);
bloch_waves = bloch_waves.*plane_q_waves;
%offset
offset_factor = exp(sqrt(-1)*q*R*pi)*ones(1,length(x));
%Wannier is the integral over the bloch_waves
Wannier = sum(bloch_waves.*offset_factor,1)*(q(2)-q(1))/2;
end
%=====

```

The Wannier functions can be used to calculate the Hubbard t and U parameters using the following function.

```

%=====
%-----
%Inputs:
%s:  lattice depth
%a:  scattering length
%m:  mass
%d:  lattice spacing
%dq: Delta q
%Outputs:
%t:  Hubbard tunneling parameter (in Er)
%U:  Hubbard interaction parameter (in Er)
%Er: Recoil energy
%-----
function [t U Er] = Hubb_t_U(s,a,m,d,dq)

hbar = 6.626068E-24/2/pi;
gs = 4*pi*hbar^2*a/m;
Er = hbar^2*(pi/d)^2/2/m;

q = -1:dq:1;

for i = 1:length(s)
    %solve for the band structure
    [Eq Psiq] = solve_lattice_bloch(s(i), q, 0);

    %t is defined from the bandwidth
    t(i) = (Eq(1)-Eq((length(q)-1)/2))/4;

    %get Wannier function
    x = -1:0.01:1;
    [PsiWann] = construct_wannier(Psiq,q,x,0);

    %U is an integral over the Wannier function
    U(i) = real(sum(PsiWann.^4*(x(2)-x(1))).^3/d^3*gs/Er);
end

```

```

end
end
%=====

```

I.6 Atomic Physics

This section contains a set of functions for calculating atomic properties.

I.6.1 Atomic Special Functions

These are a collection of useful functions for atomic physics calculations. First, this is the Clebsch-Gordan coefficient used to decompose an alkali ground state $|F, m_F\rangle$ into $|S, m_S; I, m_I\rangle$ states.

```

%=====
%-----
%Inputs:
%F: Total hyperfine number
%mF: Hyperfine projection
%mS: Spin projection
%mI: Nuclear spin projection
%I: Total nuclear spin
%Outputs:
%y: Clebsch-Gordan coefficient
%-----

function y = S_State_CG_coeff(F,mF,mS,mI,I)

%Clebsch-Gordan Coefficient for the alkali (nS) ground state
%i.e., <SS,mS;I,mI|F,mF> where |F-I|=0.5 and SS=1/2 (total spin)

if (mS+mI-mF~=0)
    y = 0;
    return;
end

if (F-I)==0.5
    y = 1/2^(0.5)*(-1)^(1 - 2*I + mS + mI + mF)*(1 + 2*(1/2 + ...
    I))^(0.5)*((factorial(1/2 + mS)*factorial(I - mI)*factorial(1/2 + I - mF)*...
    factorial(1/2+I+mF))/((1+3*I+2*I^2)*factorial(1/2 - mS)*factorial(I+mI)*...
    (factorial(I + mS - mF))^2*(factorial(1/2 - mI + mF))^2))^^(0.5);
else if (F-I)==-0.5
    y = -1/2^(0.5)*(-1)^(0.5+mS)*...
    (1 + 2*(-0.5+I))^^(0.5)*...
    ((factorial(0.5-mS)*factorial(I-mI)*factorial(I+mI)*factorial(-0.5+I-mF))/...
    ((I+2*I^2)*factorial(0.5+mS)*(factorial(-0.5+I-mS-mI))^2*...
    (factorial(0.5+mI-mF))^2*factorial(-0.5+I+mF)))^^(0.5);
end

end

%=====

```

This code calculates the Wigner-3J symbol.

```
%=====
%-----
%Inputs:
%F1: Total hyperfine number of state 1
%mF1: Hyperfine projection of state 1
%F2: Total hyperfine number of state 2
%mF2: Hyperfine projection of state 2
%q: Polarization (-1,1,0)
%Outputs:
%y: Wigner 3J coefficient
%-----
function y = wigner3J(F1,mF1,F2,mF2,q)

if ~(-mF1+mF2+q==0)
    y = 0;
    return;
end
if abs(F2-F1)>1
    y = 0;
    return;
end
if (F1==0 && F2==0)
    y = 0;
    return;
end
if (F1==F2)

    F = F1;
    y = -1/2*(-1)^(-F+q+mF2)*(q+F*q-mF1+F*mF1+q*mF1-mF1^2-mF2-F*mF2+mF1*mF2)*...
        sqrt((factorial(1+q)*factorial(F-mF1)*factorial(F+mF1)*...
            factorial(F-mF2))/(F*(1+3*F+2*F^2)*factorial(1-q)*...
            (factorial(F+q-mF1))^2*(factorial(1+mF1-mF2))^2*factorial(F+mF2)));

elseif (F2==F1+1)

    F = F1;
    y = -(-1)^(F+q-mF1)*...
        sqrt((factorial(1-q)*factorial(F-mF1)*factorial(1+F-mF2)*factorial(1+F+mF2))/...
            ((3+11*F+12*F^2+4*F^3)*factorial(1+q)*factorial(F+mF1)*...
            (factorial(F-q-mF2))^2*(factorial(1-mF1+mF2))^2));

elseif (F2==F1-1)

    F = F1;
    y = (-1)^(-F+q+mF2)*...
        sqrt((factorial(1+q)*factorial(F-mF1)*factorial(F+mF1)*factorial(-1+F-mF2))/...
            (F*(-1+4*F^2)*factorial(1-q)*factorial(-1+F+mF2)*...
            (factorial(-1+F+q-mF1))^2*(factorial(1+mF1-mF2))^2));

end
%=====
```


This code calculates the Wigner-6J symbol.

```
%=====
%-----
%Inputs:
%J: Total electronic angular momentum of state 1
%F: Total hyperfine number of state 1
%J2: Total electronic angular momentum of state 2
%F2: Total hyperfine number of state 2
%I: Total nuclear spin
%Outputs:
%y: Wigner 6J coefficient
%-----
function y = wigner6J(J,F,J2,F2,I)

if abs(F-F2)>1
    y = 0;
end
if abs(J-J2)>1
    y = 0;
end

if (J==J2)

    if (F==F2)
        if F==0
            y = 0;
        else
            y = -1/2*(-1)^(-F-I-J)*sqrt(1/(F*(1+F)*(1+2*F)*J*(1+J)*(1+2*J)))*(F+F^2-I-I^2+J+J^2);
        end
    elseif (F+1==F2)
        y = -1/2*(-1)^(-F-I-J)*sqrt((1+F+I-J)*(1+F-I+J)*(-F+I+J)*...
            (2+F+I+J)/((1+F)*(1+2*F)*(3+2*F)*J*(1+J)*(1+2*J)));
    elseif (F-1==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt((F+I-J)*(F-I+J)*(1-F+I+J)*...
            (1+F+I+J)/(F*(-1+2*F)*(1+2*F)*J*(1+J)*(1+2*J)));
    end

elseif (J+1==J2)

    if (F==F2)
        y = -1/2*(-1)^(-F-I-J)*sqrt(((F+I-J)*(1+F-I+J)*(1-F+I+J)*...
            (2+F+I+J))/(F*(1+F)*(1+2*F)*(1+J)*(1+2*J)*(3+2*J)));
    elseif (F+1==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt((1+F-I+J)*(2+F-I+J)*(2+F+I+J)*...
            (3+F+I+J)/((1+F)*(1+2*F)*(3+2*F)*(1+J)*(1+2*J)*(3+2*J)));
    elseif (F-1==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt((-2+F-I-J)*(-1+F-I-J)*(-1+F+I-J)*...
            (F+I-J)/(F*(-1+2*F)*(1+2*F)*(1+J)*(1+2*J)*(3+2*J)));
    end

elseif (J-1==J2)
```

```

    if (F==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt(((1+F+I-J)*(F-I+J)*(-F+I+J)*...
            (1+F+I+J))/(F*(1+F)*(1+2*F)*J*(-1+2*J)*(1+2*J)));
    elseif (F+1==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt((F-I-J)*(1+F-I-J)*(1+F+I-J)*...
            (2+F+I-J)/((1+F)*(1+2*F)*(3+2*F)*J*(-1+2*J)*(1+2*J)));
    elseif (F-1==F2)
        y = 1/2*(-1)^(-F-I-J)*sqrt((-1+F-I+J)*(F-I+J)*(F+I+J)*...
            (1+F+I+J)/(F*(-1+2*F)*(1+2*F)*J*(-1+2*J)*(1+2*J)));
    end

end
end
%=====

```

This code calculates a prefactor, which contains the Wigner-3J and Wigner-6J identities, for writing out the electric-dipole matrix element between states.

```

%=====
%-----
%Inputs:
%J1: Total electronic angular momentum of state 1
%F1: Total hyperfine number of state 1
%mF1: Hyperfine projection of state 1
%J2: Total electronic angular momentum of state 2
%F2: Total hyperfine number of state 2
%mF2: Hyperfine projection of state 2
%I: Total nuclear spin
%q: Polarization (-1,1,0)
%Outputs:
%y: Wigner coefficient
%-----
function y = wig.coefficient(J1,F1,mF1,J2,F2,mF2,I,q)

%this is the coefficient in front of the J->J decay rate
%#1 is the EXCITED (higher energy) state

y = 0;
y1 = (2*F1+1)*(2*J1+1)*(2*F2+1)*...
    (wigner6J(J1,F1,J2,F2,I))^2;

if nargin < 8
    for i = -1:1
        y = y + y1*(wigner3J(F1,mF1,F2,mF2,i))^2;
    end
else
    y = y1*(wigner3J(F1,mF1,F2,mF2,q))^2;
end

end
%=====

```

I.6.2 Zeeman and Hyperfine Shift

This function calculates the eigenvalues of the full Zeeman plus hyperfine interaction in the $|J, m_J; I, m_I\rangle$ basis.

```
%=====
%-----
%Inputs:
%I: Nuclear spin
%J: Electronic angular momentum
%L: Orbital angular momentum
%AHF: Hyperfine A term
%BHF: Hyperfine B term
%B: Magnetic field (Gauss)
%Outputs:
%y: vector of the eigenvalues for the (2J+1)(2I+1) states
%-----

function y = zeemanshift(I,J,L,AHF,BHF,B)
%ALL energies in Hz and B is in Gauss
%returns the energies in order of largest to smallest

num_states = (2*J+1)*(2*I+1);

%do the calculation in the |J,mj;I,mI> basis
for i = -J:J
    start_index = (i+J)*(2*I+1)+1;
    end_index = (i+J+1)*(2*I+1);
    mJ(start_index:end_index) = i;
    mI(start_index:end_index) = -I:I;
end

%gyromagnetic ratio
gJ = 3/2 + (3/4-L*(L+1))/(2*J*(J+1));

%denominator of the B hyperfine term;
b.denom = 2*I*(2*I-1)*J*(2*J-1);

if BHF == 0
    b.denom = 1;
end

%create eigenstates
Hmatrix = zeros(num_states,num_states);

for i = 1:num_states
    for j = 1:num_states

        %The hyperfine energies are on diagonal
        if i==j
            %magnetic field energy
            Hmatrix(i,i) = mJ(i)*B*gJ*1.399624555E6;
            aa = mJ(i)*mI(i);
```

```

        %hyperfine interaction
        Hmatrix(i,i) = Hmatrix(i,i) + AHF*aa + BHF/2/b_denom*...
            ((3*mJ(i)^2-J*(J+1))*(3*mI(i)^2-I*(I+1)));
    end

    %off diagonal hyperfine terms
    if (mJ(j)==(mJ(i)-1) && mI(j)==(mI(i)+1))
        Hmatrix(i,j) = 1/2*(AHF+3*BHF/b_denom*(1/2+mJ(i)*mI(i)+(mJ(i)-1)*...
            (mI(i)+1)))*alpha(-1,J,mJ(i))*alpha(1,I,mI(i));
    end

    if (mJ(j)==(mJ(i)+1) && mI(j)==(mI(i)-1))
        Hmatrix(i,j) = 1/2*(AHF+3*BHF/b_denom*(1/2+mJ(i)*mI(i)+(mJ(i)+1)*...
            (mI(i)-1)))*alpha(1,J,mJ(i))*alpha(-1,I,mI(i));
    end

    if (mJ(j)==(mJ(i)-2) && mI(j)==(mI(i)+2))
        Hmatrix(i,j) = 1/4*3*BHF/b_denom*alpha(-1,J,mJ(i))*alpha(-1,J,mJ(i)-1)*...
            alpha(1,I,mI(i))*alpha(1,I,mI(i)+1);
    end

    if (mJ(j)==(mJ(i)+2) && mI(j)==(mI(i)-2))
        Hmatrix(i,j) = 1/4*3*BHF/b_denom*alpha(1,J,mJ(i))*alpha(1,J,mJ(i)+1)*...
            alpha(-1,I,mI(i))*alpha(-1,I,mI(i)-1);
    end

end

end

y = eig(Hmatrix);
y = sort(y);

%Subfunction which gives the angular momentum addition/subtraction
%prefactor
function yy = alpha(pm,x,mx)
    yy = sqrt((x-pm*mx)*(x+pm*mx+1));
end
end
%=====
%-----
%Inputs:

```

I.6.3 Create State Structure

The first step towards doing atomic OBE calculations is to construct an atomic state structure. This structure contains all the state information (e.g., the state energy, and how the states are linked via decay paths). Each state is added with a certain index which is then used as a reference for decay. The following shows example code for how to create the state structure of ^{40}K up to the 5P level.

