

CHARACTERIZATION OF A MULTIMODE CQED SYSTEM

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF APPLIED PHYSICS

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Alexander Themis Papageorge

May 2016

Abstract

Cavity quantum electrodynamics (CQED) is a proven testbed for studies of the fundamental light-matter interaction and in recent years has become increasingly relevant in quantum simulation. By engineering the electromagnetic properties of the cavity it is possible to tailor unique interatomic interactions resulting in the possibility of observing novel phase-transitions and self-organization phenomena. These transitions distinguish themselves from existing ultra-cold atom optical lattice experiments in that the order exhibited by the atoms is not externally-imposed but emergent in the atom-cavity interaction. Increasing the variety of cavity modes available for crystallization towards a complete basis leads to the possibility of observing quantum soft-condensed matter phenomena like phonons, dislocations, and frustration in this system.

This thesis introduces a novel cold-atoms apparatus designed and constructed with the aim of observing and characterizing emergent quantum phenomena. Details of the experimental setup and procedure are presented along with recent experimental and theoretical results pertaining to a novel self-organization phenomenon. Additionally, the technical aspects of an approach to couple light selectively to arbitrary cavity modes using digital holography will be presented. This tool will be useful for seeding the cavity towards, for example, the promotion of self-organization in a particular

mode.

Acknowledgements

The work presented in this thesis is not the result of a solitary effort. I have been incredibly fortunate over the course of my life to have had unparalleled support from my instructors, peers and family. I give Benjamin Lev my sincerest thanks for concocting this wonderful project and bringing me on to help build it. It is in his lab that I have grown as a scientist and as a person, and where I have cultivated a deep appreciation for the complexity of our experiments and for complexity itself. I joined Alicia Kóliar on the CQED experiment in its nascence and together we built an unprecedented machine from an empty room. Alicia's incredible fastidiousness was invaluable as we learned enough plumbing, optics, vacuum, electronics, and programming to build a unique apparatus. Her intuition for the physics we sought simplified many of our technical problems in marvelous ways. Working with her was a privilege; the success of my PhD has as much to do with her efforts and capabilities as my own. It has been a wonderful journey that has culminated in novel science; I am fortunate to have been instrumental in the birth and youth of a brilliant cold atoms experiment.

The best of my education came at the hands of the members of the Lev Lab. Kris Baumann, Jack DiSciaccia, Matt Naides, Will Turner, Mingwu Lu, Nathaniel Burdick, and Yijun Tang have all been wonderful compatriots. Some of my best days in lab were spent working through technical problems with these individuals. In each

case I learned something about problem solving or electro-optics or ground noise that I didn't know before. This is of course only a small subset of the intellectual resources that have been at my disposal. I have had the pleasure of interacting and learning from many other talented scientists and engineers, too many to name.

That I was able to spend any time working at all is a testament to the administrative staff that supports our endeavors. Susan Ayres is spectacularly helpful in all aspects of Lev Lab operations. Paula Perron and Claire Nicholas have looked out for me my entire time at Stanford, and have had my best interests at heart when I was inattentive. Days spent machining under the guidance of Mehmet Solyali and Karlheinz Merkle were truly pleasurable; the time they spent teaching me and the trust they had that I wouldn't maim myself will not be forgotten.

Finally, I give my heartfelt love and sincerest gratitude to my family. My mother, father, and my sister Marianna have been pillars in my life. Their support, love, and encouragement give me reason every day to be. αἰὲν ἀριστεύειν

Contents

1	Introduction	1
2	The Experiment	6
2.1	Rubidium	6
2.2	BEC Production	8
2.3	Cooling Lasers	16
2.4	Laser Locks	17
2.4.1	Transport ODT	23
2.4.2	Cross Optical Dipole Trap	29
2.5	Cavity	34
3	Digital Micromirror Device	44
3.1	Overview	44
3.2	Theory of the Paraxial Cavity	52
3.3	DMD Calibration and Usage	58
3.4	Experimental outline	60
3.5	Designing a DMD mask for a particular HG mode	61
3.6	Mode-coupling procedure	64

3.7	Single-mode spectroscopy	67
3.8	Multimode Spectroscopy	68
4	Cavity Quantum Electrodynamics	71
4.1	Overview	71
4.2	Self-Organization	76
4.3	Gross-Pitaevskii Equation	84
4.4	Numerical Results	85
5	Conclusion	96
A	DMD Hologram Code	97
A.1	CountingModes.py	97
A.2	ditherimage.py	98
A.3	DMDFunctions.py	100
A.4	Patches.py	101
A.5	DMD_Calibration.py	102
A.6	PhaseMaskPost.py	105
A.7	dmd_calibrate.py	107
A.8	ModeGeneration_example.py	109
B	GPE Simulation Code	112
B.1	GPEgs.py	112
C	Cavity Lock Filters	116
D	Cavity Spacer	119

List of Tables

List of Figures

2.1	^{87}Rb D ₂ Level Diagram	7
2.2	Vacuum Chamber Rendering	10
2.3	BEC TOF	12
2.4	Science Chamber Rendering	14
2.5	Diagram of the CQED Cooling Laser System	18
2.6	Lead Lag Filter	20
2.7	All Pass Filter	20
2.8	PDH Error Signal	36
2.9	Attocube PZT Swept Sine Response	38
2.10	Circuit Diagram of the 120 Hz Resonant Filter	39
2.11	Resonant Filter Amplitude Response	40
2.12	Resonant Filter Phase Response	41
2.13	Feedback PZT Swept Sine Response	42
2.14	Circuit Diagram of the 3.5 kHz Resonant Filter	43
3.1	Schematic of the DMD Pixels	46
3.2	Schematic of a Blazed Grating	47
3.3	DMD Delta Function	50

3.4	DMD Delta Function: Aberration Corrected	51
3.5	Raw Phase Mask	53
3.6	Fully Processed Phase Mask	54
3.7	Diagram of the DMD Experiment	56
3.8	Schematic Explaining the Calibration Procedure	57
3.9	Determination of Non-Ideal Cavity Modes	62
3.10	Single Mode Spectroscopy	63
3.11	Spectroscopy of the Near-Degenerate Even Modes	66
3.12	Multimode Spectroscopy	67
3.13	Output of a Degenerate Cavity	69
4.1	Rabi Splitting	75
4.2	Below Threshold Supermode Polariton	82
4.3	Flow Chart for Split-Step Fourier Transform Technique	87
4.4	GPE Ground State	89
4.5	GPE Ground State: Atoms Pumped in the HG_{00}	90
4.6	Coherent cavity-pump potential $\theta(x, y)$ for the $HG_{0,0}$	91
4.7	GPE Ground State: Atoms Pumped in the $HG_{0,0}$, Fourier Transform	92
4.8	GPE Time Evolution: Atoms pumped in the $HG_{0,1}$	93
4.9	Coherent cavity-pump potential $\theta(x, y)$ for the $HG_{0,1}$	94
4.10	GPE Time Evolution: Atoms pumped in the $HG_{0,1}$, Fourier Transform	95
C.1	Band Bump Filter	117
C.2	Resonant Filter	118
D.1	Machine Drawing of the Brass Cavity Spacer	120

Chapter 1

Introduction

Quantum simulation is a central aim of modern Atomic Molecular and Optical (AMO) physics [1] and became a widespread goal of the field quickly following the first observation of the Bose-Einstein Condensate (BEC) [2, 3]. When weakly interacting, BECs made from atomic gases approximate a pure quantum state (in both motional and internal degrees of freedom) to a fantastic degree, and the ability to control interactions allows for studies of many-body physics [4]. Careful selection of the quantum gas also allows for the generation of ultracold clouds of fermionic atoms and the observation of degenerate fermions [5]. These experimental techniques open the field for the study of matter at the most general and fundamental level.

For twenty years cold-atoms experiments have provided a crucial testbed for these ideas, allowing for the creation and manipulation of pure quantum states in a pristine environment free from dissipation and decoherence. The ability to generate pure quantum states in ultracold atomic clouds whose behavior is governed by a precisely tailored Hamiltonian has led to the observation of quantum phenomena long theorized to exist but never directly observed in condensed matter systems [6, 7]. The dream

of the quantum simulation enterprise is to predict and observe novel quantum phenomena through improved control of the quantum degrees of freedom along with the continued verification of well-understood physics. This endeavor is on solid theoretical ground, supported by theorems concerning the equivalence of quantum systems [8, 9]. The notion that a quantum system with a sufficiently large state-space should be isomorphic to any other suggests that a generalizable quantum simulator would provide insight into existing and proposed quantum phenomena as varied as magnetism [10, 11] or neutron stars [12]. Quantum simulation has the potential to be impactful in the study of beyond-mean-field physics, where systems that behave in a way not well characterized by existing theoretical tools can be studied in a controllable manner [13].

As an experimental tool optical lattices have provided access to the types of tight-binding Bose- and Fermi- Hubbard Hamiltonians thought to describe electrons' behavior in crystals [14]. Interfering laser light results in periodic optical potentials that trap the atoms, but more generally, modify the plane-wave states of a free quantum gas into the Bloch waves of an electron under the confinement by the ions of a crystal. A multitude of crystal geometries in 1, 2, and 3 dimensions have been studied with this approach, each towards the observation of different physics [15–18]. What is lacking in these systems, however, are many qualities possessed by ‘real’ crystals. Phonons, dislocations, and frustration are examples of emergent properties and defects that exist as a consequence of the crystal's order being spontaneously generated and thus subject to thermal excitation. When the order is imposed on the atoms by an external potential, these qualities are unavailable for study. The degree to which this is limiting depends deeply on the nature of the problem being studied, but as in the case of BCS superconductivity, the limitation is obvious. Despite this, the BCS state has

been observed in quantum gases where strong contact interactions act as an inexact replacement for phonon-mediated interactions [19]. Nonetheless, self-organization is such an ubiquitous feature of condensed matter physics and encompasses such rich phenomenology that simulating it may be necessary to gain access to the complexity of real materials. Additionally there is the possibility of studying novel Hamiltonians; systems that have no analogue in nature but nevertheless provide insight to statistical and quantum physics.

Cavity QED is a particularly attractive setting to explore many-body physics using quantum gases of neutral atoms [20]. In these systems, a laser pumps energy into cavity modes either directly via a semitransparent mirror or indirectly via coherent (Rayleigh) scattering off intracavity atoms. This light leaks out of the cavity by the same semitransparent mirrors and may serve as a record of the collective dynamics driven by light-matter coupling. Atoms that successively scatter the same light are effectively interacting and under the right conditions can self-organize. Interaction timescales can be orders of magnitude faster than the superexchange or dipolar exchange interactions exploited in traditional optical lattice-based quantum simulators [21, 22]. A non-equilibrium steady state may be reached before spontaneous emission heats the gas. Despite the driven, dissipative character of this intrinsically non-equilibrium system, statistical mechanical treatments may be employed to analyze phase transitions [23–25].

Experiments of this variety have in the recent past made use of a BEC positioned near the center of, and thus coupled to, a Fabry-Pérot cavity. Pumping the atoms results in the generation of a self-consistent optical lattice made from the interference

of the cavity mode and the pump light. The atomic superfluid occupies this lattice in the process breaking a Z_2 symmetry; the atoms occupy one of two equivalent sub-lattices of the total interference potential. An example of spontaneous symmetry breaking, this experimental system is a stunning example of self-organization in a cold-atoms system, one whose threshold, excitation spectrum, entanglement, and thermodynamic properties can be studied in a clean, isolated environment. Extending the experiment to include a cavity with many, orthogonal, degenerate modes has the potential to allow for studies of phase transitions that exhibit continuous symmetry breaking; a necessary condition for observing the aforementioned emergent condensed matter phenomena. To this end, a novel multimode cavity quantum electrodynamics (CQED) apparatus was constructed with a degenerate near-confocal Fabry-Pérot cavity as its centerpiece. BECs positioned controllably with respect to the cavity modes are pumped and the resulting self-organization is studied in the time-of-flight images of the atoms and the light emitted by the cavity.

In order to study and probe the multimode cavity it is necessary to have on hand a method of mode-spectrum characterization based on driving the degenerate cavity with carefully tailored mode functions. Specifically, a technique was developed making use of a digital micromirror device (DMD), functioning as a binary spatial light modulator (SLM), which can be used to generate holographic light fields that selectively couple to each cavity mode. This capability will be important, e.g., when engineering atom-atom interactions of unusual form via shaping the pumped light fields or when measuring the multimode cQED system's dynamic response [26] to tailored drive fields. These characterization tasks can be accomplished in a manner

that is rapid and re-configurable using the DMD, enabling the versatile *in-situ* manipulation of one mode versus another. Furthermore, the DMD can be used to study the particulars of the bare cavity, i.e., how the mode structure changes as a result of aberrations and mixing when the cavity is tuned to a degenerate configuration. Iterative procedures are employed to better couple to individual cavity modes; this paves the way for the use of optimization algorithms to automatically determine the cavity mode structure.

Chapter 2

The Experiment

2.1 Rubidium

Rubidium is an alkali atom with atomic number 37. It exists primarily as two stable isotopes, ^{85}Rb and ^{87}Rb (72.2% and 27.8%), both of which have use in cold atoms experiments. The cooling line of Rubidium is 780.24 nm, an incredibly useful wavelength technologically as it coincides with the wavelength used for reading Compact Discs (CD), prompting a technological revolution in the development of diode lasers at this wavelength. With a ground state electronic configuration $1s^2 2s^2 2p^6 3s^2 3p^6 3d^{10} 4s^2 4p^6 5s^1$ Rubidium has a single $S = \frac{1}{2}$ valence electron, and as a result all relevant transitions are based on the single electron transition $5s \rightarrow 6p$, where the orbital angular momentum increases from $L = 0$ to $L = 1$. Spin-orbit coupling results in a $J = 0 \otimes \frac{1}{2} = \frac{1}{2}$ ground state and splits the excited state into $J = 1 \otimes \frac{1}{2} = \frac{1}{2} \oplus \frac{3}{2}$. The $5S_{1/2} \rightarrow 5P_{1/2}$ transition is known as the D_1 transition with a wavelength of 795 nm. The $5S_{1/2} \rightarrow 5P_{3/2}$ transition is known as the D_2 transition

with a wavelength of 780.23 nm. These transitions are further split by spin-spin interactions into manifolds defined by quantum number F . For ^{87}Rb whose nuclear spin is $I = \frac{3}{2}$, the $5S_{1/2}$ and $5P_{1/2}$ manifolds have $F = \frac{1}{2} \otimes \frac{3}{2} = 1 \oplus 2$, and the $5P_{3/2}$ manifold has $F = \frac{3}{2} \otimes \frac{3}{2} = 0 \oplus 1 \oplus 2 \oplus 3$. Being much closer to the nuclear spin the ground state hyperfine splitting is several orders of magnitude (~ 1.5) larger than the excited state hyperfine structure; the splitting between the ground states $|F = 1\rangle$ and $|F = 2\rangle$ is 6.8 GHz, while the splitting between the excited states $|F'\rangle$ is 100's of MHz. The level structure for ^{85}Rb is qualitatively the same, the only difference being that every ground and excited state $|F\rangle$ and $|F'\rangle$ has its magnetic quantum number one larger than its analogous state in ^{87}Rb , a consequence of its nuclear spin being $I = 5/2$. It turns out that this structure is conserved amongst alkali atoms, the existence of a single valence electron giving rise to a hydrogenic structure and D₁ and D₂ lines, and whose hyperfine quantum numbers are only differentiated by the atom's nuclear spin. Despite spectroscopic similarities, the two Rubidium isotopes have vastly different scattering properties, with only ^{87}Rb able to natively Bose condense (BEC) [27].




Figure 2.1: ^{87}Rb D₂ level diagram reproduced from [28].

2.2 BEC Production

In order to allow for precise and accurate control over the atomic degrees of freedom necessary to observe quantum phenomena, all cold atoms experiments take place in a pristine vacuum environment maintained at pressures as low as 10^{-12} Torr. The vacuum system used in this experiment was constructed using off-the-shelf components assembled in a unique fashion to suit the requirements and constraints of the experiment and the space in which it was built. Figure 2.2 shows a rendering of the two-part vacuum chamber. The chamber is roughly divided into three sections: the oven, the production chamber, and the science chamber. The three chambers were designed and built over time to accommodate a developing understanding of the experiment's design. Pressure is maintained by a collection of pumps; three ion pumps and two getter pumps act against leaks and out-gassing of material within the chamber. Rubidium vapor from the oven loads a magneto-optical-trap (MOT) in the center of the production chamber.

After several cooling steps the atoms are transferred to a single-beam optical dipole trap (ODT) which can be translated to the science chamber where the atoms are trapped between the mirrors of a Fabry-Pérot cavity. Here, the atom-cavity system can be addressed by light. Different self-organization phenomena can be observed by pumping the atoms, cavity, or both. The cavity is a near-confocal ($R \approx L$) optical resonator made from two medium-high-finesse ($\mathcal{F} \approx 5.7 \times 10^4$) mirrors whose radius of curvature R is 1 cm. The length L of the cavity can be modified by ± 1 mm by an Attocube nanopositioning stage built into the cavity mount. The modes of the cavity are approximately Hermite-Gauss functions with a waist of $\sqrt{\frac{\lambda}{2\pi}} (L(2R - L))^{\frac{1}{4}} \approx 35 \mu\text{m}$. The atoms can be localized within the cavity mode by much better than a waist

as the cloud's Thomas-Fermi (TF) radius is $\sim 7 \mu\text{m}$ and the passive stability of the optics prevents drift greater than this figure. Using an error signal generated by reflecting 1560 nm light off the cavity, the cavity length is stabilized by custom feedback electronics resulting in differential-mode variance of $\sim 2 \text{ pm}$. This corresponds to variation of 75 kHz in the mode frequencies; less than one-third the 260 kHz linewidth (FWHM 2κ) at 780 nm, and significantly less than Δ_c , the cavity-pump detuning used in most of the experiments.

With the atoms localized in the cavity mode they can be pumped transversely by off-resonant light. Scattered into the cavity, the light participates in self-organization phenomena that result in superradiant scattering by a cloud that has broken translational symmetry. In a single or few mode system the broken symmetry is discrete; allowing the cloud to crystallize into an increasing number of modes will result in a state that has broken a continuous phase, leading naturally to the observation of phonons, defects, partially broken symmetries, and other soft condensed matter phenomena.

The creation of ultracold gases in the production chamber follows the work of [29], with an implementation mirroring in most instances [30]. A Mini-Conflat bellows holds a glass ampoule full of ^{87}Rb . After the chamber is pumped down this bellows is flexed, breaking the ampoule and exposing the Rubidium to the vacuum. The inert gas (usually N_2) released by the ampoule is taken up by the ion pumps in the chamber. To generate a partial pressure of Rb the bellows and surrounding vacuum components are heated to over 100°C . A differential pumping tube between the oven and production chamber doubles as a velocity selector for the Rubidium gas in the oven resulting in an atomic beam directed through the Zeeman slower. The atoms




Figure 2.2: Computer rendering of experimental apparatus, excluding laser optics. Inset: photograph of multimode cavity within science chamber.

are decelerated by the Zeeman slower [31] to speeds of several cm/s, loading ^{87}Rb into a magneto-optical trap (MOT) at the center of the production chamber's trapping region. The flow of Rubidium is ceased by closing an in-vacuum shutter and the atoms are subsequently cooled to sub-doppler temperatures by polarization gradient cooling and then optically pumped to the $|1, -1\rangle$ state. Then, the MOT's magnetic quadrupole field is increased to obtain a gradient of 195 G/cm. This sequence of procedures results in approximately one-billion atoms in the $|1, -1\rangle$ state magnetically trapped at the center of the production chamber. Linearly ramping the RF frequency in coils external to the chamber from 25.5 MHz to 5.5 MHz over 5.7 s produces an evaporatively cooled cloud of 8×10^7 atoms at 40 μK .

An off-resonant ODT is tightly focused below the center of the quadrupole trap, and the atoms are directly loaded into this ODT during the course of the evaporation. We use this ODT as an optical tweezer for transporting the atoms into the cavity science chamber. It employs 9.8 W of 1064-nm laser light focused to waists of 50 μm and 73 μm in the horizontal and vertical directions, respectively. Loading into the ODT from the magnetic trap is performed by significantly reducing the quadrupole gradient over 1 s to a value just supporting the atoms against gravity. In this configuration, the quadrupole trap provides confinement in the weakly trapped, axial direction of the ODT beam. By lowering the optical power, the atoms can undergo efficient optical evaporation to high phase-space density. We have obtained a BEC of 1×10^6 atoms at 200 nK in the production chamber by further reducing the ODT power.

However, creating BECs before ODT transport to the cavity science chamber is counter-productive, as three-body collisions heat the gas during transport [32].




Figure 2.3: Typical single-shot absorption imaging density profiles of ^{87}Rb atomic clouds after 18 ms of free expansion. The top row shows the density distribution, with the colorscale equal for the two pictures. The lower row shows the integrated density along a direction perpendicular to the cavity axis. The black dots represent the data, the green curve shows the fitted bimodal profile and the red curve shows the fitted thermal portion. (a) and (b): $1.8(2) \times 10^5$ atoms at the onset of BEC, $T \sim 320$ nK. (c) and (d): Nearly pure BEC of $4.5(5) \times 10^4$ atoms at $T \approx 150$ nK. BEC fraction is $\sim 90\%$.

Rather, we evaporate to $8\text{ }\mu\text{K}$ above $T_c = 1.6\text{ }\mu\text{K}$ and transport 6×10^6 atoms to the cavity science chamber by translating the last focusing lens of the ODT. This lens is mounted on an air bearing translation stage. The stage provides a translation of the lens and of the ODT focus over a distance of 35 cm in 4.3 s. The position of this beam is adjusted to maximize loading of atoms into the ODT while maintaining a transport trajectory whose final location is the center of the multimode cavity. The motion of the stage is chosen to be a smooth function defined by three sections of constant acceleration, with the central section having zero acceleration. The parameters of this curve were chosen so as to maximize transfer of atoms into the science chamber. The overall transfer efficiency 30% is currently limited by magnetic field gradients, background gas collisions, and mechanical excitation.

The two sections of the apparatus, the ultracold atom production chamber and the cavity science chamber, are separated by a gate valve, allowing us to break vacuum in one without disturbing the other. This feature will prove useful for modifying the cavity, if for instance we want satisfy the concentric criterion $L = 2R$ or exchange the mirrors with ones of differing finesse or R . Figure 2.4 shows a cutaway of the science chamber with an emphasis on the vibration-isolation stack. The cavity mount is constructed of steel and macor and weighs altogether approximately half a kilogram. The mount is resting on viton squares on a second steel piece weighing 3.5 kg. This steel piece is resting on Viton squares on a third and final steel piece weighing 6 kg. This three-tiered stack acts as a 6-pole lowpass filter with a 3 dB point at $\sim 80\text{ Hz}$, strongly attenuating external vibrations that may otherwise shake the cavity.

