Cavity Quantum Electrodynamics with a Locally Addressable Quantum Gas

by

Justin Alan Gerber

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Dan M. Stamper-Kurn, Chair
Professor Norman Yao
Professor Daniel M. Neumark

Fall 2020

Cavity Quantum Electrodynamics with a Locally Addressable Quantum Gas

Abstract

Cavity Quantum Electrodynamics with a Locally Addressable Quantum Gas

by

Justin Alan Gerber

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Dan M. Stamper-Kurn, Chair

Motivated to develop quantum technologies and to study ever-more complex quantum systems, scientists are developing increasingly sophisticated experimental tools to control and measure quantum samples such as ensembles of ultracold atoms. A high-finesse optical cavity can be used to measure the state of atoms within its photonic mode with precision limited only by quantum uncertainty. Such a cavity can also be used to mediate interactions between different atoms within the cavity mode. High-resolution microscope objectives have been interfaced with ultracold atom experiments to allow researchers to image single atoms within optical lattices, to trap single atoms in microtweezer arrays, and to imprint arbitrary optical patterns onto atomic ensembles. Among other applications, these technologies have allowed researchers to investigate many-body quantum systems, engineer novel interactions, and realize high-fidelity quantum operations.

This dissertation presents the details of the design, construction, and operation of a new, versatile atomic physics apparatus that combines these two experimental tools. The apparatus includes a high-finesse optical cavity into which atoms are optically transported. In addition, there is a high-numerical-aperture objective aligned to image, with micron-scale resolution, atoms trapped within the center of the optical cavity. We present results demonstrating the capability of this apparatus to deliver and study ultracold atomic samples, ranging from single atoms to Bose-Einstein condensates. We demonstrate a dispersive shift of the cavity resonance due to the presence of atoms in the cavity mode and the trapping and imaging of single atoms in optical microtweezers. We also present an atomic scanning probe microscopy technique with which a single atom in a microtweezer is used to map out the spatial amplitude pattern of an optical cavity mode standing wave by monitoring the position-dependent scattering properties of the atom.

To my family

Contents

List of Figures

List of Tables

Acknowledgments

I could not have carried out the work in this thesis without the mentorship, guidance, encouragement, compassion, optimism, patience, and brilliance of so many colleagues, friends and family who supported me during this journey.

I owe an incredible thanks to my advisor Dan Stamper-Kurn for molding me into the scientist I have become, for pushing me to do my greatest work, and for being a caring mentor throughout. I'm most thankful to Dan for continuously encouraging me to stay true to myself, even when I doubted.

I have had the great pleasure of working with many incredible people on the E6 team who are talented and hardworking researchers. I'm deeply grateful to Emma Deist for her always having a song to sing or a tune to hum during all of those hours in lab together. Without her tenacity and drive I would still be at the whiteboard trying to design the 'perfect' MOT light optics layout. I'm very thankful to Johannes Zeiher for his unwavering optimism which helped us to always pursue the most interesting options, even if they were not the easiest. I want to thank Leon Lu for bravely taking over running the experimental apparatus in such short order. I'm grateful to Fang Fang for patiently and carefully teaching me so many details about laser cooling and trapping. I thank Che Liu for his eagerness and enthusiasm for E6, despite not having had the opportunity to see the lab in person yet! I also want to extend gratitude to Alec Bohnett, Rachel Tsuchiyama, Aron Lloyd, Armando Montejano, Vicky Xu, and Shantanu Debnath for all of their passion while working on E6 and for all they taught me during our time together in the lab.

I was extremely lucky to have had exceptionally kind, knowledgeable, and patient senior group members to whom I could look to as role models and tutors. Special thanks to Jonathan Kohler for so many bike rides, for all he taught me in lab, and for his patience for my crazy ideas. I'm very grateful to Claire Thomas for teaching me the importance of taking care of myself during grad school. Thanks to Tom Barter for breaking off a stuck flange on the E6 chamber with a big wrench when I was too scared to do it myself, and for what that taught me about experimental physics! Many thanks to Sydney Schreppler and Nicholas Spethmann for being the first ones to teach me about cavities and atoms. I'm grateful to Ed Marti and Ryan Olf for their advice to me about starting a new experiment. I'm especially thankful to Zephy Leung for struggling through all of the grad school milestones with me.

I want to express deep thanks to the ultracold lab members for all the help they lent me and lessons they taught me over the years. Thank you to Josh Isaacs, Julian Wolf, Kayleigh Cassella, Charles Brown, Shun Wu, Masayuki Okano, Lukas Buchmann, Severin Daiss, Scott Eustice, Shao-Wen Chang, Malte Schwarz, and Aaron Smull.

I'd like to thank Osip Schwartz for many helpful conversations about fabricating, characterizing, and cleaning superpolished cavity mirrors.

Many thanks to Warner Carlisle, Tommy Gutierrez, Abel Gonzalez, and Gordon Long in the professional machine shop for turning E6's designs into reality. Thank you also to Joseph Kant and Jesse Lopez in the student shop for helping me hone my skills and teaching me how to turn E6's designs into reality myself.

One of the hidden perks of working in the Campbell lab is getting the opportunity to catch up with Carlos Bustamante whenever you're walking between buildings! A big thanks to Carlos for helping me during the "missing mirrors saga" and for always keeping it real. I want to extend thanks to Anthony Vitan for all he does to keep the labs running. A big thanks also to Stephen Pride Raffel for so patiently answering all of my purchasing questions.

My time at Berkeley would not have been the same without so many of my friends with whom I shared this journey. I'd like to thank Erik Urban, Matt Jaffe, Sylvia Lewin, Neil Goeckner-Wald, Vinay Ramasesh, Greg Meyer, and Will Livingston especially for the time we spent on the intramural Ultimate Frisbee field. Somehow I can't avoid hyperfine lines even while playing sports. I'd like to thank Sam Kohn for being a great housemate and philosophizing with me deep into the night. I have a deep gratitude to Matt Gilbert, Halleh Balch, Robert Kealhofer and Nityan Nair for helping me enjoy a break from lab when I needed it and for always supporting and encouraging me.

I want to warmly thank my close friend Matt Bull for all of the wide-ranging physics discussions we've had over the years and for always helping me see my strengths. A special thanks to Kevin Howard for the adventures.

I want to acknowledge my family for their unending love and support for me during my time at Berkeley. To Kevin, Ann, Katie, and Allie, thank you for your unconditional support, which you lent even before you were really sure what this PhD thing was all about. To Mom, Dad, Courtney, and Amanda, thank you for teaching me to challenge myself and to be kind, and for reminding me what is important to me. To my grandparents, thank you for teaching me to see wonder in the world and always cheering for me. I dearly wish my Grandma Sandie could have been here to see me accomplish this goal. May her memory be a blessing.

Finally, to my wife, Sarah - thank you for teaching me, thank you for inspiring me, and thank you for believing in me.

Chapter 1

Introduction

1.1 From Electricity and Magnetism to Quantum Mechanics

Scientists first began systematically experimenting with electricity in the mid-18[th] century. It was at this time that the Leyden jar, an early electrical capacitor, was developed and Benjamin Franklin showed that atmospheric lightning was an electrical phenomenon. Researchers learned there were two types of electrical charge and assigned them polarities. Scientists began to understand the relationship between electricity and magnetism. One century later, James Clerk Maxwell published *On Physical Lines of Force*, codifying his famous equations, thus capturing mathematical laws that accurately describe and predict the behavior of electromagnetic phenomena [1]. By the end of the 19[th] century, electricity was changing and improving the lives of the general public with the birth and proliferation of consumer electronics such as the light bulb.

We can trace a similar historical arc regarding the development of quantum mechanics. At the turn of the 20[th] century, Planck discovered his famous empirical law resolving the ultraviolet (UV) catastrophe and enabling the dawn of quantum mechanics [2, 3]. In a comparatively, and shockingly, short period of time scientists had developed the essentially complete theory of non-relativistic quantum mechanics by the end of the 1920's and were on their way to developing relativistic quantum field theory.

Scientists views about quantum mechanics and its implications changed dramatically with the introduction of John Bell's theorem in 1964 [4] and Richard Feynman's call in 1982 for the development of engineered quantum systems as simulators of complex and computationally difficult quantum problems – quantum calculators or a quantum computers [5]. It was at this time that scientists began to realize important fundamental differences between quantum and classical mechanics. Further, they started to understand how these differences could be leveraged **in well-controlled quantum systems** to develop quantum tools to help scientists solve new problems.

Measurements of the optical spectra of atoms provided early evidence for the need for

a quantum theory beyond the classical. Throughout the decades, atoms have continued to serve as prototypical quantum systems. The first atomic clock based on a molecular beam was demonstrated in 1949 leading, over time, to the redefinition of the SI second to be based on atomic time in 1967 [6]. The first 3D magneto-optical trap (MOT), in which a gaseous cloud was cooled to sub-mK temperatures, was realized in 1987 [7]. This was followed by the realization of the first Bose-Einstein condensates (BECs) in 1995 [8, 9]. Laser and evaporative cooling allowed researchers generally to control all aspects of an atom at the quantum level: not only control of the internal atomic levels, by means of spectroscopy and optical pumping, but also the center of mass motion of the atom.

The first two decades of the new millennium have seen a blossoming in techniques both for the improved control of ultracold atomic systems and, especially, for the use of these systems as building blocks for what we can now begin to call quantum technologies. The fruits of the development of these quantum technologies are promising precursors to ever more prolific quantum technologies which may one day change and improve our lives like electromagnetic technologies did in the past.

I would highlight four pillars of quantum technology into which many quantum research efforts can be organized: fundamental quantum measurement, quantum simulation, quantum information, and quantum sensing [10–13]. Ultracold atomic systems have made major contributions to each of these four pillars, and, indeed, there is a lot of overlap between these different efforts in terms of techniques and hardware.

1.1.1   Locally Addressable Atomic Systems and Cavity QED

The above historical arc has motivated my PhD research. The work presented here is situated at a time when researchers are exploring novel techniques for the control and measurement of quantum systems in the interest of the development of new quantum technologies. The advancement of these technologies requires quantum systems which are ever more isolated from their environments, which can be measured with less noise, controlled with higher fidelity and speeds, and which are larger in scale than previously realized. To introduce this work, I will highlight three important subfields of development within ultracold atomic physics which seek to address these challenges: cavity quantum electrodynamics (cQED), quantum gas microscopy, and optical microtweezers.

Cavity QED

The Jaynes-Cummings Hamiltonian describes the coupling between a two-level system and a harmonic oscillator [14]. cQED is the experimental study of physical systems which can be modeled by the Jaynes-Cummings Hamiltonian. Often the two-level system represents two electronic levels of an atom and the harmonic oscillator represents the photonic mode of an electromagnetic resonator such as an optical cavity. Early efforts in cQED focused on the realization of basic elements within the theory, such as observations of the vacuum Rabi splitting of an optical cavity due to the presence of an atomic ensemble or sin-

gle atom within the cavity [15–18]. More recently, single-atom cQED systems have been used as prototypical elements within quantum networks [19]. Examples include the cavity-based transduction between atomic and photonic qubits or mediation of photon-photon quantum gates [20]. Alternatively, cQED systems with many atoms have been used, amongst other applications, to perform quantum simulations of lattice and other many-body Hamiltonians [21]; to prepare squeezed spin states with prospective applications for enhanced metrology, including improved atomic clock stability [22–26]; to measure quantum systems with precision fundamentally limited by quantum mechanics [27–29]; for novel laser cooling techniques [30, 31]; and to explore collective quantum phase transitions and phases such as the Dicke phase transition and the supersolid phase [32–34].

Quantum Gas Microscopes

Ultracold atomic samples often contain a large number of atoms prepared in quantum phases of matter. This makes them natural candidates to study many-body quantum physics, both to explore the fundamental physics of quantum phases of matter, but also as a sort of toy model analog for complex solid state materials. Quantum phases of matter have been explored both in bulk atomic gases, which are trapped weakly within large-volume optical or magnetic traps, and also in lattice-trapped gases, which are held in optical lattices and whose behavior resembles that of electrons in crystals. Such quantum simulator systems have been used to explore the physics of the BEC/BCS crossover, quantum magnetism models such as the Harper-Hofstadter and Haldane Hamiltonians, geometric frustration, synthetic gauge fields, topological phases, and more [35–39].

An extremely powerful tool that has been developed recently for atomic quantum simulation experiments is the quantum gas microscope [40, 41]. In such microscopes, a high-numerical aperture (NA) imaging system is integrated within the already complex experimental setup required to produce quantum gases. This microscope allows one to image individual atoms by collecting their fluorescence. The optical resolution can be sufficient to resolve the atoms in each site of an optical lattice. Such imaging provides a near-complete view of the quantum dynamics that occur in atomic many-body quantum systems. Quantum gas microscopes have led to numerous scientific breakthroughs in recent years including, but not limited to, studies of the superfluid/Mott insulator transition, anti-ferromagnetism, many-body localization and strongly interacting Fermi gases [40, 42–44].

Microtweezer Arrays

When atoms are trapped in extremely tightly focused optical dipole traps (ODTs) [45], it is possible to leverage strong light-assisted collisions, which occur due to the high densities in the small traps, to repeatably load single atoms within the trap volumes [46]. Atoms can be held with very long lifetimes and imaged with high fidelities in these tight traps, which are referred to as microtweezers. By trapping multiple atoms within an array of such tweezers and exciting them to Rydberg states, one can cause the atoms to interact. The

combination of these interactions and single atom quantum state control protocols allows researchers to perform high-fidelity quantum gates on the atoms within the array, making such systems good candidates for quantum information processing [47, 48]. Progress has been made towards decreasing the entropy in microtweezer arrays through dynamic array reconfiguration techniques, novel loading protocols, and laser cooling techniques [49–54]. Since it is possible to engineer the Hamiltonian within these arrays, it is also possible to use them for quantum simulation applications [55].

This Work

The work I present here is the culmination of an effort to build an atomic physics apparatus capable of combining these three technologies. Such a goal can be motivated in a number of ways.

On the one hand, the cavity can be seen as a tool to enhance a quantum many body system in a quantum gas microscope or microtweezer array by providing (1) the ability to engineer new, long-range interactions, (2) the ability to measure the state of the atoms within the system non-destructively and (3) the ability to introduce controlled dissipation into the closed many-body system, allowing for the exploration of dissipative many-body quantum dynamics.

On the other hand, cQED has often been restricted to focusing on symmetric and global interactions between atoms within the cavity because of the symmetric coupling of the atoms to the optical field. This has restricted cQED to looking at physics with few-degrees of freedom; either the interaction of a small number of quantum systems within the cavity [56, 57], or mean-field dynamics of a single many-body system. In this new apparatus, the quantum gas microscope or microtweezer technologies would give us the ability to address individual partitions of the quantum gas or individual atoms within a microtweezer array for either local Hamiltonian engineering or local measurement.

This apparatus then requires an in-vacuum high-finesse optical cavity to facilitate cQED as well as a high-resolution addressing system with the ability to image atoms in the focal plane and project arbitrary optical potentials with a spatial resolution at the single μm level. One of the most exciting prospects for such an apparatus would be the ability to perform real-time, quantum-limited, continuous measurements of the local state of the quantum system within the cavity, and then processing and feeding this information into the high-NA addressing systems. The high-NA addressing system could then project optical potentials onto the atoms to steer their evolution towards some desired, novel, quantum many-body state. Such an experiment would push the frontiers for research in open many-body quantum systems and quantum feedback, and could have applications for error correction protocols in quantum information processing systems.