```

%=====
%-----
%Inputs:

```

```

%B: Magnetic field
%Outputs:
%y: State structure
%-----

function y = create_state_structure(B)
%this creates a state vector y which is a structure that contains all the
%states for a certain atoms. Currently hard coded for K-40

%decay structure (for specifying the decay rates for a certain level)
decaystruct = struct('gamma',0,... %decay rate
    'index',0); %decay to which state

%speed of light (in cm/s)
c = 2.99792E10;

%K: All hyperfine splittings are for K-40. For higher states
%the energies are just for K-39. Probably the only difference are the
%standard reduced mass shifts... < 1GHz shift to energy levels.

%4S1/2
y = addstates([],1,4,1/2,4,0,-285.72E6,0,B,[],0);
decay2 = decaystruct;

%4P1/2
decay2 = decaystruct;
decay2(1).gamma = 3.75E7; %4S1/2
decay2(1).index = 1;
y = addstates(y,2,4,1/2,4,1,-34.523E6,0,B,decay2,389286.184353E9);

%4P3/2
decay2(1).gamma = 3.8E7;
y = addstates(y,3,4,3/2,4,1,-7.585E6,-3.445E6,B,decay2,391016.296050E9);

%5S1/2
decay2(1).gamma = 7.9E6; %4P1/2
decay2(1).index = 2;
decay2(2).gamma = 1.56E7; %4P3/2
decay2(2).index = 3;
y = addstates(y,4,5,1/2,4,0,68.9E6,0,B,decay2,391016.296050E9+239344.49E9);

%3D3/2
decay2(1).gamma = 2.2E7; %4P1/2
decay2(2).gamma = 4.34E6; %4P3/2
y = addstates(y,5,3,3/2,4,2,0,0,B,decay2,391016.296050E9+254647.22E9);

%3D5/2
decay2(2).gamma = 2.59E7; %4P3/2
y = addstates(y,6,3,5/2,4,2,0,0,B,decay2(2),391016.296050E9+254578.00E9);

%5P3/2
decay2(1).gamma = 1.16E6; %4S1/2
decay2(1).index = 1;

```

```

decay2(2).gamma = 4.6E6; %5S1/2
decay2(2).index = 4;
decay2(3).gamma = 1.5E5; %3D3/2
decay2(3).index = 5;
decay2(4).gamma = 1.4E6; %3D5/2
decay2(4).index = 6;
y = addstates(y,7,5,3/2,4,1,-2.45E6,-1.16E6,B,decay2,741091.33E9);

%5P1/2
decay2(1).gamma = 1.07E6; %4S1/2
decay2(1).index = 1;
decay2(2).gamma = 4.5E6; %5S1/2
decay2(2).index = 4;
decay2(3).gamma = 1.5E6; %3D3/2
decay2(3).index = 5;
y = addstates(y,8,5,1/2,4,1,-46.3E6,0,B,decay2(1:3),740529.15E9);

end
%=====

```

This is the addstates function used above.

```

%=====
%-----
%Inputs:
%curstates: Current state structure
%index: index to identify the state we are adding
%n: Principle quantum number of new state
%J: Electronic angular momentum
%I: Nuclear spin
%L: Orbital angular momentum
%A: Hyperfine A
%B: Hyperfine B
%Bfield: Magnetic Field
%decays: Decay structure
%energy: Bare energy of the state
%Outputs:
%y: Updated state structure with the new states
%-----

function y = addstates(curstates,index,n,J,I,L,A,B,Bfield,decays,energy)

%double check that index hasn't already been used
for i = 1:length(curstates)
    if curstates(i).index == index
        error('Duplicate Index');
    end
end

%Add the individual F states
y = [curstates createFstates(index,n,J,I,L,A,B,Bfield,decays,energy)];

end

```

```
%=====
```

The `addstates` function adds F states using (note that `gF` and `hfshift` are functions that calculate the gyromagnetic factor and the hyperfine energy shift) the following function.

```
%=====
%-----
%Inputs:
%index:  index to identify the state we are adding
%n:  Principle quantum number of new state
%J: Electronic angular momentum
%I: Nuclear spin
%L: Orbital angular momentum
%A: Hyperfine A
%B: Hyperfine B
%Bfield:  Magnetic Field
%decays:  Decay structure
%energy:  Bare energy of the state
%Outputs:
%y:  state structure with the new states
%-----
function y = createFstates(index,n,J,I,L,A,B,Bfield,decays,energy)

y = [];
%go through all possible F manifolds
for i = abs(J-I):abs(J+I)
    F = i;
    %calculate gF
    gF2 = gF(J,I,L,1/2,F);
    %calculate F energy splitting
    Fenergy = hfshift(A,B,I,J,F);
    %create mF states
    y = [y createmFstates(index,n,F,J,I,L,gF2,Bfield,decays,energy+Fenergy)];
end
%=====
```

The `createFstates` function adds individual m_F states using the following function.

```
%=====
%-----
%Inputs:
%index:  index to identify the state we are adding
%n:  Principle quantum number of new state
%F: Hyperfine number
%J: Electronic angular momentum
%I: Nuclear spin
%L: Orbital angular momentum
%gF: Gyromagnetic factor
%Bfield:  Magnetic Field
%decay:  Decay structure
%energy:  Bare energy of the state
```

```

%Outputs:
%y:  state structure with the new states
%-----
function y = createmFstates(index,n,F,J,I,L,gF,B,decay,energy)

%This is the structure for representing an individual |F,mF> state
y = repmat(struct('index',index,...
    'energy',0,...
    'n',n,...
    'F',F,...
    'J',J,...
    'I',I,...
    'L',L,...
    'mF',0,...
    'decay',decay),1,2*F+1);

%Create a state for each mF with the appropriate energy shift due to the
%field
for i = 1:(2*F+1)
    y(i).mF = F-i+1;
    y(i).energy = energy+1.4E6*y(i).mF*gF*B;
end
end
%=====

```

I.6.4 Multi-Level Stark Shift

Using the state structure created in §I.6.3 and the special functions from §I.6.1, we can compute the multi-level stark shift using the following code.

```

%=====
%-----
%Inputs:
%drive_wavelength:  Wavelength of field
%polarization:  Polarization of the field [sigma-,pi,sigma+], e.g.
%unpolarized is [1 1 1]/3
%state_index:  State index to find shift (e.g., "9" is 4S 9/2,9/2)
%Outputs:
%energy_shift:  AC Stark shift
%heating_rate:  Heating Rate
%-----
function [energy_shift heating_rate] = ACstark_shift(drive_wavelength, polarization, state_index)
%For a given atomic structure determine the AC Stark shift

%constants
c = 299792458;
hbar = 1.05457148E-34;
m = 40*1.67262158E-27;
kB = 1.3806503E-23;

%create state

```



```

y = create_state_structure(0);

drive_freq = c./drive_wavelength;

%sum over all states
for i = 1:length(y)

    connect_flag = 0;
    if i~=state_index

        %check if state "i" decays to the state we want the energy shift
        %for
        for j = 1:length(y(i).decay)
            if y(i).decay(j).index == y(state_index).index
                state2 = i;
                state1 = state_index;
                decayrate = y(i).decay(j).gamma;
                connect_flag = 1;
                break
            end
        end

        %check if the state we want the energy shift for decays to state
        %"i"
        if ~connect_flag
            for j = 1:length(y(state_index).decay)

                if y(state_index).decay(j).index == y(i).index

                    state2 = state_index;
                    state1 = i;
                    connect_flag = 1;
                    decayrate = y(state_index).decay(j).gamma;
                    break;
                end
            end
        end

        if connect_flag %these two states are connected

            %state2 is always the higher energy state (state we decay from)
            energydifference = abs(y(state2).energy-y(state1).energy);

            %recoil energy between the states
            Er = (hbar*2*pi*drive_freq/c).^2/(2*m);

            new_energy_shift = 0;

            %assume that "energydifferent" and "drive_freq" are both in
            %frequency units (so add in the 2*pi to make them angular
            %frequencies

            prefactor = -3*pi*c^2*decayrate/2/(2*pi*energydifference)^3*...

```

```

        (1./(energydifference-drive_freq)+1./(energydifference+drive_freq))/2/pi;

%if we are trying to find the shift in the excited state make
%the prefactor negative
if (state2==state_index)
    prefactor = -prefactor;
end

for j = -1:1
    new_energy_shift = new_energy_shift + ...
        prefactor*wig_coefficient(y(state2).J,y(state2).F,y(state2).mF,...
        y(state1).J,y(state1).F,y(state1).mF,4,j)*polarization(j+2);
end

energy_shift = energy_shift + new_energy_shift;

%don't calculate a heating if the state we are calculate
%the shift for is the "excited" state because I don't think it
%makes sense
if ~(state2==state_index)

    %calculate the scattering rate for this state just assuming
    %simple two level to get ball park
    heating_rate = heating_rate - 2*Er./kB.*new_energy_shift/hbar*decayrate.*...
        (1./(energydifference-drive_freq)+1./(energydifference+drive_freq))/2/pi.*...
        (drive_freq/energydifference).^3;

end
end
end
end
end
%=====

```

I.6.5 OBE Equations

We can also use the state structure from §I.6.3 to construct optical Bloch equations. First we construct a rate equation matrix for the diagonal density matrix terms (i.e., matrix \mathbf{A} of the equation $\dot{\rho} = \mathbf{A}\rho$) using the following code.

```

%=====
%-----
%Inputs:
%states: state structure
%Outputs:
%y: Rate matrix for the on-diagonal density matrix terms given the decays
%in states
%-----
function y = createdecaymatrix(states)
%the decay matrix is off the form dN_i = A_(ij)*N_j

```

```

y = zeros(length(states));
for i = 1:length(states)

    if ~isempty(states(i).decay)
        for jj = 1:length(states(i).decay)

            for kk = 1:length(states)
                if (states(kk).index==states(i).decay(jj).index)

                    %calculate clebsch-gordon coefficient
                    wig_coeff = wig_coefficient(states(i).J,states(i).F,states(i).mF,...
                        states(kk).J,states(kk).F,states(kk).mF,states(i).I);

                    y(i,i) = y(i,i)-states(i).decay(jj).gamma*wig_coeff;
                    y(kk,i) = y(kk,i)+states(i).decay(jj).gamma*wig_coeff;

                end
            end
        end
    end
end

end

%=====

```

We can add a driving field between certain sets of states (see considerations for multiple driving fields in §B.6.2) to the decay matrix using the following code.