Once the atomic cloud is positioned between the mirrors in the science chamber, a second 1064-nm laser beam crosses it to confine the atoms. This beam is circular




Figure 2.4: Computer rendering of a section-view of the science chamber showing the cavity mount, vibration isolation stages, and in-vacuum coupling lens. Note, colors do not reflect actual material color. Inset: photograph of cavity holder outside of vacuum, including single-crystal PZT, mirrors, and slip-stick translation stage.

with a Gaussian waist ($1/e$ radius) of $48\text{ }\mu\text{m}$. Forced evaporation in this crossed ODT produces a nearly pure BEC of $4.5(4) \times 10^4$ atoms at density $4.7(5) \times 10^{14}\text{ cm}^{-3}$, as shown in figure 2.3. The trap frequencies are $145(5)\text{ Hz}$, $115(1)\text{ Hz}$, and $167(2)\text{ Hz}$ in $\hat{x}$, $\hat{y}$, and $\hat{z}$, respectively. We expect to achieve pure condensates with 10^5 atoms after performing additional evaporation optimization. The total time to achieve a BEC within the cavity, starting from MOT loading, is 16 s.

For ease of coupling to the cavity mode at 780-nm, we can also load a thermal gas into an intracavity ODT formed with the 1560-nm locking light in the $\text{TEM}_{0,0}$ mode. To do so, we increase the power of the locking laser until the steady state circulating power provides a trap depth of $5.2(1.3)\text{ }\mu\text{K}$. The transport ODT brings atoms into the science chamber and holds the cloud within the cavity mode. Addition of a second ODT beam crossing the first increases the mode matching between the 1064-nm trap and the 1560-nm cavity mode. Optical evaporation with the crossed 1064-nm beams loads atoms in the intracavity ODT, resulting in $2.5(2) \times 10^5$ atoms at $1.8(5)\text{ }\mu\text{K}$ confined in the 1560-nm standing-wave $\text{TEM}_{0,0}$ cavity mode. The temperature and atom number uncertainties, as well as the relatively high final temperature, are dominated by the lack of intensity stabilization and dynamical intensity adjustment of the intracavity ODT power, all of which will be added in the near future.

The atoms in the intracavity ODT are thermal and the in-trap density distribution obeys Boltzmann statistics with $\rho(\mathbf{x}) = \exp(-V(\mathbf{x})/T)$, where $V(\mathbf{x})$ is the standing-wave optical potential of the intracavity ODT. The resulting atomic density distribution is Gaussian, assuming a harmonic approximation to the potential. Such an approximation is valid as long as the atoms are sufficiently cold so as to localize close to the potential minima. We choose to operate at larger trap depth, though

lower atom number, in the experiments discussed below. Since the atoms are loaded into the intracavity ODT from a separate ODT with waists larger than $\lambda/2$, the atoms will occupy multiple wells of the intracavity ODT. The waist of the 1560-nm $\text{TEM}_{0,0}$ cavity mode is $\sqrt{2}$ larger than the 780-nm $\text{TEM}_{0,0}$ mode waist: the intracavity ODT waist is $35\sqrt{2} = 50 \mu\text{m}$. For the trap parameters used below— $V_{\text{FORT}} = 22(6) \mu\text{K}$ and $T = 5(1) \mu\text{K}$ —the atomic density distribution has a transverse waist of $\sim 17 \mu\text{m}$, which is smaller than the $35 \mu\text{m}$ waist of the 780-nm mode [33].

2.3 Cooling Lasers

When cooling Rb or other Alkali atoms, two lasers are necessary for the cooling processes [34]. A Newport TLB-6900 Vortex II used as the master-cooler (MC) is tuned red of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ and performs the actual cooling by addressing the Rb cycling transition $|F = 2, m_F = 2\rangle \rightarrow |F' = 3, m_F = 3\rangle$. Low entropy monochromatic light that is coherently absorbed and spontaneously emitted by the atoms results in cooling of the atoms [35]. The limit to Doppler cooling is the atomic linewidth Γ , which corresponds to a temperature $T_D = \frac{\hbar\Gamma}{2k_B} \approx 150 \mu\text{K}$, and then the ultimate limit to laser cooling is the recoil limit $T_{\text{recoil}} = \frac{1}{k_B} \frac{\hbar^2 k^2}{2m} \approx 180 \text{ nK}$, a temperature difficult to reach through non-evaporative means due to technical limitations. The existence of a cycling transition implies that the atom can absorb and scatter light ad infinitum, with each event returning the atom to its original state, ready to participate in another scattering event. Unfortunately Raman transitions and technical imperfections in the polarization of the light allow for the MC to move atomic population into the $|F' = 2\rangle$ state, from which it can decay into the $|F = 1\rangle$ ground state, a dark state that does not participate in cooling. This process quickly depumps

the atoms and terminates the cooling. A solution to this problem is to introduce a second laser, a Toptica DL100 used as the repumping (RP) laser, which drives atomic population out of the $|F = 1\rangle$ so the atoms can continue participating in cooling interactions with the MC. The laser system is designed in such a way that the RP light is detuned 165 MHz red of the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition; this is the light the atoms see while traveling in the Zeeman slower. Atoms pumped into $|F' = 2\rangle$ will invariably fall back to $|F = 2\rangle$, where they will be pumped into the stretched state $|F = 2, m_F = 2\rangle$ and cycle with the MC. For the MOT the RP light is passed through an acousto-optic-modulator (AO) driven by 160 MHz, putting the light 5 MHz red of the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. The MC light is tuned 21.5 MHz red of the cycling transition for the MOT, and passed through an AO driven by 140 MHz to be 161.5 MHz red of the cycling transition for the Zeeman slower. The frequency of the RP light is never changed for the rest of the experiment, but the MC is shifted to accommodate the needs of sub-Doppler cooling, optical pumping, and imaging; this is accomplished by adjusting the lock, as will be discussed. All imaging and optical pumping lines are controlled by AO switches driven by 80 MHz; this frequency shift is taken into account when the MC frequency is changed. A schematic of the laser system is shown in Fig. 2.5.

2.4 Laser Locks

The two cooling lasers must be frequency stabilized in order to reliably be used in their role of cooling, manipulating, and imaging the atoms. Stabilization is accomplished by first obtaining an error signal for each laser frequency; the error signal is an electronic voltage which is roughly linear as a function of the laser frequency and




Figure 2.5: Two lasers are necessary for laser cooling Rb, one to address the Rb cycling transition (MC), and one to pump the atoms out of the dark $|F = 1\rangle$ state (RP). In this experiment these two diode lasers are amplified by tapered amplifiers before being distributed to enact all the required cooling and trapping processes. Acousto-optic-modulators (AOM) are used as switches and, by virtue of a four-wave-mixing process, change the frequency of the out-coupled beam. Frequency stability is established in two stages: the a sideband of the RP created by a fiber electro-optic-modulator (EOM) is locked to the ^{85}Rb cycling transition, and the MC is locked directly to the RP through a beat-note lock. Since both the RP sideband and the set frequency offset between the RP and the MC are computer controlled parameters, the laser frequencies can be set arbitrarily.

changes sign at the point where we want the frequency to remain fixed. This signal is then fed into homebuilt feedback electronics that essentially apply a proportional-integral (PI) filter to the signal. The PI-filtered signal is used to feedback on the frequency of the laser through modulation of the current and PZT voltage. The specific shape of the PI loops is determined in-situ for each lock, dependent on the specifics of the laser's phase and amplitude response in each channel. Feeding back on the current of a diode laser it is advisable to include a lead-lag filter as shown in Fig. 2.6 or an all-pass filter as shown in Fig. 2.7. The principal mechanism for changing the laser frequency by modulating the current at low frequencies is by affecting the temperature; at high frequencies changes to the carrier density are responsible for changes to the laser frequency. As a consequence the current response at low frequency is actually out of phase with the response at high frequency; measuring this phase response and correcting it with one of the aforementioned filters will improve lock quality.

The MC is locked to the RP by a beatnote lock; both lasers are incident on a EOT-4000 photodetector possessing a bandwidth of 12.5 GHz. The resulting beat signal at ~ 6.4 GHz is amplified and sent into feedback port of a AD4108 PLL. A ~ 64 MHz signal synthesized by a AD9958 Direct Digital Synthesizer (DDS) is provided as the reference to the PLL, and the charge pump output is used as an error signal. The PLL divides down the beat signal ~ 100 times and compares it to the reference signal. The charge pump integrates this comparison and produces an extremely sharp turn-on when one frequency is above the other. This is used as the error signal to a loop-filter that feeds back on the current and PZT voltage of the




Figure 2.6: A lead lag filter with transfer function $\frac{V_o}{V_i} = \frac{1}{50} \left(200 + \frac{1}{\frac{1}{300} + \frac{i\omega}{3125000}} \right)$.




Figure 2.7: An all pass filter (phase-shifter) with transfer function $\frac{V_o}{V_i} = \frac{i\omega RC - 1}{i\omega RC + 1}$.

MC, locking it to the RP. Changing the reference to the PLL changes the MC frequency accordingly; a fact we use to great effect in order to move the MC frequency significantly for its various cooling and imaging roles. A straightforward calculation using well-known spectroscopic tables of ^{87}Rb shows that the beat frequency given the above stipulations should be 6.425 GHz. In practice this is adjusted to optimize the trapping and cooling efficiency.

$$(4.272\text{ GHz}-72.9\text{ MHz}-165\text{ MHz})-(-2.563\text{ GHz}+193.7\text{ MHz}-21.5\text{ MHz}) = 6.425\text{ GHz}$$

The RP is locked to a Rb vapor cell and acts as the ultimate frequency reference in the lab. Traditionally the error signal has been derived by a lock-in technique. By passing a beam twice through a vapor cell the resulting photodiode signal will show a Doppler-free absorption spectrum as the frequency of the laser is scanned [36], a consequence of the forward and backward passing beams interrogating the same stationary velocity class of atoms. Since the marginal difference in the amount of scattered light decreases as the intensity of the incident light increases (i.e. the atoms are a saturable medium) the atomic resonances will appear as resonant features on top of the Doppler broadened spectrum, features that have approximately the native linewidth of the atomic transition Γ . Additionally, there are Γ -wide features for every pair of transitions at exactly the pair's average frequency. These crossover peaks correspond to atoms from a non-stationary velocity class experiencing both beams, resulting in reduced absorption at that frequency. As an aside, the crossover peaks may very well be as large or larger than the absorption peaks, making them perfectly usable for locking. Spectroscopy of the D_2 line from either the $|F = 1\rangle$

or $|F = 2\rangle$ ground state yields six different Doppler-free peaks; three absorption peaks and three crossover peaks. By modulating the frequency of the beam (using, for instance a weakly driven electro-optic-modulator (EO)) and mixing the drive frequency out of the photodetector signal, it is possible to obtain a lock-in signal which is approximately the derivative of the spectroscopy signal. The sharp Doppler-free resonances become dispersive features whose approximately linear slopes make for good error signals. The problem with this method is that the dispersive features sit on top of a broad Doppler-broadened absorptive feature; changes in the temperature modify the width of this feature and in turn move the zero-crossing of the dispersive feature, in effect forcing the laser to lock to a different non-resonant frequency.

This problem can be circumvented by making use of modulation transfer spectroscopy [37–42], a technique that takes advantage of the non-linear susceptibilities of the atoms. The setup is not drastically different from Doppler-free saturated absorption spectroscopy, but now only one beam is modulated by the EO: the beam that is not measured by the photodetector. In this way the effect of the modulation is transferred from one beam to another, hence the name. This effect is nonlinear in the intensity of the light; specifically the size of the signal scales with the intensity of the light squared. For beams of fixed power passing through a vapor cell some compromise must be made between the number of atoms interrogated, proportional to the size of the beam, and the intensity squared of the beams, proportional to the inverse area squared. Like lock-in, this technique results in dispersive curves centered at each resonance usable as error signals, but unlike Doppler-free saturated absorption spectroscopy, there are no crossover features or Doppler-broadened background. It is in this latter detail that the technique finds its use; the lock-point is much more stable

to temperature variation. The cycling transitions provide the strongest response and are thus best used for error signals, in direct contrast to Doppler-free saturated absorption spectroscopy where the cycling transitions give the weakest response. Since ^{85}Rb is the more abundant isotope ($\sim 8:3$) its cycling transition was selected for use in the modulation transfer lock. A straightforward calculation using well-known spectroscopic tables of ^{87}Rb and ^{85}Rb shows that the difference in frequency between the desired frequency for the RP and the ^{85}Rb is 5.277 GHz. This was accomplished by passing the RP through a fiber EO driven by a ~ 5.3 GHz signal. The sideband on the light was used to generate the modulation transfer signal on the ^{85}Rb cycling transition while the RP was at the appropriate frequency for addressing ^{87}Rb . The microwave signal was generated by locking a Mini-Circuits VCO with a PLL. Incidentally, this lock scheme allows for the controllable shift of the RP frequency during the experimental run by changing the parameters of the PLL.

$$\begin{aligned} & (384.230484 \text{ THz} + 4.272 \text{ GHz} - 72.9 \text{ MHz} - 165 \text{ MHz}) \\ & - (384.230406 \text{ THz} - 1.265 \text{ GHz} + 100.2 \text{ MHz}) = 5.277 \text{ GHz} \end{aligned}$$

2.4.1 Transport ODT

The design of the transport ODT is constrained in several ways. The 50 W 1064 nm IPG fiber laser is focused at the center of the magnetic trap in the production chamber where it is involved in the trapping and cooling of the atoms during RF-evaporation. Here it has the effect of pulling the atoms away from the magnetic field zero, mitigating Majorana losses. Once this evaporation is completed the atoms are

released from the magnetic trap and held solely in the ODT, which is used to transport the atoms to the center of the cavity, 35 cm away. This transport is accomplished by the translation of a single lens mounted on an air-bearing stage. The stage can move the lens up to .5 m, which is more than enough to accommodate the separation between the centers of the production and science chambers. The simplest design would have a collimated beam with waist w_c focused by a single lens on the stage possessing focal length f . This results in a focused spot with waist $w_0 = \sqrt{\frac{f^2 w_c^2 \lambda^2}{\pi^2 w_c^4 + f^2 \lambda^2}} \approx \frac{f \lambda}{\pi w_c}$ a distance $\sim f$ after the lens. Translating the lens then has the effect of translating the location of the spot in a 1:1 correspondence. Placing a pair of relay optics on the stage is also an option, one which allows for a great deal of flexibility. However, relay optics multiply transverse deviation of the beam, so that if the optics were selected to multiply the translation of the focus per translation of the stage, that same multiplicative factor would increase the deviation of the focal spot in plane per deviation of the lens in its plane. As such, this susceptibility is minimized by limiting the translated elements to a single lens. In selecting this lens there are limitations in what is readily commercially available. A Thorlabs AC508-750-C-ML, $f = 750$ mm, is selected to allow for enough distance between the production chamber center and the lens, situating the lens at one end of the stage before translation.

In trapping the atoms in the production chamber, it is important to consider what effect the beam parameters at the trap center (the beam waist) have on the atoms. Decreasing the beam waists will increase the trap depth $\propto \frac{P}{\pi w_{0x} w_{0y}}$ for the same power P . A deeper trap will catch more atoms in the hand-off at the end of RF evaporation, allowing for better loading of the ODT. Additionally, the trap frequencies depend on the beam waists $\omega_i \propto \sqrt{P}/w_{0i}$, and high trap frequencies are preferable in order to

encourage good collisional properties during evaporation. All this implies that a tight trap is preferable, something that can be accomplished by having a large collimated beam focused by the translating lens. A beam with a 1 cm waist possesses a Rayleigh range of 295 m, meaning it can be ‘collimated’ before being focused by the translation lens to a waist of 25 μm . It turns out with these beam parameters, 10 W of laser power is more than enough to trap sufficient atoms for a sizable BEC ($> 10^5$) in the science chamber. However, the smaller the waists, the more rapid the expansion of the beam as it travels to the the cavity, increasing the irradiation in a potentially destructive manner. The amount of power incident on the mirrors and mount can be estimated by integrating the Gaussian beam at the plane of the cavity while the focus is at the center of the production chamber

$$2\sqrt{\frac{2}{\pi}}\sqrt{\frac{z_r^2}{w_0^2(z^2 + z_r^2)}}\int_{\delta}^{\infty} e^{-\frac{2x^2 z_r^2}{w_0^2(z^2 + z_r^2)}} dx$$

where here, z is the direction along the beam and x is the transverse dimension. δ is the height above the center of the cavity where the mirror edge begins, which is ~ 2.5 mm. Under these assumptions, a waist of 25 μm would irradiate the cavity and mount with as much as 2.8 W of laser power. The danger in this is overtly obvious, as the cement holding the mirrors in place could potentially soften, ruining the cavity coupling. The PZTs are also at risk, possessing a Curie temperature of 150 $^{\circ}\text{C}$, above which they cease functioning permanently. A simple solution is to increase the beam waist. To this end, a beam waist of 100 μm would correspond to 151 μW of power incident on the mirrors, a totally manageable number, as the thermal mass of the two mirrors (Quartz glass cylinders, 6.35 mm long, 6.35 mm radius) is 2.6 J/ $^{\circ}\text{C}$ implying a heating rate of 0.21 $^{\circ}\text{C/hr}$ if the laser were constantly on. In actuality, the ODT

is on less than half the experimental cycle, so even if the experiment were being operated continuously, the 40% duty cycle of the ODT would reduce the amount of power incident on the mirrors. Increasing the waist by a factor of four would require an increase in the power by a factor of 16 to obtain the same trapping parameters. This would increase the heating rate of the mirrors to as high as 3.4 °C/hr, which is not dangerous, but more importantly, would require a 160 W laser to generate the ODT, which is an impossibility. 10-15 W is a safe and comfortable level at which to trap the atoms, and all further calculations will assume an incident beam of 15 W maximum.

A crucial detail is that the cavity makes for a small angular obstruction in the plane perpendicular to the cavity axis. Previous calculations assumed that all the light not passing between the mirrors was incident (and absorbed) by the cavity mirrors and mount. However, since the waist of the beam is several times the width of the cavity mirrors, a good amount of the light will simply pass to the sides of the mount, and this is another reason that a smaller waist in that direction is preferable. Thus, the final working design is an elliptical beam with $w_{0x}=50\text{ }\mu\text{m}$, and $w_{0y}=70\text{ }\mu\text{m}$. If all the laser power not passing between the cavity mirrors were absorbed by the mirrors, the heating rate would be 20 °C/hr, arguably too high for stable or safe operation of the cavity. At this rate, the thermal wave will have propagated across the mirror in $L^2/D \sim 48\text{ s}$, allowing the thermal mass of the mount to be considered in addition to that of the mirrors. Including the quarter kilogram or so of stainless steel used in the mount's construction, the heating rate of the mount-mirror system can be as low as 0.4 °C/hr, and with the thermal diffusion of steel being a hundred-fold greater than that of glass, the conduction through the material is not the limiting factor. Even if

the cavity mount and mirrors were absorbing all the light not passing between the mirrors, the heating rate would be negligible. This conclusion bears out in practice as running the ODT only marginally affects the cavity lock, an indication of the low heating of the mirrors, which should have each increased in length by 70 nm in an hour under these conditions, well within the range of the feedback PZT.

Working backwards from the atoms, a ~ 5.1 mm beam at the focusing lens would be necessary to create the desired $50\text{ }\mu\text{m}$ waist. At this point, the simplest approach to obtain the desired waist would be to (again, going backwards) shrink the beam with a 7:1 telescope so it fits through the AO, and then expand it again with a 3:1 telescope, ‘resulting’ in a beam with a 1.65 mm waist, exactly what emerges from the IPG fiber. The difference in the beam waists is created with a pair of cylindrical lenses forming a telescope. The lenses have focal lengths 60 mm and 50 mm and the first of the two is placed after the AO. Shrinking the beam in the vertical direction has the effect of creating a larger waist in this dimension at the focus. The problem with this approach is that the focus occurs 75 cm after the final focusing lens, requiring it to be less than 35 cm from the end of the stage. Iterating through the available optics, it became clear that an 850 mm achromat would ideally replace the 750 mm focusing lens previously selected. Not finding one from Thorlabs, a compromise was reached by relaxing the requirement that the beam focused by the final lens be collimated. A diverging beam will focus to a smaller spot than one which is collimated, and at a farther distance as well. The only caveat is that the ratio of focal translation to stage translation is no longer unity, and the stage must move ~ 39 cm to move the atoms 35 cm. Additionally, the waists get smaller during this translation, resulting in a trap in the science chamber with the same ellipticity, but with waists 80% what they were

in the production chamber, increasing the trap frequencies and actually improving evaporation.

A straightforward ABCD matrix simulation of the beam parameters is as follows

$$\begin{aligned}\vec{q}_x &= \text{L} [.75] \text{FS} [x_t + .6] \text{L} [.5] \text{FS} [.44] \text{L} [.05] \text{FS} [.79] \text{L} [.05] \text{FS} [.2] \text{L} [.15] \text{FS} [.49] \vec{q}_0 \\ \vec{q}_y &= \text{L} [.75] \text{FS} [x_t + .6] \text{L} [.5] \text{FS} [.44] \text{L} [.05] \text{FS} [.03] \text{L} [.05] \text{FS} [.11] \text{L} [.06] \text{FS} [.65] \text{L} [.05] \text{FS} [.2] \text{L} [.15] \text{FS} [.49] \vec{q}_0\end{aligned}$$

Where $\vec{q}_x$ and $\vec{q}_y$ are the complex q parameters for the two directions after the last focusing lens, x_t is the translation distance of the lens, either 0 or .39 depending on whether the atoms are in the production or science chamber, and $\vec{q}_0 = \left(i\pi w_{\text{IPG}}^2 / \lambda \right)$ is the complex q parameter for the beam emerging from the fiber laser with $w_{\text{IPG}} = 1.65$ mm. $\text{FS}[d]$ and $\text{L}[f]$ are the ABCD matrices for free space propagation a distance d and transmission through a lens of focal length f respectively.