This dissertation will be a detailed discussion of the design, construction and operation of the apparatus described above, named E6. My hope is that this work will contribute, if in a small way, towards the collective knowledge about how to control, measure, and understand quantum systems to help bring us closer to the day when we talk about, like the improvements
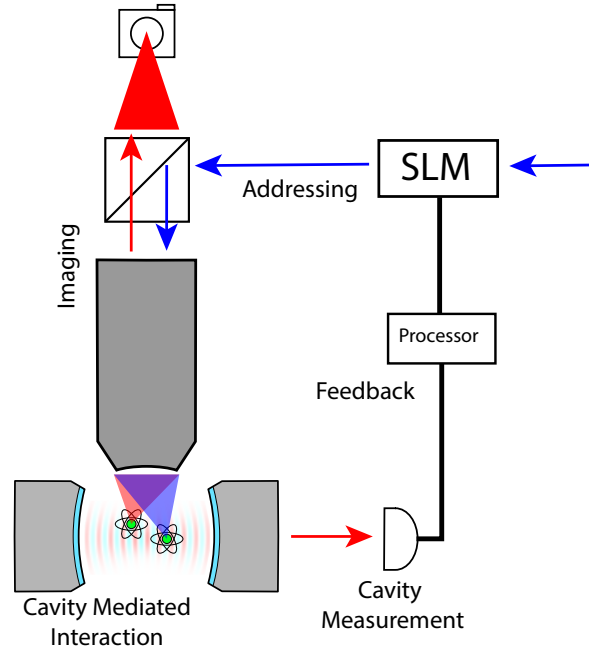
Figure 1.1: High level cartoon describing the E6 apparatus. Atoms are confined in the center of a high-finesse optical cavity. We are able to deliver a number of different types of atomic samples ranging from single atoms to Bose-Einstein condensates. The cavity can be used to mediate interactions between the atoms and to continuously and non-destructively read out the quantum state. A high-NA objective oriented transverse to the cavity axis can image the atoms with high resolution and project arbitrary optical patterns onto the atomic sample using some type of spatial light modulator. At present we control the spatial modes of light entering the objective using an acoust-optic deflector. One vision for this experiment is that we can close a feedback loop involving the quantum atoms, the cavity measurement stream, and the high-NA addressing system to explore quantum feedback applications.

to society brought on by physicists' understanding of electricity and magnetism, the great benefits to society brought on by quantum science and quantum technologies.

In the rest of this chapter, I will first give a personal account of my time in the Stamper-Kurn ultracold atoms research group at UC Berkeley, then I will give an introduction to the Jaynes-Cummings Hamiltonian and cQED generally, followed by a summary of what is to follow in the remainder of this dissertation.

1.2 My Time in the Ultracold Group

The Early Days

I joined the Stamper-Kurn Ultracold Atomic Physics research group in 2013, near the end of my first semester as a graduate student in the Berkeley physics department. I had completed my bachelors degree at the University of Colorado, Boulder the previous Spring where I had worked on superresolution infrared near-field scanning probe microscopy where I gained a curiosity and appreciation for the optical diffraction limit which has served me through my PhD as well [58].

Joining the ultracold group, I initially began a rotation with the E3[1] cavity experiment working on a small project involving improving the group's homebuilt external cavity diode laser (ECDL) design. After a short time, an opportunity opened up to work on a bigger project, the so-called E6 experiment.

E6 was to be a successor experiment to E3; a new cavity experiment that would study many-body quantum systems within the cavity using a high-resolution objective to probe and read out that quantum system. At that time, however, the full E6 apparatus was very much still a distant dream. What was known about that experiment was that it would utilize a 2D MOT for rapid loading of the 3D MOT rather than a Zeeman slower as had been used in E1, E2, E4, and E5 or a background vapor as had been used in E3. Work to explore this new technology had begun years prior to my joining the group with the work of an ambitious undergraduate, Maryrose Barrios, to develop a prototype 2D MOT system. This system included a ultra-high vacuum (UHV) vacuum chamber, magnetic coils, and optics to generate a 2D MOT in a glass cell. Originally, this system was built in the E4 labroom and used the E4 laser systems to generate the cooling and repump light needed for the 2D MOT. In early 2014 students were working to build an independent laser system for that prototype system in a small labspace. However, all of those students graduated or rotated away from the project.

I was presented with the opportunity to switch to working on this prototype 2D MOT system which I accepted. This was a great chance for me to learn independently about many of the critical components for an atomic physics apparatus — optics, spectroscopy, electronics, magnetic coils, and vacuum systems — in a setting less elaborate than a full-blown atomic physics experiment. During this time, with the support of E3 senior grad student Sydney Schreppler and E4 senior grad student Ed Marti, I was able to realize a 2D MOT in this prototype system over the course of the next year.

After reaching that milestone I faced a decision about whether I wanted to continue working on E6, committing to building out an entire atomic physics apparatus. I decided to accept this task, realizing I would learn a lot and that it would be challenging, yet fulfilling. Indeed my experience has proven those predictions to be true.

---

[1]The experiments within the group are labeled by the chronological order in which they were assembled. E1 and E2 were Dan's first experiments and had been taken offline well before I joined the group. E3, E4 and E5 were all operational when I joined the group. Now even E8 and E9 are under construction!

I was joined in late 2014 by a rotating graduate student Vicky Xu. We commenced design for the 2D and 3D MOT UHV chambers for E6 which will be described in Chapter 2.

We faced another difficult decision about where to build the new experiment. One natural choice was to build E6 in the same room as the E3 apparatus and figure out a scheme to initially share lasers with the E3 apparatus (as E3 had shared with E2 when it was first built) and eventually take over the E3 laser system. Another choice was to build the experiment in a lab space in the new Campbell Hall astronomy building. This building had been demolished in 2012, rebuilt for earthquake compliance, and re-opened in January 2015. The basement of that building was specially designed for physics labs and our group happened to have space in one of the empty rooms in that basement.

We had hesitations about starting our lab in the new building because we understood that we would need to rebuild a lot of the infrastructure that was already in place in the Birge labs. We also understood that we would be physically distanced from the other labs making it more inconvenient to pop into the Birge labs to check out spare equipment or chat with another lab member. In the end, however, we were excited to build a new robust, clean, and versatile apparatus in the new labspace where we would have the space we needed to make everything to our liking and to the required specifications for a modern apparatus. Fig. 1.2 shows a comparison between the lab as it looked the first day I entered it and as it looks now at the end of my PhD.

During 2015 Vicky, a new rotating graduate student Emma Deist, undergraduate Armando Montejano, and I continued the design for the 2D and 3D MOT vacuum chambers and optics and assembled and baked out the chambers at the end of that year.

Spending Time on E3

In the beginning of 2016, Vicky and Emma rotated out the group and, rather than continue building E6, we decided that it would be best for me to temporarily switch experiments to work on E3 with graduate student Jonathan Kohler, who had become an expert on that experiment over his years within the group. This would give me the opportunity to learn how to operate an active atomic physics experiment and get some science results before continuing the endeavor to build the new apparatus! After a few months, very fortuitously for E6 and me, Emma, excited to build the new E6 experiment, decided to rejoin the research group.

Jonathan trained Emma and me to run the E3 apparatus. We learned about laser cooling, atom chips, intensity and Pound-Drever-Hall (PDH) feedback loops, Cicero experimental control,[2] data taking, data analysis and much more! We worked on an experiment in which a thermal gas of ^{87}Rb was harmonically trapped within an optical cavity and with a magnetic field transverse to the cavity axis. In this configuration the cloud underwent collective mechanical motion within the trap and spin precession about the transverse magnetic field. Under the appropriate conditions, we were able to use the optical cavity mode (1) to monitor

---

[2]Cicero is the GUI-based experimental control software our group uses to control our experiments [59].

(a)

(b)
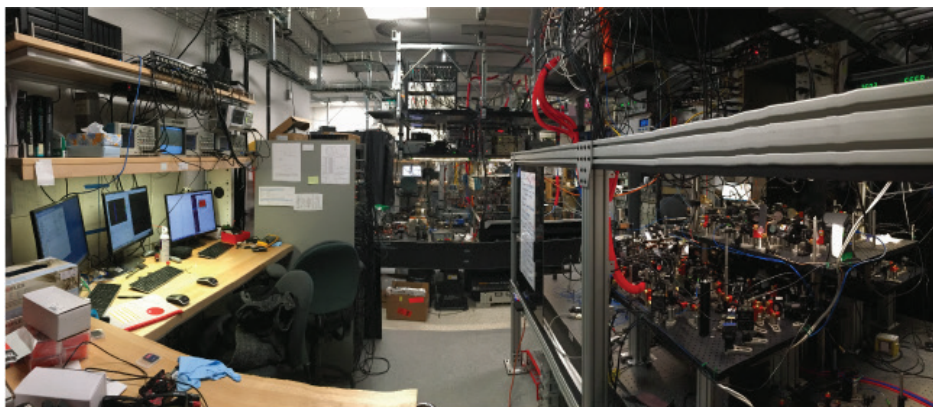
Figure 1.2: (a) Photograph of the new E6 lab space in LL104 Campbell Hall in early 2015 after completion of the Campbell Hall building. (b) Photograph of E6 lab space in Fall 2020.

both the spin and mechanical motion of the atomic cloud with very high sensitivity and (2) to couple the spin and motion of the atomic cloud. We spent a year and a half on this project, culminating in the publication of one experimental paper describing our results [56] and a second paper describing details of a tailored, matched-filter signal processing technique we had used to extract quantum information about the two oscillators within the cavity from the single photon mode exiting the optical cavity [60].

Working on the Atom Preparation System

After completing this work, Emma and I were ready to return to the E6 lab and load a 3D MOT into the vacuum chamber which had been sitting unused for over a year! We were joined by postdoc Dr. Shantanu Debnath. Shortly after moving back into E6 we did, in fact, realize a 2D MOT and our first 3D MOT before the end of 2017.

As will be explained in the following chapters, E6 utilizes a two-chamber design and we use optical transport to move the atoms from the 3D MOT chamber into the second UHV chamber called the science chamber. Our next task was to learn how to do optical transport. However, we did not yet have a second chamber into which we could transport atoms! We realized it would take us a long time to design and build the science chamber and science cavity. Rather than let this design hold up our progress with the atoms, we decided to quickly design and build what we called the 'dummy' chamber — a simple UHV chamber which we could temporarily install onto the atom preparation chamber to 'practice' optical transport while concurrently designing the science chamber. During 2018, we built the dummy chamber and optical transport system and successfully transported atoms into a secondary chamber. After success with optical transport, the science chamber was still not yet ready but we decided to press on with the atoms. We spent the final months of 2018 and first months of 2019 (with Shantanu having left the group and new postdoc Dr. Johannes Zeiher now having joined the group) working successfully to realize BEC in both the atom preparation chamber and dummy chamber.

After realizing a BEC in April 2019, Emma, Johannes, three undergraduates Aron Lloyd, Alec Bohnett, and Rachel Tsuchiyama and I turned all of our efforts towards the construction of the E6 science chamber. Our efforts were divided between finishing up the design and assembly of the science chamber, designing various optical systems such as the important high-NA system, working on various computer control systems, and building the science cavity. My main tasks during this time were the computer-assisted-design (CAD) for the science chamber, science cavity vibration isolation system, communicating with the fabrication companies for the science cavity mirrors, and building the science cavity.

The Science Cavity Mirrors

The procurement of the science cavity mirrors was a saga unto itself. The mirror fabrication involved two vendors. Perkins Precision Developments (PPD) made the custom superpolished glass substrates for the mirrors and FiveNine Optics (FNO) performed the

optical coatings. The first purchase order for substrates was made by Shantanu back in 2017. The substrate fabrication pushed the state of the art so it took months to make this first batch of mirrors. At first, this delay was ok because we were occupied with other tasks in the lab and we were still deliberating about which coatings we wanted. In early-mid 2018 we did decide which coatings we wanted and we had the substrates sent to FNO for coating. Finally at the end of summer 2018 the mirrors were being coated.

In August 2018 the first batch of cavity mirrors were shipped to us. Tragically, however, we discovered that they had been lost in the mail. UPS informed us that the mirrors had been delivered but the receiving dock at the Space Sciences Lab (where the mirrors had been accidentally shipped instead of LeConte Hall) did not receive the mirrors. An investigation was launched but nothing was turned up. These mirrors were very expensive and were the result of a year-long design and fabrication process. This taught us hard lessons about the importance of shipping terms,[3] shipping insurance, and the importance of requiring signatures upon receipt.

We restarted the fabrication with PPD. This time, PPD had problems with their superpolishing process, which extended that phase months beyond the expected lead time. We finally received this batch of mirrors at the end of summer 2019, at which time we were completing designs for the science chamber and beginning the cleaning and assembly process for all of the parts.

The Final Assembly and First Results

E6 had a massively productive push in the Fall of 2019. The science chamber and science cavity were completed, the vacuum was closed before Christmas, and the bakeout was completed early January 2020. The following months were again very productive, with us successfully transporting atoms into the new chamber with ease. In February and March, we built out the optical systems for the cavity, including input optics and heterodyne detection, allowing us to make the first measurement of a dispersive shift of the cavity resonance frequency due to the presence of the atoms.

We also completed the high-NA imaging system and acousto-optical deflector (AOD) mi-crotweezer delivery system allowing us to realize our first fluorescence images of atoms in the transport ODT and the first signatures of trapping of single atoms using the microtweezers.

Unfortunately, this progress was abruptly halted in mid-March by the COVID-19 global pandemic when the county of Alameda and UC Berkeley issued shelter-at-home orders instructing us to stay out of the lab. During quarantine, my most productive contributions were beginning to write the text for this dissertation and upgrading our image acquisition software, JKam,[4] to make it more modular, easier to read and maintain, and adding new

---

[3]'FOB destination' means the seller is liable for the goods until they reach the buyer's dock. These are favorable terms for buying important parts. A responsible vendor would make sure to insure an expensive shipment under such terms and may pass this cost onto the buyer.

[4]This software was written by Johathan Kohler for use in E3 and was inspired by the older Edcam software, written by Ed Marti, which had been the main imaging software used in the group for years.

features necessary for the E6 experimental flow.

During this time, E6 was very fortunate to have postdoc Dr. Fang Fang join the E6 effort after she finished her PhD on E4. Fang helped us greatly by sharing her expertise trapping atoms in ODTs, doing a number of numerical simulations of various system parameters and effects, and helping run the experiment remotely. We were also joined, remotely, by the eager undergraduate Che Liu who worked diligently on machine learning algorithms to help us analyze our microtweezer fluorescence images and numerical simulations of cavity induced spin-squeezing.

I was able to come back into the lab at the start of July 2020, at which time I resuscitated the experiment and began to perform more quantitative calibrations. We were able to realize single atoms trapped in multiple microtweezers and begin to see an interaction between atoms trapped in the microtweezers and optical modes supported by the cavity. Before I transitioned out of the lab, we were able to perform a measurement in which a single atom in a microtweezer was used as a local probe of the spatial intensity pattern of the cavity ODT which is described in Chapter 8.

During my final months in lab, I transferred what knowledge I could about running our experiment over to the new graduate student Leon Lu who has learned how to run the apparatus with astonishing alacrity. I am pleased to be leaving the experiment in the very capable hands of Emma, Leon, and Che and I cannot wait to see what great ideas they come up with and explore with the E6 apparatus!

1.3 Cavity QED

The two simplest quantum mechanical systems are the two-level system, which has two energy levels, $|g\rangle$ and $|e\rangle$, and the quantum harmonic oscillator, which has levels linearly spaced in energy. The Jaynes-Cummings Hamiltonian describes these two systems and their interactions as follows:

$$\hat{H}_{\mathrm{JC}} = \hbar\omega_c \hat{a}^\dagger \hat{a} + \hbar\omega_a \hat{\sigma}^\dagger \hat{\sigma} + \hbar g \left(\hat{a}^\dagger \hat{\sigma} + \hat{a}\hat{\sigma}^\dagger \right). \tag{1.1}$$

Here $\hbar$ is Planck's constant, ω_c is the harmonic oscillator resonance frequency, $\hat{a}$ is the harmonic oscillator amplitude operator, ω_a is the two-level-system energy splitting, $\hat{\sigma} = |g\rangle \langle e|$ is the two-level-system lowering operator and g is the coupling parameter.

This Hamiltonian can be realized in physical systems which include an electromagnetic resonator, such as an optical cavity, and an atom which can act as a two level system. The field of research studying such systems is the field of cQED. In the following sections I will demonstrate the basic mathematical formalism which shows the parallel between an atom in a cavity and the Jaynes-Cummings Hamiltonian.

1.3.1 The Quantization of the Electric Field: A Crash Course

In this section, I will give a rapid account of how to write down the formula for the quantized electric field based largely on a dimensional analysis with the goal of building intuition about the relationship between the amplitude of a classical electric field in an optical cavity $\boldsymbol{E}(\boldsymbol{r}, t)$ and the quantum operator $\hat{a}$ that describes the amplitude of a quantum electric field in the cavity. Much more thorough treatments can be found in Refs. [61–65].