```

%=====
%-----
%Inputs:
%curmatrix: Current OBE matrix
%states: State structure
%m: Mass
%n1: Index of the ground state
%n2: Index of the excited state
%drivingfields: Drivingfields is structures of the form
%.intensity
%.freq
%.q (vector [sigma+ sigma- pi])
%gamma1: Decay along the driving path
%atom_velocity: Velocity of the atom
%Outputs:
%y: Updated OBE matrix
%-----
function y = drivingtransitions(curmatrix,states,n1,n2,drivingfields,gamma1,atom_velocity)

y = curmatrix;
c = 299792458;
hbar = 1.05457148E-34;

for i = 1:length(n1)
    for j = 1:length(n2)

```

```

ground_state = states(n1(i));
excited_state = states(n2(j));

%energy separation between the states
omega0 = 2*pi*abs(excited_state.energy-ground_state.energy);

%prefactor
Omega = (6*pi*c^2*gamma1*drivingfields.intensity/hbar/omega0^3)^(0.5);

%calculate matrix element for the transition
matrix_element = (2*excited_state.J+1)*(2*ground_state.F+1)*(2*excited_state.F+1)*...
    (wigner6J(ground_state.J,ground_state.F,...
        excited_state.J,excited_state.F,ground_state.I))^2;

matrix_element = matrix_element*((wigner3J(excited_state.F,excited_state.mF,...
    ground_state.F,ground_state.mF,1)*drivingfields.q(1))^2+...
    (wigner3J(excited_state.F,excited_state.mF,ground_state.F,ground_state.mF,-1)*...
    drivingfields.q(2))^2+...
    (wigner3J(excited_state.F,excited_state.mF,ground_state.F,ground_state.mF,0)*...
    drivingfields.q(3))^2);

%matrix element needs to be square root
matrix_element = sqrt(matrix_element);

%add off-diagonal elements of the density matrix
if ~sum(matrix_element)==0

    %increase the size of the matrix to include the coherences
    y(:,end+2) = 0;
    y(end+2,:) = 0;

    %detuning
    detuning = excited_state.energy-ground_state.energy-...
        drivingfields.freq*(1+atom.velocity/c);

    %add off-diag decay
    %NOTE: this is the total decay out of the excited state
    y(end,end) = y(n2(j),n2(j))/2;
    y(end-1,end-1) = y(n2(j),n2(j))/2;

    %add Rabi oscillation terms
    y(end,end) = y(end,end)-sqrt(-1)*2*pi*detuning;
    y(end,n1(i)) = y(end,n1(i))-sqrt(-1)*Omega*matrix_element/2;
    y(end,n2(j)) = y(end,n2(j))+sqrt(-1)*Omega*matrix_element/2;

    y(end-1,end-1) = y(end-1,end-1)+sqrt(-1)*2*pi*detuning;
    y(end-1,n1(i)) = y(end-1,n1(i))+sqrt(-1)*Omega*matrix_element/2;
    y(end-1,n2(j)) = y(end-1,n2(j))-sqrt(-1)*Omega*matrix_element/2;

    y(n1(i),end) = y(n1(i),end)-sqrt(-1)*Omega*matrix_element/2;
    y(n1(i),end-1) = y(n1(i),end-1)+sqrt(-1)*Omega*matrix_element/2;