The two ABCD matrices are reproduced here for convenience

$$\begin{aligned}\text{FS}[d] &= \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \\ \text{L}[f] &= \begin{pmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{pmatrix}\end{aligned}$$

Finally, knowing the waists it is straightforward to calculate the trap frequencies assuming a harmonic potential [43], which is a useful assumption so long as $\eta = U_d/k_B T > 10$, where U_d is the trap depth. For ^{87}Rb when the frequency of the trap beam is far enough from the D_1 and D_2 resonances so as to average over the hyperfine

structure (as is the case for 1064 nm light), the expression for the trap depth of a Gaussian beam is

$$U_d = \frac{2P}{\pi w_{0x} w_{0y}} \left(\frac{\pi c^2}{2\omega_2^3} \Gamma_2 \frac{2}{\omega_2 - \omega_d} + \frac{\pi c^2}{2\omega_1^3} \Gamma_1 \frac{1}{\omega_2 - \omega_d} \right)$$

Where $\omega_{1,2}$ and $\Gamma_{1,2}$ are the resonant frequencies (in rad/s) and spontaneous emission rates of the $D_{1,2}$ lines respectively. ω_d is the frequency of the trapping light, and c is the speed of light. The trap frequencies (in rad/s) are given in terms of the trap depth by

$$(\omega_x, \omega_y, \omega_z) = \left(\sqrt{\frac{4U_d}{m w_{0x}^2}}, \sqrt{\frac{4U_d}{m w_{0y}^2}}, \sqrt{\frac{\lambda^2 U}{\pi^2 m} \left(\frac{1}{w_{0x}^4} + \frac{1}{w_{0y}^4} \right)} \right)$$

For a two-beam crossed dipole trap the trap frequencies are given by $\omega_i = \sqrt{\lambda_i/m}$ where λ_i are the eigenvalues of the following matrix

$$\begin{pmatrix} \frac{4U_c \cos^2 \theta}{w_{c1}^2} + \frac{4U_t}{w_{t1}^2} & 0 & \frac{4U_c \cos \theta \sin \theta}{w_{c1}^2} \\ 0 & \frac{4U_c}{w_{c2}^2} + \frac{4U_t}{w_{t2}^2} & 0 \\ \frac{4U_c \cos \theta \sin \theta}{w_{c1}^2} & 0 & \frac{4U_c \sin^2 \theta}{w_{c1}^2} \end{pmatrix}$$

Here $U_{t,c}$ is the trap depth of each beam, $w_{t,c}$ are the waists of each beam along their tight axes, and θ is the angle between the beams. This is assuming that one principal axis of each beam coincides.

2.4.2 Cross Optical Dipole Trap

Once the atoms are successfully transported between the cavity mirrors, a second ODT is used in conjunction with the transport ODT in the final evaporation step.

A single beam ODT is too loose along one direction to efficiently evaporate the gas; the evaporation rate is severely limited by what is almost always no more than a confinement of several Hz. Trapping the atoms in a second beam is an extremely effective approach to countering this problem. The difference in $\bar{\omega}$ between a single beam trap and a crossed trap can be several orders of magnitude, vastly enhancing the evaporation procedure. This technique has been used to great effect in producing a BEC entirely through optical cooling, trapping, and evaporation [44–46], eliminating the need for high-current magnetic traps. Evaporation from an all-optical trap is also quite fast; runaway evaporation cannot occur in a single beam ODT because of the specific way in which the trap depth is coupled to the trap frequencies [47]. This constraint is not present in a crossed or multi-beam OD, allowing for the achievement of very fast and efficient evaporation [48, 49]. A 50 W Nufern Fiber Amplifier is used for this purpose.

The cross ODT is designed to have a secondary utility as well; by passing the beam through an acousto-optic-modulator (AO) and correctly selecting the downstream optics it is possible to transform the angular deviation of the beam by the AO into parallel deviations of the beam at the location of the atoms. This has the effect of translating the atoms along the direction of the transport ODT in a controllable and reproducible manner; simply by changing the frequency of the signal driving the AO the location of the trap can be adjusted. As long as this translation is not larger than the Rayleigh range (~ 4 mm) of either beam the trapping will not be changed significantly and the same evaporation ramp can be used. Furthermore, by modulating the frequency driving the AO at a rate faster than the trap frequencies it is possible to create wider traps, or even a series of separated traps, by ‘smearing’

the trap over a wider area in a time-averaged manner. The means by which this is accomplished is by positioning a single lens downstream of the AO, a focal length away. Understood through geometric optics, the lens has the effect of transforming angular deviations from the AO into parallel, translational deviations. Additionally, it has the effect of focusing the beam a focal length away. If the atoms were located at this plane then changing the frequency of the signal driving the AO would have the desired effect, sliding the focus of the beam transversely. If, however, the atoms are not located at this plane then the focus has to be relayed to the atoms, which is accomplished by a series of telescopes. The introduction of a telescope into the beam path downstream of the focusing lens maintains the geometric condition while refocusing the beam in a plane past the telescope.

In the case of the ODT a number of different criteria had to be met. The focus of the beam had to be located ~ 30 cm past the final lens in order to accommodate additional optics as well as the distance between the viewport and the atoms. The waist of the beam was chosen to be $40\text{ }\mu\text{m}$, mimicking the transport ODT. Each lens in the design had to be chosen specifically to meet these conditions. What was settled on is as follows. In order to improve stability and improve beam quality the Nufern was coupled into a Thorlabs single mode polarization maintaining fiber with a quoted numerical aperture (NA) of 0.12. Exiting the fiber the light is collimated by an 8 mm aspherical lens and aligned into the AO. The first order mode of the AO is centered on a 75 mm lens, the lens which transforms angular deviations into parallel translations. The beam is then centered on a 2:3 telescope made of a 50 mm lens preceding a 75 mm lens. The result is a beam with a predicted and measured $40\text{ }\mu\text{m}$ waist. The result of the geometric optics assessment is that the beam will

translate 11 cm for every radian of angular deviation at the AO. Using the standard formula for AO deflection, $\theta_D = \lambda f_a / v$, where λ is the optical wavelength, f_a is the frequency of the signal driving the AO, and v is the acoustic velocity in a quartz AO, this implies that 50 kHz of modulation is necessary to translate the beam by $1 \mu\text{m}$.

Evaporation from the cross is accomplished by decreasing the power in the two beams according to power-law ramps $P_i(t) = \frac{P_i}{(1-t/\tau_i)^{\gamma_i}}$, with ramp parameters τ_i and γ_i . This function has the property that it forces evaporation to occur quickly at the onset and more slowly at later times in order to accommodate the decreasing trap frequencies of the trap. The ramp parameters are adjusted iteratively in order to maximize BEC purity and number, which are assessed by time-of-flight (TOF) absorption images. A good approximation when the gas isn't too dense [50] is to treat the gas as a non-interacting two-component system. The BEC is in thermal equilibrium with a non-condensed fraction and the two expand independently in TOF. In a harmonic trap under this approximation the density of the thermal component is Gaussian $\rho(\mathbf{x}) \propto \exp\left(-\beta(\frac{1}{2}m \sum_{i=1}^3 \omega_i^2 x_i^2)\right)$ and condensation occurs when the temperature $T = \frac{1}{\beta k_B}$ reaches the critical temperature in the trap $T_c = \frac{\hbar \bar{\omega}}{k_B} \left(\frac{N}{\zeta(3)}\right)^{1/3}$, where $\bar{\omega}$ is the geometric mean of the trap frequencies (rad/s), N is the total number of atoms in the trap, and $\zeta(3) \approx 1.2$ is the Riemann-Zeta function [51]. Lowering the temperature below T_c increases the condensate fraction according to $f_0 = 1 - \left(\frac{T}{T_c}\right)^3$. This analysis comes out of a statistical treatment of the Bose gas; interactions amongst the atoms renormalize both T_c and f_0 in a well-understood way. Free from interactions the thermal component of the gas remains Gaussian during TOF, and thus appears Gaussian in absorption imaging since the column integral of a Gaussian cloud is also Gaussian. Finally, the line integral in 1D is also Gaussian.

The atoms of the BEC interact primarily through s-wave contact interaction, which is a manifestation of van der Waals forces between the atoms. p-wave and higher order scattering terms are frozen out by a centrifugal barrier at such low temperatures, and dipole-dipole (and other magnetic multipole interactions) play a significantly reduced role in Rubidium scattering, but are dominant in dipolar gases like dysprosium [52]. Contact interactions like this can be treated in a mean-field way by the Gross-Pitaevskii equation, which includes an interaction term $NU|\psi(\mathbf{x})|^2$ in the Hamiltonian. Here N is the number of atoms and the interaction parameter U is related to the s-wave scattering length by $U = 4\pi\hbar^2 a_s/m$. In ^{87}Rb $a_s \approx 100a_0$ where $a_0 = \frac{\hbar}{m_e c \alpha}$ is the Bohr radius. For a sufficiently dense cloud of atoms the mean-field interaction energy dominates the kinetic energy of the gas, allowing the use of the Thomas-Fermi (TF) approximation to describe the density of the gas. The TF approximation provides an especially simple treatment for the density of a trapped gas, $\rho(\mathbf{x}) \propto \mu - V(\mathbf{x})$ where the chemical potential μ constrains the number of atoms in the trap and $V(\mathbf{x})$ is the trapping potential. The TF density has an unnaturally sharp change at the edges of the cloud; in reality this sharpness is softened by the kinetic energy and rather than ceasing discontinuously, the density falls off exponentially in spherical coordinates $\sim \exp(-\frac{2}{3}s^{2/3})$ where $s = (r - R)/(a_{ho}^4/2R)^{1/3}$ is a non-dimensional length [53]. Here $R = a_{ho}(15Na/a_{ho})^{1/5}$ and a_{ho} is the harmonic oscillator length. For traps used in this experiment the extent of this region is only a few hundred nanometers and implies a tiny modification on the TF density. Absorption imaging produce a two dimensional image which in a low-intensity limit is the three dimensional density integrated along the imaging axis $\int \rho(\mathbf{x})dz$. For a TF density integrated along one of its principal axes in imaging, and then numerically (again

along a principal axis) to produce a one dimensional line density, expressions for the TF density are as follows

$$\begin{aligned}\rho(\mathbf{x}) &= \frac{\mu - \frac{1}{2}m \sum_{i=1}^3 \omega_i^2 x_i^2}{U} \\ \rho_{2D}(\mathbf{x}) &= \frac{4 \left(\mu - \frac{1}{2}m \sum_{i=1}^2 \omega_i^2 x_i^2 \right)^{3/2}}{3U \sqrt{\frac{m\omega_z^2}{2}}} \\ \rho_{1D}(\mathbf{x}) &= \frac{\pi \left(\mu - \frac{1}{2}m\omega_x^2 x^2 \right)^2}{2U \sqrt{\frac{m\omega_y^2}{2}} \sqrt{\frac{m\omega_z^2}{2}}}\end{aligned}$$

If the imaging or numerical integration is not performed along a principal axis the forms of the equations are the same, e.g. $\rho_{2D}(\mathbf{x}) = c_1 (c_2 - c_x x^2 - c_y y^2)^{3/2}$ but the coordinates x and y are a new set of projected principal axes, and the constants c_i are functions of the trap frequencies, chemical potential, and Euler angles describing the cloud's orientation. A partially condensed cloud will possess a density that is the sum of a thermal Gaussian and a TF distribution; fitting to this bimodal distribution is the hallmark of the onset of condensation.

2.5 Cavity

In extending the study of self-organization phenomena in cavity-QED systems to include phase transitions marked by continuous symmetry breaking, it is important to control the number and density of available cavity modes into which the atoms can self organize, from a single mode to a continuum of modes. While there are many types of degenerate optical resonators, a confocal cavity ($L = R$) was selected for its mechanical stability; if the mirrors drift with respect to each other the mode of the

cavity is shifted on the mirror's surface by only a marginal amount. In contrast, for a concentric cavity (defined by $L = 2R$), not only is the cavity only marginally stable with regards to its length, but any drift of the mirrors will result in a cavity mode which quickly falls off the mirrors. Mechanical stability of the cavity is especially important for this particular implementation as the cavity length is adjustable in-vacuum. This cavity length is changed by activating a Attocube nanopositioning built into the cavity mount. Imperfections in the assembly of the cavity and mount make it impossible to guarantee that the axis along which the cavity mirror is translated is the same as the cavity axis along which the mode lives. The mechanical stability of confocal cavities obviates the need for any further (coarse) control of the orientation or position of the cavity mirror; when the length is changed the cavity mode may reorient itself, but by so little as to be irrelevant. While the nanopositioning stage determines the cavity's length coarsely, a single-crystal PZT attached to the other mirror is used to length-stabilize the cavity through active feedback. An error signal is obtained by the Pound-Drever-Hall (PDH) technique [54]. The PDH error signal, as depicted in Fig. 2.8, actually has three usable lock points; the cavity can be locked such that a resonance is at the carrier frequency ω_c or at either sideband frequency $\omega_c \pm \omega_m$. This is advantageous as locking on the carrier results in a large cavity field; 100 μW of light incident on the cavity is enhanced by as much as $\mathcal{F}/\pi = 18000$ times, producing a 2 W trapping potential in the cavity. This potential can be used to determine the center of the cavity mode, to assist in evaporation, or to hold the atoms for self-organization studies. Alternatively, locking on a sideband puts almost no light in the cavity whatsoever; the strength of the PDH signal is proportional to the product of the carrier and sideband intensities, meaning that a weak sideband

can be compensated for by a strong carrier. Locking at the sideband frequency, there is so little power in the sideband as to not affect the atoms in any noticeable manner.




Figure 2.8: A simulation of a Pound-Drever-Hall error signal showing three lock points. The three dispersive features are extremely linear around the zero-crossings and thus are effective as error signals.

The cavity is locked by filtering the PDH error signal and driving the single crystal PZT holding the bottom mirror. If the response of the PZT were flat, a simple PI control scheme might be sufficient. However, in trying to lock the cavity it was discovered that the plant response possessed a number of resonant features that would destabilize a conventional lock scheme. These features were discovered through interferometric measurement of the mirrors' response. By using the flat backside of each mirror in turn as one port of a Michelson-Morley interferometer and

then driving the associated piezo with a swept-sine, it was possible to measure the electro-mechanical response of the system. To reduce noise in the measurement one of the two interferometer beams was the first order of an AO driven by 60 MHz, which was then mixed out of the photodetector signal. Additionally, the PD response was fed to the measured piezo in order to loosely lock it to the laser frequency, further improving data quality.

Two important features were discovered in these measurements. In the response spectrum of the Attocube piezo there is an enormous feature located at ≈ 120 Hz, and this response curve is shown in Fig. 2.9. Despite being suspiciously close to twice the mains frequency, e.g. 2×60 Hz, no other harmonics are visible in the spectrum and it is almost certain that the feature is a mechanical resonance of the mount or Attocube. This feature implies a frightening ability of the system to pick-up noise at this frequency, and as a consequence the controller should exhibit a similarly enhanced amount of gain at this frequency to suppress the noise.

An active resonant filter, whose basic structure is shown in Fig. 2.10, was developed to enhance the gain in this band. This filter exhibits the Sallen-Key topology with an independently adjustable resonance and quality factor, and in this implementation increases the gain at 120 Hz by ~ 20 dB.

The transfer function of the filter is given by the following expression. While somewhat complicated, the form of the transfer function is that of a damped resonator.

$$\frac{iC_1R_2R_3\omega(R_4 + R_5)}{i\omega(C_1R_2R_5(R_1 + R_3) - C_1R_1R_3R_4 + C_2R_3R_5(R_1 + R_2)) - C_1C_2R_1R_2R_3R_5\omega^2 + R_5(R_1 + R_2)}$$




Figure 2.9: Interferometric response of a cavity mirror, mounted on an Attocube nanopositioning stage, whose piezo is driven by a swept-sine. There is a strong susceptibility at 120 Hz that requires enhanced gain at that frequency to suppress.




Figure 2.10: The Sallen-Key topology can be used to create a wide variety of active resonant filters with independently tunable resonant frequencies and quality factors.

Useful for the design of such a filter are the resonant frequency ω_0 , and the quality factor Q are

$$\omega_0 = \frac{\sqrt{R_1 + R_2}}{\sqrt{C_1 C_2 R_1 R_2 R_3}}$$

$$Q = \frac{\sqrt{C_1} \sqrt{C_2} \sqrt{R_1} \sqrt{R_2} \sqrt{R_3} \sqrt{R_1 + R_2} R_5}{C_1 R_1 R_3 R_4 - R_5 (C_1 R_2 (R_1 + R_3) + C_2 R_3 (R_1 + R_2))}$$

A simulated phase and amplitude response of the filter is shown in Fig. 2.11 and Fig. 2.12.




Figure 2.11: The amplitude response of the filter used to enhance the gain at 120 Hz. In this simulation the response is almost 80 dB, showing the potential for this topology to provide extremely strong and narrow amplification.

Performing the interferometry on the feedback PZT results in the curve shown in Fig. 2.13. While it is perhaps unreasonable to expect a large single-crystal PZT to have a flat response beyond 1 kHz, the dispersive features at 3.5 kHz and 6.5 kHz may be significantly limit the efficacy of the feedback. These features are the result of electromechanical resonances in the PZT and prevent PI control from being effective at these frequencies. One solution would be to limit the bandwidth to below the features, but it is highly preferable to extend the bandwidth and cancel these features,




Figure 2.12: The phase response of the filter used to enhance the gain at 120 Hz. As expected from a resonant feature the filter has 180° phase shift across the resonance.

an approach which proved feasible. A resonant circuit was built to have the same dispersive lineshape as the 3.5 kHz feature observed in the response spectrum; putting this filter in front of the PZT has the effect of flattening the plant gain. Only the 3.5 kHz feature was cancelled, as it was assumed (and borne out to be true) that 5 kHz was sufficient to adequately lock the cavity.

The filter used to produce this response is shown in some detail in Fig. 2.14.

The filter's transfer function is written out below.

$$\frac{R_8}{R_5} \frac{i(R_3 + R_4)}{C_1 R_1 R_3 \omega + R_4 (-i C_1 C_2 R_1 R_2 \omega^2 - C_2 \omega (R_1 + R_2) + i)} - \frac{R_6}{R_5}$$




Figure 2.13: Interferometric response of a cavity mirror, mounted on a single-crystal PZT which is driven by a swept-sine. There is a dispersive feature at 3.5 kHz, probably of electro-mechanical origin, that makes feedback at that frequency impossible.




Figure 2.14: The dispersive feature is created by adding a resonance to a flat background, turning the 180° phase shift into an amplitude shift.

Chapter 3

Digital Micromirror Device

3.1 Overview

The Digital Micromirror Device (DMD) used in this project, a Texas Instruments DLP4500NIR, is an array of 912×1140 individually addressable aluminum mirrors. These mirrors are arranged on a rectangular substrate on which I designate a right handed coordinate system wherein the unit vector pointing down the long edge of the rectangle is $\hat{x} = (1, 0)$, and the unit vector pointing along the short edge of the rectangle is $\hat{y} = (0, 1)$. Each mirror is square and measures approximately $10.8 \mu\text{m}$ on the diagonal. The mirrors are arranged in a square lattice with a lattice constant of $a = 7.6 \mu\text{m}$ whose lattice vectors $\vec{u}_1 = \frac{a}{\sqrt{2}}(1, 1)$ and $\vec{u}_2 = \frac{a}{\sqrt{2}}(1, -1)$ are rotated 45° with respect to $\hat{x}$ and $\hat{y}$. With incredible improvements in nano-fabrication technology over the past decade perhaps it is no surprise that each tiny mirror be actuated and tilted $\pm 12^\circ$ with respect to the DMD surface. The axis of rotation for the mirrors is $\hat{y}$. A binary bitmap file (.bmp) with the resolution of the DMD is used to simultaneously address all the mirrors at once, where a value of ‘1’ tilts a mirror to one position,

and a value of ‘0’ tilts it to the other. Because the lattice vectors of the DMD array are tilted with respect to the DMD surface, indexing the mirrors is done in a highly non-intuitive manner.

Fig. 3.1 shows the numbering of the pixels; pixels in the same column offset by 1 row are separated by alternating lattice vectors. Pixels in the same row offset by 1 column are separated horizontally by a combination of the lattice vectors $\vec{u}_1 + \vec{u}_2$. This confusing system leads to the surprising fact that the $\hat{y}$ direction, which corresponds to the shorter side of the DMD, contains 1140 rows, while the longer side has fewer columns. The vertical separation between mirrors in adjacent rows is $a\sqrt{2}$, while the separation between mirrors in adjacent columns is twice that. As a consequence, the vertical extent of the DMD is ~ 6.1 mm, while the horizontal DMD of the mirrors is ~ 9.8 mm. Once a mapping between the pixels of the .bmp image and the mirrors of the DMD is made, it is straightforward to convert between the two, but dealing with this fact can be confusing at first.

The mirrors, angled as they are and regularly arranged on a square lattice, constitute what is known as a blazed grating. A blazed grating scatters light by two mechanisms; incident light is diffracted by the regular array of scatterers, and reflected specularly by each scatterer which in this case is a small mirror. If one of the diffraction orders is in the same direction as the reflected light, the fraction of light scattered in this order will be maximized. In this case, the DMD is a 2D grating, and the blazing condition can only be met if the incident and scattered light is in a plane orthogonal to the mirrors’ axis of rotation $\hat{y}$.

A more in-depth analysis of the blazing condition begins with a coordinate system; the coordinates parallel to the surface of the DMD are $\hat{x}$ and $\hat{y}$, and the third




Figure 3.1: A schematic of the DMD showing how the pixels are numbered and arranged. Reproduced from [55].