Maxwell's equations are solved by spatio-temporal vector fields of the form

$$\boldsymbol{E}_{\boldsymbol{k}}^{(+)}(\boldsymbol{r}, t) = E_{\boldsymbol{k}}^{(+)} \boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r}) e^{-i\omega t}. \tag{1.2}$$

Here I have represented the positive rotating component of the electric field corresponding to temporal frequency ω and spatial wavevector $\boldsymbol{k}$. Maxwell's equations ensure $k = |\boldsymbol{k}| = \frac{\omega}{c}$. $E_{\boldsymbol{k}}^{(+)}$ is the complex amplitude for the field, and $\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})$ is the complex spatial mode shape, including the local polarization vector, for the solution. I take a Fourier transform convention for which $e^{-i\omega t}$ is a positive frequency function justifying the $(+)$ notation. See Appendix A for the definitions of Fourier transforms used in this thesis.

The total *real* electric field[5] can be written as

$$\boldsymbol{E}_{\boldsymbol{k}}(\boldsymbol{r}, t) = \boldsymbol{E}_{\boldsymbol{k}}^{(+)}(\boldsymbol{r}, t) + \boldsymbol{E}_{\boldsymbol{k}}^{(-)}(\boldsymbol{r}, t) = 2\mathrm{Re}\left(\boldsymbol{E}_{\boldsymbol{k}}^{(+)}(\boldsymbol{r}, t)\right). \tag{1.3}$$

The function $\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})$ satisfies the vector Helmholtz equation:

$$\left(\nabla^2 + k^2\right) \boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r}) = 0. \tag{1.4}$$

In a finite volume system, such as an optical cavity, $\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})$ is one element of a countable family of orthogonal mode functions that span the space of possible solutions to the Helmholtz equation. We are most familiar with plane wave solutions which and arise in free space and that look like

$$\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r}) = \boldsymbol{\epsilon} e^{i\boldsymbol{k}\cdot\boldsymbol{r}}, \tag{1.5}$$

where $\boldsymbol{\epsilon}$ is a possibly complex polarization vector satisfying $\boldsymbol{\epsilon} \cdot \boldsymbol{k} = 0$. In this work, we will see that $\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})$ might be a confined Hermite-Gaussian transverse mode of an optical cavity as described in Sec. 3.2.1. For confined mode functions $\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})$ we normalize so that

$$\max\left(|\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})|^2\right) = 1,$$
$$\int |\boldsymbol{f}_{\boldsymbol{k}}(\boldsymbol{r})|^2 \, d^3\boldsymbol{r} = V_{\mathrm{mode}}, \tag{1.6}$$

---

[5]Throughout this text the $(\pm)$ superscript will indicate the complex positive or negative rotating component of the field, whereas the lack of such a subscript will indicate the total real electric field. An alternative convention, $E = \mathrm{Re}\left(E^{(+)}\right)$, is also regularly encountered. This factor of 2 can lead to differences in equations for power and energy.

where I've defined the mode volume V_{mode}, which quantifies the spatial volume within which the electromagnetic field has appreciable amplitude and energy.

The total time-averaged energy in an electromagnetic field mode is given by

$$\langle U_{\text{EM}} \rangle = 2\epsilon_0 \int \left| \boldsymbol{E}_{\boldsymbol{k}}^{(+)}(\boldsymbol{r}, 0) \right|^2 d^3\boldsymbol{r} = 2\epsilon_0 V_{\text{mode}} \left| E_{\boldsymbol{k}}^{(+)} \right|^2. \tag{1.7}$$

Quantum mechanically, we know that the energy stored in the electromagnetic field is quantized in units of $\hbar\omega$ [61, 65]. This energy quantization means that the magnitude of the electric field is quantized as well. The minimum field magnitude (assuming a purely real field amplitude) is given by

$$2\epsilon_0 V_{\text{mode}} \left| E_{\omega,\text{SP}}^{(+)} \right|^2 = \hbar\omega,$$
$$E_{\omega,\text{SP}}^{(+)} = \sqrt{\frac{\hbar\omega}{2\epsilon_0 V_{\text{mode}}}}, \tag{1.8}$$

where the SP subscript indicates the electric field of a single photon. We see that the field amplitude per photon scales with $1/\sqrt{V_{\text{mode}}}$. The intuition here is that, since the energy per photon is fixed, if you reduce the volume over which the energy of a single photon is distributed, the local field amplitude must increase.

With this in mind we can write down the form for the quantum electric field by taking

$$E_{\boldsymbol{k}}^{(+)} e^{-i\omega t} \to E_{\omega,\text{SP}}^{(+)} \hat{a}_{\boldsymbol{k}}(t). \tag{1.9}$$

Here $\hat{a}_{\omega,\boldsymbol{k}}(t)$ is a dimensionless quantum operator that carries information about the quantum amplitude of the electric field and satisfies (dropping subscripts) the bosonic commutation relation

$$\left[\hat{a}, \hat{a}^\dagger \right] = 1. \tag{1.10}$$

With this we have that $\hat{n} = \hat{a}^\dagger \hat{a}$ is a number operator with integer eigenvalues that count the number of energy excitations (photons) in the electromagnetic field.

Putting this all together we have, summing over all electric field modes,

$$\hat{\boldsymbol{E}}(\boldsymbol{r}, t) = \sum_{\boldsymbol{k},s} \left(E_{\omega,\text{SP}}^{(+)} \boldsymbol{f}_{\boldsymbol{k},s}(\boldsymbol{r}) \hat{a}_{\boldsymbol{k},s}(t) + h.c. \right),$$
$$\hat{H}_{\text{EM}} = \sum_{\boldsymbol{k},s} \hbar\omega \hat{a}_{\boldsymbol{k},s}^\dagger \hat{a}_{\boldsymbol{k},s}. \tag{1.11}$$

I've added an index s to indicate there are two possible polarization for each mode function. Note the consistency between these expressions, which arise in a quantized treatment, and the similar classical expressions.

1.3.2 Electric Dipole Interaction

Photons interact with atoms by applying electric forces that redistribute the valence electron wavefunction about the atomic nucleus. In particular, for electric dipole transitions, the electric field induces a dipole moment between the free electrons and the nucleus. I'll consider single valence electron atoms here.

Classically the interaction energy for a dipole in an electric field is given by $H_{\text{int}} = -\boldsymbol{E} \cdot \boldsymbol{d}$. The corresponding quantum Hamiltonian is

$$\hat{H}_{\text{int}} = -\hat{\boldsymbol{E}} \cdot \hat{\boldsymbol{d}}, \tag{1.12}$$

where $\hat{\boldsymbol{d}} = -e\hat{\boldsymbol{r}}$ is the dipole moment resulting from the separation between the nucleus and electron. In atomic systems, because of the anharmonicity of the atomic energy spectrum, it is often the case that light of a particular frequency only strongly couples two specific states with an energy splitting close the optical frequency. In this case we can consider the atom to be a a two level system with states $|g\rangle$, $|e\rangle$. Because the dipole operator is odd under a parity transformation, it only couples states of opposite parity. If we assume $|e\rangle$ and $|g\rangle$ are atomic states of opposite parity, then we have

$$\begin{aligned} \hat{\boldsymbol{d}} &= \langle g| \, \hat{\boldsymbol{d}} \, |e\rangle \, |g\rangle \, \langle e| + \langle e| \, \hat{\boldsymbol{d}} \, |g\rangle \, |e\rangle \, \langle g| \\ &= \boldsymbol{d}_{ge}\hat{\sigma} + \boldsymbol{d}_{eg}\hat{\sigma}^{\dagger} \\ &= \hat{\boldsymbol{d}}^{(+)} + \hat{\boldsymbol{d}}^{(-)}. \end{aligned} \tag{1.13}$$

The $(\pm)$ superscripts indicate whether the Heisenberg version of the respective operator under the free atomic Hamiltonian is positive or negative frequency according to our convention. For a classical plane wave drive field under the rotating wave approximation and in the rotating frame, the interaction Hamiltonian for an atom at position $\boldsymbol{r}$ reduces to

$$\hat{H}_{\text{int}} = -E^{(-)}\boldsymbol{f}^{*}(\boldsymbol{r}) \cdot \boldsymbol{d}_{ge}\hat{\sigma} - E^{(+)}\boldsymbol{f}(\boldsymbol{r}) \cdot \boldsymbol{d}_{eg}\hat{\sigma}^{\dagger}. \tag{1.14}$$

With this we define the Rabi frequency as

$$\Omega(\boldsymbol{r}) = -2E^{(-)}\frac{\boldsymbol{f}^{*}(\boldsymbol{r}) \cdot \boldsymbol{d}_{ge}}{\hbar}, \tag{1.15}$$

leading to the interaction Hamiltonian

$$\hat{H}_{\text{int}} = \frac{\hbar\Omega(\boldsymbol{r})}{2}\hat{\sigma} + \frac{\hbar\Omega^{*}(\boldsymbol{r})}{2}\hat{\sigma}^{\dagger}. \tag{1.16}$$

If the electric field polarization vector is constant as a function of space then $\boldsymbol{f}(\boldsymbol{r}) = f(\boldsymbol{r})\boldsymbol{\epsilon}$ and we can define $d_{ge} = \boldsymbol{\epsilon}^{*} \cdot \boldsymbol{d}_{ge}$. Recalling that $f(r)$ is normalized so that its peak magnitude

is unity, we can write down the formula for the maximal coupling strength within the mode volume

$$\Omega_0 = -2\frac{E^{(-)}d_{ge}}{\hbar},$$
$$|\Omega_0| = \frac{|E||d_{ge}|}{\hbar},$$
$$\Omega(\boldsymbol{r}) = \Omega_0 f(\boldsymbol{r}). \tag{1.17}$$

The factor of 2 in the definition of Ω is conventional. Also, for many problems, we are free to choose phases for the quantum states and the electric field so that Ω is purely real.

For the case of a quantum drive field, making the transformation in Eq. (1.9), the interaction Hamiltonian becomes

$$\hat{H}_{\text{int}} = -E^{(-)}_{\omega,\text{SP}}\boldsymbol{f}^*(\boldsymbol{r}) \cdot \boldsymbol{d}_{ge}\hat{a}^\dagger\hat{\sigma} - E^{(+)}_{\omega,\text{SP}}\boldsymbol{f}(\boldsymbol{r}) \cdot \boldsymbol{d}_{eg}\hat{a}\hat{\sigma}^\dagger. \tag{1.18}$$

We define the quantum coupling parameter:[6]

$$g(\boldsymbol{r}) = -E^{(-)}_{\omega,\text{SP}}\frac{\boldsymbol{f}^*(\boldsymbol{r}) \cdot \boldsymbol{d}_{ge}}{\hbar} = -\sqrt{\frac{\omega}{2\epsilon_0\hbar V_{\text{mode}}}}\boldsymbol{f}^*(\boldsymbol{r}) \cdot \boldsymbol{d}_{ge}. \tag{1.19}$$

There is no factor of 2 here as there was in the definition of Ω, again, by convention.

For a constant polarization vector we can define d_{ge} as above and write down the peak coupling parameter:

$$g_0 = -\frac{E^{(-)}_{\omega,\text{SP}}d_{ge}}{\hbar} = -\sqrt{\frac{\omega}{2\epsilon_0\hbar V_{\text{mode}}}}d_{ge},$$
$$|g_0| = 2\frac{|E_{\omega,\text{SP}}||d_{ge}|}{\hbar},$$
$$g(\boldsymbol{r}) = g_0 f(\boldsymbol{r}). \tag{1.20}$$

We see that g_0 quantifies the dipole energy associated with the electric field of a single photon. The dipole interaction Hamiltonian is then given by

$$\hat{H}_{\text{int}} = \hbar g(\boldsymbol{r})\hat{a}^\dagger\hat{\sigma} + \hbar g^*(\boldsymbol{r})\hat{a}\hat{\sigma}^\dagger, \tag{1.21}$$

the interaction term in the Jaynes-Cummings Hamiltonian Eq. (1.1)!

Note that the transformation in Eq. (1.9) is equivalent to

$$\frac{\Omega}{2} \to g\hat{a}. \tag{1.22}$$

---

[6]Avoid confusion between the coupling strength g and the ground state $|g\rangle$.

The discussion so far has assumed we are working with a two-level atom. In practice we work with real atoms which have many energy levels, but, which can, under certain circumstances, be approximated as two-level atoms. In E6, we work with ^{87}Rb atoms and typically work on the $5\,S_{1/2} \to 5\,P_{3/2}$ transition, called the D_2 transition. This transition couples two hyperfine ground-state manifolds (and a total of eight magnetic sublevels) to four excited-state hyperfine manifolds (and a total of 16 magnetic sublevels). See Fig. 2.3 for a level diagram. Details about how to calculate transition dipole elements for such a multilevel atom can be found in Appendix B.

For our purposes here, it will suffice to utilize two transition dipole elements. First, we define d_0 to be the transition dipole element for the D_2 cycling transition $|F = 2, m_F = +2\rangle \to |F' = 3, m'_F = +3\rangle$ which has a numerical value of

$$d_0 = 2.534 \times 10^{-29}\,\mathrm{C\,m} = 2.989\ ea_0. \tag{1.23}$$

We will use $d_{ge} = d_0$ when defining g_0 for E6 so that

$$g_0 = -\frac{E_{\omega,\mathrm{SP}}^{(-)} d_0}{\hbar} = -\sqrt{\frac{\omega}{2\epsilon_0 \hbar V_{\mathrm{mode}}}} d_0. \tag{1.24}$$

Next, we define $d_{\mathrm{eff}} = \sqrt{\frac{2}{3}} d_0$, the so-called effective far-detuned transition dipole element [66] which has a numerical value of

$$d_{\mathrm{eff}} = 2.069 \times 10^{-29}\,\mathrm{C\,m} = 2.441\ ea_0. \tag{1.25}$$

The far-detuned transition dipole element is useful for determining average energy shifts due to linearly polarized far-detuned ODTs.

1.3.3 Energy Levels in the Jaynes-Cummings Hamiltonian

The Jaynes-Cummings Hamiltonian in the rotating frame is given by

$$\frac{\hat{H}_{\mathrm{JC}}}{\hbar} = \Delta \hat{a}^\dagger \hat{a} + g \left(\hat{a}^\dagger \hat{\sigma} + \hat{a} \hat{\sigma}^\dagger \right), \tag{1.26}$$

where $\Delta = \Delta_{CA} = \omega_C - \omega_A$ is the detuning between the cavity and atomic resonance frequencies. For a single atom in the cavity, the Hilbert space for this system is spanned by states of the form $|g, n\rangle$ and $|e, n\rangle$ where the first label indicates the state of the atom and the second label indicates the number of photons in the cavity.

In the single excitation subspace spanned by $|e, 0\rangle$ and $|g, 1\rangle$, the matrix for this interaction is given by

$$\frac{\boldsymbol{H}_{\mathrm{JC}}}{\hbar} = \begin{pmatrix} \Delta & g \\ g & 0 \end{pmatrix}. \tag{1.27}$$
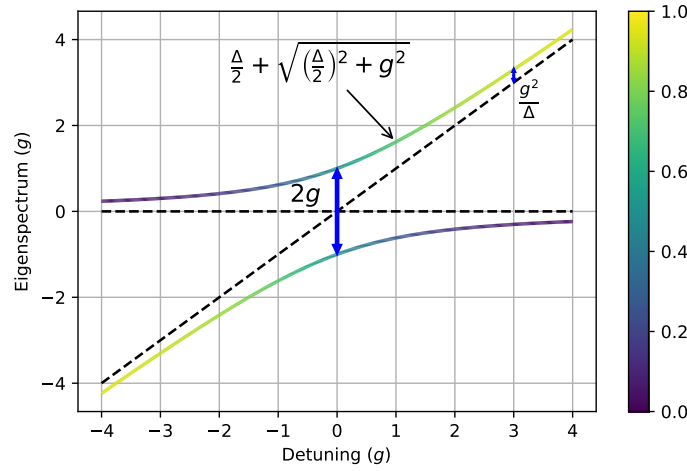
Figure 1.3: Eigenfrequencies of the cavity photon system as a function of the detuning, Δ, between the cavity photon and the atomic transition frequency. Uncoupled energy levels are shown as dashed black lines. The horizontal dashed line represents the state with an excited atom and no photons and the sloped dashed line represents the state with a ground state atom and a single photon. The full functional form for the eigenfrequencies is indicated (colored lines) as well as the on-resonance vacuum Rabi splitting of $2g$ and the off-resonance dispersive shift, $\frac{g^2}{\Delta}$ valid for large detuning. The colormap indicates the squared amplitude of the overlap of the eigenstates with the single photon excited state.