    y(n2(j),end) = y(n2(j),end)+sqrt(-1)*Omega*matrix_element/2;

```

```

        y(n2(j),end-1) = y(n2(j),end-1)-sqrt(-1)*Omega*matrix_element/2;
    end
end
end
end
%=====

```

Appendix J

Technical Drawings





Figure J.1: Technical drawing of the ^{40}K apparatus top vacuum flange.






Figure J.3: Technical drawing of the ^{40}K science chamber with the top and bottom flanges attached.





Figure J.4: Technical drawing of the ^{40}K MOT Cell as sent to the glassmaker.

Plot: BBAR @ 700-1100 NM(SiO₂): Reflectance




Table: BBAR @ 700-1100 NM(SiO₂): Performance

Design		BBAR @ 700-1100 NM(SiO ₂)		
Reference Wavelength (nm)		790.00		
Incident Angle (deg)		0.00		
Wavelength (nm)	Reflectance (%)	Transmittance (%)	Reflectance-Phase (deg)	Transmittance-Phase (deg)
405	4.741030	94.902868	165.082105	15.983297
425	8.816627	90.878817	159.998282	52.021952
445	13.712363	86.030285	171.818347	83.303933
465	16.070400	83.717394	-173.655988	110.727770
485	15.883809	84.181657	-159.401072	135.872620
505	13.518134	86.386932	-146.660551	158.719777
525	10.628258	89.306404	-135.843867	-179.658672
545	7.843082	92.103997	-127.315638	-158.269090
565	5.699820	94.362730	-121.247180	-140.103798
585	4.012625	95.956431	-117.300654	-122.273113
605	2.933593	97.053900	-114.882819	-105.568394
625	2.195437	97.803593	-113.101926	-89.968917
645	1.665252	98.344748	-111.268062	-75.454389
665	1.229318	98.770682	-109.436773	-61.840243
685	0.880601	99.119369	-107.712310	-49.055680
705	0.596706	99.403295	-106.580909	-37.018262
725	0.378672	99.623328	-107.033496	-25.788715
745	0.221059	99.778941	-110.635365	-15.244531
765	0.124958	99.875042	-119.630247	-5.270550
785	0.078123	99.921877	-136.100413	4.192309
805	0.067445	99.932565	-156.782714	13.180053
825	0.079230	99.920770	-173.216819	21.726323
845	0.101092	99.896908	177.215193	29.862705
865	0.123188	99.876812	172.289151	37.618808
885	0.138793	99.861207	169.798845	45.022237
905	0.144360	99.855640	168.335133	52.098536
925	0.139205	99.860795	166.964014	58.871158
945	0.124998	99.875002	164.883945	65.361480
965	0.105165	99.894835	161.198511	71.588872
985	0.084296	99.915704	154.711023	77.570788
1005	0.067595	99.932405	143.995483	83.322900
1025	0.050422	99.936678	128.604254	88.859267
1045	0.037918	99.932082	110.973523	94.192482
1065	0.024742	99.905258	96.252575	99.333857
1085	0.144991	99.855009	86.379042	104.293589

Figure J.5: Coating specifications for ⁴⁰K science chamber windows.




Figure J.6: Coating specifications for ^{40}K sapphire imaging window at (a) 0 degrees (b) 25 degrees and (c) 50 degrees incident angle and assuming unpolarized light.




Figure J.7: (a) Dimensions of the Special Optics objective. (b) Dimensions of the Special Optics tube lens. (c) Objective and tube lens imaging configuration. Rays are just a guide and not representative of the actual ray traces (e.g., the objective is not infinity conjugated). Working distance of the objective is 2.43mm and the effective focal length is 3.3mm.





Figure J.8: (a) MTF of the Special Optics objective at $NA = 0.6$ and $\lambda = 405\text{nm}$. (b) MTF of the Special Optics objective at $NA = 0.45$ and $\lambda = 767\text{nm}$. The objective is not diffraction limited at higher NA for $\lambda = 767\text{nm}$.

Appendix K

References

- [1] D. McKay, M. White, M. Pasienski, and B. DeMarco. Phase-slip-induced dissipation in an atomic Bose-Hubbard system. *Nature*, 453:76, 2008. [ii](#), [8](#), [129](#), [134](#)
- [2] S. S. Natu, D. C. McKay, B. DeMarco, and E. J. Mueller. Evolution of condensate fraction during rapid lattice ramps. *Phys. Rev. A*, 85:061601(R), 2012. [ii](#), [8](#), [137](#)
- [3] D. McKay, M. White, and B. DeMarco. Lattice thermodynamics for ultracold atoms. *Phys. Rev. A*, 79:063605, 2009. [ii](#), [9](#), [10](#), [142](#), [146](#), [150](#)
- [4] D. McKay and B. DeMarco. Thermometry with spin-dependent lattices. *New. J. Phys.*, 12:055013, 2010. [iii](#), [10](#), [52](#), [56](#), [168](#), [188](#), [226](#)
- [5] P. J. Mohr, B. N. Taylor, and D. B. Newell. CODATA recommended values of the fundamental physical constants: 2010. 2012, arXiv:1203.5425. [xiv](#), [189](#)
- [6] C. F. Fischer. Atomic structure: Multiconfiguration Hartree-Fock theories. In G. W. F. Drake, editor, *Springer Handbook of Atomic, Molecular and Optical Physics*, chapter 21, page 307. Springer, 2006. [1](#)
- [7] P. W. Anderson. More is different. *Science*, 177:393, 1972. [1](#)
- [8] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Harcourt Inc., 1976. [2](#), [29](#), [133](#)
- [9] G. Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91:147902, 2003. [3](#)
- [10] D. M. Ceperley. Path integrals in the theory of condensed helium. *Rev. Mod. Phys.*, 67:279, 1995. [3](#)
- [11] D. A. Bonn. Are high-temperature superconductors exotic? *Nature Physics*, 2:129, 2006. [4](#), [7](#)
- [12] P. A. Lee, N. Nagaosa, and X.-G. Wen. Doping a Mott insulator: Physics of high-temperature superconductivity. *Rev. Mod. Phys.*, 78:17, 2006. [4](#), [5](#), [6](#), [48](#), [111](#)
- [13] R. P. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21:467, 1982. [3](#)
- [14] D. P. Divincenzo. The physical implementation of quantum computation. *Fortschr. Phys.*, 48:771, 2000. [3](#)
- [15] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. *Nature*, 464:45, 2010. [3](#)
- [16] S. Lloyd. Universal quantum simulators. *Science*, 273:1073, 1996. [3](#)
- [17] I. Buluta and F. Nori. Quantum simulators. *Science*, 326:108, 1009. [3](#), [5](#)
- [18] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos. Universal digital quantum simulation with trapped ions. *Science*, 334:57, 2011. [3](#)

- [19] P. Hauke, F. M. Cucchietti, L. Tagliacozzo, M. Lewenstein, and I. Deutsch. On the robustness of quantum simulators. 2011, arXiv:1109.6457. [5](#)
- [20] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999. [5](#)
- [21] W. Ketterle and N. J. van Druten. Evaporative cooling of atoms. In B. Bederson and H. Walther, editors, *Advances in Atomic, Molecular and Optical Physics*, volume 37, page 181. Elsevier, 1996. [5](#), [15](#)
- [22] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269:198, 1995. [5](#), [24](#)
- [23] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969, 1995. [5](#), [59](#), [78](#)
- [24] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. *Science*, 285(5434):1703–1706, 1999. [5](#), [250](#)
- [25] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. Cold bosonic atoms in optical lattices. *Phys. Rev. Lett.*, 81:3108, 1998. [6](#), [7](#), [36](#)
- [26] W. Hofstadter, J. I. Cirac, P. Zoller, E. Demler, and M. D. Lukin. High-temperature superfluidity of fermionic atoms in optical lattices. *Phys. Rev. Lett.*, 89:220407, 2002. [6](#), [7](#)
- [27] J. Hubbard. Electron correlations in narrow energy bands. *Proc. R. Soc. Lond. A*, 276:238, 1963. [6](#)
- [28] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher. Boson localization and the superfluid-insulator transition. *Phys. Rev. B*, 40:546, 1989. [6](#), [41](#), [43](#), [44](#), [48](#)
- [29] B. G. Orr, H. M. Jaegar, A. M. Goldman, and C. G. Kuper. Global phase coherence in two-dimensional granular superconductors. *Phys. Rev. Lett.*, 56:378, 1986. [6](#), [47](#)
- [30] H. S. J. van der Zant, F. C. Fritschy, W. J. Elion, L. J. Geerligs, and J.E. Mooij. Field-induced superconductor-to-insulator transitions in Josephson-junction arrays. *Phys. Rev. Lett.*, 69:2971, 1992. [6](#), [47](#)
- [31] J.K. Freericks and H. Monien. Phase diagram of the Bose-Hubbard model. *Europhys. Lett.*, 26:545, 1994. [6](#), [44](#), [48](#)
- [32] N. F. Mott. The basis of the electron theory of metals, with special reference to the transition metals. *Proc. Phys. Soc. A*, 62:416, 1949. [6](#)
- [33] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415:39, 2002. [7](#), [43](#), [188](#)
- [34] U. Schneider, L. Hackermüller, S. Will, Th. Best, I. Bloch, T. A. Costi, R. W. Helmes, D. Rasch, and A. Rosch. Metallic and insulating phases of repulsively interacting fermions in a 3D optical lattice. *Science*, 322:1520, 2008. [7](#), [45](#), [142](#)
- [35] R. Jördens, N. Strohmaier, K. Günter, H. Moritz, and T. Esslinger. A Mott insulator of fermionic atoms in an optical lattice. *Nature*, 455:204, 2008. [7](#), [45](#)
- [36] D. C. McKay and B. DeMarco. Cooling in strongly correlated lattices: prospects and challenges. *Rep. Prog. Phys.*, 74:054401, 2011. [7](#), [39](#), [102](#), [110](#), [144](#), [145](#), [190](#), [237](#)
- [37] H. Pichler, A. J. Daley, and P. Zoller. Nonequilibrium dynamics of bosonic atoms in optical lattices: Decoherence of many-body states due to spontaneous emission. *Phys. Rev. A*, 82:063605, 2010. [7](#), [39](#), [57](#), [110](#), [166](#), [237](#)

- [38] A. Polkovnikov, K. Sengupta, A. Silva, and M. Vengalattore. Colloquium: Nonequilibrium dynamics of closed interacting quantum systems. *Rev. Mod. Phys.*, 83:863, 2011. [7](#), [138](#)
- [39] B. Capogrosso-Sansone, N. V. Prokof'ev, and B. V. Svistunov. Phase diagram and thermodynamics of the three-dimensional Bose-Hubbard model. *Phys. Rev. B*, 75:134302, 2007. [7](#), [41](#), [42](#), [43](#)
- [40] N. E. Bickers, D. J. Scalapino, and S. R. White. Conserving approximations for strongly correlated electron systems: Bethe-Salpeter equation and dynamics for the two-dimensional Hubbard model. *Phys. Rev. Lett.*, 62:961, 1989. [7](#), [47](#)
- [41] Matthew White. *Ultracold Atoms in a Disordered Optical Lattice*. PhD thesis, University of Illinois, 2009. [7](#), [19](#), [49](#), [51](#), [52](#), [53](#), [54](#), [55](#), [85](#), [186](#), [239](#), [294](#)
- [42] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Simplified system for creating a Bose-Einstein condensate. *J. Low Temp. Phys.*, 132:209, 2002. [7](#)
- [43] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In M. Inguscio, S. Stringari, and C. E. Wieman, editors, *Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics "Enrico Fermi", Course CXL*, page 67. IOS Press, 1999. [7](#), [17](#)
- [44] L. De Sarlo, L. Fallani, J. E. Lye, M. Modugno, R. Saers, C. Fort, and M. Inguscio. Unstable regimes for a Bose-Einstein condensate in an optical lattice. *Phys. Rev. A*, 72:013603, 2005. [8](#), [132](#)
- [45] L. Fallani, L. De Sarlo, J. E. Lye, M. Modugno, R. Saers, C. Fort, and M. Inguscio. Observation of dynamical instability for a Bose-Einstein condensate in a moving 1D optical lattice. *Phys. Rev. Lett.*, 93:140406, 2004. [8](#)
- [46] A. Polkovnikov, E. Altman, E. Demler, B. Halperin, and M. D. Lukin. Decay of superfluid currents in a moving system of strongly interacting bosons. *Phys. Rev. A*, 71:063613, 2005. [8](#), [130](#), [135](#), [136](#)
- [47] J. S. Lander and V. Ambegaokar. Intrinsic resistive transition in narrow superconducting channels. *Phys. Rev.*, 164:498, 1967. [8](#), [130](#)
- [48] M. Tinkham. *Introduction to Superconductivity*. McGraw Hill, 2 edition, 1996. [8](#), [111](#), [130](#)
- [49] A. Bezryadin, C. N. Lau, and M. Tinkham. Quantum suppression of superconductivity in ultrathin nanowires. *Nature*, 404:971, 2000. [8](#), [130](#), [137](#)
- [50] A. M. Goldman. Superconductor-insulator transitions in the two-dimensional limit. *Physica E*, 18:1, 2003. [8](#)
- [51] C. D. Fertig, K. M. O'Hara, J. H. Huckans, S. L. Rolston, W. D. Phillips, and J. V. Porto. Strongly inhibited transport of a degenerate 1D Bose gas in a lattice. *Phys. Rev. Lett.*, 94:120403, 2005. [8](#), [132](#)
- [52] J. Mun, P. Medley, G. K. Campbell, L. G. Marcassa, D. E. Pritchard, and W. Ketterle. Phase diagram for a Bose-Einstein condensate moving in an optical lattice. *Phys. Rev. Lett.*, 99:150604, 2007. [8](#), [132](#)
- [53] A. J. Leggett. Bose-Einstein condensation in the alkali gases: Some fundamental concepts. *Rev. Mod. Phys.*, 73:307, 2001. [8](#), [94](#), [154](#)