Figure 3.2: A schematic of a blazed grating showing the blazing angle θ_B as well as the path traced by a ray of light reflecting off the surface. Reproduced from [56].

orthogonal coordinate is $\hat{z}$. The light incident on the surface of the DMD has a k-vector $\vec{k} = k(\sin \theta_i, 0, \cos \theta_i)$ and can be divided into a part perpendicular to the DMD surface and a part parallel to the DMD surface, $\vec{k} = \vec{k}_{\parallel} + \vec{k}_{\perp}$. After reflecting off the mirrors' surface, the resulting k-vector is $\vec{k}_r = \vec{k} - 2(\hat{n} \cdot \vec{k})\hat{n}$, where $\hat{n}$ is the mirrors' surface normal $\hat{n} = (\sin \beta, 0, \cos \beta)$ where $\beta = \pm 12^\circ = \pm \pi/15$. The resulting reflected beam is

$$\vec{k}_r = -k \sin(2\beta - \theta_i)\hat{x} - k \cos(2\beta - \theta_i)\hat{z}$$

Diffraction off a grating can be understood as adding an integer number of reciprocal lattice vectors to the in-plane component of the incident light beam, effecting the transformation $\vec{k}_{\parallel} \rightarrow \vec{k}_{\parallel} + m\vec{K}_1 + n\vec{K}_2$ for $m, n \in \mathbb{Z}$. The perpendicular component must then be modified to satisfy energy conservation, i.e. $|k_d| = |k|$. In the case of a 2D square lattice the reciprocal vectors are $\frac{\sqrt{2}\pi}{a}(1, 1, 0)$ and $\frac{\sqrt{2}\pi}{a}(1, -1, 0)$, and the locations of the diffraction orders are given by

$$\vec{k}_d = \left(\frac{\sqrt{2}\pi(m+n)}{a} + k \sin \theta_i \right) \hat{x} + \frac{1}{a} \sqrt{2}\pi(m-n) \hat{y} + \frac{1}{a} \sqrt{a^2 k^2 \cos^2 \theta_i - 2\pi \left(\sqrt{2}ak \sin \theta_i (m+n) + 2\pi (m^2 + n^2) \right)} \hat{z}$$

To keep the diffracted beam in the $x - z$ plane, m must equal n , and the above expression simplifies to

$$\vec{k}_d = \left(\frac{2\sqrt{2}\pi n}{a} + k \sin \theta_i \right) \hat{x} + \sqrt{k^2 - \left(\frac{2\sqrt{2}\pi n}{a} + k \sin \theta_i \right)^2} \hat{z}$$

The blazing condition is met whenever $\vec{k}_d = \vec{k}_r$, which occurs for 780 nm light when $n = 0, 1, 2$, corresponding to angles of $\theta_i = 78^\circ, 98.3^\circ, 122^\circ$ with respect to vertical. Understanding this is not necessary to operating the DMD; simply positioning it in such a way that it is uniformly illuminated and the light from the ± 12 mirrors can be separated is sufficient. However, considering the laser power lost into all the other orders it may be prudent to establish the blazing condition at first. Once the DMD has been placed, a Fourier lens is selected and put approximately a focal length's distance down the path traversed by the desired ('on') DMD light. Placing a camera a focal length's distance from the lens and imaging the light, the field at the camera plane should be approximately the Fourier Transform of the field at the DMD plane. The waist of a Gaussian beam in the image plane is $2f/kw_0$ when starting with a waist w_0 in the DMD plane. In this project the test cavity was a confocal cavity made from 2 cm mirrors whose waist is $49 \mu\text{m}$ at the center of the cavity. Since the largest waist possible on the DMD is approximately 2 mm, a 200 mm focal length lens would result in a $25 \mu\text{m}$ waist, on the order of the cavity mode.

Since the DMD is to an excellent approximation acting as a binary amplitude grating, algorithms from the field of digital holography can be used to design masks for the DMD. However, it will quickly be discovered that aberrations in the optical system as well as errors made in placing the components will prevent the planes from truly being Fourier Transform pairs. Figure 3.3 shows the result of an attempt to produce a tight spot (an approximate delta function) in the camera plane by virtue of putting a grating in the DMD plane. A calibration procedure is necessary to measure the optical aberrations and cancel them, enforcing the Fourier reciprocity between the two planes. Figure 3.4 shows the result of including the aberration cancellation into the Hologram design. The spot is extremely sharp on the CCD camera.

The calibration procedure yields a phase mask $\phi_{\Delta}(x, y)$ that can be understood as the optical phase difference between light reflected from different parts of the DMD. This mask represented in the real-space coordinates of the DMD surface can be included in hologram design to cancel aberrations. Figure 3.5 shows one such mask derived from raw data. Figure 3.6 is an unwrapped version of the same mask and demonstrates the scale over which the aberrations occur. The resolution of the calibration procedure is certainly high enough to capture the aberrative features. The pixels on these masks represent superpixels, or groups of pixels on the DMD. The mask is processed further to be full-sized over the surface of the DMD so that each mask pixel corresponds to a single DMD mirror, and then that phase mask is incorporated into the design of every hologram used in this thesis. After implementing this calibration procedure the DMD can be used for full phase and amplitude control of the field in the imaging plane. In the methodology chosen for this study a grating is necessary to impart phase control; using different algorithms and methodologies it




Figure 3.3: Attempt to produce a sharp spot with by writing a grating on the DMD and imaging the Fourier plane. Placing the Fourier lens approximately in the correct location with respect to the image plane and the DMD plane results in a poorly focused spot suffering from aberrations.




Figure 3.4: After incorporating the aberration-canceling phase mask into the hologram design, the resulting spot is extremely sharp. The Image is saturated to showcase the higher order peaks, but the maximum intensity of the central peak is $> 100\times$ the maximum intensity of the satellite peaks.

is possible to neglect either phase or amplitude in favor of the other, a quality useful for the construction of off-resonant light traps [57]. Closed-loop optimization schemes can readily be applied by measuring the amplitude of the field [58], but feeding back on the phase of the desired field is a more involved process. Performing feedback on a mask coupled to a cavity mode is certainly possible if the cavity is locked and repeatable transmission measurements can be obtained.

3.2 Theory of the Paraxial Cavity

The spectroscopic properties of an ideal Fabry-Pérot cavity are well described by paraxial wave theory [59]. The results of such analyses are two-fold: first, the modes themselves are Hermite-Gauss (HG) beams. Indeed in practice, the fundamental mode of a cavity is nearly indistinguishable from an ideal Gaussian beam, even when using spherical, rather than parabolic, mirrors. Second, the frequencies (in Hz) of these modes are given by

$$f_{qmn} = \frac{c}{2L} \left(q + \frac{m + \frac{1}{2}}{\pi} \arccos \sqrt{g_{1x}g_{2x}} + \frac{n + \frac{1}{2}}{\pi} \arccos \sqrt{g_{1y}g_{2y}} \right), \quad (3.1)$$

where c is the speed of light in the cavity, L is the length of the two-mirror resonator, $g_{ij} = 1 - \frac{L}{R_{ij}}$, and R_{ij} is the radius of curvature (ROC) of the i^{th} mirror in the j^{th} direction. These denote orthogonal Cartesian axes. This formula captures many of the salient features of optical resonators. In a symmetric cavity where all ROCs are equal, the cavity may exhibit degenerate configurations at particular ratios of cavity length to ROC wherein a class of spatially incommensurate modes have identical resonant frequencies. Specifically, this degeneracy can occur when the Gouy phase




Figure 3.5: The DMD calibration procedure measures the optical phase difference between every pair of super-pixels on the surface of the DMD. This information is encapsulated in a 48×30 .bmp image (top) where each pixel is real-valued $\in [-\pi, \pi]$. An unwrapped version of the phase mask is shown as well (bottom), clearly showing that the phase variation is smooth at this resolution. In principle either image can be decomposed in the Zernike polynomials, giving some insight into the nature of the aberration.




Figure 3.6: The fully processed phase mask embedded in a rectangular domain where locations in the array have a one-to-one spatial correspondence with the centers of the DMD pixels.

$\arccos(\pm g)$ is, e.g., 0 for planar cavities, $\pi/2$ for confocal cavities, and π for concentric cavities.

The ideal confocal cavity is one well-known example of a degenerate two-mirror resonator; a confocal cavity's length is equal to the ROC of its mirrors $L = R$. The mode-spectrum separates into two degenerate families alternating every half free spectral range (FSR), where an FSR is equal to $\frac{c}{2L}$. One degenerate family consists of all the even modes with $m + n \equiv 0 \pmod{2}$, while the other consists of the odd modes with $m + n \equiv 1 \pmod{2}$. It is worth noting that the azimuthal invariance in ideal confocal cavities means that the eigenmodes are also well-described by Laguerre-Gauss functions, complex admixtures of the real-valued HG functions.¹

The practicalities of mirror fabrication and mounting often break the mode degeneracy²; even if the mirrors are perfect paraboloids, they are unlikely to have the same four curvatures and same axial alignments. In these realistic cases, the spectroscopy of the near-confocal cavity is marked by two energy scales for the HG modes. The first, $\Delta E_1 = c(\delta_L - \delta_R)/\pi R^2$, is the frequency difference between families of different $m + n$, where $\delta_L = L - R$ and $\delta_R = \bar{R}_i - \bar{R}_j$. Here, $\bar{R}_i = (R_{1i} + R_{2i})/2$, where R_{1i} is the ROC of mirror 1 along direction i , etc. The second energy scale is $\Delta E_2 = \frac{c\delta_R}{2\pi R^2}$, which is the energy offset between modes within the same family. This picture is complicated by the fact that the cavity mirrors may not be mounted in such a way that the principal axes of the two mirrors are co-planar. Moreover, mixing of the ideal HG modes into a non-uniform ΔE_2 spectrum within modes of a family arises from aberrations in the shape of the mirror, e.g., spherical aberration, beyond-paraxial

¹More generally, the Ince-Gauss modes provide a smooth crossover between these two families and are useful in describing real cavity modes where mode-mixing is present.

²Additionally, the two polarizations of each mode can be frequency-split by birefringence in the mirror and its coatings exacerbated by strain caused by the mounting supports, glue, etc.

effects in the propagation of the light, and defects on the mirror surface. This mixing increases as the modes begin to overlap in frequency as degeneracy is reached.




Figure 3.7: Overview of the experiment. Laser light at 780.23 nm is incident on the DMD. The ‘on’ pixels direct the light into the cavity via the first-order diffraction peak of the programmed holographic grating. The light reflected from the ‘off’ pixels is caught by a beam block (not shown). The light emitted by the cavity is analyzed by a photodetector and a CCD camera. The computer can analyse the images for mode content or generate the aberration-canceling phase mask. Software communicates a binary mask file to the DMD, called the DMD mask, for arbitrary field generation.




Figure 3.8: An example of the window functions used in calibrating the DMD along with the associated interference patterns in the far field. a) Schematic of two activated regions containing a phase-coherent grating. In the actual DMD, the 912×1140 pixels are grouped into 1,440 rectangular regions, each of which is 19×38 pixels in extent. b) The interference pattern generated in the Fourier plane by a DMD programmed with the above mask. The phase of the fringe pattern reflects the $\phi_{\Delta}(\Delta x, \Delta y)$ accumulated due to aberrations. c) The real and d) imaginary parts of the numerical Fourier transform of the interference pattern shown in panel b). The function describing the fringes in panel b) is $\cos^2(k_1 x' + k_2 y')$, meaning that in the Fourier space of panels c) and d), the spots are at $\pm(k_1, k_2)$ and $(0, 0)$. The orientation of the fringes in each spot in c) and d) reflects the real space orientation of the DMD grating in panel a). For ease, the phase $\phi_{\Delta}(\Delta x, \Delta y)$ is found not by analysing the fringes in panel b), but by taking the difference in phase between the pair of fringes marked by white circles in panel c) and d).

3.3 DMD Calibration and Usage

The DMD is positioned at one plane of a $2f$ imaging system; in the paraxial limit, the electric field in this plane is connected to that in another plane by a Fourier transform. See Fig. 3.7. This transform scales one set of spatial dimensions by k/f with respect to the other, where $k = 2\pi/\lambda$ is the wavenumber of the laser field. As such, a delta function in one plane $\delta(x - x_0)\delta(y - y_0)$ corresponds to a plane wave in the other plane $\exp[ik(x'x_0 + y'y_0)/f]$, where (x, y) and (x', y') are pixel coordinates in the DMD and image planes, respectively. It is this duality which allows us to use the DMD to generate arbitrary field profiles in the image plane using one of several algorithms known from digital holography [60, 61].

A necessary step before the DMD can be used to accurately produce field profiles is the measurement of aberrations in the optical system, including the DMD itself [62]. The protective glass of the DMD is known to be under stress from manufacturing and thus exhibits non-uniform strain. It is possible to characterize these aberrations by measuring the phase shift imparted onto a set of orthogonal basis functions generated at the DMD plane [63]. Locally adding phase shifts in the DMD plane provides full phase control of the holographic image for either aberration removal, phase-sensitive image construction, or both simultaneously.

For a binary amplitude mask like a DMD, one way to add phase to the field in the image plane is to include a grating into the mask design. This is accomplished by using the DMD mask to modulate the amplitude of the field at the DMD plane via the grating function $\cos^2(k_x x + k_y y + \phi_m(x, y))$, where ϕ_m contains all the phase information of the mode and $k_x = kx_0/f$ and $k_y = ky_0/f$. For simplicity, we choose approximate delta functions as our set of nearly-orthogonal basis functions since they are both

orthogonal in the DMD plane and in any downstream plane. We approximate delta functions by partitioning the 912×1140 mirror-pixels of the DMD into 1,440 individual 48×30 rectangular regions. Examples of two such windows are shown in Fig. 3.8a. The finite size of these window regions provides space within which to overlay phase-coherent grating functions. The specific size is chosen as a balance between making them small enough to provide many partitions of the full DMD array—thereby providing high resolution for aberration correction—while also ensuring that each window is not too small, such that the transform-limited spot-size is not too large in the Fourier-image plane. Setting $\phi_m(x, y) = 0$ for the present discussion, each orthogonal basis function has the approximate form of $\cos^2\left(\frac{k}{f}(xx_0 + yy_0)\right)\Pi\left(\frac{x-x_c}{w}\right)\Pi\left(\frac{y-y_c}{h}\right)$ at the DMD plane, where $\Pi(x)$ is the rectangle function, (x_c, y_c) is the center of the region, and w and h are the width and height of the region, respectively. Turning on any two windows—defined by differing (x_c, y_c) —yields two interference patterns in the image plane centered at $\pm(2x_0, 2y_0)$, as shown in Fig. 3.8b. Each of these has the form $\cos^2\left(\frac{k}{f}(x'\Delta x + y'\Delta y)\right)$, where Δx and Δy are the separation in x and y in the DMD plane between the two centers of the window regions.

Aberrations in the imaging system are determined by measuring the phase shift in the image plane acquired by light reflected from these window regions in the DMD plane. In the presence of aberrations in the optical path, the interference pattern shown in Fig. 3.8b is modified to be $\cos^2\left(\frac{k}{f}(x'\Delta x + y'\Delta y) + \phi_\Delta(\Delta x, \Delta y)\right)$, where the phase $\phi_\Delta(\Delta x, \Delta y)$ is that which is accumulated due to the particular optical aberration encountered by the light propagating from the window region to the image plane. The phase of this pattern is easily obtained by taking the Fourier transform of the interference pattern, which produces three complex-valued intensity features.

These can be split into a real-valued image and an imaginary-valued image, as shown in Figs. 3.8c and d. The phase $\phi_{\Delta}(\Delta x, \Delta y)$ is found by taking the difference in phase between the pair of fringes marked by white circles in Figs. 3.8c and d. The $\phi_{\Delta}(\Delta x, \Delta y)$ is then measured for all windows tiling the DMD pixel area. We employ this procedure by fixing one window in the center of the DMD pixel array while another window is sequentially tiled over all other window regions. Together this allows one to determine a phase mask that is a full representation of the linear aberrations of the optical system. The inverse of such a mask can be superimposed onto any hologram mask programmed into the DMD for the purpose of cancelling these aberrations. In this fashion we are able to produce high-fidelity representations of desired field mode profiles.

3.4 Experimental outline

The DMD used is a Texas Instruments DLP4500NIR consisting of 912×1140 individually addressable square aluminum mirrors, which measure $\sim 10.8 \mu\text{m}$ on the diagonal. They are arranged on a square lattice and can be independently controlled by sending an array of 1's ('on') and 0's ('off') to the controller. Physically, these values correspond to the mirrors changing their angle by $\pm 12^\circ$ with respect to the substrate on which they are fabricated. Light reflected from 'off' mirrors is directed into a beam-blocker, while light reflected from the 'on' mirrors is used to form the desired hologram and cavity in-coupling light. We generate a set of DMD masks for each experiment we wish to perform; the rate at which we can update holograms is presently limited by the 60-Hz refresh rate of the software driver. The laser beam incident on the DMD surface is magnified to a $1/e$ radius of 8 mm, quite a bit larger

than the 7-mm extent of the 45° -projected DMD surface. This insures that the DMD is exposed to an homogeneous light intensity across its surface, which we verify by imaging the surface.

The 2-cm cavity mirrors are glued onto a rotationally symmetric brass mount for temperature and vibration stability. The cavity length can be adjusted via the turning of an 80-turns-per-inch screw that holds one of the two mirrors. This screw provides course control of the cavity length: $4.4 \mu\text{m}/5^\circ$. Glued in between the screw and the mirror is a single-crystal PZT that can be voltage-tuned through 210 nm. The lens connecting the two Fourier planes has a focal length of $f = 200$ mm, chosen to maximize the DMD area used in producing the HG fields that will be mode-matched to the cavity. Figure 3.7 presents a diagram of the setup used to couple holograms generated by the DMD into the cavity. The finesse of the cavity is $\sim 5 \times 10^4$ at the wavelength we use for this work, 780 nm, resulting in a cavity linewidth of ~ 300 kHz.

3.5 Designing a DMD mask for a particular HG mode

The HG modes are a good basis set for DMD studies because they are eigenfunctions of the Fourier transform. That is, if the field created by the mask at the DMD plane has the form of an HG mode, then the field in the associated Fourier image plane will be a scaled version of the same HG mode [64].

To couple light to the cavity, it is necessary to create HG modes with the appropriate scale using the DMD. We first choose a target HG mode $\Phi_{HG}(x, y)$ normalized such that $\max |\Phi_{HG}(x, y)| = 1$. The phase of the mode $\phi_{HG}(x, y)$ is superimposed




Figure 3.9: Cavity transmission versus pump frequency. a) Coupling to the cavity (3,1) mode. Native mixing in the cavity results in mixing of the ideal cavity modes such that the (3,1) does not equal the ideal $\text{HG}_{3,1}$. The red transmission curve results from in-coupling an ideal $\text{HG}_{3,1}$ mode into the cavity. The imperfect overlap induces a coupling to the adjacent modes in addition to the targeted (3,1). The blue curves show the result of increasing optimization of the DMD mask, resulting in a hologram that couples exclusively to the cavity (3,1) mode, which is shown in the inset. We determine that the cavity (3,1) mode roughly consists of 8% $\text{HG}_{2,2}$, 8% $\text{HG}_{4,0}$, and 84% $\text{HG}_{3,1}$. b) Same as in panel a), but with respect to the (2,1) mode. We determine the cavity (2,1) mode roughly consists of 4% $\text{HG}_{1,2}$, 4% $\text{HG}_{3,0}$, and 92% $\text{HG}_{2,1}$.




Figure 3.10: Cavity transmission versus pump frequency for different DMD masks. Each trace corresponds to a distinct mask generating a unique and orthogonal hologram on the DMD. a) Even mode families. Insets show the cavity transmission for the 2nd and 4th-order families. The feature above the noise floor marked ‘reference’ was created with a superimposed Gaussian beam in order to align the transmission curves to one another. b) Same as in panel a), but odd mode families. Insets show the cavity transmission for the 3rd and 5th-order families of modes.

with the aberration-correcting phase such that the phase of the hologram to be programmed into the DMD is $\phi_m(x, y) = \frac{1}{2}(\phi_{HG}(x, y) - \phi_{\Delta}(x, y))$. Both phases are normalized to run between 0 and 2π . Including the grating structure necessary to create the HG mode in the Fourier image plane, the final field to be produced by the DMD mask is

$$E_{DMD} = |\Phi_{HG}(x, y)| \cos \left(\frac{k}{f}[xx_0 + yy_0 + \frac{z_0}{2f}C(x, y)] + \phi_m(x, y) \right)^2, \quad (3.2)$$

where kx_0/f and ky_0/f determine the periodicity of the grating in the DMD plane and thus the center location of the mode in the image plane at $\pm(2x_0, 2y_0)$. z_0 is a focusing parameter and $C(x, y)$ is a focusing field that is in the form of a paraboloid that corrects for the 45° -angle of the DMD face. The function in Eqn. 3.2 is sampled using Floyd-Steinberg dithering with a two-color palette to numerically develop a pixel-by-pixel representation of this field at the DMD surface for programming into the controller. A final reshaping is necessary to generate the file that correctly maps pixel values to DMD mirrors. This file is then loaded into the DMD controller. The whole process takes only 30 s per mode-image using a commercial laptop running code written in Python.

3.6 Mode-coupling procedure

To couple a mode into the cavity, we first choose x_0 and y_0 to be large enough to sufficiently separate the grating orders in the Fourier image plane. The first-order peak contains the mode image to be coupled and is directed into the cavity while the other orders are blocked. Transmission through the cavity is maximized using two

mirrors placed immediately before the cavity. The Gaussian waist of the incoming mode and z_0 are optimized through manual iteration to maximize the coupling of the mode. For simplicity, we typically use the fundamental $\text{HG}_{0,0}$ mode for this procedure, which fixes all but one of the cavity parameters. The indices are the number of nodes in the i and j axes, respectively.

The remaining parameter, the azimuthal angle of the nodal axes in the transverse plane, is determined manually by observing which modes in a set of non-azimuthally symmetric modes maximize coupling. For simplicity, we use a collection of $\text{HG}_{1,0}$ beams at different azimuthal angles θ , finally choosing the θ corresponding to the $\text{HG}_{1,0}$ mode exhibiting the best coupling efficiency.

This procedure would, in principle, fully determine the HG mode-basis of the cavity if the cavity conformed to the ideal paraxial model. However, intrafamily mixing is present, as mentioned above, and the actual cavity modes do not always correspond to ideal HG-mode shapes. An example of this can be seen in the insets of Fig. 3.9. To label the modes comprising the actual mode spectrum, we number each mode of an $n + m = N$ family by indices (α, β) , where α and β no longer necessarily correspond to the number of nodes along a direction i and j , respectfully. However, the condition $\alpha + \beta = N$ remains fixed for each mode family N .³ Moreover, we choose to denote the highest-frequency mode in a family the $(0, N)$ mode, rather than the $(N, 0)$ mode.⁴

³We note that the 210-MHz splitting between families ensures that inter-family mixing is negligible, thereby ensuring that $\alpha + \beta = n + m = N$ is a good approximation.

⁴Among the mode families we explore here, the modes with N nodes along only one direction typically possess the highest resonance frequency.




Figure 3.11: Spectroscopy of the cavity in a near-degenerate configuration wherein the first even mode becomes degenerate with the fundamental. Inset: zoomed-in spectrum of the first clump of degenerate modes.




Figure 3.12: Cavity transmission when the cavity is pumped by fields created by DMD masks of ideal HG-modes. a) Even mode spectra. b) Odd mode spectra.

3.7 Single-mode spectroscopy

The single-mode character of the cavity is probed by setting the length of the cavity to be 1-mm-longer than the confocality condition. While most modes in a given family possess a spatial structure very close to an ideal HG-mode, we do find at least one per family that significantly deviates from the HG-modal structure. For example, in the $N = 3$ and $N = 4$ families, the $(2, 1)$ and $(3, 1)$ modes both show significant deviation from the ideal $HG_{2,1}$ and $HG_{3,1}$ modes, respectively, as can be seen from the insets of Fig. 3.9. We determine the HG-modal decompositions through a 2D fitting routine, and the fraction of other intrafamily HG modes is listed in the captions.