The eigenvalues, or energy levels, for this Hamiltonian are given by

$$\frac{E_{\pm}}{\hbar} = \frac{\Delta}{2} \pm \sqrt{\left(\frac{\Delta}{2}\right)^2 + g^2}. \tag{1.28}$$

This spectrum is depicted in Fig. 1.3.

In absence of the cQED interaction, for $\Delta = 0$, the two excited states, $|e, 0\rangle$ and $|g, 1\rangle$ are degenerate. The Jaynes-Cummings interaction lifts this degeneracy, leaving two states with energy levels split by the Rabi frequency $2g$, and strongly mixing the photonic and atomic excitations. I will refer to the regime for which $|\Delta| \ll g$ as the regime of resonant cavity quantum electrodynamics (QED). If $g \gg \kappa, \Gamma$ then if the cavity probe is swept across resonance, the transmission spectrum will exhibit two peaks corresponding to the two energy branches [16, 18].

The regime for which $|\Delta| \gg g$ is the dispersive regime of cQED. In this regime the energy of the photon-like branch is given by

$$\frac{E}{\hbar} \approx \Delta + \frac{g^2}{\Delta} = \Delta + g_C, \tag{1.29}$$

where $g_C = g^2/\Delta$ reflects the dispersive shift of the cavity resonance due to the presence of a single atom. In this regime, the admixture of the atomic excited state is $g^2/\Delta^2 \ll 1$. The physics here is analogous to the ac Stark effect and can be derived using second order perturbation theory.

In this case, if the cavity probe is swept across resonance, the spectrum would exhibit a single resonance which is slightly shifted in frequency by an amount g_C. We can approximate the Hamiltonian as

$$\hat{H}_{\rm JC} \approx \hbar \left(\Delta + g_C\right) \hat{a}^\dagger \hat{a}. \tag{1.30}$$

In the dispersive regime, as with the AC Stark effect, we can provide a semi-classical explanation for the shift in the cavity frequency. The atom can be thought of as a small, polarizable medium that exhibits an induced polarization when driven by the oscillating electric field of the probe beam. The electric field radiated by the driven atomic polarization interferes with the incident field to give a total field with a phase shift related to the atomic polarizability. The effect of this phase shift is to alter the cavity round trip phase resulting in a shift of the cavity resonance frequency as will be explored in Sec. 3.2.3. In this picture, the atom acts like a small, dispersive element within the cavity.

If there is an ensemble of N_A atoms in the cavity mode, the total energy shift will be $N_A g_C$. This shift in cavity resonance can be used as a sensitive measurement of the atom number. More generally, if g is position- or spin-dependent, then the cavity frequency shift can be used as measure of average position or spin of the atomic ensemble [22, 25, 56, 67–71]. Above I have described a scheme in which the cavity frequency shift is measured by sweeping the probe across cavity resonance. However, the cavity frequency shift can also be measured by monitoring the amplitude and/or phase of the transmitted cavity light because this also depends on the detuning of the incident probe light from cavity resonance.

1.3.4 The Cooperativity Parameter
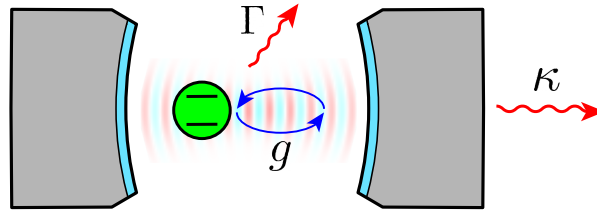
Figure 1.4: A schematic diagram showing the coupling g, between a two-level atom and an optical cavity photonic mode. The atom can decay via spontaneous emission at rate Γ, and the photonic mode can decay at rate κ due to a small fraction of light which exits, or leaks out of, the cavity. The cooperativity parameter is given by $C = 4g^2/(\kappa\Gamma)$.

Above I have shown how an atom placed within the optical mode of an optical cavity with a quantized electric field is modeled by the Jaynes-Cummings Hamiltonian, Eq. (1.1). This Hamiltonian describes the unitary, closed-system dynamics of a two-level system coupled to a harmonic oscillator. However, the real physical system involving the atom and cavity is not a closed system. There are two important loss channels which are very important for cQED: spontaneous emission from the atomic excited state $|e\rangle$ into free space with rate Γ, and leakage of the optical mode out of the optical cavity with rate κ. Here Γ and κ both refer to the energy (as opposed to amplitude) decay rates for the atom and cavity field, respectively.

The strong-coupling regime of cQED is defined as the regime in which the rate of coherent dynamics, g, exceeds the rates of incoherent dynamics κ and Γ. The cooperativity[7] is a dimensionless parameter which quantifies this comparison [17]:

$$C = 4\frac{g^2}{\kappa\Gamma}. \tag{1.31}$$

If the cQED system begins with an atomic or photonic excitation, then it will undergo Rabi oscillations at a frequency of $2g$. In the limit that $g \gg \kappa, \Gamma$, the number of Rabi oscillations between atomic and photonic excitations that can be completed, before the energy decays out of the system via κ and Γ, is proportional to $\sqrt{C}$. Purcell enhancement, another import cQED effect, is an effect whereby an atom in the excited state will exhibit enhanced spontaneous emission at a rate $(1 + C)\Gamma$ rather than Γ [74]. Importantly, this enhanced emission will be into the cavity mode. It turns out that the C is the appropriate figure of merit for a very wide range of other quantum control and quantum measurement tasks so cQED apparatuses are designed to realize a large cooperativity.

Following Ref. [75, 76], we can re-express C for a two-level atom by noting that both g^2 and Γ scale with d_{ge}^2:

$$g^2 = \frac{\omega d_{ge}^2}{2\epsilon_0 \hbar V_{\text{mode}}}, \tag{1.32}$$

$$\Gamma = \frac{\omega^3 d_{ge}^2}{3\pi\epsilon_0 \hbar c^3}, \tag{1.33}$$

leaving us with

$$\frac{g^2}{\Gamma} = \frac{3}{8\pi}\frac{\lambda^2 c}{V_{\text{mode}}} = \frac{3}{\pi^2}\frac{\lambda^2}{w_0^2}\frac{c}{2L_{\text{cav}}}. \tag{1.34}$$

We have used that $\omega/c = 2\pi/\lambda$ and, for a standing wave Gaussian mode with waist w_0 in a cavity with length L_{cav}, that $V_{\text{mode}} = \frac{\pi}{4}w_0^2 L_{\text{cav}}$ as shown in Eq. (3.74). For the ^{87}Rb D$_2$ transition $\Gamma = 2\pi \times 6.07\,\text{MHz}$ [66].

---

[7]A warning: different references include different numerical pre-factors in the definition of C. In particular, the E3 literature uses a definition of C which differs from this one by a factor of 2: $C_{E6} = 2C_{E3}$ [72, 73].

In Sec. 3.2.3, I will introduce the cavity free spectral range as $f_{\mathrm{FSR}} = c/2L_{\mathrm{cav}}$ and show that the cavity energy decay rate is given by $\kappa = 2\pi f_{\mathrm{FSR}}/\mathcal{F}$ where the cavity finesse $\mathcal{F}$, defined in Sec. 3.4, is inversely proportional to the cavity round trip survival probability for a photon. Plugging these expressions in we find

$$C = \frac{6}{\pi^3} \frac{\lambda^2}{w_0^2} \mathcal{F} = \frac{12}{\pi^2} \frac{\sigma_0}{w_0^2} \mathcal{F}, \tag{1.35}$$

where I've recalled that $\sigma_0 = \frac{\lambda^2}{2\pi}$ is the far-field resonant scattering cross section for the atomic transition. We see that the cooperativity is given by the ratio of the resonant scattering cross section to the mode waist of the beam with an enhancement by the cavity finesse. This lends a geometric interpretation to the cooperativity as the ratio of the free-space to cavity emission probability for the atom [71, 75]. I will explore the implications of this interpretation for cavity design in Chapter 4.

The calculation above used g^2 and Γ as defined for a two level atom. The results would vary by a geometric factor for a multilevel atom. However, if we are working on the cycling transition then both g_0^2 and Γ scale with $d_{JJ'}^2(2J+1)/(2J'+1)$ so the same result holds true:

$$C = \frac{4g_0^2}{\kappa\Gamma}. \tag{1.36}$$

The cooperativity does not depend on any details of the atomic transition other than the wavelength or scattering cross section. This tells us that it would be easier to reach the strong coupling regime for cQED using longer wavelength transitions. Perhaps, then, it is no surprise that some early success was found in cQED systems which operated on a microwave transition [77].

1.4 Outline

Most of my PhD was spent designing and building the E6 apparatus. As a result, most of this dissertation will also be about the design and assembly of the apparatus for the benefit of future E6 lab workers and for future builders of new experiments who are seeking clarification or inspiration. However, along the way, I spent a good deal of time identifying general design principles and 'clear ways of thinking about certain problems' that I think are generally useful for an atomic physicist, so, likewise, I will point out these insights along the way through this text.

In Chapter 2, I describe our atom preparation system. This system includes a 2D and 3D MOT UHV chamber, an optical transport scheme, and results on the realization of a BEC. Chapters 3, 4, and 6 are heavily focused on the science cavity, which is the focal point of the apparatus and which was, in many ways, the focal point of my PhD. Chapter 3 presents the general theory and formalism for the physics of an optical cavity. This chapter should provide the reader with good intuition for how to think about cavities generally and serve

as a reference for various definitions and derivations. Chapter 4 details the science cavity design, fabrication, assembly, and characterization. Chapter 6 gives details about the cavity probe and cavity ODT laser systems including various frequency locks and our heterodyne detection system. Chapter 5 describes, in detail, the design and construction of all of the components within the science chamber including the cavity vibration isolation system (excluding the science cavity which is addressed in the previously mentioned chapters). Chapter 7 describes the designs and characterizations for the high-NA addressing system including the microtweezer generation system and the fluorescence imaging systems. Chapter 8 lays out and discusses the main experimental results we realized upon completion of the assembly. These results include:

- a measurement of a dispersive shift of the cavity resonance due to the presence of atoms within the cavity indicating the realization of a cQED interaction,

- high-fidelity imaging of single atoms trapped in microtweezers,

- and a measurement in which single atoms trapped in microtweezers are scanned in position to map out the spatial amplitude distribution of a mode of the optical cavity.

These results demonstrate simultaneous operation of the two key technologies for the E6 apparatus: the high-finesse cavity and the high-NA addressing system. Finally, Chapter 9 will outline future science targets for this experiment.

Chapter 2

Atom Preparation

Ultracold atomic physics experiments rely on multiple stages of laser and evaporative cooling to prepare an ultra-low entropy atomic sample for subsequent quantum control and measurement operations. There are as many schemes for producing such samples as there are research groups working within the field.

The atom preparation scheme required depends on the number of atoms and their states that will be used in the final stages of the experiment. Early on in the design phase for E6, we considered the following possibilities for the types of atomic samples we would like to prepare and investigate in our apparatus.

- Thermal atomic clouds trapped in ODTs containing 10^2 to 10^5 atoms [21, 22, 56, 78].

- Scalar or spinor BECs in loose harmonic traps containing up to 10^6 atoms [79–81].

- BECs loaded into optical lattices such as those studied in quantum gas microscopes to allow us to study Hubbard physics. Such a lattice may include few to hundreds of sites with zero to few atoms each [40, 41, 82–84].

- Arrays of single atoms trapped in tightly confining optical microtweezers [46, 49–51, 53].

Each option enables a certain range of science to be targeted but also comes with trade-offs in terms of experimental complexity, cycle time and build time. For example, small-scale microtweezer arrays require only a 3D MOT and small cycle times, whereas BEC experiments require very low vacuum pressures and additional steps of laser cooling and evaporation, leading to long cycle times. While we had interest in microtweezer arrays and quantum gas microscopes when we began designing E6, these were both new technologies to the group compared to more traditional BEC machines.

In the end, we designed the apparatus so as to not exclude any of these possibilities. While this may have resulted in a more complex and costly apparatus in terms of money and build time, it has resulted in a very versatile apparatus with which we can explore

almost any type of single species bosonic atomic sample ranging from single or few atoms up to BECs of hundreds of thousands of atoms.

To facilitate the production of BECs, an optical cavity, and a high-NA objective, we elected to use a two-chamber design. The initial stages of atom preparation, including the 3D MOT, are realized in a first chamber, called the atom preparation or 3D MOT chamber. After atoms are trapped in the 3D MOT, they undergo additional cooling and are subsequently trapped in the focus of an ODT and translated into a second chamber, the science chamber. In the science chamber the atoms interface with the optical cavity and high-NA objective. This two-chamber design allowed us to realize the most geometric flexibility in the design of the science chamber by obviating the need for MOT beams.

Throughout the design process we were aware of the importance of experimental cycle time, an increasingly important specification for cold atoms apparatuses. Such experiments are typically operated as follows. First a room temperature, or hotter, source of atoms supplies a portion of the vacuum chamber with a gas of atoms of the appropriate species. These atoms are then manipulated and cooled in subsequent stages. Finally the science experiment, utilizing custom-designed optical, magnetic, radio frequency (rf) and microwave fields, and measurement schemes, is performed. This whole cycle from atom preparation to measurement has historically lasted about 10s of seconds [8, 9]. However, in recent years, advances have been made to decrease experimental cycle times to the level of a few seconds or less. These advances include improved evaporation techniques, all-optical trapping and cooling methods, atom chips, and rapid loading of microtweezer arrays [49, 50, 85–90].

Short cycle times allow for many rapid repetitions of experimental sequences, so that researchers can quickly collect sufficient statistics to quantify subtle quantum mechanical effects and characterize very high-fidelity quantum gates [91]. Additionally, a short cycle time allows for more agile testing and reconfiguration of the apparatus. These rapid adjustments allow for fast prototyping of new ideas, and also for an effective means of identifying and removing sources of systematic bias in measurements.

2.1 Atom Preparation Vacuum Chamber

Ultracold atoms are delicate. A collisions with a room temperature molecule will impart enough energy to eject a neutral atom from any trap that holds it. Such collisions then limit the experimental duration. For this reason, ultracold atomic physics experiments are evacuated as much as possible of residual room-temperature gas, achieving ultimate vacuum pressures in the range of 1×10^{-11} torr.

The UHV chamber must be designed with considerations for ultimate pressure, optical access, proximity of various field coils (magnetic, rf, microwave), atomic transport, and any additional considerations that may be necessary in a particular experiment. In this section I will describe the design for the E6 3D and 2D MOT chambers.
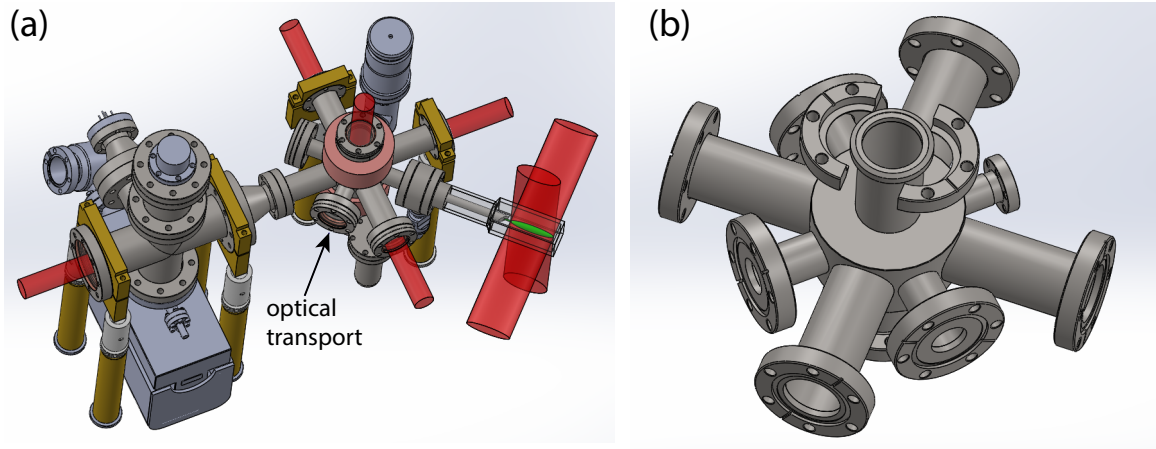
Figure 2.1: 3D MOT vacuum chamber. (a) Red beams represent the 3D and 2D MOT beams. The optical transport beam enters through the indicated viewport and exits through the gate valve at the top of the image into the science chamber (not shown). The 3D MOT chamber port opposite the 2D MOT chamber can be used for fluorescence imaging of the 3D MOT and/or to couple in microwaves for forced microwave evaporation. Also seen at left is the pumping chamber which includes an ion pump (with attached angle valve for pump down), an ion gauge, and a TSP. (b) Dedicated view of the bare 3D MOT chamber. Note the split flange on the upper and lower ports. The split flange allows the ID of the anti-Helmholtz coils, seen in (a), to be smaller than the OD of the flange, limited only by the diameter of the CF knife edge. Note that this design makes it impossible to remove the anti-Helmholtz coils without breaking vacuum.