- [54] K. V. Krutitsky and P. Navez. Excitation dynamics in a lattice Bose gas within the time-dependent Gutzwiller mean-field approach. *Phys. Rev. A*, 84:033602, 2011. [8](#), [44](#)
- [55] D. A. Butts and D. S. Rokhsar. Trapped fermi gases. *Phys. Rev. A*, 55:4346, 1997. [9](#)
- [56] N. Gemelke, X. Zhang, C.-L. Hung, and C. Chin. In situ observation of incompressible Mott-insulating domains in ultracold atomic gases. *Nature*, 460:995, 2009. [9](#), [10](#), [43](#), [102](#), [152](#)
- [57] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature*, 467:68, 2010. [9](#), [10](#), [12](#), [43](#), [81](#), [82](#), [152](#)

- [58] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner. Probing the superfluid-to-Mott insulator transition at the single-atom level. *Science*, 329:547, 2010. [9](#), [10](#), [12](#), [43](#), [82](#), [140](#), [152](#)
- [59] D. M. Weld, P. Medley, H. Miyake, D. Hucul, D. E. Pritchard, and W. Ketterle. Spin gradient thermometry for ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 103:245301, 2009. [10](#), [152](#)
- [60] Y. Kato, Q. Zhou, N. Kawashima, and N. Trivedi. Sharp peaks in the momentum distribution of bosons in optical lattices in the normal state. *Nature Physics*, 4:617, 2008. [10](#), [153](#), [156](#), [164](#)
- [61] Q. Zhou, Y. Kato, N. Kawashima, and N. Trivedi. Direct mapping of the finite temperature phase diagram of strongly correlated quantum models. *Phys. Rev. Lett.*, 103:085701, 2009. [10](#)
- [62] T.-L. Ho and Q. Zhou. Obtaining the phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases. *Nature Physics*, 6:131, 2010. [10](#), [43](#), [101](#), [151](#)
- [63] P. N. Ma, L. Pollet, and M. Troyer. Measuring the equation of state of trapped ultracold bosonic systems in an optical lattice with in-situ density imaging. *Phys. Rev. A*, 82:033627, 2010. [10](#)
- [64] G.-D. Lin, W. Zhang, and L.-M. Duan. Characteristics of Bose-Einstein condensation in an optical lattice. *Phys. Rev. A*, 77:043626, 2008. [10](#), [123](#)
- [65] C. Schori, T. Stöferle, H. Moritz, M. Köhl, and T. Esslinger. Excitations of a superfluid in a three-dimensional optical lattice. *Phys. Rev. Lett.*, 93:240402, 2004. [10](#)
- [66] I. B. Spielman, W. D. Phillips, and J. V. Porto. Condensate fraction in a 2D Bose gas measured across the mott-insulator transition. *Phys. Rev. Lett.*, 100:120402, 2008. [10](#), [43](#), [150](#)
- [67] S. Trotzky, L. Pollet, F. Gerbier, U. Schnorrberger, I. Bloch, N.V. Prokof'ev, B. Svistunov, and M. Troyer. Suppression of the critical temperature for superfluidity near the Mott transition. *Nature Physics*, 6:998, 2010. [10](#), [43](#), [110](#), [153](#), [164](#), [166](#)
- [68] K. Sheshadri, H. R. Krishnamurthy, R. Pandit, and T. V. Ramakrishnan. Superfluid and insulating phases in an interacting-boson model: Mean-field theory and the RPA. *Europhys. Lett.*, 257:22, 1993. [10](#), [44](#), [124](#)
- [69] T.-L. Ho and Q. Zhou. Squeezing out the entropy of fermions in optical lattices. *Proc. Nat. Acad. Sci.*, 106:6916, 2009. [11](#), [181](#)
- [70] A. Griessner, A. J. Daley, S. R. Clark, D. Jaksch, and P. Zoller. Dark-state cooling of atoms by superfluid immersion. *Phys. Rev. Lett.*, 97:220403, 2006. [11](#), [187](#)
- [71] L. J. LeBlanc and J. H. Thywissen. Species-specific optical lattices. *Phys. Rev. A*, 75:053612, 2007. [11](#)
- [72] J. Catani, G. Barontini, G. Lamporesi, F. Rabatti, G. Thalhammer, F. Minardi, S. Stringari, and M. Inguscio. Entropy exchange in a mixture of ultracold atoms. *Phys. Rev. Lett.*, 103:140401, 2009. [11](#), [166](#)
- [73] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hänsch, and I. Bloch. Coherent transport of neutral atoms in spin-dependent optical lattice potentials. *Phys. Rev. Lett.*, 91:010407, 2003. [11](#), [54](#), [55](#), [229](#)
- [74] M. Karski, L. Förster, J.-M. Choi, A. Steffen, W. Alt, D. Meschede, and A. Widera. Quantum walk in position space with single optically trapped atoms. *Science*, 325:174, 2009. [11](#), [229](#)
- [75] B. Gadway, D. Pertot, R. Reimann, and D. Schneble. Superfluidity of interacting bosonic mixtures in optical lattices. *Phys. Rev. Lett.*, 105:045303, 2010. [11](#), [55](#)
- [76] L. Förster, M. Karski, J.-M. Choi, A. Steffen, W. Alt, D. Meschede, A. Widera, E. Montano, J. H. Lee, W. Rakreungdet, and P. S. Jessen. Microwave control of atomic motion in optical lattices. *Phys. Rev. Lett.*, 103:233001, 2009. [11](#), [55](#), [107](#)

- [77] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Phys. Rev. A*, 56:R1095–R1098, 1997. [11](#), [175](#), [233](#)
- [78] R. Jáuregui, N. Poli, G. Roati, and G. Modugno. Anharmonic parametric excitation in optical lattices. *Phys. Rev. A*, 64:033403, 2001. [11](#)
- [79] T. Müller, S. Fölling, A. Widera, and I. Bloch. State preparation and dynamics of ultracold atoms in higher lattice orbitals. *Phys. Rev. Lett.*, 99:200405, 2007. [11](#), [175](#)
- [80] J. H. Challis, S. M. Girvin, and L. I. Glazman. Theory of preparation and relaxation of a p -orbital atomic Mott insulator. *Phys. Rev. A*, 79:043609, 2009. [11](#)
- [81] J.-P. Martikainen. Dynamical instability and loss of p -band bosons in optical lattices. *Phys. Rev. A*, 83:013610, 2011. [11](#)
- [82] B. Gadway, D. Pertot, J. Reeves, and D. Schneble. Probing an ultracold-atom crystal with matter waves. *Nature Physics*, 2012. [12](#), [183](#)
- [83] S. Palzer, C. Zipkes, C. Sias, and M. Köhl. Quantum transport through a Tonks-Girardeau gas. *Phys. Rev. Lett.*, 103:150601, 2009. [12](#), [183](#)
- [84] G. Ferrari, M. Inguscio, W. Jastrzebski, G. Modugno, G. Roati, and A. Simoni. Collisional properties of ultracold K-Rb mixtures. *Phys. Rev. Lett.*, 89:053202–1, 2002. [12](#), [183](#)
- [85] P. Maddaloni, M. Modugno, C. Fort, F. Minardi, and M. Inguscio. Collective oscillations of two colliding Bose-Einstein condensates. *Phys. Rev. Lett.*, 85:2413, 2000. [12](#), [183](#), [184](#)
- [86] B. Horstmann, S. Dürr, and T. Roscilde. Localization of cold atoms in state-dependent optical lattices via a Rabi pulse. *Phys. Rev. Lett.*, 105:160402, 2010. [12](#)
- [87] B. Paredes, F. Verstraete, , and J. I. Cirac. Exploiting quantum parallelism to simulate quantum random many-body systems. *Phys. Rev. Lett.*, 95:140501, 2005. [12](#)
- [88] D. C. McKay, D. Jervis, D. J. Fine, J. W. Simpson-Porco, G. J. A. Edge, and J. H. Thywissen. Low-temperature, high-density magneto-optical trapping of potassium using the open 4s-5p transition at 405 nm. *Phys. Rev. A*, 84:063420, 2011. [12](#), [264](#)
- [89] A. S. Tychkov, J. C. J. Koelemeij, T. Jeltès, W. Hogervorst, and W. Vassen. Two-color magneto-optical trap for metastable helium. *Phys. Rev. A*, 69:055401, 2004. [13](#), [251](#), [268](#)
- [90] J. C. J. Koelemeij, R. J. W. Stas, W. Hogervorst, and W. Vassen. Magneto-optical trap for metastable helium at 389 nm. *Phys. Rev. A*, 67:053406, 2003. [13](#), [251](#), [268](#)
- [91] P. M. Duarte, R. A. Hart, J. M. Hitchcock, T. A. Corcovilos, T. L. Yang, A. Reed, and R. G. Hulet. All-optical production of a lithium quantum gas using narrow-line laser cooling. *Phys. Rev. A*, 84:061406(R), 2011. [13](#), [251](#)
- [92] C. Sanner, E. J. Su, W. Huang, A. Keshet, J. Gillen, and W. Ketterle. Correlations and pair formation in a repulsively interacting Fermi gas. *Phys. Rev. Lett.*, 108:240404, 2012. [14](#), [29](#)
- [93] W. D. Phillips and H. Metcalf. Laser deceleration of an atomic beam. *Phys. Rev. Lett.*, 48:596, 1982. [15](#)
- [94] C. Monroe, W. Swann, H. Robinson, and C. Wieman. Very cold trapped atoms in a vapor cell. *Phys. Rev. Lett.*, 65:1571, 1990. [15](#)
- [95] P. Böhi, M. F. Riedel, J. Hoffrogge, J. Reichel, T. W. Hänsch, and P. Treutlein. Coherent manipulation of Bose-Einstein condensates with state-dependent microwave potentials on an atom chip. *Nature Physics*, 5:592, 2009. [15](#)

- [96] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer. Adiabatic radio-frequency potentials for the coherent manipulation of matter waves. *Phys. Rev. A*, 73:033619, 2006. [15](#)
- [97] R. Gerritsma, S. Whitlock, T. Fernholz, H. Schlatter, J. A. Luigjes, J.-U. Thiele, J. B. Goedkoop, and R. J. C. Spreeuw. Lattice of microtraps for ultracold atoms based on patterned magnetic films. *Phys. Rev. A*, 76:033408, 2007. [17](#), [188](#)
- [98] T.-L. Ho. Spinor Bose condensates in optical traps. *Phys. Rev. Lett.*, 81:742, 1998. [17](#), [196](#)
- [99] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. In B. Bederson and H. Walther, editors, *Advances In Atomic, Molecular, and Optical Physics*, volume 42, page 95. Academic Press, 2000. [18](#), [209](#)
- [100] P. W. Milonni and J. H. Eberly. *Lasers*. Wiley, 1989. [18](#), [230](#), [231](#)
- [101] D. S. Naik and C. Raman. Optically plugged quadrupole trap for Bose-Einstein condensates. *Phys. Rev. A*, 71:033617, 2005. [18](#), [59](#), [78](#)
- [102] C.-L. Hung, X. Zhang, N. Gemelke, and C. Chin. Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps. *Phys. Rev. A*, 78:011604, 2008. [18](#), [51](#)
- [103] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto. Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential. *Phys. Rev. A*, 79:063631, 2009. [18](#), [51](#), [200](#)
- [104] P. Soltan-Panahi, J. Struck, P. Hauke, A. Bick, W. Plenkers, G. Meineke, C. Becker, P. Windpassinger, M. Lewenstein, and K. Sengstock. Multi-component quantum gases in spin-dependent hexagonal lattices. *Nature Physics*, 7:434, 2011. [19](#)
- [105] C. Becker, P. Soltan-Panahi, J. Kronjäger, S. Dörscher, K. Bongs, and K. Sengstock. Ultracold quantum gases in triangular optical lattices. *New. J. Phys.*, 12:065025, 2010. [19](#)
- [106] G.-B. Jo, J. Guzman, C. K. Thomas, P. Hosur, A. Vishwanath, and D. M. Stamper-Kurn. Ultracold atoms in a tunable optical Kagome lattice. *Phys. Rev. Lett.*, 108:045305, 2012. [19](#)
- [107] L. Tarruell, D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger. Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice. *Nature*, 483:302, 2012. [19](#)
- [108] S. Fölling, S. Trotzky, P. Cheinet, M. Feld, R. Saers, A. Widera, T. Müller, and I. Bloch. Direct observation of second-order atom tunnelling. *Nature*, 448:1029, 2007. [19](#)
- [109] J. Sebby-Strabley, M. Anderlini, P. S. Jessen, and J. V. Porto. Lattice of double wells for manipulating pairs of cold atoms. *Phys. Rev. A*, 73:033605, 2006. [19](#)
- [110] Matthew Pasienski. *Transport Properties of Ultracold Atoms in a Disordered Optical Lattice*. PhD thesis, University of Illinois, 2011. [19](#)
- [111] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga. Feshbach resonances in ultracold gases. *Rev. Mod. Phys.*, 82:1225, 2010. [19](#), [21](#), [193](#)
- [112] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, 1994. [19](#), [89](#), [199](#), [202](#)
- [113] J. Dalibard. Collisional dynamics of ultra-cold atomic gases. In M. Inguscio, S. Stringari, and C. E. Wieman, editors, *Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics “Enrico Fermi”, Course CXL*, page 321. IOS Press, 1999. [21](#)
- [114] K. Huang. *Statistical Mechanics*. Wiley, 1987. [22](#)
- [115] M. Lu, N. Q. Burdick, S. H. Youn, and B. L. Lev. Strongly dipolar Bose-Einstein condensate of dysprosium. *Phys. Rev. Lett.*, 107:190401, 2011. [22](#)

- [116] J. J. McClelland and J. L. Hanssen. Laser cooling without repumping: A magneto-optical trap for erbium atoms. *Phys. Rev. Lett.*, 96:143005, 2006. [22](#)
- [117] C. C. Bradley, J. J. McClelland, W. R. Anderson, and R. J. Celotta. Magneto-optical trapping of chromium atoms. *Phys. Rev. A*, 61:053407, 2000. [22](#)
- [118] J. P. D’Incao and B. D. Esry. Scattering length scaling laws for ultracold three-body collisions. *Phys. Rev. Lett.*, 94:213201, 2005. [22](#)
- [119] D. S. Petrov, C. Salomon, and G. V. Shlyapnikov. Weakly bound dimers of fermionic atoms. *Phys. Rev. Lett.*, 93:090404, 2004. [22](#), [29](#)