Figure 3.9 shows the results of an iterative procedure used to probe the $(2, 1)$ and $(3, 1)$ modes. In each case, attempting to couple to the mode with an ideal HG resulted in cross-coupling to the immediately adjacent modes. The optimized masks resulted primarily of the HG mode in question with small (5%-10%) admixtures of the

adjacent HG modes. Each mask was rotated to an angle θ optimal for mode coupling. Though there can be much mixing of HG modes, we are able to successfully generate holograms that couple to the (2,1) and (3,1) cavity modes with high specificity. As seen in Fig. 3.9, all transmission in non-targeted modes are below the level of our detection signal-to-noise. Although optimizing by hand is quite straightforward, an automated procedure can be applied to determine the eigenmodes of the cavity more quickly.

Other than the (2,1) and (3,1) modes, all other modes in these families were efficiently coupled to by DMD masks employing the ideal HG mode functions. The result of a spectroscopic scan of the low-order odd and even mode families is shown in Fig. 3.10. Each trace corresponds to a separate DMD mask designed to couple to a specific single mode of the cavity. Our manual optimization routine ensures that unwanted cross-coupling is minimal and indeed not apparent in these traces. All modes emitted out of the cavity are nearly identical to ideal HG modes, with the exception of the two aforementioned mixtures. We note that the rapidity of the frequency sweep results in an asymmetric lineshape for each mode due to ringdown effects.

3.8 Multimode Spectroscopy

We experimentally probe the basic properties of the near-degenerate cavity at hand, because a numerical characterization [65] is impractical due to the need to minutely characterize the mode mixing arising from mirror defects, etc. To proceed, the cavity is tuned to a near-degenerate configuration wherein the length was adjusted until the nearest second-order mode (0,2) became degenerate with the fundamental




Figure 3.13: a) Four transmission images of the near-degenerate cavity when pumped at a particular frequency by a mask consisting of an increasing amount of an $HG_{2,0}$ mode added to an $HG_{0,0}$ mode. Each image is the light emitted from the cavity and is the coherent addition of the two modes at a single probe frequency. b) Same as panel a), but with an increasing amount of an $HG_{3,0}$ mode added to an $HG_{1,0}$ mode.

(0,0).⁵ Figure 3.11 illustrates the character of the degeneracy in this configuration. To allow the observation of the frequency spacing of modes, the cavity was pumped with a broad beam slightly off-center. The modes are seen to arrange into clusters in this near-confocal configuration. It is possible to study some aspects of cavity degeneracy in this configuration since many modes appear with separations of less than a few linewidths.

The transmission of the cavity was measured for each ideal HG mask as before. Figure 3.12 shows the cavity transmission for each mask. As expected, there is a now a greater degree of mode mixing, especially in the higher order modes, as can be seen by the prevalence of side-peaks at many distinct probe frequencies. Nevertheless, each ideal HG mask couples to a small collection of cavity modes within its spectral vicinity. We note that the two lowest-order even and odd modes—blue traces at the bottom of each figure—can be selectively addressed by distinct, orthogonal HG masks. Finally, to demonstrate that the near-degenerate cavity can be pumped by tailored field modes, we show in Fig. 3.13 that any linear combination of, e.g., $\text{HG}_{0,0}$ with $\text{HG}_{2,0}$ or $\text{HG}_{1,0}$ with $\text{HG}_{3,0}$, can be coupled into the near-degenerate cavity.

⁵Due to misalignment of the cavity mounts, this cavity length was the point at which the maximum mode density was achieved. Other cavities can achieve far greater mode density, e.g., see Ref. [66].

Chapter 4

Cavity Quantum Electrodynamics

4.1 Overview

In 1946 Edward Purcell performed a simple experiment monitoring the excited state population of a collection of nuclear spins [67]. Placing the spin in the vicinity of a near-resonant radio-frequency cavity had the effect of reducing the particle's excited state lifetime. The eponymous Purcell effect is due to a modification of the electromagnetic mode density in the vicinity of the emitter, in effect changing the collection of modes to which the emitter is coupled and thereby affecting the rate of spontaneous emission. Broadly speaking, incoherent processes are the result of dephasing among a multitude of coherent processes. An atom's ability to coherently exchange energy in the form of photons with any of a multitude of electromagnetic modes results in a wash-out of the Rabi oscillations that are the hallmark of a quantum process. The effective absence of coherence effects a crossover where rate equations become descriptive of the atom's excited and ground state populations.

Wigner and Weiskopf [68] calculated the incoherent decay rate of a two-level system coupled to a radiative bath using Fermi's golden rule. The rate is proportional to $\sum_{ks} \left| \sqrt{\frac{\omega_k}{2\hbar\epsilon_0 V}} (\mathbf{d} \cdot \boldsymbol{\epsilon}_{ks}) \right|^2$, where the sum is over all free-space electromagnetic modes, where ω_k is the frequency of the mode, V is the mode volume, $\mathbf{d}$ is the coupling operator, and $\boldsymbol{\epsilon}_{ks}$ is the polarization vector of the light field. For a dipole emitter in three dimensions, this summation yields the spontaneous decay rate Γ . Alkali atoms behave approximately as two level systems, and as such an atom in its excited state is stimulated to decay into its ground state by vacuum fluctuations in the bath of electromagnetic modes to which it is coupled. The rate of spontaneous emission for ^{87}Rb in any of its $5P_{3/2}$ excited states is $\sim \Gamma = 2\pi \times 6.2 \text{ MHz}$.

Any modification of the free-space electromagnetic spectrum, in the frequency or spatial structure of the modes, will lead to a modification of the spontaneous emission rate. If the electromagnetic environment can be engineered to include a single mode more strongly coupled than any other mode to the emitter, that single mode can be involved in fundamental studies of the light-matter interaction, and can be used to engineer novel quantum devices [69]. In practice this is accomplished by placing the atom near the field maximum of the mode of an optical resonator, e.g. a Fabry-Pérot cavity. The resonator has been engineered to harbor a mode whose frequency is near the atomic resonance and whose loss rate κ is manageable. The field-enhancement exhibited by the cavity concomitantly enhances the coupling of vacuum fluctuations to the atom's induced dipole moment, and its coupling to the cavity mode is $g(\mathbf{x}) = \sqrt{\frac{3c\Gamma\lambda^2}{8\pi V}} \Xi(\mathbf{x})$, where c is the speed of light, λ is the wavelength of the two-level transition, and $\Xi(\mathbf{x})$ is the cavity mode normalized to have an extremal value of magnitude 1. For an atom optimally coupled to the cavity mode, it is conventional

to refer to the coupling in terms of the vacuum Rabi splitting $2g_0 = g(\mathbf{x}_{max})$, where $\mathbf{x}_{max}$ is the position that maximizes the magnitude of $g(\mathbf{x})$.

The Hamiltonian that describes the coupling of an atom's dipole moment to a mode of the electromagnetic field is the quantized version of $H_{cl} = \vec{d} \cdot \vec{E}$, the classical dipole interaction. The dipole moment operator is $\propto e\mathbf{r}$ [70], a function of odd parity. In the basis of an atom's ground and excited state wavefunctions (which are themselves eigenstates of parity), the representation of this operator will only have off-diagonal elements and can be simply expressed in this two-level scheme in terms of the Pauli matrices σ_x and σ_y . Expressing the electric field operator in terms of its annihilation and creation operators, the interaction Hamiltonian for a dipole interacting with a single electromagnetic mode is $H_{int} = \frac{g_0}{2}\sigma_x(a + a^\dagger)$. For induced dipole moments in atoms, the separation in energy between the ground and excited states will be many orders of magnitude greater than the coupling g_0 , allowing for the application of the rotating wave approximation (RWA), and the elimination of the counter-rotating or energy non-conserving terms. Finally, including the energy of the bare atomic and cavity excitations yields the celebrated Jaynes-Cummings (JC) Hamiltonian

$$H_{JC} = \omega_c a^\dagger a + \omega_a \sigma^\dagger \sigma + \frac{g_0}{2}(a\sigma^\dagger + a^\dagger\sigma) \quad (4.1)$$

In the strong-coupling regime when $g \gg \kappa, \Gamma$, the spectrum of this Hamiltonian becomes visible in experiment [71, 72], and provides theoretical support for the first evidence of electric field quantization [73]. The Hamiltonian can be block-diagonalized in the basis of collective excitations n of the atom and light field. First, the diagonal matrix $\omega_a(a^\dagger a + \sigma^\dagger \sigma)$ is subtracted and the JC Hamiltonian is rewritten

$$H_{JC} = \Delta_c a^\dagger a + \frac{g_0}{2}(a\sigma^\dagger + a^\dagger\sigma) \quad (4.2)$$

Here $\Delta_c = (\omega_c - \omega_a)$. In each n -excitation manifold, where $n \geq 1$, the dressed states are $|n, +\rangle = \cos \theta_n |e, n-1\rangle + \sin \theta_n |g, n\rangle$ and $|n, -\rangle = \sin \theta_n |e, n-1\rangle - \cos \theta_n |g, n\rangle$, with $\tan \theta_n = \sqrt{\frac{\sqrt{ng_0^2 + \Delta_c^2} + \Delta_c}{\sqrt{ng_0^2 + \Delta_c^2} - \Delta_c}}$. The eigenvalues for these states are $E_n^\pm = n\omega_c - \frac{1}{2}\Delta_c \pm \frac{1}{2}\sqrt{ng_0^2 + \Delta_c^2}$, with respect to the energy of $|g, 0\rangle$. Compared to the bare-cavity spectrum $E_n = \hbar\omega_c n$ the atom provides a significant nonlinearity for low energy manifolds, leading to the vacuum Rabi splitting observed in the low-power response of the cavity-atom system. This nonlinearity disappears at higher manifolds as the separation in energy between $|n, +\rangle$ and $|n+1, +\rangle$ becomes $\omega_c + g_0/2\sqrt{n}$, and thus independent of n as $n \rightarrow \infty$. In practice the nonlinearity disappears at high n as the laser or atomic linewidths become on the order of $g_0/\sqrt{n}$. Figure 4.1 shows a model transmission spectrum for a cavity coupled strongly to a two-level single atom on resonance. The hybridization of the cavity and atom degrees of freedom results in eigenstates $|+\rangle$ and $|-\rangle$ that each have cavity character, and can thus be probed independently. The transmission peaks occur roughly at $\pm g$ with respect to the atom and cavity resonance, reflecting the JC eigenvalues. The transmission peaks are broadened with respect to the bare cavity since the atomic loss channel (the rate of spontaneous emission is perhaps on order of magnitude greater than the rate of cavity leakage) is now available to the cavity states.

There are two natural ways to extend the JC model: increasing the number of atoms and increasing the number of cavity modes. Increasing the number of atoms




Figure 4.1: Modeled cavity transmission showing vacuum Rabi splitting (red), against bare cavity transmission (blue). The cavity and atom parameters are $\kappa = g/50$ and $\Gamma = g/5$.

leads to the Dicke model with the counter-rotating terms in place, and the Tavis-Cummings model without. The Tavis-Cummings model displays similar phenomenology as the JC model in the lowest manifold by replacing g_0 with $g_0\sqrt{N}$, and $|e\rangle$ with $\sum_i \frac{1}{\sqrt{N}}|e_i\rangle$, the atomic excitation coherently distributed over the entire cloud of N atoms. However, the observed splitting is not a quantum phenomenon in that it can be explained by classical dispersion theory [74]. The Dicke model has long been known to harbor a superradiant phase transition [75] when $g_0/\sqrt{\omega_c\omega_a} > \frac{1}{2}$, wherein the atoms exhibit a collective dipole moment. Not only is g_0 usually many orders of magnitude (~ 9) below ω_c and ω_a in atom-cavity systems, but it is in fact limited by a dipole sum rule to lie below the threshold condition [76]. Relevant here is the subtle fact that a term $\propto A^2$ has been ignored, where $\vec{A}$ is the vector potential of the electromagnetic field. This term is necessary to fully describe the light-matter interaction, and stems from a gauge transformation that is made when treating a quantum particle in an electromagnetic field, namely $\vec{\nabla} \rightarrow \vec{\nabla} - \frac{e}{c}\vec{A}$ [77]. This is an aside, and the relevant fact is that several schemes have been developed in recent years that decouple g_0 from ω_a , thereby circumventing the sum rule [78–81] and making observations of the mean-field Dicke transition feasible.

4.2 Self-Organization

Remarkably, increasing the available cavity modes such that they form a complete basis is theorized to result in a system with a quantum phase transition [82]. Our analysis of this cavity-atom system begins with the system Hamiltonian which consists of the following parts written in second-quantized form [83, 84]

$$\begin{aligned}
H_a &= \int d^3\mathbf{x} \psi_e^\dagger(\mathbf{x}) \left(\frac{\mathbf{p}^2}{2m} + \omega_a + V_e(\mathbf{x}) \right) \psi_e(\mathbf{x}) + \int d^3\mathbf{x} \psi_g^\dagger(\mathbf{x}) \left(\frac{\mathbf{p}^2}{2m} + V_g(\mathbf{x}) \right) \psi_g(\mathbf{x}) \\
H_c &= \sum_{\alpha} \omega_{c,\alpha} a_{\alpha}^{\dagger} a_{\alpha} \\
H_{cp} &= \sum_{\alpha} \Omega_{c,\alpha} (a_{\alpha} e^{i\mu_{\alpha} t} + a_{\alpha}^{\dagger} e^{-i\mu_{\alpha} t}) \\
H_{aa,g} &= \frac{U_{gg}}{2} \int d^3\mathbf{x} \psi_g^\dagger(\mathbf{x}) \psi_g^\dagger(\mathbf{x}) \psi_g(\mathbf{x}) \psi_g(\mathbf{x}) \\
H_{aa,e} &= \int d^3\mathbf{x} \left(\frac{U_{ge}}{2} \psi_g^\dagger(\mathbf{x}) \psi_g^\dagger(\mathbf{x}) \psi_e(\mathbf{x}) \psi_e(\mathbf{x}) + h.c. + \frac{U_{ee}}{2} \psi_e^\dagger(\mathbf{x}) \psi_e^\dagger(\mathbf{x}) \psi_e(\mathbf{x}) \psi_e(\mathbf{x}) \right) \\
H_{ap} &= \int d^3\mathbf{x} \left(\psi_e^\dagger(\mathbf{x}) \eta(\mathbf{x}) \psi_g(\mathbf{x}) e^{-i\nu t} + \psi_g^\dagger(\mathbf{x}) \eta^*(\mathbf{x}) \psi_e(\mathbf{x}) e^{i\nu t} \right) \\
H_{ac} &= \sum_{\alpha} \int d^3\mathbf{x} \left(\psi_e^\dagger(\mathbf{x}) g_{\alpha}(\mathbf{x}) \psi_g(\mathbf{x}) a_{\alpha} + \psi_g^\dagger(\mathbf{x}) g_{\alpha}^*(\mathbf{x}) \psi_e(\mathbf{x}) a_{\alpha}^{\dagger} \right)
\end{aligned}$$

Here we assume the atom is a two-level system with ground and excited states separated in energy by ω_a . In each expression $\psi_i(\mathbf{x})$ is the annihilation operator for an atom in its ground or excited state located at $\mathbf{x}$, and obeys the standard commutation relation $[\psi_i(\mathbf{x}), \psi_j^\dagger(\mathbf{y})] = \delta(\mathbf{x} - \mathbf{y}) \delta_{ij}$ where $i \in g, e$. a_{α} is the photon annihilation operator for cavity mode α which has spatial structure $\propto g_{\alpha}(\mathbf{x})$. H_a is the atoms-only terms, which include the kinetic and potential energy of the ground and excited state atoms. H_c is the cavity-only terms, consisting of the bare energy of the orthogonal cavity modes. H_{cp} is a semi-classical description of a set of external fields with Rabi frequency $\Omega_{c,\alpha}$ and frequency μ_{α} pumping each cavity mode. H_{aa} describes the contact interaction among the atoms, and is separated into two parts; the $H_{aa,g}$ is the mean-field energy due to contact interactions among the ground state atoms, and $H_{aa,e}$ is all interactions involving the excited state atoms. H_{ap} is the semi-classical pump of the atoms, where the monochromatic pump field has frequency ν and spatial profile $\eta(\mathbf{x})$. Finally, H_{ac} is the atom-cavity interaction with each cavity

mode. This term exchanges energy between the cavity and atomic excited state, and as a second order process allows for pump light directed at the atoms to scatter into the cavity modes.

In our studies we focus on pumping either the atoms or the cavity with a monochromatic light field. Future work will see the cavity pumped longitudinally and seeded with light while the pump is brought above threshold in order to study self-organization in the presence of a bias field. For now, we set either $\Omega_{c,\alpha}$ or $\eta(\mathbf{x})$ to 0. Transforming to a frame rotating with the cavity or atom pump frequency is accomplished by applying a unitary transformation to the Hamiltonian $H \rightarrow -iU^\dagger \dot{U} + U^\dagger H U$, with $U = \exp\left(i\nu t(\int d^3\mathbf{x} \psi_e^\dagger(\mathbf{x})\psi_e(\mathbf{x}) + \sum_\alpha a_\alpha^\dagger a_\alpha)\right)$, and has the effect of changing the frequencies ω_a and $\omega_{c,\alpha}$ into detunings Δ_a and $\Delta_{c,\alpha}$. Finally, by putting the cavity resonances and pump far away from the atomic resonance we can adiabatically eliminate the excited state, yielding an approximate algebraic relationship $\psi_e(\mathbf{x}) = -\frac{1}{\Delta_a}(\eta(\mathbf{x}) + \sum_\alpha g_\alpha(\mathbf{x})a_\alpha)\psi_g(\mathbf{x})$. Replacing $\psi_e(\mathbf{x})$ this way leads to the effective Hamiltonian for our problem:

$$H = \sum_\alpha \Delta_{c,\alpha} a_\alpha^\dagger a_\alpha + \int d^3\mathbf{x} \psi^\dagger(\mathbf{x}) \left(\frac{\mathbf{p}^2}{2m} + V(\mathbf{x}) + \frac{U}{2} \psi^\dagger(\mathbf{x})\psi(\mathbf{x}) \right) \psi(\mathbf{x}) \\ - \frac{1}{\Delta_a} \int d^3\mathbf{x} \psi^\dagger(\mathbf{x}) \left(\sum_{\alpha,\beta} g_\alpha^*(\mathbf{x})g_\beta(\mathbf{x})a_\alpha^\dagger a_\beta + \sum_\alpha \left(\eta(\mathbf{x})g_\alpha^*(\mathbf{x})a_\alpha^\dagger + \eta^*(\mathbf{x})g_\alpha(\mathbf{x})a_\alpha \right) + \eta^*(\mathbf{x})\eta(\mathbf{x}) \right) \psi(\mathbf{x})$$

Far away from resonance the pump light scattered by atoms is almost entirely in phase with the pump, with a dipole spatial profile whose node is parallel with the polarization of the pump light. The atoms behave like classical scatterers in this Rayleigh scattering regime, that is they behave like a classical refractive material.

From here on we will drop the ‘g’ subscript that distinguishes ground from excited state terms in the Hamiltonian, there being no reference to the excited state atoms whatsoever. It is important to realize that this Hamiltonian describes a single atom; entanglement between atoms in a cloud cannot be captured by this model. Having framed the problem we can proceed to obtain some intuition for the underlying phenomena by considering simple cases. For instance, consider a static atomic distribution and no atom pump $\eta(\mathbf{x}) = 0$; in the absence of atomic motion the atoms-only terms provide a constant offset to the Hamiltonian and do not contribute to any observable property. This situation arises when the system is pumped below threshold and the atoms are held in place by an external potential, for instance an optical dipole trap. The relevant part of the Hamiltonian is the dispersive mode-mixing term

$$H_{\text{mix}} = \sum_{\alpha} \Delta_{c,\alpha} a_{\alpha}^{\dagger} a_{\alpha} - \frac{1}{\Delta_a} \sum_{\alpha,\beta} a_{\alpha}^{\dagger} a_{\beta} \int d^3\mathbf{x} \rho(\mathbf{x}) g_{\alpha}^*(\mathbf{x}) g_{\beta}(\mathbf{x})$$

Here $\rho(\mathbf{x}) = \psi_g^{\dagger}(\mathbf{x})\psi_g(\mathbf{x})$ is the density of the gas, and will be given by the solution to the atoms-only Hamiltonian with the external trapping potential $V(\mathbf{x})$. For a harmonic trap the density is described to an extremely high degree by the Thomas-Fermi approximation, $\rho(\mathbf{x}) = \frac{\mu - \frac{m}{2} \sum_i \omega_i^2 x_i^2}{U}$. The chemical potential μ is given in this approximation by $\mu = \left(\frac{225m^3 N^2 U^2 \bar{\omega}^6}{512\pi^2} \right)^{\frac{1}{5}}$ [85], where m is the atomic mass, and $\bar{\omega}$ is the geometric mean of the trap frequencies. The atoms couple the cavity modes, causing population in one mode to scatter into the others. An equivalent picture is to note that the cavity eigenstates are now admixtures of the bare states, and occupation of a bare eigenstate constitutes in this context occupation of many actual eigenstates, leading to sloshing in the dynamics. We can more clearly appreciate this by expressing this Hamiltonian in the basis of single excitations of the vacuum $|\Omega\rangle$, e.g. $|\alpha\rangle = a_{\alpha}^{\dagger}|\Omega\rangle$,

wherein the matrix elements are

$$H_{\text{mix}}^{\alpha\beta} = \Delta_{c,\alpha} \delta_{\alpha\beta} - \frac{1}{\Delta_a} \int d^3\mathbf{x} \rho(\mathbf{x}) g_\alpha^*(\mathbf{x}) g_\beta(\mathbf{x})$$

The bare cavity frequencies $\Delta_{c,\alpha}$ are given to lowest order by paraxial wave theory. Coupling of the ideal HG modes leading to renormalization of the mode resonances is caused by a number of factors including complications in mounting the mirrors, mirror aberrations, beyond-paraxial effects, and the existence of dirt or other scattering centers on the mirror surface. In any case, the atomic density acting as a dispersive element can mix the ideal modes considerably. This mixing can be observed in low power spectroscopy of the cavity-atom system as a shift of the cavity wavelengths, and as a modification of the spatial profile of the cavity eigenstates. A pathological but pedagogically useful case to consider is that a delta function atomic density $\rho(\mathbf{x}) = N\delta(\mathbf{x} - \mathbf{x}_0)$, and all the cavity modes degenerate, $\Delta_{c,\alpha} = \Delta_c$. In this case the Hamiltonian reduces to

$$H_{\text{mix},\delta} = -\frac{N}{\Delta_a} \vec{g}^*(\mathbf{x}_0) \otimes \vec{g}(\mathbf{x}_0)$$

where $\vec{g}(\mathbf{x}_0)$ is the vector of mode functions $g_\alpha(\mathbf{x})$ evaluated at the location of the atoms $\mathbf{x}_0$. Notable about this matrix is the existence of a single non-zero eigenvalue of magnitude $N\vec{g}^*(\mathbf{x}_0) \cdot \vec{g}(\mathbf{x}_0)/\Delta_a$, with associated eigenvector $\vec{g}^*(\mathbf{x}_0)$ [86]. The resulting field profile $\propto \sum_\alpha g_\alpha^*(\mathbf{x}_0) g_\alpha(\mathbf{x})$ maximizes the amplitude at the location of the atoms, almost as though the cavity is trying to envelop and protect the atoms. Unsurprisingly, increasing the number of modes increases the localization of the cavity field, and if the cavity were to harbor a complete basis of modes the resulting cavity mode would

also be a delta function. Such a tight supermode can never be formed in practice; the finite size of the cavity mirrors effectively cuts off the supported modes above a certain order, and the finite width of the trap prevents a true point-localization of the atoms. Figure 4.2 shows the increasing localization of the supermode as the number of cavity modes increases. Considering only even modes, when only modes up to 4th order are included, the supermode consists almost exclusively of the $\text{HG}_{4,0}$ since the atoms lie just outside the range of the 4th order modes. Including up to 8th order the resulting supermode is clearly localized on the location of the atoms, and this localization only improves when modes up to 16th order are included.