2.1.1   3D MOT Chamber

As will be described in Sec. 2.2, the main geometric constraint for a conventional 3D MOT is the requirement for six pairs of counter-propagating cooling beams to enter a volume of space along three orthogonal axes. Another important geometric requirement is the presence of a spherical quadrapole magnetic field centered at the intersection of these beams. In designing our 3D MOT vacuum chamber there were additional requirements for a port through which we could pump out the vacuum chamber, a port through which we deliver a high-flux, laser-cooled atomic beam and a port out which atoms from the MOT could be optically transported into a neighboring science chamber.

Fig. 2.1(b) illustrates the 3D MOT chamber we designed to satisfy these geometric constraints. The 3D MOT chamber is a 10-way cross with eight ports within the horizontal plane and two ports along the vertical axis. One of the ports in the horizontal plane is a short 1.33" flange, chosen to minimize the total optical transport distance. The two vertical ports utilize split conflat (CF) flanges, which can be entirely removed from the assembly. This design allowed for the installation of small inner diameter (ID) anti-Helmholtz coils

(a)

(b)

Custom Bored Double Flange Adapter

Dual Size Differential Pumping tube

102 mm

1.5 mm

7.5 mm

31 mm

Rb Vapor Side Port

Angle Valve
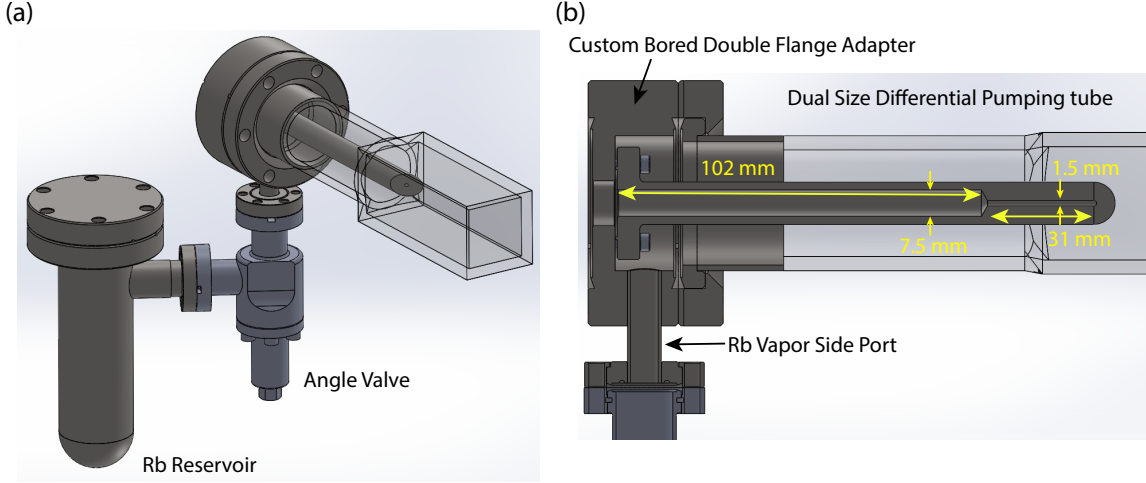
Rb Reservoir

Figure 2.2: 2D MOT vacuum chamber design. (a) Overview including double flange adapter, glass cell, differential pumping tube, angle valve and Rb reservoir. (b) Section view showing differential pumping tube dimensions. The differential pumping tube is attached to the double flange adapter using 6 vented bolts. The mating surface is also only a small rim rather than a large flat surface to avoid capturing air pockets with the mated surfaces that could lead to virtual leaks.

above and below the central chamber prior to closing the vacuum chamber.

The 3D MOT chamber is attached to a pump down chamber, which includes an ion pump (TiTan 75S ion pump, Gamma Vacuum), titanium sublimation pump (TSP) (up to $500\,\mathrm{L/s}$ pump speed for H_2, Gamma Vacuum) and ion gauge (T-NUDE-F, Duniway).

2.1.2   2D MOT Chamber

Our 2D MOT design was motivated by Refs. [92–98]. The 2D MOT subchamber is based around a custom thick 2.75" CF double flange, one side of which is attached onto the 3D MOT chamber, and the other side onto which the 2D MOT glass cell is mounted. See Fig. 2.2. The double flange includes a central bore through which the atomic beam can pass, as well as a bolt ring to which a stainless steel (SSL) differential pumping tube is bolted. The differential pumping tube increases the pressure differential between the (relatively) high pressure 2D MOT cell and the 3D MOT chamber while allowing the atomic beam to pass between chambers. See Appendix C for details about the differential pumping calculation used to design the differential pumping tube. Finally, the double flange also has a small port into which gaseous Rb can flow from the Rb reservoir, through an angle valve[1] into the 2D MOT glass cell.

---

[1]The purpose of the angle valve is to, in principle, allow the Rb supply to be replenished without breaking vacuum on the MOT chamber. In practice this would require pumping down the reservoir while it

The Rb reservoir has inside of it a steel ball and an ampoule containing 5 g of solid ^{87}Rb and buffer argon gas. After the chamber is pumped out, the SSL ball is physically manipulated using a strong out-of-vacuum magnet so that it crushes the ampoule to expose the solid Rb to the 2D MOT chamber. The solid ^{87}Rb sublimates to fill the chamber with a partial pressure of 3×10^{-7} torr of gaseous ^{87}Rb [66]. Enhanced 2D MOT fluxes can be realized by heating the reservoir and increasing this vapor pressure [98] but we have not found that to be necessary.

The differential pumping tube consists of two sections: one small diameter section which provides the conductance choke and one longer, larger diameter section which permits a larger atomic flux cone to pass. The differential pumping tube protrudes into the rectangular section of the glass cell. The geometry of the differential pumping tube presents an angular opening with a full opening angle of about 70 mrad. The end of the differential pumping tube is cut at a 45° angle and electro-polished so that it can serve as an in-vacuum mirror to enable a 2D$^+$ MOT [97, 98]. However, we have not needed to use either a pusher or retroreflected 2D$^+$ beam to realize sufficient fluxes.

2.2 Laser Systems

2.2.1 Background

A 3D MOT is a configuration of laser beams and magnetic fields within a volume of space (about the size of a few cm^3) that imparts viscous and trapping forces onto any atom which enters this volume [7]. To load atoms into our 3D MOT, we generate a cold, high-flux, beam of atoms using a 2D MOT. We chose to load our 3D MOT with a 2D MOT rather than a Zeeman slower —as has traditionally been used in our research group— due to the more compact and simple design and promises of large atomic fluxes offered by a 2D MOT [93, 97, 98]. The compact design leaves space on the table for hardware required to achieve ever more advanced science targets, and the high flux increases cycle time.

See Ref. [99] and other introductory atomic physics references for more details on MOT physics than are presented here. A MOT requires laser cooling light which is red-detuned from the $|F = 2, m_F = +2\rangle \rightarrow |F' = 3, m_{F'} = +3\rangle$ cycling transition combined with a spherical (3D MOT) or radial (2D MOT) quadrapole magnetic field. In the case of a typical 3D MOT, there are three pairs of counter-propagating laser beams along three orthogonal axes to effect cooling of all motional degrees of freedom. In a 2D MOT, two pairs of counter-propagating beams cool atoms in two orthogonal directions, leaving an atomic beam traveling with some velocity along the final axial direction. Additionally repump light tuned to the $|F = 1\rangle \rightarrow |F = 2\rangle$ transition is required.

Naively, one would expect a 3D MOT to realize trapped samples as cold as the Doppler temperature, $T_{\text{Doppler}} = \hbar\Gamma/2k_B \approx 140\,\mu\text{K}$. However, because of sub-Doppler molasses

---

is attached to the angle valve. The pump down arm could then perhaps be cut from the reservoir using a vacuum pinch-off tube. The reservoir would likely need to be redesigned in this event.

cooling which occurs near the center of the 3D MOT, temperatures below $100\,\mu$K can be realized [100].

2.2.2 Optics Overview

The E6 apparatus includes two optics tables. One table, the experiment or science table, has mounted on it the vacuum chamber apparatus and beam injection optics and is where the experiments take place. The other table is the laser table and has mounted on it most of our laser sources and optics. We generate the laser beams on the laser table and control their frequency and power levels before sending them to the science table where their polarizations, spatial mode structures, and beam pointing directions are controlled for optimal interactions with the atoms.

We prepare both the cooling and repump light on the laser table. The cooling and re-pump beams are separated in frequency by about 6.8 GHz, the ^{87}Rb ground state hyperfine splitting. We choose to generate these two frequencies using two independent lasers. The cooling light is generated by a 780 nm distributed Bragg reflector (DBR) diode laser (Photodigm) which is driven by a Vescent current and temperature controller. The repump light is generated by a 780 nm distributed feedback laser (DFB) diode laser from Eagleyard, again driven by a Vescent current and temperature controller.

See Fig. 2.3 for a schematic diagram of the cooling path. The cooling laser is used to derive cycling transition light for both the 2D and 3D MOTs, imaging light for a few absorption imaging paths, and optical molasses in the science chamber There is some light leftover that is currently used for experimental test setups but which, in the future, could be used for optical pumping or generating a pusher beam for the 2D MOT.

Along both the cooling and repump paths, we pick off a small amount of light to perform saturated absorption spectroscopy on Rb vapor within a heated vapor cell. The saturated absorption signal is used to stabilize the frequencies of the cooling and repump lasers to known frequencies with respect to the appropriate atomic transitions.

We utilize two slightly non-standard saturated absorption spectroscopy schemes to help us get our lasers to the appropriate frequencies while minimizing the loss of various resources such as acousto-optical modulators (AOMs) or laser power. These schemes allow us to lock each laser at an offset from its final desired frequency so that we can, downstream, make up the frequency difference using an AOM. This gives us the ability to control the frequency and intensity of the beams as needed. The resultant offset frequencies are indicated in Fig. 2.3. Details about these slight modifications to our saturated absorption spectroscopy are found in Fig. 2.4. Our scheme was also designed to take advantage of the strongest spectroscopic features we observed.

Fig. 2.3 shows how the light from the two lasers is split into multiple paths, each passing through an AOM that shifts the frequency of the beam by the indicated frequency shift. Additionally, in order to produce the large optical powers that are required to operate the 2D and 3D MOT setups above saturation intensity, we employ two tapered amplifiers (TAs) to boost the powers of the 2D and 3D MOT cooling beams, respectively.
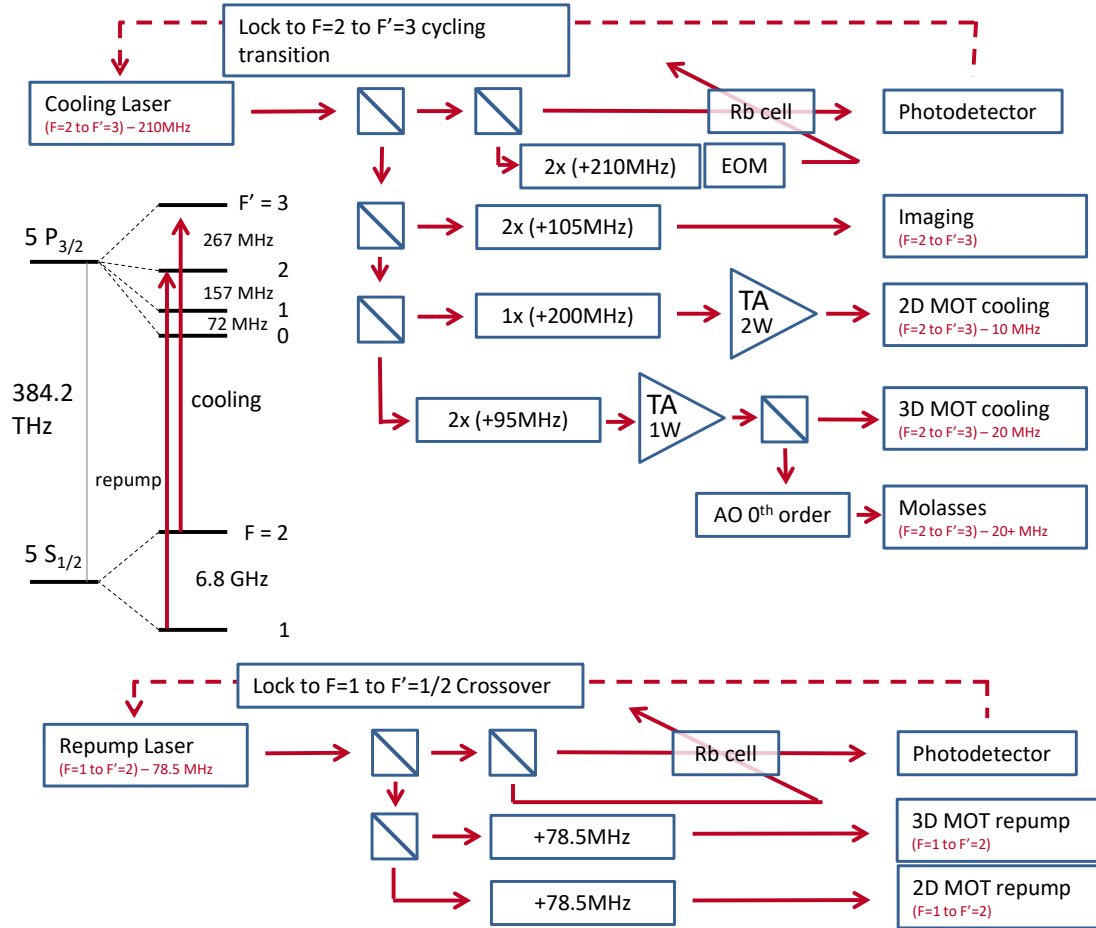
Figure 2.3: Cooling and repump laser frequency schemes. For the cooling spectroscopy setup the EOM is modulated and the lock error signal is fed into the DBR current. For the repump spectroscopy, the laser current is directly modulated and we also feedback to the laser current. We pick off some light after the 3D MOT TA and pass it through an AOM and use the zeroth order output light for our science chamber molasses light. This configuration allows us to intensity stabilize this light without changing its frequency. This figure was modified from earlier figures made by Emma Deist and Vicky Xu.
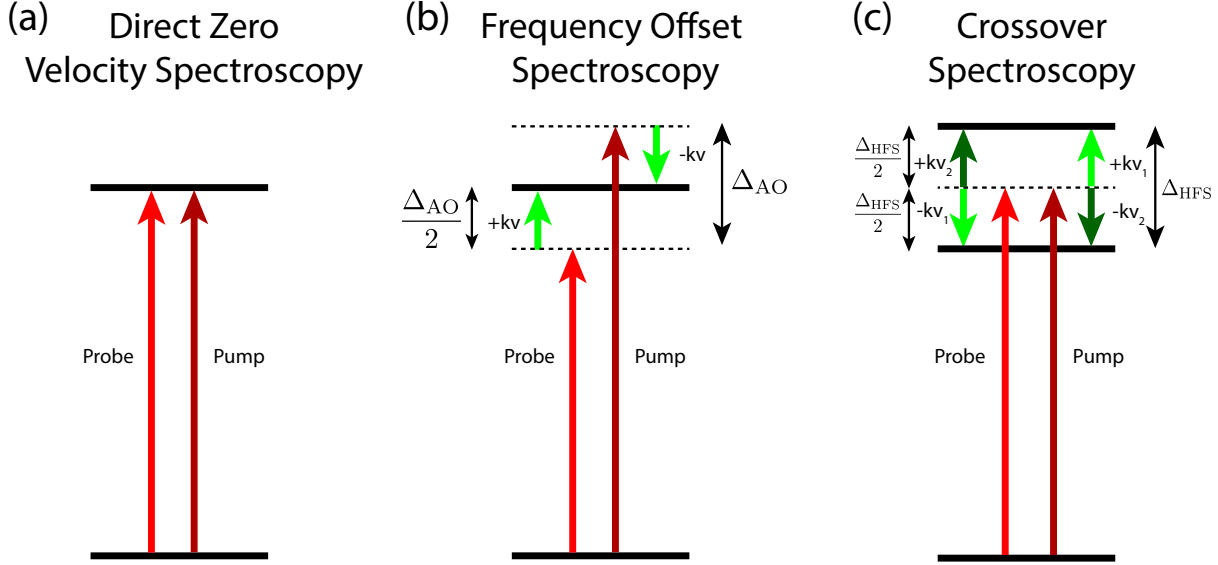
Figure 2.4: Saturated absorption spectroscopy techniques. (a) In direct saturated absorption spectroscopy, atoms moving with zero velocity are optically excited by the pump beam so that the probe experiences less absorption as it propagates through the medium than it would in the absence of the pump. This creates a sub-Doppler feature in the probe spectroscopy signal to which a laser can be locked. (b) In frequency offset spectroscopy, the pump and the probe are detuned by Δ_{AO} from each other so that they are never simultaneously resonant with any transition. However, when the probe is red detuned by $-\Delta_{AO}/2$ from a transition then the pump is blue detuned by $\Delta_{AO}/2$. There is, then, a certain subset of atoms moving towards the probe with velocity v such that the Doppler shift $kv = \Delta_{AO}/2$ blue (red) detunes the probe (pump) onto resonance. (c) In the simplest form of crossover spectroscopy, the pump and the probe are at the same frequencies and detuned halfway between two atomic transitions. There is, then, a class of atoms whose Doppler shift red or blue detunes the pump into resonance with one transition and blue or red detunes the probe onto resonance with another transition. A saturated absorption feature again appears in the probe spectroscopy signal because of the depletion of the ground state by the pump. Crossover spectroscopy can also be performed with the pump and probe detuned from each other if the pump and probe are also respectively detuned from two different transitions by equal and opposite amounts.