- [120] R. Shankar. *Principles of Quantum Mechanics*. Plenum Press, 1994. [23](#)
- [121] B. DeMarco, H. Rohner, and D. S. Jin. An enriched 40K source for fermionic atom studies. *Rev. Sci. Instrum.*, 70:1967, 1999. [24](#), [59](#), [65](#)
- [122] A. L. Fetter and J. D. Walecka. *Quantum Theory of Many-Particle Systems*. Dover, 2003. [25](#)
- [123] S. Giorgini, L. P. Pitaevskii, and S. Stringari. Thermodynamics of a trapped Bose-condensed gas. *J. Low Temp. Phys.*, 109:309, 1997. [25](#), [26](#), [28](#), [121](#), [122](#)
- [124] A. Griffin, T. Nikuni, and E. Zaremba. *Bose-Condensed Gases at Finite Temperatures*. Cambridge University Press, 2009. [26](#), [111](#), [121](#), [122](#), [132](#), [181](#)
- [125] C. J. Pethick and H. Smith. *Bose-Einstein condensation in dilute gases*. Cambridge University Press, 2002. [26](#), [169](#), [181](#)
- [126] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer, and A. Aspect. Experimental study of the thermodynamics of an interacting trapped Bose-Einstein condensed gas. *Phys. Rev. A*, 70:013607, 2004. [28](#), [95](#), [154](#), [285](#)
- [127] M. Naraschewski and D. M. Stamper-Kurn. Analytical description of a trapped semi-ideal Bose gas at finite temperature. *Phys. Rev. A*, 58:2423, 1998. [28](#)
- [128] R. P. Smith, N. Tammuz, R. L. D. Campbell, M. Holzmann, and Z. Hadzibabic. Condensed fraction of an atomic Bose gas induced by critical correlations. *Phys. Rev. Lett.*, 107:190403, 2011. [28](#)
- [129] S. Giorgini, L. P. Pitaevskii, and S. Stringari. Thermodynamics of a trapped Bose-condensed gas. *Rev. Mod. Phys.*, 80:1215, 2008. [29](#)
- [130] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:1175, 1957. [29](#)
- [131] W. Ketterle and M. W. Zwierlein. Making, probing and understanding ultracold Fermi gases. In M. Inguscio, W. Ketterle, and C. Salomon, editors, *Ultracold Fermi Gases, Proceedings of the International School of Physics “Enrico Fermi”, Course CLXIV*. IOS Press, 2008. [29](#)
- [132] W. Zwerger. Mott-Hubbard transition of cold atoms in optical lattices. *J. Opt. B*, 5:S9, 2003. [33](#), [36](#)
- [133] A. M. Rey, G. Pupillo, C. W. Clark, and C. J. Williams. Ultracold atoms confined in an optical lattice plus parabolic potential: A closed-form approach. *Phys. Rev. A*, 72:033616, 2005. [34](#), [35](#), [116](#), [147](#)
- [134] X.-F. Zhou, Z.-X. Chen, Z.-W. Zhou, Y.-S. Zhang, and G.-C. Guo. Frustrated tunneling of ultracold atoms in a state-dependent optical lattice. *Phys. Rev. A*, 81:021602(R), 2010. [38](#)
- [135] T. D. Kühner, S. R. White, and H. Monien. One-dimensional Bose-Hubbard model with nearest-neighbor interaction. *Phys. Rev. B*, 61:12474, 2000. [38](#)
- [136] S. Will, T. Best, U. Schneider, L. Hackermüller, D.-S. Lühmann, and I. Bloch. Time-resolved observation of coherent multi-body interactions in quantum phase revivals. *Nature*, 465:197, 2010. [38](#)

- [137] F. Gerbier. Boson Mott insulators at finite temperatures. *Phys. Rev. Lett.*, 99:120405, 2007. [42](#), [43](#), [120](#), [151](#)
- [138] I. B. Spielman, W. D. Phillips, and J. V. Porto. Mott-insulator transition in a two-dimensional atomic Bose gas. *Phys. Rev. Lett.*, 98:080404, 2007. [43](#)
- [139] M. Köhl, H. Moritz, T. Stöferle, C. Schori, and T. Esslinger. Superfluid to Mott insulator transition in one, two, and three dimensions. *J. Low Temp. Phys.*, 138:635, 2005. [43](#), [114](#), [142](#)
- [140] S. Fölling, A. Widera, T. Müller, F. Gerbier, and I. Bloch. Formation of spatial shell structure in the superfluid to Mott insulator transition. *Phys. Rev. Lett.*, 97:060403, 2006. [43](#), [105](#)
- [141] G. K. Campbell, J. Mun, M. Boyd, P. Medley, A. E. Leanhardt, L. G. Marcassa, D. E. Pritchard, and W. Ketterle. Imaging the Mott insulator shells by using atomic clock shifts. *Science*, 313(5787):649, 2006. [43](#), [105](#)
- [142] T. D. Kühner and H. Monien. Phases of the one-dimensional Bose-Hubbard model. *Phys. Rev. B*, 58:R14741, 1998. [43](#)
- [143] B. Capogrosso-Sansone, Ş. G. Söyler, N. Prokof'ev, and B. Svistunov. Monte Carlo study of the two-dimensional Bose-Hubbard model. *Phys. Rev. A*, 77:015602, 2008. [43](#)
- [144] P. Anders, E. Gull, L. Pollet, M. Troyer, and P. Werner. Dynamical mean field solution of the Bose-Hubbard model. *Phys. Rev. Lett.*, 105:096402, 2010. [43](#), [128](#)
- [145] P. Anders, E. Gull, L. Pollet, M. Troyer, and P. Werner. Dynamical mean-field theory for bosons. *New. J. Phys.*, 13:07013, 2011. [43](#), [128](#)
- [146] K. Sengupta and N. Dupuis. Mott-insulator-to-superfluid transition in the Bose-Hubbard model: A strong-coupling approach. *Phys. Rev. A*, 71:033629, 2005. [44](#), [120](#), [128](#)
- [147] D. van Oosten, P. van der Straten, and H. T. C. Stoof. Quantum phases in an optical lattice. *Phys. Rev. A*, 63:053601, 2001. [44](#), [123](#)
- [148] M. J. Rozenberg, R. Chitra, and G. Kotliar. Finite temperature Mott transition in the Hubbard model in infinite dimensions. *Phys. Rev. Lett.*, 83:3498, 1999. [45](#), [47](#)
- [149] A. B. Kuklov and B. V. Svistunov. Counterflow superfluidity of two-species ultracold atoms in a commensurate optical lattice. *Phys. Rev. Lett.*, 90:100401, 2003. [45](#)
- [150] F. Werner, O. Parcollet, A. Georges, and S. R. Hassan. Interaction-induced adiabatic cooling and antiferromagnetism of cold fermions in optical lattices. *Phys. Rev. Lett.*, 95:056401, 2005. [45](#)
- [151] D. Greif, L. Tarruell, T. Uehlinger, R. Jördens, and T. Esslinger. Probing nearest-neighbor correlations of ultracold fermions in an optical lattice. *Phys. Rev. Lett.*, 106:145302, 2011. [45](#)
- [152] R. Jördens, L. Tarruell, D. Greif, T. Uehlinger, N. Strohmaier, H. Moritz, T. Esslinger, L. De Leo, C. Kollath, A. Georges, V. Scarola, L. Pollet, E. Burovski, E. Kozik, and M. Troyer. Quantitative determination of temperature in the approach to magnetic order of ultracold fermions in an optical lattice. *Phys. Rev. Lett.*, 104:180401, 2010. [47](#)
- [153] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.*, 68:13, 1996. [47](#), [48](#), [128](#)
- [154] R. Staudt, M. Dzierzawa, and A. Muramatsu. Phase diagram of the three-dimensional Hubbard model at half filling. *Euro. Phys. J. B*, 17:411, 2000. [47](#)
- [155] T. Paiva, R. Scalettar, M. Randeria, and N. Trivedi. Fermions in 2D optical lattices: Temperature and entropy scales for observing antiferromagnetism and superfluidity. *Phys. Rev. Lett.*, 104:066406, 2010. [47](#)

- [156] E. Kozik, K. Van Houcke, E. Gull, L. Pollet, N. Prokof'ev, B. Svistunov, and M. Troyer. Diagrammatic Monte Carlo for correlated fermions. *Europhys. Lett.*, 90:10004, 2010. [47](#), [128](#)
- [157] P. R. C. Kent, M. Jarrell, T. A. Maier, and Th. Pruschke. Efficient calculation of the antiferromagnetic phase diagram of the three-dimensional Hubbard model. *Phys. Rev. B*, 72:060411(R), 2005. [47](#)
- [158] D. F. B. ten Haaf and J. M. J. van Leeuwen. High-temperature series expansions for the Hubbard model. *Phys. Rev. B*, 46:6313, 1992. [47](#), [128](#)
- [159] L. De Leo, J.-S. Bernier, C. Kollath, A. Georges, and V. W. Scarola. Thermodynamics of the three-dimensional Hubbard model: Implications for cooling cold atomic gases in optical lattices. *Phys. Rev. A*, 023606:83, 2011. [47](#)
- [160] C. J. Jia, B. Moritz, C.-C. Chen, B. Sriram Shastry, and T. P. Devereaux. Fidelity study of the superconducting phase diagram in the two-dimensional single-band Hubbard model. *Phys. Rev. B*, 84:125113, 2011. [47](#)
- [161] E. H. Lieb and F. Y. Wu. Absence of Mott transition in an exact solution of the short-range, one-band model in one dimension. *Phys. Rev. Lett.*, 20:1445, 1968. [47](#)
- [162] M. Rigol, A. Muramatsu, G. G. Batrouni, and R. T. Scalettar. Local quantum criticality in confined fermions on optical lattices. *Phys. Rev. Lett.*, 91:130403, 2003. [47](#)
- [163] R. W. Helmes, T. A. Costi, and A. Rosch. Mott transition of fermionic atoms in a three-dimensional optical trap. *Phys. Rev. Lett.*, 100:056403, 2008. [47](#)
- [164] Y. Nagaoka. Ferromagnetism in a narrow, almost half-filled s band. *Phys. Rev.*, 147:392, 1966. [47](#)
- [165] D. J. Scalapino. The case for $d_{x^2-y^2}$ pairing in the cuprate superconductors. *Phys. Rep.*, 250:329, 1995. [47](#)
- [166] T. A. Maier, M. Jarrell, T. C. Schulthess, P. R. C. Kent, and J. B. White. Systematic study of d-wave superconductivity in the 2D repulsive Hubbard model. *Phys. Rev. Lett.*, 95:237001, 2005. [47](#)
- [167] T. Gericke, F. Gerbier, A. Widera, S. Fölling, O. Mandel, and I. Bloch. Adiabatic loading of a Bose-Einstein condensate in a 3D optical lattice. *J. Mod. Opt.*, 54:735, 2007. [53](#), [109](#), [166](#)
- [168] J. P. Gordon and A. Ashkin. Motion of atoms in a radiation trap. *Phys. Rev. A*, 21:1606, 1980. [57](#), [218](#)
- [169] F. Gerbier and Y. Castin. Heating rates for an atom in a far-detuned optical lattice. *Phys. Rev. A*, 82:013615, 2010. [57](#)
- [170] J. F. O'Hanlon. *A User's Guide to Vacuum Technology*. Wiley, 3 edition, 2003. [60](#)
- [171] Amir Mazouchi. Feasibility of imaging single atoms in an optical lattice. Master's thesis, University of Toronto, 2007. [61](#)
- [172] G. Modugno, C. Benkō, P. Hannaford, G. Roati, and M. Inguscio. Sub-doppler laser cooling of fermionic ^{40}K atoms. *Phys. Rev. A*, 60:R3373, 1999. [63](#), [250](#)
- [173] Brian DeMarco. *Quantum Behavior of an Atomic Fermi Gas*. PhD thesis, University of Colorado - Boulder, 2001. [63](#), [262](#), [285](#)
- [174] Jonathan Goldwin. *Quantum Degeneracy and Interactions in the 87Rb - 40K Bose-Fermi Mixture*. PhD thesis, University of Colorado, 2005. [63](#)
- [175] Joshua Zirbel. *Ultracold Fermionic Feshbach Molecules*. PhD thesis, University of Colorado, 2008. [63](#)
- [176] Silke Ospelkaus-Schwarzer. *Quantum Degenerate Fermi-Bose Mixtures of 40K and 87Rb in 3D Optical Lattices*. PhD thesis, Universität Hamburg, 2006. [63](#)

- [177] Tim Rom. *Bosonische und fermionische Quantengase in dreidimensionalen optischen Gittern: Präparation, Manipulation und Analyse*. PhD thesis, Ludwig-Maximilians-Universität München, 2009. [63](#)
- [178] Kenneth Günter. *Interacting Fermi gases and Bose-Fermi mixtures in optical lattices*. PhD thesis, Swiss Federal Institute of Technology Zurich, 2007. [63](#)
- [179] Marcius Extavour. *Fermions and bosons on an atom chip*. PhD thesis, University of Toronto, 2009. [63](#)
- [180] Giacomo Roati. *Quantum Degenerate Potassium-Rubidium Mixtures*. PhD thesis, Università Degli Studi De Trento, 2003. [63](#)
- [181] D.-Z. Xiong, H.-X. Chen, P.-J. Wang, X.-D. Yu, F. Gao, and J. Zhang. Quantum degenerate Fermi-Bose mixtures of 40K and 87Rb atoms in a quadrupole-Ioffe configuration trap. *Chin. Phys. Lett.*, 25:843, 2007. [63](#)
- [182] Carsten Klempt. *Wechselwirkung in Bose-Fermi-Quantengasen*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, 2007. [63](#)
- [183] Matthias Taglieber. *Quantum Degeneracy in an Atomic Fermi-Fermi-Bose Mixture*. PhD thesis, Ludwig-Maximilians-Universität München, 2008. [63](#)
- [184] Eric Wille. *Preparation of an Optically Trapped Fermi-Fermi Mixture of 6Li and 40K Atoms and Characterization of the Interspecies Interactions by Feshbach Spectroscopy*. PhD thesis, Universität Innsbruck, 2009. [63](#)
- [185] C.-H. Wu, I. Santiago, J. W. Park, P. Ahmadi, and M. W. Zwierlein. Strongly interacting isotopic Bose-Fermi mixture immersed in a Fermi sea. *Phys. Rev. A*, 84:011601(R), 2011. [63](#)
- [186] S. S. Kondov, W. R. McGehee, J. J. Zirbel, and B. DeMarco. Three-dimensional localization of ultracold matter. *Science*, 334:66, 2011. [63](#)
- [187] B. Fröhlich, M. Feld, E. Vogt, M. Koschorreck, W. Zwerger, and M. Köhl. Radio-frequency spectroscopy of a strongly interacting two-dimensional Fermi gas. *Phys. Rev. Lett.*, 106:105301, 2011. [63](#)