The self-organization can be studied in greater detail by first focusing on the single mode case of the effective Hamiltonian when the atoms are pumped transversely. The pump light is scattered into the cavity by the atoms, and the combination of the pump field and the cavity field results in a self-sustaining light-shift-induced potential that maintains the atoms in a configuration that increases the rate at which they scatter light into the cavity field, deepening the potential that holds them there, etc. There are three terms that involved in this positive-feedback phenomenon.

The term $\int d^3\mathbf{x} \frac{\eta^*(\mathbf{x})\eta(\mathbf{x})}{\Delta_a} \psi^\dagger(\mathbf{x})\psi(\mathbf{x})$ is due to the pump field, and for a standing wave pump field results in a 1D optical lattice along the direction of the pump beam, which is orthogonal to the cavity axis. $\int d^3\mathbf{x} \frac{g^*(\mathbf{x})g(\mathbf{x})}{\Delta_a} a^\dagger a \psi^\dagger(\mathbf{x})\psi(\mathbf{x})$ is a trapping potential due to the cavity field caused by light scattered into the cavity by the atoms. Finally, $\int d^3\mathbf{x} \frac{\eta(\mathbf{x})g^*(\mathbf{x})}{\Delta_a} a^\dagger \psi^\dagger(\mathbf{x})\psi(\mathbf{x})$ is a potential with spatial profile $\propto \eta(\mathbf{x})g^*(\mathbf{x})$ caused by the interference of the cavity and pump fields. In our experiments the pump $\eta(\mathbf{x})$ is approximately a standing wave $\propto \cos(kx)$ created by retroreflecting a single beam launched into the atomic cloud. The beam is wide enough that the atoms confined




Figure 4.2: The cavity component of the supermode generated by a delta function atomic density. Increased localization is observed by incorporating ideal HG cavity modes up to increasing order.

in an approximately harmonic optical dipole trap do not experience the transverse structure of the pump beam.

In addition to acting as a 1D optical lattice that partially confines the atoms, the pump light is scattered into the cavity generating the interference potential. Again due to the limited extent of the atoms, to a high degree the cavity field looks like $g(\mathbf{x}) \propto \cos(kz)$. The potential due to all three terms is effectively a 2D optical lattice $\propto (A \cos(kx) \pm B \cos(kz))^2$, where the sign is self-consistently determined by the configuration of the atoms.

Consider the lattice sites of the aforementioned potential. If atoms were to inhabit every site, or equivalently if the atomic gas was homogenous over the lattice, the pump light would not be scattered into the cavity; like a phased array consisting of dipole emitters, the atoms' dipole fields would destructively interfere along the cavity axis preventing the buildup of light in the cavity. Seen another way, the atoms couple to the cavity with amplitude $g(\mathbf{x}_i)$, and the sum of the cavity field at all lattice sites is $\sum_n \cos(k\lambda n) = \cos(2\pi n) = 0$. If, however, the atoms inhabited every other lattice site in an alternating fashion they would constructively interfere and scatter light into the cavity, deepening the potential they experienced, etc. The occupation of the 'even' or 'odd' sublattice is an example of a discrete symmetry breaking associated with the Dicke phase transition. This positive feedback leads to a superradiant configuration of the atomic density and cavity light field, where the atoms quickly move to occupy one of the two sublattices at random, in the process rapidly filling the cavity with light coherently scattered by the atoms [87]. The atoms are organized into a 2D square lattice with lattice vectors $\frac{\lambda}{2}(1, 0, 1)$ and $\frac{\lambda}{2}(1, 0, -1)$. The atoms are unconfined in the third direction but for the external potential. This intuitive picture doesn't change

much if the cavity mode changes, but new phenomena arise as the number of cavity modes increases from one to a few, and then from a few to many.

4.3 Gross-Pitaevskii Equation

Due to the mean-field scattering interaction the Hamiltonian needs to be solved self-consistently, and is an example of the Gross-Pitaevskii equation (GPE) [53]. The self-organization can be studied in greater depth by deriving and solving the equations of motion for the atom and cavity fields. The equations of motion for each operator are determined in standard fashion by commuting the operator with the Hamiltonian, the generator of time translations, e.g. $\frac{d}{dt}\psi(\mathbf{x}) = i[H, \psi(\mathbf{x})]$.

$$i\partial_t a_\alpha = \Delta_{c,\alpha} a_\alpha - \frac{1}{\Delta_a} \int d^3\mathbf{x} \psi^\dagger(\mathbf{x}) \left(\eta(\mathbf{x}) g_\alpha^*(\mathbf{x}) + g_\alpha^*(\mathbf{x}) \sum_\beta g_\beta(\mathbf{x}) a_\beta \right) \psi(\mathbf{x}) - i\kappa a_\alpha$$

$$i\partial_t \psi(\mathbf{x}) = \left(\frac{\mathbf{p}^2}{2m} + V(\mathbf{x}) - \frac{\eta^*(\mathbf{x})\eta(\mathbf{x})}{\Delta_a} + U\psi^\dagger(\mathbf{x})\psi(\mathbf{x}) \right) \psi(\mathbf{x}) \\ - \frac{1}{\Delta_a} \left[\sum_{\alpha,\beta} g_\alpha^*(\mathbf{x}) g_\beta(\mathbf{x}) a_\alpha^\dagger a_\beta + \sum_\alpha \left[\eta(\mathbf{x}) g_\alpha^*(\mathbf{x}) a_\alpha^\dagger + \eta^*(\mathbf{x}) g_\alpha(\mathbf{x}) a_\alpha \right] \right] \psi(\mathbf{x})$$

The time evolution of the wavefunction is governed by an effective Hamiltonian $i\partial_t \psi = H_{\text{eff}}(a, a^\dagger)\psi$, and κ , the cavity linewidth, has been added to model the leakage of light from the cavity. Since the time scale of the atomic motion is much longer than the cavity dynamics $\sim 1/\kappa$, we can assume that the cavity field adiabatically follows the atom wavefunction, in effect setting $\partial_t a_\alpha = 0$. Furthermore, it is assumed

that the cavity field is at all times a coherent field, in which case the operator a_α is demoted to a c -number [88]. This assumption prevents the theoretical study of quantum correlations in the light field but allows for a straightforward and quantitatively accurate model of the atomic self-organization threshold, spectrum, and dynamics. Under these assumptions for the light field, the coherent state amplitudes are determined from the atomic density and fields using the following relationships, one for each coherent state amplitude a_α .

$$\frac{1}{\Delta_a} \int d^3\mathbf{x} \rho(\mathbf{x}) \eta(\mathbf{x}) g_\alpha^*(\mathbf{x}) = \Delta_{c,\alpha} a_\alpha - \frac{1}{\Delta_a} \int d^3\mathbf{x} \rho(\mathbf{x}) g_\alpha^*(\mathbf{x}) \sum_\beta g_\beta(\mathbf{x}) a_\beta - i\kappa a_\alpha$$

The vector of coherent state amplitudes can be determined by matrix inversion $a_\alpha = A_{\alpha\beta}^{-1} B_\beta$, where the matrix $A_{\alpha\beta}$ and the vector B_β are defined as follows.

$$\begin{aligned} A_{\alpha\beta} &= (\Delta_{c,\alpha} - i\kappa) \delta_{\alpha\beta} - \frac{1}{\Delta_a} \int d^3\mathbf{x} \rho(\mathbf{x}) g_\alpha^*(\mathbf{x}) g_\beta(\mathbf{x}) \\ B_\alpha &= \frac{1}{\Delta_a} \int d^3\mathbf{x} \rho(\mathbf{x}) \eta(\mathbf{x}) g_\alpha^*(\mathbf{x}) \end{aligned}$$

4.4 Numerical Results

Solving the dynamics of the problem numerically involves stepping through the evolution of the system iteratively, whereby the cavity field and atomic wavefunction are updated sequentially, the former according to the above matrix inversion. The time evolution of the atomic wavefunction is determined by the effective Hamiltonian from earlier, $H_{\text{eff}}(a, a^*)$, and a very simple scheme would have

$\psi(t_n) = \exp(-i\delta t H_{\text{eff}})\psi(t_{n-1})$, where $\delta t = t_n - t_{n-1}$ is the time increment for the simulation that must be small enough to assure the Hamiltonian commutes with itself approximately at adjacent time steps. The problem with this approach is that the Hamiltonian contains functions of canonically conjugate operators $\mathbf{p}$ and $\mathbf{x}$. If the wavefunction is expressed in real-space where $\mathbf{x}$ is a diagonal operator, then the application of the momentum operator is a computationally difficult task that would significantly slow the simulation, and vice-versa.

Decades of widespread utility have made the Fast Fourier Transform an extremely fast and reliable algorithm written into probably every computer language known. The idea of the split-step Fourier Transform technique is to separate the Hamiltonian into momentum and real-space components $H_{\text{eff}} = H_{\mathbf{p}} + H_{\mathbf{x}}$ and to apply each independently, Fourier Transforming the wavefunction at each half step. This assures that the relevant time-evolution operator $\exp(-i\delta t H_{\mathbf{p}})$ or $\exp(-i\delta t H_{\mathbf{x}})$ is diagonal in the space of the wavefunction it is applied to. This amounts to applying the operators in sequence, which is mathematically disallowed since $[H_{\mathbf{p}}, H_{\mathbf{x}}] \neq 0$. By the Baker-Hausdorff-Campbell formula we can expect the error in this approach to be $\propto [\delta t H_{\mathbf{p}}, \delta t H_{\mathbf{x}}] \sim \mathcal{O}(\delta t^2)$. It turns out that the error can be made even smaller by splitting the Hamiltonian judiciously. The particular splitting used in these calculations is as follows

$$\exp(-i\delta t H_{\text{eff}}) = \exp\left(-i\frac{\delta t}{2} H_{\mathbf{p}}\right) \exp(-i\delta t H_{\mathbf{x}}) \exp\left(-i\frac{\delta t}{2} H_{\mathbf{p}}\right) \exp(\mathcal{O}(\delta t^3))$$

Time evolution proceeds by initializing $\psi(\mathbf{x})$ and a_α , and iteratively applying the split-step operators interspersed with a FFT of $\psi(\mathbf{x})$, as conveyed in Fig.4.3.




Figure 4.3: A flow chart describing the split-step Fourier Transform technique for the cavity-atom system.

The results of numerical simulations are presented here. All simulations occur in 2D and involve only a single cavity mode. Care must be taken in 2D to correctly account for the scattering interaction; in 3D $U = 4\pi\hbar^2 a_s/m$, but the dimensionality of

this constant makes it inapplicable in lower dimensions. One approach to circumvent this problem is to assume the atoms are confined in the free direction by a harmonic potential. Performing all other calculations in 2D, this has the effect of renormalizing by the Harmonic oscillator length $l_y = \sqrt{\frac{\pi\hbar}{m\omega_y}}$, i.e. $U \rightarrow U/l_y$. The approach taken in this thesis is to set Thomas-Fermi expressions for the chemical potential of a 2D or 1D gas equal to the expression for a 3D gas and solve for an effective interaction parameter. This approach yields the following expressions for the interaction parameters in 2D and 1D

$$\begin{aligned} U_{2D} &= \left(\frac{50625}{262144} \frac{\pi m U^4 \bar{\omega}^{12}}{N(\omega_x \omega_z)^5} \right)^{1/5} \\ U_{1D} &= \left(\frac{125}{18} \frac{m^2 U^3 \bar{\omega}^9}{\pi^3 N^2 \omega_x^5} \right)^{1/5} \end{aligned}$$

Here $\bar{\omega}$ is the geometric mean of the trap frequencies ω_x , ω_y , and ω_z for the harmonically confined gas. In addition to time-evolution, this approach can be used to calculate the ground state of the Hamiltonian via the imaginary time formalism; replacing the time increment δt with $i\delta t$ effects a further suppression of all states but the ground state after each loop of the split-step algorithm. When this calculation converges, the ground state has been determined. Done in this fashion, Fig. 4.4 shows the GPE ground state, where the pump has been eliminated. This numerically generated field is extremely close to the 2D Thomas-Fermi solution except near the edges where the mean-field energy is overcome by the kinetic energy of the gas.

Pumping the atoms with a standing wave $\eta(\mathbf{x}) \propto \cos(kx)$ above-threshold, the ground state of the cavity-atom system exhibits one of the two self-organized states




Figure 4.4: The simulated ground state atomic density of a 2D atomic gas obeying the GPE.

as seen in Fig. 4.5. The atoms are arranged on a square with lattice constant $\lambda/\sqrt{2}$. Below-threshold the atoms remain unchanged by the pump although some light is scattered into the cavity as a consequence of the inevitable asymmetry between the occupation of the ‘even’ and ‘odd’ sites for a randomly placed cloud. This asymmetry amounts to a bias field, in principle predetermining into which lattice sites the atoms self-organize. For a sufficiently large cloud this bias field is unnoticeable in experiment and simulation [89]. The difference between the two possible ground states is further clarified in Fig. 4.6 where the checkerboard interference potential for this simulation is shown. The atoms are induced to organize into either the red or blue spots in a self-consistent fashion. Superposing the self-organized atomic distribution onto the potential clearly shows that the atoms are occupying one of the sublattices.

The Fourier Transform of this atomic wavefunction is shown in Fig. 4.7, where it is used to approximate the time-of-flight density of the gas. A better method of modeling the gas’s expansion from the trap would be to time evolve the self-organized wavefunction in the absence of a pump or external potential. This numerical




Figure 4.5: The simulated self-consistent ground state of a 2D atomic gas with contact interactions coupled to the $\text{HG}_{0,0}$ mode of a cavity and pumped transversely. The pump is above threshold and the gas has self-organized into the cavity mode.




Figure 4.6: (left) The effective potential generated by the interference of the pump and the light scattered off the atoms into the $\text{HG}_{0,0}$ mode into which the atoms self-organize above-threshold. (right) With the atoms overlaid on the potential it is clear that the ‘even’ sites are filled, while the ‘odd’ sites are left empty; a necessary condition to ensure constructive interference in the light scattered from the atoms.

experiment would best mimic the hydrodynamic expansion of the gas but would demonstrate qualitatively the same density as the Fourier Transform. Unsurprisingly, the Fourier Transform shows distinct Bragg peaks at momenta that correspond to the reciprocal lattice vectors, as would be expected from any scattering experiment involving a crystal.

The result of a similar simulation is shown in Fig. 4.8, but instead of the $\text{HG}_{0,0}$, the atoms are coupled to the $\text{HG}_{0,1}$. At this scale the transverse structure of the cavity modes is not apparent, and the $\text{HG}_{0,1}$ like the $\text{HG}_{0,0}$ appears indistinguishable from $\propto \cos(kz)$. However, what is clearly visible in the first order mode is its nodal $y-z$ plane. The phase of the mode changes by π across this plane, and thus the ‘odd’ and ‘even’ sites reverse roles crossing the plane. Again, the atoms must self-organize into the red or blue sites of the original checkerboard, but in the first order mode




Figure 4.7: The Fourier Transform, and thus approximate time-of-flight density, of the atomic gas self-organized into the $\text{HG}_{0,0}$.

the checkerboard switches colors across the nodal plane. This is clear as the atom density exhibits a mirror symmetry across a plane (the nodal plane) that does not contain any atoms. This would be impossible if the atoms were forced to remain ‘on the same checkerboard.’ This is further evidenced in Fig. 4.9 where the first order interference potential exhibits a flipped checkerboard, and where the atomic density overlaid on the zeroth order interference potential shows that the locations of the atoms correspond to different checkerboards on either side of the nodal plane. A consequence of this phase discontinuity is that the atomic wavefunction must mirror the cavity in this way; the relative phase of the atomic distributions localized on either side of the nodal plane is π .

The Fourier transform of atomic density looks qualitatively similar to the zeroth order self-organization, but with one very telling difference. The first order Bragg




Figure 4.8: A simulation of the cavity-atom system where the initial atomic density (GPE ground state) is coupled to the $\text{HG}_{0,1}$ and pumped under time-varying conditions. The pump power is quadratically ramped for $100 \mu\text{s}$ and then held constant for another $100 \mu\text{s}$. In this time the atoms scatter light into the cavity and self-organize into the cavity mode.




Figure 4.9: (left) The effective potential generated by the interference of the pump and the light scattered off the atoms into the $HG_{0,1}$ mode into which the atoms self-organize above-threshold. (right) With the atoms overlaid on the potential it is clear that they occupy ‘even’ sites on one side of the mode’s nodal plane and the ‘odd’ sites on the other. This can happen because there is a π phase jump in the atomic wavefunction across this nodal plane.

peaks exhibit a node which is a relic of the phase discontinuity of the atomic wavefunction, a feature that has been observed clearly in experiment.




Figure 4.10: The Fourier Transform, and thus approximate time-of-flight density, of the atomic gas self-organized into the $\text{HG}_{0,1}$.

Chapter 5

Conclusion

This thesis has described the design, construction, and characterization of a novel multimode CQED apparatus for observing and studying beyond mean-field self-organization phenomena. We designed an apparatus that allows us to place a microscopic Bose-Einstein condensate of rubidium 87 within the mode volume of a 1-cm-long degenerate near-confocal Fabry-Pérot cavity. Pumping the atoms transversely allows for the possibility of observing and studying quantum phase transitions marked by the emergence of quantum states that cannot be described by mean-field theories. Developing DMD technology was necessary to properly and flexibly address the cavity, whose myriad degrees of freedom are involved in the crystallization phenomena.

Appendix A

DMD Hologram Code

This appendix provides all the code necessary to

1. Generate the calibration masks used in the calibration procedure.
2. Perform the calibration and store the interference patterns.
3. Generate an aberration-correcting phase mask from the interference patterns.

The phase mask is then incorporated into the generation of every hologram. A sample script has been provided that generates a $\text{HG}_{0,0}$ and $\text{LG}_{0,0}$ mode. Running the calibration requires that `pyflycapture2` is installed on a computer running Linux. The module is in the QED Github repository

<https://github.com/mmcqed/pyflycapture2>

A.1 CountingModes.py

This is a quick bit of code to enumerate the modes, even, odd, or both, up to a maximum order.

```

from __future__ import division
import scipy

# All modes
def print_all_modes(max_mode):
    num_modes = int((max_mode+1)*(max_mode+2)/2)
    for k in range(num_modes):
        order = int(scipy.floor(.5*(-1+scipy.sqrt(1+8*k))))
        l = int(k - order*(order+1)/2)
        m = int(order - 1)

        print (l,m)

# Even modes
# max_mode should be even
def print_even_modes(max_mode):
    num_modes = int((max_mode/2+1)**2)
    for k in range(num_modes):
        order = 2*int(scipy.sqrt(k))
        l = int(k-(order/2)**2)
        m = order-1

        print (l,m)

# Odd modes
# max_mode should be odd
def print_odd_modes(max_mode):
    num_modes = int((1+max_mode)*(3+max_mode)/4)
    for k in range(num_modes):
        order = 2*int(scipy.floor(.5*(-1+scipy.sqrt(1+4*k))))+1
        l = int(k-(order**2-1)/4)
        m = order-1
        print (l,m)

```

A.2 ditherimage.py

This is a collection of helper functions used to generate the calibration masks and the final phase mask.

```

from __future__ import division
import scipy
import copy
from scipy.constants import pi,c

```

```

def FS_Dither(img,palette):
    '''FS_Dither(img,palette) produces a dither image of img using the Floyd-Steinberg algorithm on the basis of color
        palette palette'''
    A = copy.deepcopy(img)
    yt,xt = scipy.shape(A)
    newA = [[0 for x in range(xt)] for y in range(yt)]

    for j in range(yt):
        for i in range(xt):
            temp = []
            pix = A[j][i]
            temp[:] = [abs(x-pix) for x in palette]
            newA[j][i] = palette[temp.index(min(temp))]
            qerr = A[j][i] - newA[j][i]

            if ((i+1>=0) & (j>=0) & (i+1<xt) & (j<yt)):
                A[j][i+1] += qerr*7/16
            if ((i-1>=0) & (j+1>=0) & (i-1<xt) & (j+1<yt)):
                A[j+1][i-1] += qerr*3/16
            if ((i>=0) & (j+1>=0) & (i<xt) & (j+1<yt)):
                A[j+1][i] += qerr*5/16
            if ((i+1>=0) & (j+1>=0) & (i+1<xt) & (j+1<yt)):
                A[j+1][i+1] += qerr*1/16

    return newA

def rgb2gray(rgb):
    return scipy.dot(rgb[:, :, 0:3], [0.299, 0.587, 0.144])

def Lanczos(x,a):
    if x == 0:
        return 1
    if ((0 < abs(x)) & (abs(x) < a)):
        return a*scipy.sin(pi*x)*scipy.sin(pi*x/a)/(pi**2 * x**2)
    else: return 0

def PadImg(small,h,w):
    m,n = scipy.shape(small)
    ml = h-m
    nl = w-n
    big = scipy.vstack((scipy.zeros((int(ml/2)+ml%2,n)),small,scipy.zeros((int(ml/2),n))))
    big = scipy.hstack((scipy.zeros((h,int(nl/2)+nl%2)),big,scipy.zeros((h,int(nl/2)))))

    return big