All of the light generated on the laser table is sent to the experimental table via single mode optical fibers. We used fiber splitters extensively to easily combine multiple laser beams for delivery to the atoms. The fiber splitters are advantageous because they allow us to combine beams without requiring additional optics on the table. This saves on cost, table space, optical alignment maintenance, and possibly beam power. The downsides are that there may be slight and unadjustable power imbalances coming out of the different ports, and, also, that the fiber splitter is potentially a single point of failure. For example, if one port of the fiber splitter becomes irreparably damaged, then the entire fiber splitter would need to be replaced. Such a replacement would be expensive and inconvenient. That said, we have improved and had to deploy our skills at cleaning and polishing fiber tips to increase the lifetime of each fiber port.[2] Overall we have been happy with our choice to use fiber splitters throughout the laser system.

The fiber routing diagram is shown in Fig. 2.5. The 2D (3D) MOT system uses a 2×4 (2×6) fiber splitter to combine repump and cooling beams and deliver them to all of the respective MOT ports. We deliver about $400\,\mathrm{mW}$ of cooling light to the 3D MOT fiber splitter and get about $25\,\mathrm{mW}$ of output power at each port. We deliver about $650\,\mathrm{mW}$ of cooling light to the 2D MOT fiber splitter and get about $100\,\mathrm{mW}$ of output power at each port. These insertion losses for the fiber splitters of $-4.3\,\mathrm{dB}$ (38% transmitted) and $-2.1\,\mathrm{dB}$ (62% transmitted) for the 3D and 2D MOT fiber splitters seem to be due to limited input coupling efficiency. In addition, we also couple repump light into the second input port of each of these fibers. We have about $2.5\,\mathrm{mW}$ of repump light emitted from each 2D MOT fiber port and about $1.5\,\mathrm{mW}$ emitted from each 3D MOT fiber port. We use a 2×2 fiber splitter to combine absorption imaging light, resonant with the cycling transition, with a small amount of repump light and deliver these beams to two absorption imaging beampaths.

One challenge we have faced in delivering light through the 2D MOT fiber splitter is the coupling of high power light emitted from the TA. The spatial mode of light emitted from our TA is very poor. We have been concerned that trying to fiber couple it at high power would damage the fiber splitter input port, because the large fraction of light that is not mode matched into the fiber would be absorbed within the surrounding fiber cladding, causing the cladding and fiber to heat up and burn. To mitigate this risk, we installed a pinhole spatial mode filter after the TA and before the fiber. The spatial mode filter cleans up the optical mode, thus improving the fiber coupling efficiency into the fiber splitter port, but at the cost of about 30% of the output power from the TA.

---

[2]Our 2D MOT fiber ports have about $100\,\mathrm{mW}$ of power coming out. We have observed that if we unplug the fiber when power is coming out there is a high likelihood of damaging the fiber tip. We suspect this may be due to either high-power reflections off of the metallic fiber coupler body during removal or the high power beam optically trapping and/or melting dust in the air against the fiber tip. We're able to recover from this damage by re-polishing the fiber end face. This is all not to mention the laser safety concerns that come with unplugging this high power fiber with power coming out. The lesson is do **not** unplug the fibers when light is coming out!
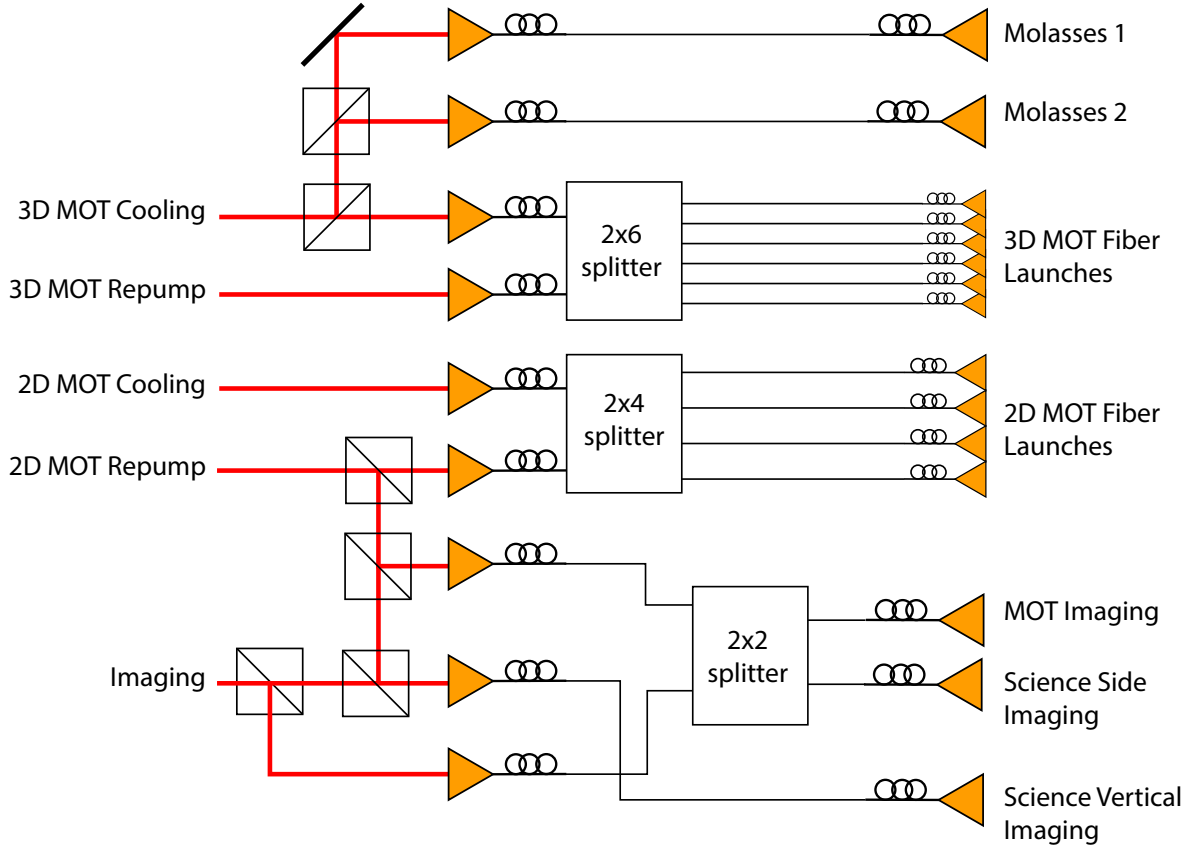
Figure 2.5: Fiber routing for the cooling laser system.

2.2.3   Fiber Launches

For both the MOTs, we use what we call fiber-launch telescopes to deliver the light from the optical fibers into the vacuum chamber; see Fig. 2.6. One of the challenges of delivering MOT beams is making a large beam with the appropriate polarization. We handle the polarization first and then expand the beam. Each fiber-launch begins with a fiber collimator which creates a 3.75 mm beam waist radius. The output then passes through a cage-mounted $\lambda/2$ waveplate, polarization cleanup polarizing beamsplitter (PBS), and then a $\lambda/4$ to generate the circularly polarized light needed for MOT laser cooling. Next each beam needs to be telescoped to the appropriate size.

The 2D MOT beams fill a rectangular glass cell. Each beam passes through a circular diverging lens with $f = -9$ mm. One axis of the diverging beam is collimated by a $f = 50.8$ mm cylindrical lens and the other axis is collimated by a large $f = 150$ mm cylindrical lens, resulting in MOT beams with waists of 20×60 mm. With 100 mW of MOT light and

$I_{\mathrm{sat}} = 1.669\,\mathrm{mW/cm^2}$ for the D_2 cycling transition, this gives us a large saturation parameter of $s = I/I_{\mathrm{sat}} \approx 30$ per beam [66].

These telescope systems are mounted to vertical 1.5" posts that we call 2D MOT towers. Some of the beams are then redirected using metal mirrors[3] for alignment before passing through the glass cell. We chose to use (4x) counter-propagating beams with no retroreflection out of an abundance of caution since this was the first 2D MOT within the group and we weren't sure if the beam imbalance from retro reflecting would have been too problematic. In hindsight I think we would have been fine if we had gone with a power-saving retro reflecting approach.

The 3D MOT fiber launches use the same polarization optics as the 2D MOT fiber launch. The beam then passes through a $f = -15\,\mathrm{mm}$ focal length[4] lens and is then collimated by a $f = 60\,\mathrm{mm}$ focal length lens, giving us a 3D MOT beam with a waist of $14\,\mathrm{mm}$ and a saturation parameter of $s \approx 4.8$. Depending on the particular 3D MOT arm under consideration, the fiber-launch either attaches directly to the chamber with a cage-CF flange adapter or reflects off of a cage mirror prior to entering the chamber. We use (6x) counter-propagating beams with no retroreflection.

Some fiber launches have additional ports for the injection and detection of absorption imaging light or for monitoring the power of the MOT beams.

We enjoy working with these fiber launches because we rarely, if ever, have to do any alignment of the MOT lasers. Further, the optical setups are very compact, leaving plenty of room for other important optics on the table.

---

[3]We use metal mirrors because, unlike dielectric mirrors, they do not introduce any birefringence into the beams which would degrade the polarization. However, one must remember that the handedness of circularly polarized light reverses upon reflection from a metal mirror.

[4]In both telescope systems the diverging lenses allow us to realize larger magnifications with shorter optical paths to save space.
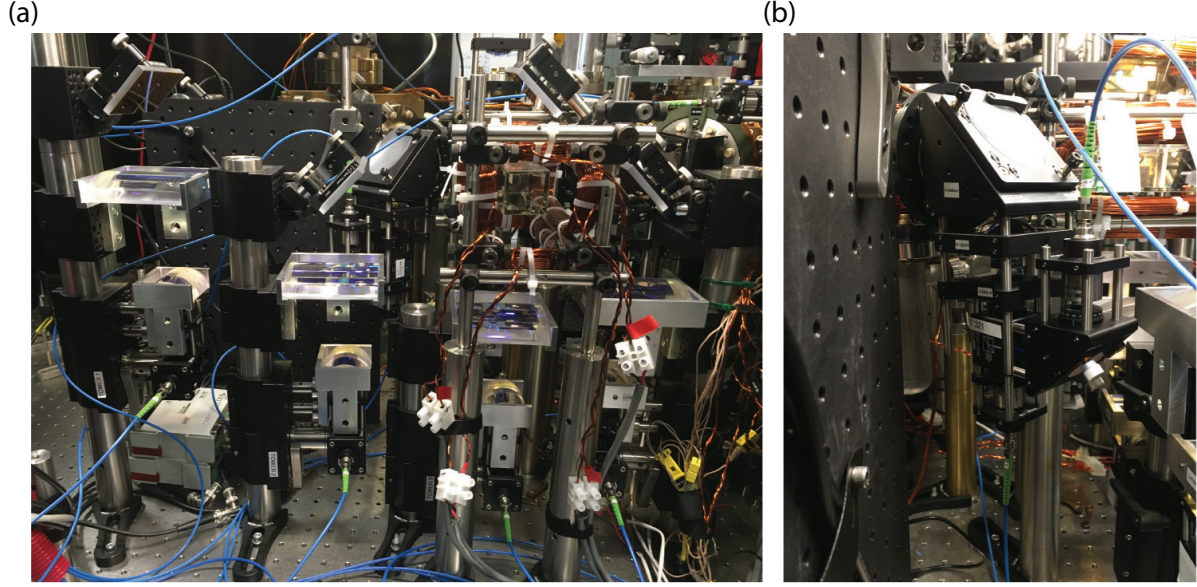
Figure 2.6: (a) 2D MOT fiber launch towers. (b) 3D MOT fiber launch. This fiber launch includes an extra port to inject the MOT absorption imaging light. The fiber at bottom carries 3D MOT cooling and repump light and the fiber pointing down carries absorption imaging and repump light.

2.3 Magnetic Coils

The 3D MOT anti-Helmholtz coils were wound from square hollow insulated magnet wire with a cross-section side length of about 3 mm. Each coil has about 65 turns, an ID of 50 mm, an outer diameter (OD) of 100 mm and a height of 45 mm. Each coil is offset by 22 mm from the atoms. The coils provide an estimated gradient of $3\,\mathrm{G\ cm^{-1}A^{-1}}$. We run the coils at low amperage for the 3D MOT to realize about $15\,\mathrm{G\ cm^{-1}}$. At these currents we do not require water cooling. However, plumbing is in place to water cool these coils for applications requiring higher gradients and currents such as magnetic trapping.

We wound three pairs of bias coils around the arms of 3D MOT chamber. We use the bias coils to zero the magnetic field at the atoms for polarization gradient (PG)-cooling, adjust the position of the field-zero of the spherical quadrapole field to overlap the 3D MOT with the transport ODT, and to supply a bias field to improve the fidelity of absorption imaging. In the horizontal direction, these coils have about 100 turns, an approximate radius of 3 cm and a distance to the atoms of 5 cm. In the vertical direction, the coils are about 50 turns with a radius of 5 cm and distance to the atoms of 4 cm.

The 2D MOT 'racetrack' coils have dimensions of about 5×18 cm and are about 50 turns of AWG 12 square magnet wire. We run them at a few A to generate a gradient of about 35 G/cm.
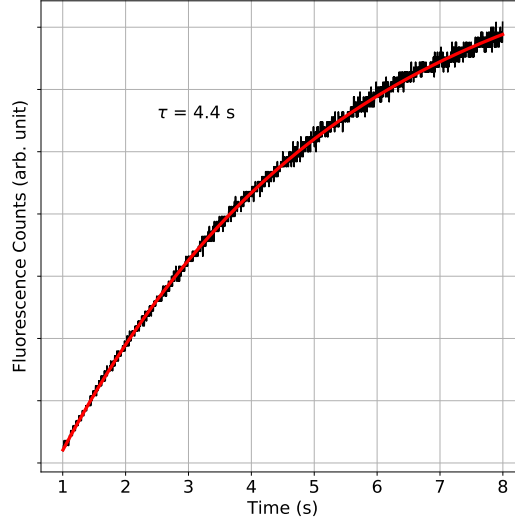
Figure 2.7: 3D MOT fluorescence loading measurement. We extract a loading time of 4.4 s.