- [188] A. Ridinger, S. Chaudhuri, T. Salez, U. Eismann, D. R. Fernandes, K. Magalhães, D. Wilkowski, C. Salomon, and F. Chevy. Large atom number dual-species magneto-optical trap for fermionic ^6Li and ^{40}K atoms. *Euro. Phys. J. D*, 65:223, 2011. [63](#), [268](#)
- [189] Tobias Tiecke. *Feshbach resonances in ultracold mixtures of the fermionic quantum gases ^6Li and ^{40}K* . PhD thesis, University of Amsterdam, 2009. [63](#), [268](#)
- [190] B. Shirinzadeh and C. Wang. Accurate determination of the vapor pressure of potassium using optical absorption. *Applied Optics*, 22:3265, 1983. [65](#), [189](#), [255](#)
- [191] Alma Bardon. Private Communication, 2011. [67](#)
- [192] J. E. Sansonetti. Wavelengths, transition probabilities, and energy levels for the spectra of potassium (K I through K XIX). *J. Phys. Chem. Ref. Data*, 37:7, 2008. [68](#), [189](#), [190](#), [279](#)
- [193] J. E. Sansonetti. Wavelengths, transition probabilities, and energy levels for the spectra of rubidium (Rb I through Rb XXXVII). *J. Phys. Chem. Ref. Data*, 35:301, 2006. [69](#), [189](#), [191](#), [194](#), [239](#)
- [194] Felix Stubenrauch. Frequency modulation saturation spectroscopy laser lock of an interference-filter-stabilized external-cavity diode laser. Research report, University of Toronto, 2010. [69](#)
- [195] Matthias Scholl. Interference filter stabilized external cavity diode laser. Research report, University of Toronto, 2010. [69](#)

- [196] William Cairncross and Matthias Scholl. Interference filter stabilized external cavity diode lasers. Research report, University of Toronto, 2011. [69](#)
- [197] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard. High densities of cold atoms in a dark spontaneous-force optical trap. *Phys. Rev. Lett.*, 70:2253, 1993. [72](#)
- [198] Josefine Metzkes. Design of a magnetic transport system. Research report, University of Toronto, 2008. [76](#)
- [199] Michael Yee. Magnetic trapping and transport of ultracold atoms. BAsC thesis, University of Toronto, 2009. [76](#)
- [200] M. Greiner, I. Bloch, T. W. Hänsch, and T. Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Phys. Rev. A*, 63:031401(R), 2001. [76](#), [142](#)
- [201] Michael Sprague. Frequency offset lock and microwave evaporation for fermions in an optical lattice. Master’s thesis, University of Toronto, 2010. [78](#)
- [202] B. DeMarco, J. L. Bohn, J. P. Burke, M. Holland, and D. S. Jin. Measurement of p-wave threshold law using evaporatively cooled fermionic atoms. *Phys. Rev. Lett.*, 82:4208, 1999. [78](#), [195](#)
- [203] Brian DeMarco. Private Communication, 2010. [78](#)
- [204] K. Jiménez-García, R. L. Compton, Y.-J. Lin, W. D. Phillips, J.V. Porto, and I. B. Spielman. Phases of a two-dimensional Bose gas in an optical lattice. *Phys. Rev. Lett.*, 105:110401, 2010. [81](#), [142](#)
- [205] M. Born and E. Wolf. *Principles of Optics, 7th (expanded) Edition*. Cambridge University Press, 2002. [83](#), [103](#), [274](#), [276](#)
- [206] Gabriello Presenza-Pitman. Determination of the contrast and modulation transfer functions for high resolution imaging of individual atoms. NSERC summer report, University of Toronto, 2009. [83](#)
- [207] H. Im, S. H. Lee, N. J. Wittenberg, T. W. Johnson, N. C. Lindquist, P. Nagpal, D. J. Norris, and S.-H. Oh. Template-stripped smooth ag nanohole arrays with silica shells for surface plasmon resonance biosensing. *ACS Nano*, 5:6244, 2011. [83](#)
- [208] Tout Wang. Interferometric optical testing for high resolution imaging in an optical lattice. NSERC summer report, University of Toronto, 2008. [83](#)
- [209] F. Henkel, M. Krug, J. Hofmann, W. Rosenfeld, M. Weber, and H. Weinfurter. Highly efficient state-selective submicrosecond photoionization detection of single atoms. *Phys. Rev. Lett.*, 105:253001, 2010. [86](#)
- [210] T. Gericke, P. Würtz, D. Reitz, T. Langen, and H. Ott. High-resolution scanning electron microscopy of an ultracold quantum gas. *Nature Physics*, 4:949, 2008. [86](#)
- [211] H. Miyake, G. A. Siviloglou, G. Puentes, D. E. Pritchard, W. Ketterle, and D. M. Weld. Bragg scattering as a probe of atomic wave functions and quantum phase transitions in optical lattices. *Phys. Rev. Lett.*, 107:175302, 2011. [86](#)
- [212] C. Weitenberg, P. Schauß, T. Fukuhara, M. Cheneau, M. Endres, I. Bloch, and S. Kuhr. Coherent light scattering from a two-dimensional Mott insulator. *Phys. Rev. Lett.*, 106:215301, 2011. [86](#), [103](#)
- [213] T. A. Corcovilos, S. K. Baur, J. M. Hitchcock, E. J. Mueller, and R. G. Hulet. Detecting antiferromagnetism of atoms in an optical lattice via optical Bragg scattering. *Phys. Rev. A*, 81:013415, 2010. [86](#)
- [214] S. Fölling, F. Gerbier, A. Widera, O. Mandel, T. Gericke, and I. Bloch¹. Spatial quantum noise interferometry in expanding ultracold atom clouds. *Nature*, 434:481, 2005. [90](#)

- [215] U. Schneider, L. Hackermüller, J. P. Ronzheimer, S. Will, S. Braun, T. Best, I. Bloch, E. Demler, S. Mandt, D. Rasch, and A. Rosch. Fermionic transport and out-of-equilibrium dynamics in a homogeneous Hubbard model with ultracold atoms. *Nature Physics*, 8:213, 2012. [90](#)
- [216] M. Greiner, I. Bloch, O. Mandel, T. W. Hänsch, and T. Esslinger. Exploring phase coherence in a 2D lattice of Bose-Einstein condensates. *Phys. Rev. Lett.*, 87:160405, 2001. [91](#)
- [217] J. H. Denschlag, J. E. Simsarian, H. Häffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston, and W. D. Phillips. A Bose-Einstein condensate in an optical lattice. *J. Phys. B*, 35:3095, 2002. [91](#)
- [218] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto. Controlled exchange interaction between pairs of neutral atoms in an optical lattice. *Nature*, 448:452, 2007. [93](#)
- [219] R. P. Smith, R. L. D. Campbell, N. Tammuz, and Z. Hadzibabic. Effects of interactions on the critical temperature of a trapped Bose gas. *Phys. Rev. Lett.*, 106:250403, 2011. [94](#)
- [220] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels, and P. van der Straten. Thermodynamics of Bose-Einstein-condensed clouds using phase-contrast imaging. *Phys. Rev. A*, 81:053632, 2010. [94](#)
- [221] Y. Castin and R. Dum. Bose-Einstein condensates in time dependent traps. *Phys. Rev. Lett.*, 77:5315, 1996. [95](#), [286](#)
- [222] F. Ferlaino, P. Maddaloni, S. Burger, F. S. Cataliotti, C. Fort, M. Modugno, and M. Inguscio. Dynamics of a Bose-Einstein condensate at finite temperature in an atom-optical coherence filter. *Phys. Rev. A*, 66:011604, 2002. [95](#), [134](#)
- [223] B. D. Busch, C. Liu, Z. Dutton, C. H. Behroozi, and L. Vestergaard Hau. Observation of interaction dynamics in finite-temperature bose condensed atom clouds. *Europhys. Lett.*, 51:485, 2000. [97](#), [285](#)
- [224] J. Szczepkowski, R. Gartman, M. Witkowski, L. Tracewski, M. Zawada, and W. Gawlik. Analysis and calibration of absorptive images of Bose-Einstein condensate at nonzero temperatures. *Rev. Sci. Instrum.*, 80:053103, 2009. [97](#)
- [225] S. Nascimbène, N. Navon, K. J. Jiang, F. Chevy, and C. Salomon. Exploring the thermodynamics of a universal Fermi gas. *Nature*, 463:1057, 2010. [101](#), [108](#)
- [226] K. Van Houcke, F. Werner, E. Kozik, N. Prokof'ev, B. Svistunov, M. J. H. Ku, A. T. Sommer, L. W. Cheuk, A. Schirotzek, and M. W. Zwierlein. Feynman diagrams versus Fermi-gas Feynman emulator. *Nature Physics*, XXX:XXX, 2012. [101](#)
- [227] X. Zhang, C.-L. Hung, S.-K. Tung, and C. Chin. Observation of quantum criticality with ultracold atoms in optical lattices. *Science*, 335:1070, 2012. [101](#), [111](#)
- [228] S. L. Winoto, M. T. DePue, N. E. Bramall, and D. S. Weiss. Laser cooling at high density in deep far-detuned optical lattices. *Phys. Rev. A*, 59:R19, 1999. [102](#)
- [229] M. T. DePue, C. McCormick, S. L. Winoto, S. Oliver, and D. S. Weiss. Unity occupation of sites in a 3D optical lattice. *Phys. Rev. Lett.*, 82:2262, 1999. [102](#)
- [230] M. J. H. Ku, A. T. Sommer, L. W. Cheuk, and M. W. Zwierlein. Revealing the superfluid lambda transition in the universal thermodynamics of a unitary Fermi gas. *Science*, 335:563, 2012. [102](#)
- [231] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr. Single-spin addressing in an atomic Mott insulator. *Nature*, 471:319, 2011. [102](#)
- [232] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Ba nuls, L. Pollet, I. Bloch, and S. Kuhr. Observation of correlated particle-hole pairs and string order in low-dimensional Mott insulators. *Science*, 334:200, 2011. [102](#)

- [233] S. W. Hell. Far-field optical nanoscopy. *Science*, 316:1153, 2007. [102](#)
- [234] W. S. Bakr, P. M. Preiss, M. E. Tai, R. Ma, J. Simon, and M. Greiner. Orbital excitation blockade and algorithmic cooling in quantum gases. *Nature*, 480:500, 2011. [103](#)
- [235] X. Li, T. A. Corcovilos, Y. Wang, and D. S. Weiss. 3D projection sideband cooling. *Phys. Rev. Lett.*, 108:103001, 2012. [107](#)
- [236] C.-L. Hung, X. Zhang, N. Gemelke, and C. Chin. Slow mass transport and statistical evolution of an atomic gas across the superfluid Mott-insulator transition. *Phys. Rev. Lett.*, 104:160403, 2010. [110](#), [138](#), [151](#)
- [237] S. Sachdev. Quantum magnetism and criticality. *Nature Physics*, 4:173, 2008. [111](#)
- [238] S. Stringari. Collective excitations of a trapped Bose-condensed gas. *Phys. Rev. Lett.*, 77:2360, 1996. [111](#)
- [239] H. L. Stormer. Nobel lecture: The fractional quantum Hall effect. *Rev. Mod. Phys.*, 71:875, 1999. [113](#)
- [240] W. F. Vinen. Critical velocities in liquid He ii. In G. Careri, editor, *Liquid Helium, Proceedings of the International School of Physics “Enrico Fermi”, Course XXI*, page 336. Academic Press, 1963. [113](#)
- [241] M. Greiner, O. Mandel, T. W. Hänsch, and I. Bloch. Collapse and revival of the matter wave field of a Bose-Einstein condensate. *Nature*, 419:51, 2002. [113](#), [178](#)
- [242] D. Chen, M. White, C. Borries, and B. DeMarco. Quantum quench of an atomic Mott insulator. *Phys. Rev. Lett.*, 106:235304, 2011. [114](#), [144](#)
- [243] E. M. Lifshitz and L. P. Pitaevskii. *Statistical Physics Part 2*. Pergamon, 1980. [116](#)
- [244] M. Köhl. Thermometry of fermionic atoms in an optical lattice. *Phys. Rev. A*, 73:031601, 2006. [119](#)
- [245] A. M. Rey, G. Pupillo, and J. V. Porto. The role of interactions, tunneling, and harmonic confinement on the adiabatic loading of bosons in an optical lattice. *Phys. Rev. A*, 73:023608, 2006. [120](#)
- [246] R. P. Feynman. Application of quantum mechanics to liquid helium. In C. J. Gorter, editor, *Progress in Low Temperature Physics, Volume 1*, chapter 2, page 17. Elsevier, 1955. [130](#)
- [247] J. S. Lander and M. E. Fisher. Intrinsic critical velocity of a superfluid. *Phys. Rev. Lett.*, 19:560, 1967. [130](#)
- [248] P. W. Anderson. Considerations on the flow of superfluid Helium. *Rev. Mod. Phys.*, 38:298, 1966. [130](#)
- [249] J. Steinhauer, K. Schwab, Y. Mukharsky, J. C. Davis, and R. E. Packard. The determination of the energy barrier for phase slips in superfluid 4He. *J. Low Temp. Phys.*, 100:281, 1995. [130](#)
- [250] B. Wu and Q. Niu. Superfluidity of Bose-Einstein condensate in an optical lattice: Landau-Zener tunnelling and dynamical instability. *New. J. Phys.*, 5:104.1, 2003. [132](#)
- [251] A. O. Caldeira and A. J. Leggett. Quantum tunnelling in a dissipative system. *Ann. Phys.*, 149:374, 1983. [134](#)
- [252] W. F. Vinen. Mutual friction in a heat current in liquid helium ii. iii. Theory of the mutual friction. *Proc. R. Soc. Lond. A*, 242:493, 1957. [136](#)
- [253] S. Coleman. Fate of the false vacuum: Semiclassical theory. *Phys. Rev. D*, 15:2929, 1977. [136](#)
- [254] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. Vortex formation in a stirred Bose-Einstein condensate. *Phys. Rev. Lett.*, 84:806, 2000. [136](#)
- [255] D. V. Freilich, D. M. Bianchi, A. M. Kaufman, T. K. Langin, and D. S. Hall. Real-time dynamics of single vortex lines and vortex dipoles in a Bose-Einstein condensate. *Science*, 329:1192, 2010. [137](#)