```

A.3 DMDFunctions.py

This is a set of helper functions used to generate the holograms.

```
import scipy

def RealSpace(maxLen,N):
    space = scipy.linspace(-maxLen,maxLen,N,endpoint=False)
    return space

def MomSpace(maxLen,N):
    maxK = N/4./maxLen
    space = scipy.linspace(-maxK,maxK,N,endpoint=False)
    return space

def CFFT(Ax):
    Ap = scipy.fftpack.fftshift(scipy.fftpack.fft(scipy.fftpack.ifftshift(Ax)))
    return Ap

def iCFFT(Ap):
    Ax = scipy.fftpack.fftshift(scipy.fftpack.ifft(scipy.fftpack.ifftshift(Ap)))
    return Ax

def RealSpace2D(maxLen,N):
    space = scipy.linspace(-maxLen,maxLen,N,endpoint=False)
    xspace,yspace = scipy.meshgrid(space,space)    #x,y coordinates
    return xspace, yspace

def RealSpace2DAspect(maxLen,asp,N):
    spacex = scipy.linspace(-asp*maxLen,asp*maxLen,asp*N,endpoint=False)
    spacey = scipy.linspace(-maxLen,maxLen,N,endpoint=False)
    xspace,yspace = scipy.meshgrid(spacex,spacey)    #x,y coordinates
    return xspace, yspace

def MomSpace2D(maxLen,N):
    maxK = N/4./maxLen
    space = scipy.linspace(-maxK,maxK,N,endpoint=False)
    kspace,lspace = scipy.meshgrid(space,space)
    return kspace, lspace

def CFFT2D(Ax):
    Ap = scipy.fftpack.ifftshift(scipy.fftpack.fft2(scipy.fftpack.fftshift(Ax)))
    return Ap

def iCFFT2D(Ap):
    Ax = scipy.fftpack.ifftshift(scipy.fftpack.ifft2(scipy.fftpack.fftshift(Ap)))
    return Ax

def RandOneArray(l,n1,d=1):
```

```

if d==1:
    oneArray = scipy.zeros(1)
    while sum(oneArray)<n1:
        oneArray[int(scipy.floor(1*scipy.rand(1))[0])] = 1

if d==2:
    oneArray = scipy.zeros((1,1))
    while scipy.sum(oneArray)<n1:
        oneArray[int(scipy.floor(1*scipy.rand(1))),int(scipy.floor(1*scipy.rand(1)))] = 1
return oneArray

```

A.4 Patches.py

This script generates the calibration masks.

```

from __future__ import division
import scipy
import matplotlib.pyplot as plt
import scipy.misc
from scipy.constants import c, pi

import ditherimage

import DMDFunctions as DF

# Generate the DMD masks used for calibration. Each mask turns on two
# superpixels, one of which is always the same reference.

# These three lines set parameters and generate a rectangular grid in which the
# DMD pixels lives. This is to account for the pixels being square, but the
# array being effectively rotated 45deg.
N = 1481
m = 910
DMD = [[ ((i+j > m) & (i+j < 2*N-1-m) & (i-j < m+2) & (j-i < m+3)) for i in range(N)] for j in range(N+1)]

# These are the parameters for the diffraction grating. x0 and y0 determine the
# location of the first order, while z0 determines the overall focus, all in m.
x0 = .5*15.e-4
y0 = -.5*15.e-4
z0 = 0

dx = 7.64e-6 # dx is the real of a square DMD pixel in meters
lam = 780.e-9 # laser wavelength in meters

```

```

f = 150e-3 # focal length in meters

x,y = DF.RealSpace2DAspect(741*dx,1481/1482,1482) # x and y values over the DMD

# The grating
DMDfield = DMD*scipy.cos( (2*pi)/(f*lam)*(x*x0 + y*y0 + (z0/2/f)*(x**2 + y**2)) )**2

# Perform Floyd-Steinberg dithering on the field to binarize it
DMDfieldBin = ditherimage.FS_Dither(DMDfield,[1,0])

# Create a canvas for the .bmp that represents the DMD
MasterPatch = scipy.zeros((1140,912))

# Load the DMD field into the .bmp format
for l in range(912):
    for j in range(2):
        for i in range(570):
            MasterPatch[2*i+j,l] = DMDfieldBin[i+j+1][911+i-1]

# The size of the super pixel is 38x19, which looks square on the DMD by virtue
# of its rotated orientation.

x1 = 38
x2 = 19

# Size of the superpixel grid

xS1 = int(1140/x1)
xS2 = int(912/x2)

# Choose pairs of superpixels to make up the xS1 x xS2 calibration images
savedir = '/Patches/x15y-15/'
for i in range(xS1):
    for j in range(xS2):
        patch = scipy.zeros((1140,912))
        patch[i*x1:(i+1)*x1,j*x2:(j+1)*x2] = MasterPatch[i*x1:(i+1)*x1,j*x2:(j+1)*x2]
        patch[(xS1/2)*x1:((xS1/2)+1)*x1,(xS2/2)*x2:((xS2)/2+1)*x2] = MasterPatch[(xS1/2)*x1:((xS1/2)+1)*x1,(xS2/2)*x2:((
            xS2/2)+1)*x2]
        bmparray = scipy.array((patch,patch,patch))
        scipy.misc.imsave(savedir+'Patch'+str(j + i*xS2)+' .bmp', bmparray)

```

A.5 DMD_Calibration.py

This script analyzes the fringe images and produces the super-pixel phase mask.

```

from __future__ import division
import scipy.misc
import scipy.fftpack as ft
import matplotlib.pyplot as plt
import os, os.path
import copy

# Performs DMD calibration using a sequence of interference images generated by
# pairs of 'superpixels' activated at a time. In the limit that the superpixels
# are infinitesimally small, each is a delta function field in the DMD plane.
# Placing a lens of focal length f halfwafer between the camera and DMD, which are
# separated by 2f ideally connects the two fields by Fourier Transform. Thus, a
# delta function in the DMD plane corresponds to a plane wave in the camera
# plane. The interference pattern resulting from the superposition of two plane
# waves is recorded for N pairs of superpixels, where one of the pair is held
# constant as a reference.

caldir = '/dmd_calibration/cal_x15y-15_1440_real/'

CalImg = scipy.misc.imread(caldir + 'cal250.bmp')
CalImg_ft = ft.fftshift(ft.fft2(ft.ifftshift(CalImg)))
CalImg_phi = scipy.arctan2(scipy.imag(CalImg_ft), scipy.real(CalImg_ft))

dimX, dimY = scipy.shape(CalImg)
center = scipy.array([dimX/2, dimY/2])

# A collection of masks blotting out the 0th order mode (the DC component of the
# FT) and the unwanted 1st order mode.

WinSize = 5
Win0 = scipy.ones((960, 1280))
Win0[480 - WinSize:480 + 1 + WinSize, 640 - WinSize:640 + 1 + WinSize] = 0

Win0u = copy.copy(Win0)
Win0u[480:960, :] = 0

Win0d = copy.copy(Win0)
Win0d[0:480, :] = 0

Win0l = copy.copy(Win0)
Win0l[:, 640:1280] = 0

Win0r = copy.copy(Win0)
Win0r[:, 0:640] = 0

# The size of the super pixel is 38x19, which looks square on the DMD by virtue
# of its rotated status.

x1 = 38

```

```

x2 = 19

# Size of the superpixel grid

xS1 = int(1140/x1)
xS2 = int(912/x2)

Phase1 = scipy.zeros((int(1140/x1),int(912/x2)))
Phase2 = scipy.zeros((int(1140/x1),int(912/x2)))
Amp1 = scipy.zeros((int(1140/x1),int(912/x2)))
Amp2 = scipy.zeros((int(1140/x1),int(912/x2)))

numfiles = len([entry for entry in os.listdir(caldir) if '.bmp' in entry])

# Watch the video, is the top-most peak moving L2R or R2L?
L2R = False

for i in range(xS1):
    for j in range(xS2):
        print int(i*xS2 + j)
        # Pull up the interference pattern, find its FT, identify the location
        # of the 1st order mode, measure the phase there.

        CalImg = scipy.misc.imread(caldir + 'cal' + str(int(i*xS2 + j)) + '.bmp')
        CalImg_ft = ft.fftshift(ft.fft2(ft.ifftshift(CalImg)))
        CalImg_phi = scipy.arctan2(scipy.imag(CalImg_ft),scipy.real(CalImg_ft))

        # Identify the correct peak, 1 of 3

        if i<xS1/2:
            peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0u) == abs(CalImg_ft*Win0u).max()))[:,0]
        if i>xS1/2:
            peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0d) == abs(CalImg_ft*Win0d).max()))[:,0]
        if i==xS1/2:
            if L2R:
                if j<xS2/2:
                    peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0l) == abs(CalImg_ft*Win0l).max()))[:,0]
                if j>xS2/2:
                    peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0r) == abs(CalImg_ft*Win0r).max()))[:,0]
            else:
                if j>xS2/2:
                    peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0l) == abs(CalImg_ft*Win0l).max()))[:,0]
                if j<xS2/2:
                    peak1 = scipy.array(scipy.where(abs(CalImg_ft*Win0r) == abs(CalImg_ft*Win0r).max()))[:,0]

        #peak1 = scipy.array(scipy.where(abs(CalImg_ft) == abs(CalImg_ft*Win0).max()))[:,0]
        # The location of the second 1st order peak is reliably the inversion of
        # the first 1st order peak.
        peak2 = -(peak1 - center) + center

```

```

Phase1[i,j] = CalImg_phi[peak1[0],peak1[1]]
Phase2[i,j] = CalImg_phi[peak2[0],peak2[1]]

Amp1[i,j] = 1/abs(CalImg_ft[peak1[0],peak1[1]])
Amp2[i,j] = 1/abs(CalImg_ft[peak1[0],peak1[1]])

plt.imshow(Phase1,interpolation='none')
plt.show()

# Save Phase1 in some context, as that is the result of this script

scipy.save('/dmd_calibration/Phase1_real.npy',Phase1)

```

A.6 PhaseMaskPost.py

This script takes the super-pixel phase mask and performs several post-processing steps that ultimately result in a phase mask that can be used to cancel out the optical aberrations when developing holograms.

```

from __future__ import division
import scipy
import scipy.linalg as ln
import copy
from scipy.constants import pi

import ditherimage

# Develop the phase mask for use in producing aberration-corrected fields

# Load the calibration data generated by DMD_Calibration.py
PhaseMask = scipy.load('/dmd_calibration/Phase1_real.npy')

# Interpolate the 'bad region' in the middle

testRegion = PhaseMask[13:18,22:27]

IntPnts = [(i,j) for i in range(5) for j in range(5) if ((abs(2-i)>1) | (abs(2-j)>1))]
Vals = scipy.array([testRegion[i[0],i[1]] for i in IntPnts])

t = scipy.floor([(1+scipy.sqrt(1+8*i))/2 for i in range(10)])
tn = t*(t-1)/2
t1 = range(10) - tn
t0 = t-t1-1

```

```

Amat_t = scipy.array([[i[0]**t0[j] * i[1]**t1[j] for j in range(10)] for i in IntPnts])

Amat = scipy.array([[1, i[0], i[1], i[0]**2, i[0]*i[1], i[1]**2, i[0]**3, i[0]**2*i[1], i[0]*i[1]**2, i[1]**3] for i in
    IntPnts])

bvec = scipy.dot(ln.inv(scipy.dot(Amat.transpose(),Amat)),scipy.dot(Amat.transpose(),Vals))

NewPnts = [(i,j) for i in range(5) for j in range(5) if ((abs(2-i)<2) & (abs(2-j)<2))]
NewAmat = scipy.array([[1, i[0], i[1], i[0]**2, i[0]*i[1], i[1]**2, i[0]**3, i[0]**2*i[1], i[0]*i[1]**2, i[1]**3] for i
    in NewPnts])
NewVals = scipy.dot(NewAmat,bvec)

for inum,i in enumerate(NewPnts):
    testRegion[i[0],i[1]] = NewVals[inum]

# Bad region is replaced
PhaseMask[13:18,22:27] = testRegion

w,l = scipy.shape(PhaseMask)

PhaseMaskUnwrap = copy.copy(PhaseMask)

# This confusing snippet unwraps the mask, first in columns, and again in rows

for i in range(l-1):
    PhaseMaskUnwrap[:,i+1:] = PhaseMaskUnwrap[:,i+1:] - scipy.tile(2*pi*scipy.sign(PhaseMaskUnwrap[:,i+1]-
        PhaseMaskUnwrap[:,i])*(abs(PhaseMaskUnwrap[:,i+1]-PhaseMaskUnwrap[:,i])>pi), (scipy.shape(PhaseMaskUnwrap)[1]-i-
        -1,1)).transpose()
for i in range(w-1):
    PhaseMaskUnwrap[i+1::,:] = PhaseMaskUnwrap[i+1::,:] - scipy.tile(2*pi*scipy.sign(PhaseMaskUnwrap[i+1,:]-
        PhaseMaskUnwrap[i,:])*(abs(PhaseMaskUnwrap[i+1,:]-PhaseMaskUnwrap[i,:])>pi), (scipy.shape(PhaseMaskUnwrap)[0]-i-
        -1,1))

# Each superpixel is actually a region of 38x19 actual pixels. This resampling
# needs be performed to produce the phase profile in real space.

a = 3

magX = 38
magY = 19

PhaseMask_enl = scipy.zeros((1140,912))

sumspace = scipy.linspace(-a+1,a,2*a)

for x in range(1140):
    print x
    for y in range(912):
        counter = 0

```

```

    for i in sumspace:
        for j in sumspace:
            posX = (scipy.floor(x/magX)+i)
            posY = (scipy.floor(y/magY)+j)
            if ((0<=posX<w) & (0<=posY<1)):
                counter += PhaseMaskUnwrap[posX,posY]*ditherimage.Lanczos(x/magX-(scipy.floor(x/magX)+i),a)*
                           ditherimage.Lanczos(y/magY-(scipy.floor(y/magY)+j),a)
            PhaseMask_enl[x,y] = counter

# After resampling to the new resolution, the phase can be put on a real space
# grid.

PhaseMask_RS = scipy.zeros((1482,1481))

for l in range(912):
    for j in range(2):
        for i in range(570):
            PhaseMask_RS[i+j+1,911+i-1] = PhaseMask_enl[2*i+j,l]

# Save the real-space phase calibration to be included in all further masks
scipy.save('PhaseMaskDMD_real.npy',PhaseMask_RS)

```

A.7 dmd_calibrate.py

This script performs the DMD calibration by successively selecting and projecting a calibration mask onto the DMD and taking a picture of the resulting fringe pattern.

```

from __future__ import division

import flycapture2 as fc2
import numpy as np
import scipy.misc
import sys
sys.path.append('/usr/local/lib/python2.7/site-packages')
import cv2
import time

def main():

    c = fc2.Context()
    c.connect(*c.get_camera_from_index(0))
    c.set_video_mode_and_frame_rate(fc2.VIDEOMODE_1280x960Y16,fc2.FRAMERATE_7_5)

```

```

p = c.get_property(fc2.BRIGHTNESS)
p['abs_control'] = True
p['abs_value'] = 0 # %
c.set_property(**p)

p = c.get_property(fc2.AUTO_EXPOSURE)
p['on_off'] = False
p['abs_control'] = True
p['abs_value'] = 0 # EV
p['auto_manual_mode'] = False
p['one_push'] = False
c.set_property(**p)

p = c.get_property(fc2.SHARPNESS)
p['on_off'] = True
p['abs_value'] = 0
p['auto_manual_mode'] = False
c.set_property(**p)

p = c.get_property(fc2.GAMMA)
p['abs_value'] = 1
c.set_property(**p)

p = c.get_property(fc2.IRIS)
p['on_off'] = True
p['abs_value'] = 256
p['auto_manual_mode'] = False
c.set_property(**p)

p = c.get_property(fc2.PAN)
p['on_off'] = False
p['auto_manual_mode'] = False
c.set_property(**p)

p = c.get_property(fc2.TILT)
p['on_off'] = False
p['auto_manual_mode'] = False
c.set_property(**p)

p = c.get_property(fc2.SHUTTER)
p['auto_manual_mode'] = False
p['abs_value'] = 20 # ms
p['one_push'] = False
c.set_property(**p)

p = c.get_property(fc2.GAIN)
p['abs_value'] = 0 # dB
p['auto_manual_mode'] = False
p['one_push'] = False
c.set_property(**p)

```

```

# c.start_capture()
# im = fc2.Image()

for i in range(72,100):
    img = cv2.imread('Patches/x15y-15_1440/patch' + str(i) + '.bmp')
    cv2.namedWindow("HGmn", cv2.WND_PROP_FULLSCREEN)
    cv2.setWindowProperty("HGmn", cv2.WND_PROP_FULLSCREEN, cv2.cv.CV_WINDOW_FULLSCREEN)
    cv2.imshow("HGmn",img)
    key = cv2.waitKey(500)
    if key == -1:
        c.start_capture()
        im = fc2.Image()
        c.retrieve_buffer(im)
        c.stop_capture()
        a = np.array(im)
        scipy.misc.imsave('Patches/real_cal_x15y-15_1440/cal' + str(i) + '.bmp',a)

# c.stop_capture()
c.disconnect()

if __name__ == '__main__':
    main()

```

A.8 ModeGeneration__example.py

This script creates an example $HG_{0,0}$ and $LG_{1,1}$ mode using the fully processed phase mask to cancel optical aberrations in the system.

```

from __future__ import division
import DMDFunctions as DF
import scipy
from scipy.constants import pi,c
import matplotlib.pyplot as plt
import scipy.misc
import numpy.polynomial.hermite
import scipy.fftpack as ft
import ditherimage

# This little function determines the DMD coordinates from realspace coordinates
def RowColID(x,y):
    return (x-y, ((x+y)-(x+y)%2)/2)

```

```

# Returns a HG mode
def philm_gen(l,m,x,y,w0,a,phi):
    ta = .5*((1+a)*x + (1-a)*y)*scipy.cos(phi) + ((1-a)*x + (1+a)*y)*scipy.sin(phi)
    sa = .5*((1-a)*x + (1+a)*y)*scipy.cos(phi) - ((1+a)*x + (1-a)*y)*scipy.sin(phi)
    model = numpy.polynomial.hermite.hermval(scipy.sqrt(2)*ta/w0,scipy.concatenate((scipy.zeros(l),scipy.ones(1))))
    modem = numpy.polynomial.hermite.hermval(scipy.sqrt(2)*sa/w0,scipy.concatenate((scipy.zeros(m),scipy.ones(1))))
    phi_lm = model*modem*scipy.exp(- (ta/w0)**2 - (sa/w0)**2)
    return phi_lm

# Returns a LG mode
def phiLGpl_gen(p,l,x,y,w0,alpha,phi):
    t = x*scipy.cos(phi) + y*scipy.sin(phi)
    s = -x*scipy.sin(phi) + y*scipy.cos(phi)
    rho = scipy.sqrt((alpha*t)**2+s**2)
    phi = scipy.arctan2(s,alpha*t)
    LG_pl = scipy.special.genlaguerre(p,abs(l))
    norm = scipy.sqrt(2*scipy.misc.factorial(p)/pi/scipy.misc.factorial(abs(l)+p))
    phi_pl = norm * (1/w0) * ((2*.5*rho/w0)**abs(l)) * scipy.exp(-rho**2/(w0**2) - 1j*l*phi) * LG_pl(2*rho**2/w0**2)
    return phi_pl

# N determines the size of the rectangle necessary to house the DMD space.
# m effectively sets the aspect ratio of the rectangle.
N = 1481
m = 910

# DMD is a pixel-by-pixel representation of real space.
DMD = [[ (i+j > m) & (i+j < 2*N-l-m) & (i-j < m+2) & (j-i < m+3)) for i in range(N)] for j in range(N+1)]

# x0, y0, and z0 are the parameters of the grating, where x0 and y0 set the
# position in m of the first order, and z0 sets the focus.
x0 = .5*2.e-3
y0 = -.5*2.e-3
z0 = 0

# These are fundamental parameters
dx = 7.64e-6
lam = 780.e-9
f = 200e-3

# x,y are defined as real-space coordinates over the surface of DMD
x,y = DF.RealSpace2DAspect(741*dx,1481/1482,1482)

# This is the calibrated phase, in the same domain as DMD
CalPhase = scipy.load('PhaseMaskDMD_real.npy')

#This creates each mask by performing the following steps
#1) Calculating the HG field desired on the DMD
#2) Applying the necessary grating including phase correction to the HG field
#3) Generating a binarized version of the field with FS dithering

```

```

#4) Putting the values of the actual DMD pixels onto an 1140x912 grid, which is
# saved as a .bmp to be projected on the DMD itself

# HG00
l = 0
m = 0
w0 = 0.00155268
z0 = 0.0023464
modename = 'HG_00_example.bmp'
# obliquity parameter
a = 1/.83
fieldx = philm_gen(l,m,x,y,w0/a,1/a,scipy.arctan(1482/1481))
fieldx = fieldx/scipy.amax(abs(fieldx))
Phase = (1/2)*(pi*(fieldx<0) - CalPhase)
circle = .5*((1 + a**2)*y**2 + 2*(a**2 - 1)*x*y + (1 + a**2)*x**2)
DMDfield = abs(fieldx)*scipy.cos( (2*pi)/(f*lam)*(x*x0 + y*y0 + (z0/2/f)*circle) + Phase )**2
DMDfieldBin = ditherimage.FS_Dither(DMDfield,[1,0])
DMDmat = scipy.zeros((1140,912))
for l in range(912):
    for j in range(2):
        for i in range(570):
            DMDmat[2*i+j,l] = DMDfieldBin[i+j+1][911+i-1]
scipy.misc.toimage(scipy.array((DMDmat,DMDmat,DMDmat)), cmin=0.0, cmax=1.0).save(modename)

# LG11
p = 1
l = 1
w0 = 0.00155268
z0 = 0.0023464
modename = 'LG_11_example.bmp'
# obliquity parameter
a = 1/.83
fieldx = phiLGpl_gen(p,l,x,y,w0,1/a,3*pi/4)
fieldx = fieldx/scipy.amax(abs(fieldx))
Phase = (1/2)*(scipy.arctan2(scipy.imag(fieldx),scipy.real(fieldx)) - CalPhase)
DMDfield = DMD*abs(fieldx)*scipy.cos( (2*pi)/(f*lam)*(x*x0 + y*y0 + (z0/2/f)*(x**2 + y**2)) + Phase )**2
DMDfieldBin = ditherimage.FS_Dither(DMDfield,[1,0])
DMDmat = scipy.zeros((1140,912))
for l in range(912):
    for j in range(2):
        for i in range(570):
            DMDmat[2*i+j,l] = DMDfieldBin[i+j+1][911+i-1]
scipy.misc.toimage(scipy.array((DMDmat,DMDmat,DMDmat)), cmin=0.0, cmax=1.0).save(modename)