2.4   3D MOT Loading

We perform absorption imaging along one of the MOT axes. The imaging light is combined and split from the MOT light using PBSs. We image through a 3D MOT fiber launch telescope and realize a magnification of 0.36, calibrated by a time of flight (TOF) free-fall gravity measurement. We image using a charge-coupled device (CCD) camera (Grasshopper USB3 GS3-U3-15S5M-C, FLIR) to capture 'with atoms', 'without atoms', and 'dark' images in succession. We then analyze these three images to extract atom numbers in the MOT and subsequent atom traps. We also perform real-time fluorescence imaging of the 3D MOT using an infrared (IR)-sensitive security camera that we monitor as the experiment runs to qualitatively confirm the 3D MOT is behaving as expected.

To maximize the 3D MOT loading rate, we optimized the 2D MOT flux by adjusting the 2D MOT detuning, optical polarization, beam pointing, and magnetic field gradients and bias fields. By independently controlling the current to each of the four 2D MOT racetrack coils we can adjust the position of the field-zero for the radial quadrapole field to align the 2D MOT beam through the differential pumping tube aperture. We find the loading rate to be optimized with a red-detuning of only about -10 MHz and a magnetic field gradient of 35 G/cm.

Next, we adjusted the 3D MOT parameters to optimize the 3D MOT loading rate and final atom number. We adjusted the same parameters as for the 2D MOT. We find the 3D MOT to be optimized with a detuning of about -20 MHz.

I present two data points to quantify our system performance. First, under certain optimized experimental conditions we have measured a saturated atom number of about 10^9 atoms using the absorption imaging system described above. A second measurement is shown in Fig. 2.7. Here the 3D MOT has been imaged onto a photodiode so that the output voltage is proportional to the MOT fluorescence. For these data the scaling between the photodiode voltage and atom number remains poorly calibrated. Nevertheless, by fitting to the loading curve we can extract a typical loading time $\tau = 4.4\,\text{s}$. For many of our recent experiments, we work with fewer atoms and load the MOT only for around $1\,\text{s}$.

More carefully calibrated measurements should be recorded of the MOT saturated atom number, loading time, loading rate, and lifetime to serve as benchmarks to help assess the performance of the system over time.

2.5 Optical Transport

Since we are using a two-chamber design, we require a means to transport atoms from the preparation chamber, across a distance of about $32\,\text{cm}$, into the science chamber. We chose to use optical transport with a tunable focus lens because it had been shown to be compact, versatile, fast, and efficient [101].

The transport ODT beam enters the vacuum chamber, focuses into the center of the 3D MOT and then transmits through a gate valve into the science chamber where it subsequently exits out of a viewport on the other side of the chamber to be analyzed or dumped.

To realize a high number of atoms in the transport ODT, we load the MOT for up to $2\,\text{s}$ and then perform a compressed MOT by detuning the 3D MOT light, increasing the magnetic field gradient, and decreasing the repump power. We then perform a stage of molasses cooling by increasing the detuning and reducing the power of the 3D MOT cooling light and turning off the magnetic field gradient for up to $25\,\text{ms}$. These two steps have the combined effect of decreasing the cloud temperature down to about $18\,\mu\text{K}$ and increasing the density of the atomic cloud to optimize loading into the transport ODT. The transport ODT is on during all three of these stages and captures any atoms which are cooled into its trapping volume during molasses cooling.

We generate the transport ODT using a $1064\,\text{nm}$ IPG fiber laser. We intensity stabilize the laser to about $6\,\text{W}$ for loading and transport using an AOM. The transport ODT is focused to about $50\,\mu\text{m}$ giving us a trap depth of about $230\,\mu\text{K}$ and an axial trap frequency of about $4.5\,\text{Hz}$.[5] We can load as many as 10^7 atoms in the transport ODT.

We then transport the atoms by adjusting the focal length of one of the lenses in the transport ODT optical path [101]. As shown in Fig. 2.8(a), the beam shaping lenses are positioned so that tuning the focal length of the tunable focus lens adjusts the position of the focus but not the size of the waist. The tunable focus lens we use is part number EL-C-16-40-TC from Optotune. The lens is formed by a malleable polymer sac that can be

---

[5]See Sec. 6.4.1 for useful formulas for calculating ODT trap parameters.
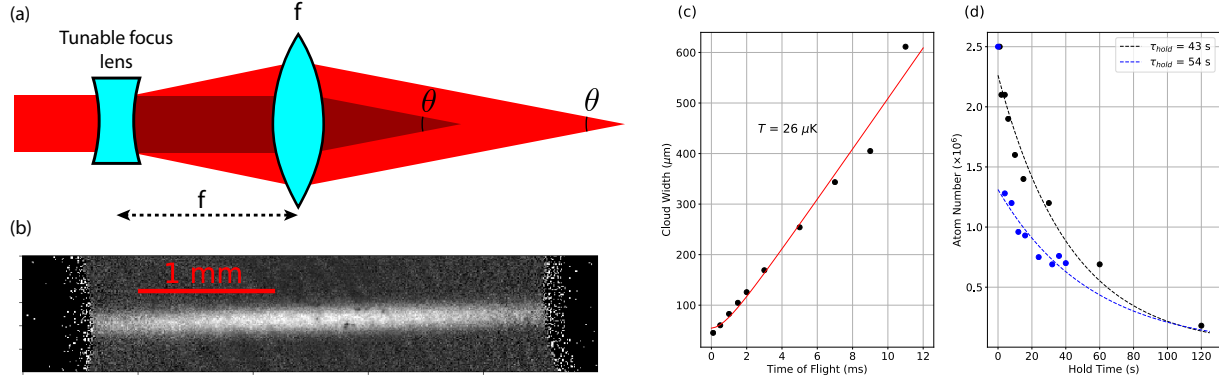
Figure 2.8: (a) We utilize a scheme described in Ref. [101] for optical transport. If the focal length of the tunable focus lens is adjusted, then the position of the focused waist will translate but the size of the waist will remain unchanged. (b) Absorption image of atoms released from the transport ODT after $500\,\mu$s of TOF. This image is taken perpendicular to the transport ODT axis and at $45°$ to the science cavity, the mirrors of which can be seen as dark regions at the left and right sides of the frame. See Fig. 5.1(a). (c) TOF measurement of the temperature of the atoms in the transport ODT after transport. (d) Lifetime of atoms in the transport ODT while holding the atoms still (black) and while continuously transporting the atoms (blue).

compressed or expanded by applying current to electromagnets within the housing which attract or repel each other and compress or expand the polymer sac. The lens can be tuned from $-10\,$dpt to $+10\,$dpt.[6] In practice we restrict the full tuning range of the lens to avoid accidentally focusing the transport ODT on the vacuum viewport and damaging the anti-reflection (AR) coating.

We transport the atoms using a simple spline position profile over about $2\,$s and are able transport as many as about 2×10^6 atoms into the science chamber at a temperature of about $26\,\mu$K. Fig. 2.8(b) shows an absorption image of atoms released from the transport ODT in the science chamber.

To determine the efficiency of our optical transport system, we performed an experiment where we transported the atoms back and forth from the atom preparation chamber to the dummy chamber many times and monitored loss as a function of time. We also measured the lifetime of the atoms in the atom preparation chamber with no transport. The results of this measurement are shown in Fig. 2.8(d). We extract a lifetime of $43\,$s when holding the atoms in place. When transporting, we see an initial loss of atoms the first time they are transported, but we otherwise measure a lifetime of $54\,$s. Surprisingly. the lifetime is longer while transporting than holding the atoms. This may be due to an improved vacuum pressure in the dummy chamber.

---

[6]Recall that a focal length is converted to diopters by dpt $= \frac{1\,\text{m}}{f}$.

At present we are able to deliver atoms into the science chamber within $3\,\mathrm{s}$. For experiments with microtweezers that require only a very small number of atoms (sometimes only a single atom is needed), we can reduce the MOT loading and transport time to improve the experiment cycle time.

2.6   Bose-Einstein Condensate

Above, I described the performance of our experiment when atoms are loaded into the optical transport system directly from laser cooling within the 3D MOT chamber. Alternatively, to produce even higher atomic phase space densities in the science chamber, we adopt a different, slower approach where we introduce an intermediate stage of magnetic trapping between the laser cooling and optical transport stages. The specific sequence is as follows.

As before, we load atoms into a 3D MOT, then compressed MOT and then turn on PG cooling. Next, we ramp on a large magnetic field gradient of over $100\,\mathrm{G/cm}$ for forced microwave evaporation. We shine in microwaves from an rf horn to flip hot atoms from the trapped $|F = 2, m_F = +2\rangle$ hyperfine states into the anti-trapped $|F = 1, m_F = -1\rangle$ states. We sweep the microwave frequency to continuously eliminate the hottest atoms within the trap. Next, we ramp on the transport ODT and ramp down the magnetic trap to form a hybrid magnetic and optical trap [89]. Finally we perform hybrid evaporation in this combined trap by reducing the magnetic field gradient and reducing the optical trap depth. At this point we have a dense cloud of atoms trapped in the transport ODT in the atom preparation chamber.
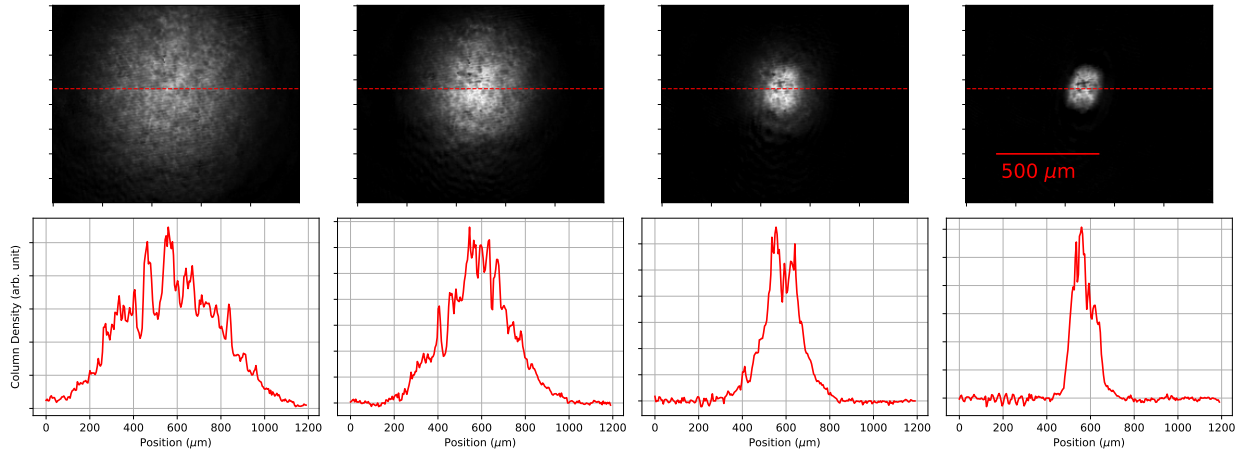
Figure 2.9: TOF images of atomic gas at various stages during the final optical evaporation. Evaporation proceeds from left to right. Notice the bimodal distribution apparent in the third panel and the complete disappearance of the thermal background in the fourth panel.

These initial conditions within the transport ODT are sufficient, following further evaporative cooling out of the optical trap, to produce Bose-Einstein condensed gases. We verified that we could produce BECs in two different configurations: 1) within the 3D MOT chamber and 2) following optical transport into the dummy chamber before the science chamber was installed.

To realize a BEC in the 3D MOT chamber we continued the hybrid evaporation described above until condensation. To realize a BEC in the dummy chamber we turned off the magnetic field gradient and transported the atoms into the dummy chamber as described above. We installed another $1064\,$nm laser with a $50\,\mu$m waist and similar power as the transport ODT at $90°$ to the transport ODT in the dummy chamber. This ODT was referred to as the cross ODT. By overlapping the cross ODT with the transport ODT in the dummy chamber and ramping up the cross ODT we were able to tightly confine the atoms in all three dimensions. We then performed a final evaporation in this joint cross ODT trap by reducing the optical power in both the cross and transport ODTs. The evaporation resulted in a sufficient increase in phase space density to again realize Bose-Einstein condensation.

I show representative TOF images of the atomic gas at various stages during the final optical evaporation in the dummy chamber in Fig. 2.9. In the third image we begin to see a bimodal structure in the atomic distribution, characteristic of the BEC transition. In the final image there is no clearly visible thermal background indicating that we have realized a high condensate fraction. We estimate that the BEC has order $\approx 5 \times 10^5$ atoms at a temperature of about $200\,$nK.

We don't have very precise quantification of the conditions that produced these BECs because, at the time, we were only briefly testing out this approach in the dummy chamber prior to installation of the science chamber. It would be beneficial for the present experimental team to recreate a BEC, now with the high-finesse science cavity in place in the science chamber, and to carefully characterize the conditions under which it is produced as an important benchmark for future work.

Chapter 3

Optical Cavity Theory

In this chapter I will outline, from very basic principles, the general theory of a high finesse optical resonator. Much of this chapter can be understood as providing the relevant background information to ease the calculation of the relative beatnote between a 780 nm and 1560 nm cavity superlattice presented in Sec. 3.6. However, I hope this chapter will also serve two additional purposes. (1) To present readers with a coherent and complete classical theoretical framework for the dynamics of the optical field inside of an optical cavity and (2) to serve as a reference for various useful definitions and formulas such as, for example, the relationship between input or output power and the intracavity photon number.

I will begin this chapter by writing down a matrix formalism for the reflection of light from a single cavity mirror, Next, I will explore the geometric, temporal, and spectral properties of a two-mirror optical cavity. Sec. 3.5 provides an experimentally useful reference discussion about power flow and energy within an optical cavity. Finally, Sec 3.6 discusses the superlattice beatnote mentioned above. Much of the material presented here is covered in more or less detail in Refs. [102, 103].

3.1   Scattering Matrix for Single Mirror

We begin by exploring a detailed description of an ideal (lossless) optical mirror as a two-port scattering network. At the end of the section, I will include a brief discussion of how loss can be included in such a model

The mirrors we consider are typically glass substrates with dielectric coatings on one surface. We imagine a single mirror as a two port network, each port consisting of an input and an output plane wave as depicted in Fig. 3.1. Each of the four fields is of the following form, for example:

$$E_A^{\text{in},(+)} e^{+ik\left(z-z_A^{\text{in}}\right)} e^{-i\omega t}. \tag{3.1}$$

Each term consists of a positive frequency temporal complex exponential as well as a positive or negative wavevector spatial complex exponential depending on whether the field is left
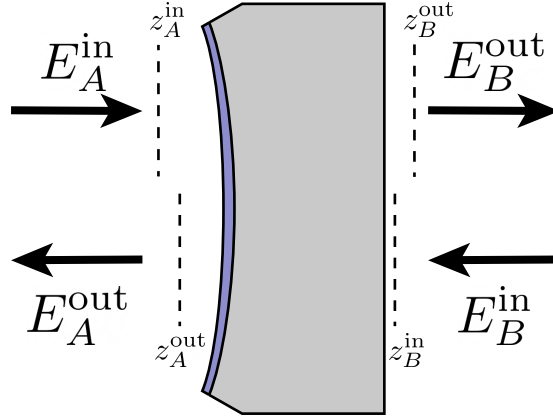
Figure 3.1: Field configuration for mirror transmission and reflection scattering matrix. Here the cavity is depicted as a glass substrate (gray) with a high reflectivity coating (blue) on one surface. Four input and output plane waves are depicted and labeled along with their respective reference planes.

or right propagating. Recall from Appendix A that, based on our choice for Fourier transformations, e^{+ikz} is a positive frequency spatial function while $e^{-i\omega t}$ is a positive frequency temporal function. Each field is defined with respect to a specific reference plane labeled by $z_{A/B}^{\text{in/out}}$. More discussion about these reference planes will follow below.