- [256] J.-P. Brantut, J. Meineke, D. Stadler, S. Krinner, and T. Esslinger. Conduction of ultracold fermions through a mesoscopic channel. 2012, arXiv:1203.1927. [137](#)
- [257] S. S. Natu, K. R. A. Hazzard, and E. J. Mueller. Local versus global equilibration near the bosonic Mott-insulator–superfluid transition. *Phys. Rev. Lett.*, 106:125301, 2011. [142](#)
- [258] R. B. Diener, Q. Zhou, H. Zhai, and T.-L. Ho. Criterion for bosonic superfluidity in an optical lattice. *Phys. Rev. Lett.*, 98:180404, 2007. [145](#), [153](#)
- [259] F. Gerbier, S. Fölling, A. Widera, and I. Bloch. Visibility of a Bose-condensed gas released from an optical lattice at finite temperatures. 2007, arXiv/cond-mat/0701420. [153](#)
- [260] F. Gerbier, A. Widera, S. Fölling, O. Mandel, T. Gericke, and I. Bloch. Phase coherence of an atomic Mott insulator. *Phys. Rev. Lett.*, 95:050404, 2005. [153](#)
- [261] F. Gerbier, S. Trotzky, S. Fölling, U. Schnorrberger, J. D. Thompson, A. Widera, I. Bloch, L. Pollet, M. Troyer, B. Capogrosso-Sansone, N.V. Prokof'ev, and B. V. Svistinnov. Expansion of a quantum gas released from an optical lattice. *Phys. Rev. Lett.*, 101:155303, 2008. [155](#), [156](#), [157](#), [164](#), [291](#)
- [262] Y. B. Band, M. Trippenbach, J. P. Burke, and P. S. Julienne. Elastic scattering loss of atoms from colliding bose-einstein condensate wave packets. *Phys. Rev. Lett.*, 84:5462, 2000. [156](#)
- [263] L. Pollet, C. Kollath, K. Van Houcke, and M. Troyer. Temperature changes when adiabatically ramping up an optical lattice. *New. J. Phys.*, 10:065001, 2008. [164](#)
- [264] D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle. Reversible formation of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 81:2194, 1998. [166](#)
- [265] M. C. Garrett, A. Ratnapala, E. D. van Ooijen, C. J. Vale, K. Weegink, S. K. Schnelle, O. Vainio, N. R. Heckenberg, H. Rubinsztein-Dunlop, and M. J. Davis. Growth dynamics of a Bose-Einstein condensate in a dimple trap without cooling. *Phys. Rev. A*, 83:013630, 2011. [166](#)
- [266] P. W. H. Pinkse, A. Mosk, M. Weidemüller, M. W. Reynolds, T. W. Hijmans, and J. T. M. Walraven. Adiabatically changing the phase-space density of a trapped Bose gas. *Phys. Rev. Lett.*, 78:990, 1997. [166](#)
- [267] E. Timmermans and R. Côté. Superfluidity in sympathetic cooling with atomic Bose-Einstein condensates. *Phys. Rev. Lett.*, 80:3419, 1998. [181](#)
- [268] F. Zambelli, L. Pitaevskii, D. M. Stamper-Kurn, and S. Stringari. Dynamic structure factor and momentum distribution of a trapped Bose gas. *Phys. Rev. A*, 61:063608, 2000. [181](#)
- [269] G. M. Kavoulakis, C. J. Pethick, and H. Smith. Relaxation processes in clouds of trapped bosons above the Bose-Einstein condensate temperature. *Phys. Rev. Lett.*, 81:4039, 1998. [183](#)
- [270] B. DeMarco and D. S. Jin. Spin excitations in a Fermi gas of atoms. *Phys. Rev. Lett.*, 88:040405, 2002. [183](#)
- [271] T. Kinoshita, T. Wenger, and D. S. Weiss. A quantum Newton's cradle. *Nature*, 440:900, 2006. [183](#)
- [272] A. Sommer, M. Ku, G. Roati, and M. W. Zwierlein. Universal spin transport in a strongly interacting Fermi gas. *Nature*, 472:201, 2011. [183](#)
- [273] A. P. Chikkatur, A. Gorlitz, D. M. Stamper-Kurn, S. Inouye, S. Gupta, and W. Ketterle. Suppression and enhancement of impurity scattering in a Bose-Einstein condensate. *Phys. Rev. Lett.*, 85:483, 2000. [183](#)
- [274] D. S. Hall, M. R. Matthews, J. R. Ensher, C. E. Wieman, and E. A. Cornell. Dynamics of component separation in a binary mixture of Bose-Einstein condensates. *Phys. Rev. Lett.*, 81:1539, 1998. [183](#), [184](#)

- [275] C. Menotti and N. Trivedi. Spectral weight redistribution in strongly correlated bosons in optical lattices. *Phys. Rev. B*, 77:235120, 2008. [185](#)
- [276] G. Audi, , A. H. Wapstra, and C. Thibault. The AME2003 atomic mass evaluation (ii). tables, graphs and references. *Nuc. Phys. A*, 729:337, 2003. [189](#)
- [277] J. K. Böhlke, J. R. de Laeter, P. De Bièvre, H. Hidaka, H. S. Peiser, K. J. R. Rosman, and P. D. P. Taylor. Isotopic compositions of the elements, 2001. *J. Phys. Chem. Ref. Data*, 34:57, 2005. [189](#)
- [278] E. Arimondo, M. Inguscio, , and P. Violino. Experimental determinations of the hyperfine structure in the alkali atoms. *Rev. Mod. Phys.*, 49:31, 1977. [189](#), [190](#), [192](#), [194](#)
- [279] A. Gallagher and E.L. Lewis. Determination of the vapor pressure of rubidium by optical absorption. *J. Opt. Soc. Am.*, 63:864, 1973. [189](#)
- [280] G. Audi, O. Bersillon, J. Blachot, and A. H. Wapstra. The NUBASE evaluation of nuclear and decay properties. *Nuc. Phys. A*, 729:3, 2003. [189](#)
- [281] H. Wang, J. Li, X. T. Wang, C. J. Williams, P. L. Gould, and W. C. Stwalley. Precise determination of the dipole matrix element and radiative lifetime of the ^{39}K 4p state by photoassociative spectroscopy. *Phys. Rev. A*, 55:R1569, 1997. [190](#)
- [282] R. W. Berends, W. Kedzierski, J. B. Atkinson, and L. Krause. The radiative lifetimes of the potassium 5P,6P and 7P states. *Spectrochim. Acta.*, 43B:1069, 1988. [190](#)
- [283] A. Mills, J. A. Behr, L. A. Courneyea, and M. R. Pearson. Lifetime of the potassium 5P $_{1/2}$ state. *Phys. Rev. A*, 72:024501, 2005. [190](#), [255](#)
- [284] R. F. Gutterres, C. Amiot, A. Fioretti, C. Gabbanini, M. Mazzoni, and O. Dulieu. Determination of the ^{87}Rb 5p state dipole matrix element and radiative lifetime from the photoassociation spectroscopy of the Rb_2 $0_g^-(P_{3/2})$ long-range state. *Phys. Rev. A*, 66:024502, 2002. [191](#)
- [285] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005. [192](#), [193](#)
- [286] G. Breit and I. I. Rabi. Measurement of nuclear spin. *Phys. Rev.*, 38:2082, 1931. [193](#)
- [287] E. G. M. van Kempen, S. J. J. M. F. Kokkelmans, D. J. Heinzen, and B. J. Verhaar. Interisotope determination of ultracold rubidium interactions from three high-precision experiments. *Phys. Rev. Lett.*, 88:093201, 2002. [195](#), [196](#)
- [288] Stephan Falke, Horst Knöckel, Jan Friebe, Matthias Riedmann, Eberhard Tiemann, and Christian Lisdat. Potassium ground-state scattering parameters and Born-Oppenheimer potentials from molecular spectroscopy. *Phys. Rev. A*, 78:012503, 2008. [195](#)
- [289] C. A. Regal, M. Greiner, and D. S. Jin. Observation of resonance condensation of fermionic atom pairs. *Phys. Rev. Lett.*, 92:040403, 2004. [195](#)
- [290] C. A. Regal and D. S. Jin. Measurement of positive and negative scattering lengths in a Fermi gas of atoms. *Phys. Rev. Lett.*, 90:230404, 2003. [195](#)
- [291] A. Simoni, M. Zaccanti, C. D’Errico, M. Fattori, G. Roati, M. Inguscio, and G. Modugno. Near-threshold model for ultracold KRb dimers from interisotope Feshbach spectroscopy. *Phys. Rev. A*, 77:052705, 2008. [195](#)
- [292] S. Aubin, S. Myrskog, M. H. T. Extavour, L. J. Leblanc, D. McKay, A. Stummer, and J. H. Thywissen. Rapid sympathetic cooling to Fermi degeneracy on a chip. *Nature Physics*, 2:384, 2006. [196](#)
- [293] D. J. Heinzen. Ultracold atomic interactions and Bose-Einstein condensation. *Int. J. Mod. Phys. B*, 11:3297, 1997. [196](#)

- [294] P. Meystre. Light-matter interaction. In G. W. F. Drake, editor, *Springer Handbook of Atomic, Molecular and Optical Physics*, chapter 21, page 307. Springer, 2006. [198](#), [209](#)
- [295] B. H. Bransden and C. J. Joachain. *Physics of atoms and molecules*. Longman Group, 1983. [198](#)
- [296] L. I. Schiff. *Quantum Mechanics*. McGraw-Hill, 1949. [199](#)
- [297] R. Loudon. *The Quantum Theory of Light*. Oxford, 1983. [200](#)
- [298] M. O. Scully and M. S. Zubairy. *Quantum Optics*. Cambridge University Press, 1997. [201](#), [213](#)
- [299] C. Zener. Non-adiabatic crossing of energy levels. *Proc. R. Soc. Lond. A*, 137:696, 1932. [205](#)
- [300] J. D. Jackson. *Classical Electrodynamics*. Wiley, 1999. [215](#), [216](#)
- [301] B. R. Mollow. Power spectrum of light scattered by two-level systems. *Phys. Rev.*, 188:1969, 1969. [217](#)
- [302] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6:2023, 1989. [221](#)
- [303] F. S. Cataliotti, E. A. Cornell, C. Fort, M. Inguscio, F. Marin, M. Prevedelli, L. Ricci, and G. M. Tino. Magneto-optical trapping of fermionic potassium atoms. *Phys. Rev. A*, 57:1136–1138, 1998. [250](#)
- [304] T. P. Dinneen, C. D. Wallace, K.-Y. N. Tan, and P. L. Gould. Use of trapped atoms to measure absolute photoionization cross sections. *Op. Lett.*, 17:1706, 1992. [255](#)
- [305] M. Anderlini, E. Courtade, D. Ciampin, J. H. Müller, O. Morsch, and E. Arimondo. Two-photon ionization of cold rubidium atoms. *J. Opt. Soc. Am. B*, 21:480, 2004. [255](#)
- [306] E. Courtade, M. Anderlini, D. Ciampini, J. H. Müller, O. Morsch, E. Arimondo, M. Aymar, and E. J. Robinson. Two-photon ionization of cold rubidium atoms with a near resonant intermediate state. *J. Phys. B*, 37:967, 2004. [255](#)
- [307] R. V. Ambartsumian, N. P. Furzikov, V. S. Letokhov, and A. A. Puretsky. Measuring photoionization cross-sections of excited atomic states. *Appl. Phys.*, 9:335, 1976. [255](#)
- [308] I. D. Petrov, V. L. Sukhorukov, E. Leber, and H. Hotop. Near threshold photoionization of excited alkali atoms Ak(np) (Ak = Na, K, Rb, Cs; n = 3–6). *Euro. Phys. J. D*, 10:53, 2000. [255](#)
- [309] N. Amina, S. Mahmooda, S.U. Haqa, M.A. Kalyara, M. Rafiq, and M.A. Baig. Measurements of photoionization cross sections from the 4p, 5d and 7s excited states of potassium. *J. Quant. Spec. Rad. Trans.*, 109:863, 2008. [255](#)
- [310] U. Gustafsson, J. Alnis, and S Svanberg. Atomic spectroscopy with violet laser diodes. *Am. J. Phys.*, 68:660, 2000. [255](#)
- [311] U. Gustafsson, G. Somesfalean, J. Alnis, and S. Svanberg. Frequency-modulation spectroscopy with blue diode lasers. *Applied Optics*, 39:3774, 2000. [255](#)
- [312] S. Uetake, K. Hayasaka, and M. Watanabe. Saturation spectroscopy of potassium for frequency stabilization of violet diode lasers. *Jpn. J. Appl. Phys.*, 42:L332, 2003. [255](#)
- [313] A. Behrle, M. Koschorreck, and M. Köhl. Isotope shift and hyperfine splitting of the 4S→5P transition in potassium. *Phys. Rev. A*, 83:052507, 2011. [255](#), [257](#)
- [314] L. J. S. Halloran, S. Fostner, E. Paradis, and J. A. Behr. Specific mass shift of potassium 5p_{1/2} state. *Op. Comm.*, 282:554, 2009. [255](#)
- [315] Y. Iwata, Y. Inoue, and M. Minowa. Development of a compact resonance ionization mass spectrometer for trace element analysis of potassium. 2010, arXiv:1010.5902. [257](#)

- [316] Robert S. Williamson III. *Magneto-Optical Trapping of Potassium Isotopes*. PhD thesis, University of Wisconsin - Madison, 1997. [257](#)
- [317] C. Y. Park and T. H. Yoon. Frequency stabilization of injection-locked violet laser diode with Doppler-free absorption signal of ytterbium. *Jpn. J. Appl. Phys.*, 42:L754–L756, 2003. [261](#)
- [318] K. Komori, Y. Takasu, M. Kumakura, Y. Takahashi, and T. Yabuzaki. Injection-locking of blue laser diodes and its application to the laser cooling of neutral ytterbium atoms. *Jpn. J. Appl. Phys.*, 42:5059, 2003. [261](#)
- [319] Josh Zirbel. Private Communication, 2010. [263](#), [264](#)
- [320] J. Weiner, V. S. Bagnato, S. Zilio, and P. S. Julienne. Experiments and theory in cold and ultracold collisions. *Rev. Mod. Phys.*, 71:1–85, 1999. [267](#)
- [321] G. Modugno, G. Roati, M. Inguscio, M. S. Santos, G. D. Telles, L. G. Marcass, and V. S. Bagnato. Comparative investigation of ^{39}K and ^{40}K trap loss rates: alternative loss channel at low light intensities. *Euro. Phys. J. D*, 23:409, 2003. [268](#)
- [322] T. Chanelière, J.-L. Meunier, R. Kaiser, C. Miniatura, and D. Wilkowski. Extra-heating mechanism in Doppler cooling experiments. *J. Opt. Soc. Am. B*, 22:1819, 2005. [270](#)
- [323] Markus Greiner. *Ultracold quantum gases in three-dimensional optical lattice potentials*. PhD thesis, Ludwig-Maximilians-Universität München, 2003. [291](#)