```

Appendix B

GPE Simulation Code

B.1 GPEgs.py

This script determines the ground state of the mean-field atom cavity system for an atomic cloud coupled to a single cavity mode pumped transversely. With a few modifications the script can perform time evolution of the GPE self-consistently. The time step is made positive real and the convergence criterion conditioning the `while` loop is changed to a `for` loop run for as many time steps as is desired. It is important to select a sensible initial wavefunction `psi0`, for example the GPE ground state, which can be determined using the regular script.

```
from __future__ import division
import scipy
import scipy.misc
import scipy.fftpack as ft
import scipy.linalg
import scipy.constants
import numpy
from scipy.constants import pi, hbar, c, m_p

def RealSpace(maxLen, N):
    space = ft.ifftshift(scipy.linspace(-maxLen, maxLen, N, endpoint=False))
```

```

xspace, yspace = scipy.meshgrid(space, space)    #x, y coordinates
return xspace, yspace

def MomSpace(maxLen, N):
    endpoint = N/4./maxLen
    space = ft.ifftshift(scipy.linspace(-endpoint, endpoint, N, endpoint=False))
    kspace, lspace = scipy.meshgrid(space, space)
    return kspace, lspace

def phin(n, x, z, w0, k):
    zr = k * w0**2 / 2
    zeta = z/zr
    w = w0*scipy.sqrt(1+zeta**2)
    S = (x*scipy.sqrt(2))/w
    normalize_factor = scipy.sqrt(1/2.**n/scipy.misc.factorial(n))
    hermvals = numpy.polynomial.hermite.hermval(S, scipy.concatenate((scipy.zeros(n), scipy.ones(1))))

    #phi_n = normalize_factor * hermvals * scipy.exp(-x**2/w**2) * scipy.cos(k*z + (2*n+1)*scipy.arctan(zeta)/2 - (zeta
        + 1/zeta)**(-1) * x**2/w0**2)
    phi_n = normalize_factor * hermvals * scipy.exp(-x**2/w**2) * scipy.cos(k*z)
    return phi_n

deltaT = -1j*1e-6
# Fundamental Constants
m = 87*m_p
Nat = 1.e5
lambda0 = 780.e-9
gam = 2*pi*6.e6
gamperp = .5*gam
DeltaA = -2*pi*102.e9
a0 = scipy.constants.value('Bohr_radius')
L = 1.e-2
r = 1.e-2
kappa = 10*2*pi*3.e5
# Derived Parameters
w0 = scipy.sqrt(lambda0/(2*pi))*scipy.sqrt(scipy.sqrt(L*(2*r-L)))
Vol = pi*w0**2*L/4
g0 = scipy.sqrt(3*c*(lambda0**2)*gamperp/(4*pi*Vol))
U0 = (g0**2)/DeltaA
k0 = 2*pi/lambda0
gc = 6.7e-14
#discretization
Nxy = 2048
omegaHar = 2*pi*150
# Initialize psi0(x, y)
#extent of x
maxLenx = 4*3.4e-6
#x, y coordinates
x, y = RealSpace(maxLenx, Nxy)
dx = x[0,1]-x[0,0]

```

```

# Initialize constant real space variables and operators
gNorm = phin(0,y,x,w0,k0)
etaNorm = scipy.cos(k0*y)
ThetaNorm = etaNorm*gNorm
UOper = U0*(gNorm**2)
VHar = (x**2 + y**2)*m*omegaHar**2/(2*hbar)
# Initialize momentum space variables and operators
P, Q = MomSpace(maxLenx,Nxy)
MomFac = hbar/(2*m)
TOper = MomFac*((2*pi*P)**2 + (2*pi*Q)**2)
expT = scipy.exp(-1j*deltaT*TOper)
expT2 = scipy.exp(-1j*(deltaT/2.)*TOper)

def FindTheta(eta,DeltaC,psi0):
    psix = scipy.array(psi0)
    # Half Momentum Step
    psik = ft.fft2(psix)
    psik0 = scipy.array(psik)
    psik = expT2*psik
    # Speedup Variables
    c1 = g0/DeltaA
    c2 = DeltaC - 1j*kappa
    c3 = Nat*gc

    #loop starts
    while ((abs(psik - psik0)**2).sum()/(abs(psik)**2).sum() > 1.e-6):
        psik0 = scipy.array(psik)
        psix = ft.ifft2(psik)
        psix = psix/(scipy.linalg.norm(psix)*dx)
        #Expectation Values for the steady state coherent field
        Theta = (ThetaNorm*abs(psix)**2).sum()*dx*dx
        B = (UOper*abs(psix)**2).sum()*dx*dx
        alpha = Nat*c1*eta*Theta/(c2-Nat*B)
        V = c3*abs(psix)**2 + (UOper*alpha*alpha.conj() + c1*eta*ThetaNorm*2*scipy.real(alpha) + (eta**2/DeltaA)*etaNorm
            **2) + VHar
        # Full Position Step
        expV = scipy.exp(-1j*deltaT*V)
        psix = expV*psix
        # Full Momentum Step
        psik = ft.fft2(psix)
        psik = expT*psik

    psix = ft.ifft2(psik)
    psix = psix/(scipy.linalg.norm(psix)*dx)
    #Expectation Values for the steady state coherent field
    Theta = (ThetaNorm*scipy.absolute(psix)**2).sum()*dx*dx
    B = (UOper*scipy.absolute(psix)**2).sum()*dx*dx
    alpha = Nat*c1*eta*Theta/(c2-Nat*B)
    V = c3*scipy.absolute(psix)**2 + (UOper*alpha*alpha.conj() + c1*eta*ThetaNorm*2*scipy.real(alpha) + (eta**2/DeltaA)*
        etaNorm**2) + VHar

```

```

# Full Position Step
expV = scipy.exp(-1j*deltaT*V)
psix = expV*psix

# Half Momentum Step
psik = ft.fft2(psix)
psik = expT2*psik

# Final position-space representation of wavefunction
psix = ft.ifft2(psik)
psix = psix/(scipy.linalg.norm(psix)*dx)
Theta = (ThetaNorm*scipy.absolute(psix)**2).sum()*dx*dx
B = (UOper*scipy.absolute(psix)**2).sum()*dx*dx
alpha = Nat*c1*eta*Theta/(c2-Nat*B)

return Theta, alpha, psix


if __name__ == '__main__':
    mu = scipy.sqrt(hbar*gc*m*Nat/pi)*omegaHar
    n2D = (mu - hbar*VHar)/(hbar*gc)*((mu - hbar*VHar)>0)
    eta = 1.05*1e9
    DeltaC = -2*pi*1.e8
    psi0 = scipy.sqrt(n2D)
    theta, alpha, psix = FindTheta(eta,DeltaC,psi0)

```

Appendix C

Cavity Lock Filters

Figure C.1




Appendix D

Cavity Spacer

The pocket for the internal threading was made by drilling through the piece with a 25/64 in drillbit, and then boring the appropriate length out to a diameter of .737 in. This number was derived from the formula for such things, $D - \frac{5}{4}p \sin(60^\circ) = .737$ in where the pitch $p = 1/80$ and the major diameter of the tap is $D = .75$ in. This formula assumes 60 ° threads and follows the standard for UN and metric threads, namely that they are truncated by 1/8 the thread height in the major diameter and 1/4 the thread height in the minor diameter, leaving 5/8 of thread height.

The external threading (single point) was required: a flat with a 60 ° inset cut normal to the edge, a cutting tool with a 60 ° point, and a fine bristly pad. After chamfering the edge, it was simply a matter of setting the gears appropriately on the lathe, and cutting one mil at a time, measuring regularly. From a major diameter of .749, cut threads .015 in deep, and then make 3 mil cuts thereafter, each time checking to see how tightly the piece mates. After each cut carefully remove the chips with the bristly pad.

Figure D.1



Bibliography

1. Bloch, I., Dalibard, J. & Nascimbène, S. Quantum simulations with ultracold quantum gases. *Nature Physics* **8**, 267–276 (2012).
2. Davis, K. B. *et al.* Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.* **75**, 3969 (1995).
3. Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. & Cornell, E. A. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science* **269**, 198–201 (1995).
4. Chin, C., Grimm, R., Julienne, P. & Tiesinga, E. Feshbach resonances in ultracold gases. *Reviews of Modern Physics* **82**, 1225 (2010).
5. O'hara, K., Hemmer, S., Gehm, M., Granade, S. & Thomas, J. Observation of a strongly interacting degenerate Fermi gas of atoms. *Science* **298**, 2179–2182 (2002).
6. Dahan, M. B., Peik, E., Reichel, J., Castin, Y. & Salomon, C. Bloch oscillations of atoms in an optical potential. *Physical Review Letters* **76**, 4508 (1996).
7. Abanin, D. A., Kitagawa, T., Bloch, I. & Demler, E. Interferometric approach to measuring band topology in 2D optical lattices. *Physical review letters* **110**, 165304 (2013).

8. Feynman, R. P. Simulating physics with computers. *International journal of theoretical physics* **21**, 467–488 (1982).
9. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073 (1996).
10. Richerme, P. *et al.* Quantum catalysis of magnetic phase transitions in a quantum simulator. *Physical review letters* **111**, 100506 (2013).
11. Jo, G.-B. *et al.* Itinerant ferromagnetism in a Fermi gas of ultracold atoms. *Science* **325**, 1521–1524 (2009).
12. Zwierlein, M. W., Abo-Shaeer, J. R., Schirotzek, A., Schunck, C. H. & Ketterle, W. Vortices and superfluidity in a strongly interacting Fermi gas. *Nature* **435**, 1047–1051 (2005).
13. Viefers, S. Quantum Hall physics in rotating Bose–Einstein condensates. *Journal of Physics: Condensed Matter* **20**, 123202 (2008).
14. Greiner, M., Mandel, O., Esslinger, T., Hänsch, T. W. & Bloch, I. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *nature* **415**, 39–44 (2002).
15. Becker, C *et al.* Ultracold quantum gases in triangular optical lattices. *New Journal of Physics* **12**, 065025 (2010).
16. Jo, G.-B. *et al.* Ultracold atoms in a tunable optical kagome lattice. *Physical review letters* **108**, 045305 (2012).
17. Paredes, B. *et al.* Tonks–Girardeau gas of ultracold atoms in an optical lattice. *Nature* **429**, 277–281 (2004).
18. Schneider, U *et al.* Metallic and insulating phases of repulsively interacting fermions in a 3D optical lattice. *Science* **322**, 1520–1525 (2008).

19. Bourdel, T. *et al.* Experimental study of the BEC-BCS crossover region in lithium 6. *Physical Review Letters* **93**, 050401 (2004).
20. Ritsch, H., Domokos, P., Brennecke, F. & Esslinger, T. Cold atoms in cavity-generated dynamical optical potentials. *Reviews of Modern Physics* **85**, 553 (2013).
21. Bloch, I., Dalibard, J. & Zwerger, W. Many-body physics with ultracold gases. *Reviews of Modern Physics* **80**, 885 (2008).
22. Cirac, J. I. & Zoller, P. Goals and opportunities in quantum simulation. *Nature Physics* **8**, 264–266 (2012).
23. Asbóth, J., Domokos, P., Ritsch, H & Vukics, A. Self-organization of atoms in a cavity field: Threshold, bistability, and scaling laws. *Physical Review A* **72**, 053417 (2005).
24. Szymańska, M., Keeling, J & Littlewood, P. Nonequilibrium quantum condensation in an incoherently pumped dissipative system. *Physical review letters* **96**, 230602 (2006).
25. Dalla Torre, E. G., Diehl, S., Lukin, M. & Strack, P. Keldysh approach for nonequilibrium phase transitions in quantum optics: Dicke model in optical cavities. *gas*] (2012).
26. Landig, R., Brennecke, F., Mottl, R., Donner, T. & Esslinger, T. Measuring the dynamic structure factor of a quantum gas undergoing a structural phase transition. *Nat. Commun.* **6**, 7046 (May 2015).

27. Cornish, S. L., Claussen, N. R., Roberts, J. L., Cornell, E. A. & Wieman, C. E. Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions. *Phys. Rev. Lett.* **85**, 1795–1798 (9 2000).
28. Moon, G. & Noh, H.-R. A comparison of the dependence of saturated absorption signals on pump beam diameter and intensity. *JOSA B* **25**, 2101–2106 (2008).
29. Lin, Y.-J., Perry, A., Compton, R., Spielman, I. & Porto, J. Rapid production of R 87 b Bose-Einstein condensates in a combined magnetic and optical potential. *Physical Review A* **79**, 063631 (2009).
30. Naides, M. *Trapping Ultracold Gases Near Cryogenic Materials with Rapid Reconfigurability* PhD thesis (Stanford University, 2014).
31. Metcalf, H. J. & van der Straten, P. *Laser Cooling and Trapping of Neutral Atoms* (Wiley-VCH Verlag GmbH & Co. KGaA, 2007).
32. Gustavson, T. *et al.* Transport of Bose-Einstein condensates with optical tweezers. *Physical Review Letters* **88**, 020401 (2001).
33. Kollár, A. J., Papageorge, A. T., Baumann, K., Armen, M. A. & Lev, B. L. An adjustable-length cavity and Bose–Einstein condensate apparatus for multimode cavity QED. *New Journal of Physics* **17**, 043012 (2015).
34. Wieman, C. E., Pritchard, D. E. & Wineland, D. J. Atom cooling, trapping, and quantum manipulation. *Rev. Mod. Phys.* **71**, S253–S262 (2 1999).
35. Cohen-Tannoudji, C. *Laser cooling and trapping of neutral atoms* in *Perspectives in Neutrinos, Atomic Physics and Gravitation: Proceedings of the XXVIIIth Rencontre de Moriond, Series: Moriond Workshops: Villars Sur Ollon, Switzerland, January 30-February 6, 1993* **76** (1993), 343.

- 36. Steck, D. A. *Rubidium 87 D line data* 2001.
- 37. Camy, G, Bordé, C. J. & Ducloy, M. Heterodyne saturation spectroscopy through frequency modulation of the saturating beam. *Optics Communications* **41**, 325–330 (1982).
- 38. Shirley, J. H. Modulation transfer processes in optical heterodyne saturation spectroscopy. *Optics Letters* **7**, 537–539 (1982).
- 39. McCarron, D., King, S. & Cornish, S. Modulation transfer spectroscopy in atomic rubidium. *Measurement science and technology* **19**, 105601 (2008).
- 40. Ma, L.-S. & Hall, J. Optical heterodyne spectroscopy enhanced by an external optical cavity: toward improved working standards. *Quantum Electronics, IEEE Journal of* **26**, 2006–2012 (1990).
- 41. Schenzle, A., DeVoe, R. G. & Brewer, R. G. Phase-modulation laser spectroscopy. *Physical Review A* **25**, 2606 (1982).
- 42. Gehrtz, M., Whittaker, E. A. & Bjorklund, G. C. Quantum-limited laser frequency-modulation spectroscopy. *JOSA B* **2**, 1510–1526 (1985).
- 43. Grimm, R., Weidemüller, M. & Ovchinnikov, Y. B. Optical dipole traps for neutral atoms. *arXiv preprint physics/9902072* (1999).
- 44. Arnold, K. & Barrett, M. All-Optical BEC in a 1.06 μ m dipole trap (2011).
- 45. Kinoshita, T., Wenger, T. & Weiss, D. S. All-optical Bose-Einstein condensation using a compressible crossed dipole trap. *Physical Review A* **71**, 011602 (2005).
- 46. Barrett, M., Sauer, J. & Chapman, M. All-optical formation of an atomic Bose-Einstein condensate. *Physical Review Letters* **87**, 010404 (2001).

47. OŌhara, K., Gehm, M., Granade, S. & Thomas, J. E. Scaling laws for evaporative cooling in time-dependent optical traps. *Physical Review A* **64**, 051403 (2001).
48. Cl  ment, J.-F. *et al.* All-optical runaway evaporation to Bose-Einstein condensation. *Physical Review A* **79**, 061406 (2009).
49. Olson, A. J., Niffenegger, R. J. & Chen, Y. P. Optimizing the efficiency of evaporative cooling in optical dipole traps. *Physical Review A* **87**, 053613 (2013).
50. Turner, R. *Microscopic Magnetometry with Degenerate 1D Bose Gases* PhD thesis (Stanford University, 2015).
51. Giorgini, S, Pitaevskii, L. & Stringari, S. Condensate fraction and critical temperature of a trapped interacting Bose gas. *Physical Review A* **54**, R4633 (1996).
52. Lu, M., Burdick, N. Q., Youn, S. H. & Lev, B. L. Strongly dipolar Bose-Einstein condensate of dysprosium. *Physical review letters* **107**, 190401 (2011).
53. Pitaevskii, L. & Stringari, S. *Bose-Einstein Condensation* (Clarendon Press, 2003).
54. Black, E. D. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics* **69**, 79–87 (2001).
55. *DLP4500NIR 0.45 WXGA Near-Infrared DMD* Texas Instruments (2014).
56. Wikipedia. *Blazed Grating* — *Wikipedia, The Free Encyclopedia* [Online; accessed 21-March-2016]. 2016.
57. Gaunt, A. L. & Hadzibabic, Z. Robust digital holography for ultracold atom trapping. *Scientific reports* **2** (2012).

- 58. Nogrette, F. *et al.* Single-atom trapping in holographic 2D arrays of microtraps with arbitrary geometries. *Physical Review X* **4**, 021034 (2014).
- 59. Siegman, A. E. (University Science Books, 1986).
- 60. Lee, W.-H. Binary Synthetic Holograms. English. *Appl. Opt.* **13**, 1677 (1974).
- 61. Stuart, D., Barter, O. & Kuhn, A. Fast algorithms for generating binary holograms. *arXiv:1409.1841*. arXiv: 1409 . 1841 [physics.optics] (Sept. 2014).
- 62. Cizmar, T., Mazilu, M. & Dholakia, K. In situ wavefront correction and its application to micromanipulation. *Nature Photon.* **4**, 388–394 (2010).
- 63. Vellekoop, I. M. & Mosk, A. P. Focusing coherent light through opaque strongly scattering media. English. *Opt. Lett.* **32**, 2309–2311 (Aug. 2007).
- 64. Zupancic, P. P. J. *Dynamic Holography and Beamshaping using Digital Micromirror Devices* MA thesis (Ludwig-Maximilians-Universität, Munich, 2013).
- 65. Vinet, J. Y. & Hello, P. Matrix Simulation of Optical Cavities. English. *Journal of Modern Optics* **40**, 1981–1993 (Mar. 2007).
- 66. Kollár, A. J., Papageorge, A. T., Baumann, K., Armen, M. A. & Lev, B. L. An adjustable-length cavity and Bose–Einstein condensate apparatus for multimode cavity QED. *New J. Phys.* **17**, 043012 (Apr. 2015).
- 67. Purcell, E. M. Spontaneous Emission Probabilities at Radio Frequencies. *Phys. Rev.* **69**, 674–674 (11-12 1946).
- 68. Weisskopf, V. & Wigner, E. Berechnung der natürlichen Linienbreite auf Grund der Diracschen Lichttheorie. *Zeitschrift für Physik* **63**, 54–73.

- 69. Bloch, I., Dalibard, J. & Nascimbene, S. Quantum simulations with ultracold quantum gases. *Nat Phys* **8**, 267–276 (Apr. 2012).
- 70. Scully, M. O. & Zubairy, M. S. *Quantum optics* (Cambridge university press, 1997).
- 71. Boca, A. *et al.* Observation of the Vacuum Rabi Spectrum for One Trapped Atom. *Phys. Rev. Lett.* **93**, 233603 (23 2004).
- 72. Wallraff, A. *et al.* Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431**, 162–167 (Sept. 2004).
- 73. Rempe, G., Walther, H. & Klein, N. Observation of quantum collapse and revival in a one-atom maser. *Phys. Rev. Lett.* **58**, 353–356 (4 1987).
- 74. Zhu, Y. *et al.* Vacuum Rabi splitting as a feature of linear-dispersion theory: Analysis and experimental observations. *Phys. Rev. Lett.* **64**, 2499–2502 (21 1990).
- 75. Dicke, R. H. Coherence in Spontaneous Radiation Processes. *Phys. Rev.* **93**, 99–110 (1 1954).
- 76. Nataf, P. & Ciuti, C. No-go theorem for superradiant quantum phase transitions in cavity QED and counter-example in circuit QED. *Nat Commun* **1** (Sept. 2010).
- 77. Vukics, A., Grieser, T. & Domokos, P. Elimination of the A-Square Problem from Cavity QED. *Phys. Rev. Lett.* **112**, 073601 (7 2014).
- 78. Baden, M. P., Arnold, K. J., Grimsmo, A. L., Parkins, S. & Barrett, M. D. Realization of the Dicke Model Using Cavity-Assisted Raman Transitions. *Phys. Rev. Lett.* **113**, 020408 (2 2014).

- 79. Dimer, F., Estienne, B., Parkins, A. S. & Carmichael, H. J. Proposed realization of the Dicke-model quantum phase transition in an optical cavity QED system. *Phys. Rev. A* **75**, 013804 (1 2007).
- 80. Domokos, P. & Ritsch, H. Collective Cooling and Self-Organization of Atoms in a Cavity. *Phys. Rev. Lett.* **89**, 253003 (25 2002).
- 81. Baumann, K., Guerlin, C., Brennecke, F. & Esslinger, T. Dicke quantum phase transition with a superfluid gas in an optical cavity. *Nature* **464**, 1301–1306 (Apr. 2010).
- 82. Gopalakrishnan, S., Lev, B. L. & Goldbart, P. M. Atom-light crystallization of Bose-Einstein condensates in multimode cavities: Nonequilibrium classical and quantum phase transitions, emergent lattices, supersolidity, and frustration. English. *Phys. Rev. A* **82**, 043612–29 (Oct. 2010).
- 83. Baumann, K. G. *Experimental Realization of the Dicke Quantum Phase Transition* PhD thesis (ETH Zürich, 2011).
- 84. Ritsch, H., Domokos, P., Brennecke, F. & Esslinger, T. Cold atoms in cavity-generated dynamical optical potentials. English. *Rev. Mod. Phys.* **85**, 553–601 (Apr. 2013).
- 85. Pethick, C. J. & Smith, H. *Bose-Einstein Condensation in Dilute Gases* first (Cambridge).
- 86. Wickenbrock, A., Hemmerling, M., Robb, G. R. M., Emary, C. & Renzoni, F. Collective strong coupling in multimode cavity QED. *Phys. Rev. A* **87**, 043817 (4 2013).

- 87. Baumann, K., Mottl, R., Brennecke, F. & Esslinger, T. Exploring Symmetry Breaking at the Dicke Quantum Phase Transition. *Phys. Rev. Lett.* **107**, 140402 (14 2011).
- 88. Nagy, D. *Collective effects of radiatively interacting ultracold atoms in an optical resonator* PhD thesis (Budapest University of Technology and Economics, 2010).
- 89. Baumann, K, Mottl, R, Brennecke, F & Esslinger, T. Exploring symmetry breaking at the Dicke quantum phase transition. *Physical review letters* **107**, 140402 (2011).

Alexander Themis Papageorge

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Benjamin Lev) Principal Advisor

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Sandy Fetter)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Amir Safavi-Naeini)

Approved for the University Committee on Graduate Studies