Maxwell's equations at and within the dielectric interfaces on the mirror constrain the complex amplitudes of these four fields to be related by a 2×2 scattering matrix with certain properties:

$$\begin{bmatrix} E_B^{\text{out},(+)} \\ E_A^{\text{out},(+)} \end{bmatrix} = \begin{bmatrix} r_B & t_A \\ t_B & r_A \end{bmatrix} \begin{bmatrix} E_B^{\text{in},(+)} \\ E_A^{\text{in},(+)} \end{bmatrix}. \tag{3.2}$$

Each element of the scattering matrix is a complex number. Conservation of energy constrains this matrix to be unitary and time-reversal symmetry constrains the magnitude of the off-diagonal elements to be equal [102]. We can implement these constraints by writing the scattering matrix as

$$\begin{bmatrix} E_B^{\text{out},(+)} \\ E_A^{\text{out},(+)} \end{bmatrix} = \begin{bmatrix} e^{-i\phi_{r_B}} r & e^{-i\phi_{t_A}} t \\ e^{-i\phi_{t_B}} t & -e^{-i\phi_{r_A}} r \end{bmatrix} \begin{bmatrix} E_B^{\text{in},(+)} \\ E_A^{\text{in},(+)} \end{bmatrix}, \tag{3.3}$$

where r and t are real with

$$r^2 + t^2 = R + T = 1. \tag{3.4}$$

The complex phases satisfy

$$e^{-i\phi_{r_B}} e^{-i\phi_{t_A}} e^{-i\phi_{t_B}} \left(-e^{-i\phi_{r_A}}\right) = -1$$
$$\phi_{r_B} + \phi_{t_A} + \phi_{t_B} + \phi_{r_A} = 2\pi n \tag{3.5}$$

with n an integer. We will assume that each phase satisfies $0 < \phi < 2\pi$.

In most optical applications these optical phases are not of importance and, in fact, the scattering matrix is typically presented with all of these phase factors set to zero and r and t simply taken to be real. We will see later, in Sec. 3.6, that these optical phases have importance when considering the spatial overlap between two simultaneously resonant optical cavity modes, thereby motivating a more thorough exploration here. The effect of mirror phases is also considered carefully in Ref. [104], in which the spatial overlap of multiple cavity modes is of importance.

Sec. 11.1 in Ref. [102] provides a thorough discussion of these scattering phases for a laser mirror. There, it is discussed that the reference plane for each of the four input and output beams can be shifted in position, with the net effect of altering the phases of the elements of the scattering matrix in a symmetric way depending on which reference plane is shifted. In fact, physically speaking, there is no single, preferred, physical plane where the entire reflection occurs in a discontinuous sense within the dielectric. The reflection of the electromagnetic wave is, rather, distributed and built up interferometrically as the beam penetrates further into the dielectric coating. This leaves the choice for the location of the reference plane open to convention.

One natural choice would be for all four reference planes to be located at the air-side physical surface of the dielectric coating. However, this choice has the disadvantage that, in general, the reflection phases $\phi_{r_{A/B}}$ are non-zero, in contrast to the case of a perfect reflector. The other phases in the scattering matrix are also, in general, non-zero. An alternative choice which is sometimes taken [105] would be to move all of the reference planes by an amount $\Delta L = \frac{\phi_{r_A}}{2\pi} \frac{\lambda}{2}$ away from the air-side surface into the dielectric coating such that the transformed scatting matrix has no reflection phase on port A. This means that we intuitively would think of the light as reflecting from a slightly different location than the physical mirror surface. This distance ΔL is the so-called penetration depth of the light into the mirror (modulo $\frac{\lambda}{2}$).

When considering multiple wavelengths of light within an optical cavity it is, in general, possible that the penetration depths differ between the two mirrors and that the penetration depths differ between the different wavelengths. This means that the relative position of the two cavity modes may not be symmetrically overlapped as one would naively expect. Here I choose appropriate conventions that allow me to capture these effects while keeping the notation simple. I will choose to enforce the following constraints:

1. The scattering matrix element for reflection on port A is the most important because this is the one that determines the location and properties of the cavity mode within the chosen coordinate system. For this reason we will take the two reference planes

for port A to be located at the air-side physical surface of the dielectric mirror. This means the location of the reference plane (and thus our mental picture of the location of the mirror) does not vary for different wavelengths but is fixed in space. Yet, we pay the price that ϕ_{r_A} may be non-zero, so we cannot assume that light simply acquires a π phase shift upon reflection.

2. We don't care about the physical location of either of the port B reference planes. Thus, we will adjust the position of these planes to simplify the scattering matrix. We will place these planes so that the two transmission phases ϕ_{t_A} and ϕ_{t_B} vanish.

The resultant scattering matrix under these conventions will be

$$\begin{bmatrix} e^{+i\phi_r}r & t \\ t & -e^{-i\phi_r}r \end{bmatrix}. \tag{3.6}$$

Note that there is just one extra phase ϕ_r that we must keep track of under this convention choice.

3.1.1   Mirror Loss

In all of the above we have neglected mirror loss. There are two sources of loss important for our optical cavity. Loss from light absorption within the dielectric coating substrate, and loss from light which is scattered off of surface impurities in the mirror surface. A more thorough discussion of the experimental aspects of mirror losses will appear in Sec. 4.3. Here I will consider only the theoretical implications for losses on the form of the scattering matrix.

A thorough theoretical treatment of such loss mechanisms might include these loss channels as additional ports in the scattering matrix. That is, the two-port network could be expanded to a three- or four- port network where the new ports indicate new loss channels. In a quantum optics formalism these loss ports could be introduced by including fictitious beam-splitters in front of the mirror which scatter some light away. This approach is slightly problematic because a beam-splitter introduces *two* new ports rather than just one as needed to include loss.

Regardless, I bring this up to point out that, in the presence of loss, the 2×2 sub-matrix describing reflection and transmission will no longer necessarily be unitary because it is only part of the full unitary scattering matrix. As a result, its elements are not constrained in the same way as above. Nevertheless, in the case that these losses are small (as in our experiment), they perturb the unitarity conditions only slightly and can be neglected for now.

We will continue to use Eq. 3.6 to describe the mirror but we will bear in mind that (1) in the presence of loss which removes the unitarity condition, the phases may be different than those that appear in that formula, and (2) the amplitudes of r and t will be slightly

adjusted so that

$$r^2 + t^2 + l^2 = R + T + L = 1. \tag{3.7}$$

3.2   Conditions for Cavity Resonance

In this section we will calculate the spatial properties of the optical modes supported by a pair of cavity mirrors. See Fig. 3.2.

The basic principle for an optical cavity is as follows. A beam of light enters the cavity through one of the mirrors. This light then bounces back and forth between the mirrors many times. The light that is continuously pumped into the cavity interferes with the light that is circulating within the cavity. If (1) the spatial mode of light is the same at the beginning and end of a round trip through the cavity and (2) the phase collected by the light on a round trip through the cavity is an integer multiple of 2π, then the light entering the cavity will constructively interfere with the circulating light resulting in a buildup of the stored optical electromagnetic energy with the cavity. If the round trip phase deviates from 2π, then destructive interference dominates in the cavity such that the stored optical energy within the cavity is minimal and the pump field is nearly entirely reflected from the cavity input mirror. In this section, we will mathematically explore the conditions under which the constructive interference, or cavity resonance, described above is achieved.

3.2.1   Gaussian Modes

An optical cavity can support a range of transverse spatial modes. To this end we write down the general formula for a Hermite-Gaussian optical beam. A standard TEM$_{00}$ Gaussian
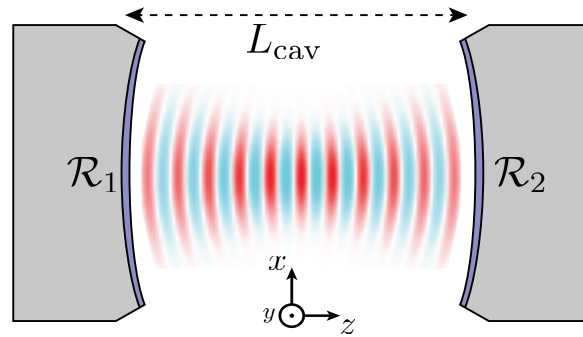
Figure 3.2: Optical cavity formed by two mirrors. The cavity mode electromagnetic field amplitude and phase is shown between the cavity mirrors. Positive phases are indicated by red shading and negative phases by blue shading. The spatial amplitude of the mode is indicated by the color saturation.

beam focused at $z = 0$ is described by

$$E_{00}^{(+)}(x, y, z, t) = E_0^{(+)} \frac{w_0}{w(z)} \exp\left(-\frac{x^2 + y^2}{w(z)^2}\right) \exp\left(ik\frac{x^2 + y^2}{2\mathcal{R}(z)}\right)$$
$$\times \exp\left(ikz - i\omega t\right) \exp\left(-i\psi(z)\right). \tag{3.8}$$

This expression represents the complex scalar field. The vector nature of the electromagnetic field has been suppressed with the understanding that there are two vector polarizations possible. The terms of this equation are as follows.

- $E_{00}^{(+)}(x, y, z, t)$ is the complex scalar electric field as a function of space and time

- $E_0^{(+)}$ is the complex field amplitude. $E_{00}^{(+)}(0, 0, 0, 0) = E_0^{(+)}$.

- w_0 is the beam waist at $z = 0$. This is the radius at which the amplitude (intensity) falls off by $1/e$ ($1/e^2$).

- $w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2}$ is the expanding beam waist as a function of z.

- $z_R = \frac{\pi w_0^2}{\lambda}$ is the beam Rayleigh range or confocal parameter.

- $\mathcal{R}(z) = z\left(1 + \left(\frac{z_R}{z}\right)^2\right)$ is the evolving wavefront radius-of-curvature (ROC) as a function of z.

- $k = \frac{2\pi}{\lambda}$ is the beam wavenumber where λ is the wavelength of light.

- $\omega = ck$ is the beam temporal frequency where c is the speed of light.

- $\psi(z) = \arctan\left(\frac{z}{z_R}\right)$ is the Gouy phase.

The standard Gaussian beam described above is the fundamental mode of an infinite family of Gaussian beams. These higher-order beams have mode functions of the form

$$E_{nm}^{(+)}(x, y, z, t) = E_{00}^{(+)}(x, y, z, t)H_n\left(\frac{\sqrt{2}x}{w(z)}\right)H_m\left(\frac{\sqrt{2}y}{w(z)}\right)e^{-i(n+m)\psi(z)}, \tag{3.9}$$

where H_n are Hermite polynomials. The n index indicates the mode order for the x-direction and the m index indicates the mode order for the y-direction. The Hermite polynomial $H_n(\cdot)$ has n nodes.

3.2.2 Fitting a Mode Between the Mirrors

The first condition for cavity resonance, as stated above, was that the spatial mode of the electromagnetic field is the same at the beginning and end of a cavity round trip. This condition is satisfied if the beam reflects exactly back onto itself at each cavity mirror. For a complicated spatial mode, this condition is met if the mirror surface is normal to the local wavevector for the optical beam. This is equivalent to the mirror surface being a surface of constant phase for the beam. For paraxial Gaussian beams, the wavefronts are approximately spherical with radii of curvature $\mathcal{R}(z)$ given above. For this reason, the cavity mirrors that we use are fabricated with spherical reflective surfaces with radii of curvature $\mathcal{R}_i$. In this section, we will explore how the choice for $\mathcal{R}_i$ for each mirror affects the geometry of the mode supported by a cavity formed from these mirrors.

We set up our coordinate system as follows. The waist of the beam is located at $z_0 = 0$, the left mirror is at z_1 and the right mirror is located at z_2 so that the cavity length is given by $L_{\text{cav}} = z_2 - z_1$. Note that these positions z_1 and z_2 correspond to the positions of the reflection reference planes for each mirror discussed in Sec. 3.1.

The condition that both mirror surfaces overlap with wavefronts of a Gaussian beam is then given by

$$\mathcal{R}(z_1) = -\mathcal{R}_1 = -z_1 \left(1 + \frac{z_R^2}{z_1^2}\right),$$

$$\mathcal{R}(L_{\text{cav}}) = \mathcal{R}_2 = z_2 \left(1 + \frac{z_R^2}{z_2^2}\right). \tag{3.10}$$

This system of equations can be solved to yield

$$z_1 = L\frac{L - \mathcal{R}_2}{\mathcal{R}_1 + \mathcal{R}_2 - 2L_{\text{cav}}},$$

$$z_2 = L\frac{\mathcal{R}_1 - L}{\mathcal{R}_1 + \mathcal{R}_2 - 2L_{\text{cav}}},$$

$$z_R = \sqrt{L_{\text{cav}}\frac{(L_{\text{cav}} - \mathcal{R}_1)(L_{\text{cav}} - \mathcal{R}_2)(\mathcal{R}_1 + \mathcal{R}_2 - L_{\text{cav}})}{(2L_{\text{cav}} - \mathcal{R}_1 - \mathcal{R}_2)^2}}. \tag{3.11}$$

For a given cavity geometry, we see that the Rayleigh range is independent of the wavelength of light used to form the optical mode.

These less-than-beautiful expressions can be written more compactly by introducing the cavity stability or degeneracy parameter:[1]

$$\mathsf{g}_i = 1 - \frac{L_{\text{cav}}}{\mathcal{R}_i}. \tag{3.12}$$

---

[1] The font for g for the cavity degeneracy parameter is distinguished from that for g which is used to notate atom-photon coupling parameters. While the distinction is notationally subtle it should be clear which quantity is intended from context.

The Rayleigh range can then be re-expressed as

$$z_R = L_{\text{cav}} \sqrt{\frac{\text{g}_1\text{g}_2(1 - \text{g}_1\text{g}_2)}{(\text{g}_1 + \text{g}_2 - 2\text{g}_1\text{g}_2)^2}}. \tag{3.13}$$

In this expression the quantity under the square root must be positive to ensure z_R is real and positive. This is equivalent to the condition that $0 < \text{g}_1\text{g}_2 < 1$; this is the famous cavity stability criterion. If this criterion is not met, it is an indication that there is no Gaussian mode that has the property that both mirrors are surfaces of constant phase for that mode. It is, then, impossible for this cavity to exhibit a resonance condition and thus we say there is no cavity.

In this work, we will consider geometrically symmetric cavities for which $\mathcal{R}_1 = \mathcal{R}_2$ so that $\text{g} = \text{g}_1 = \text{g}_2$. For such a cavity, the stability criterion reduces to $-1 < \text{g} < 1$.

In the limit that $\text{g} \to 1$ we have that $L_{\text{cav}} \to 0$. Such a cavity is referred to as near-planar because, compared to the length of the cavity, the mirror ROC is so large that the mirror surface is effectively planar. Near-planar cavities have historically been used for cQED due to their very small mode volumes which result in large cQED coupling parameters g [17, 67, 106–109].

In the limit that $\text{g} = 0$ we have that $L_{\text{cav}} = \mathcal{R}$. Such a cavity is called a confocal cavity because the focal points of the two mirrors, when thought of as focusing optics, overlap in the center of the cavity. In Refs. [80, 110], a tunable-length, near-confocal cavity is used to realize coupling of BECs to many high order optical cavity modes to realize tunable short range interactions between the atoms within the BECs.

In the limit that $\text{g} \to -1$ we have that $L_{\text{cav}} \to 2\mathcal{R}$. Such a cavity is referred to as near-concentric. A near-concentric cavity is the longest cavity that can be made with mirrors of fixed radii of curvature before the mode becomes unstable and vanishes. The optical cavity designed and assembled for E6 is in the near-concentric regime, so we will explore the properties of such cavities thoroughly. There are a few other efforts to realize cQED interactions within near-concentric cavities using both neutral atoms and ions [21, 111, 112].

It will be useful to define the distance to concentric:

$$\delta_{\text{conc}} = 2\mathcal{R} - L_{\text{cav}}. \tag{3.14}$$

With this definition in hand we can express the basic beam properties for the optical modes supported by a geometrically symmetric cavity in terms of $\mathcal{R}_i$, L_{cav} and g:

$$z_2 = -z_1 = \frac{L_{\text{cav}}}{2},$$

$$z_R = \frac{L_{\text{cav}}}{2}\sqrt{\frac{1+\text{g}}{1-\text{g}}} = \frac{\mathcal{R}}{2}\sqrt{(1-\text{g}^2)} = \frac{\delta_{\text{conc}}}{2}\sqrt{2\frac{\mathcal{R}}{\delta_{\text{conc}}} - 1} = \frac{1}{2}\sqrt{L_{\text{cav}}\delta_{\text{conc}}},$$

$$w_0^2 = \frac{\lambda}{\pi}z_R = \frac{\lambda}{2\pi}\mathcal{R}\sqrt{1-\text{g}^2} = \frac{\lambda}{2\pi}\sqrt{\delta_{\text{conc}}}\sqrt{2\mathcal{R} - \delta_{\text{conc}}}. \tag{3.15}$$

We see that, for a symmetric cavity, the beam waist appears at the center of the cavity as expected, and again, z_R depends only the mirror and cavity geometry. For reference, I have given multiple expressions for z_R to demonstrate differing dependences.

The waist of the cavity mode on the cavity mirrors is given by

$$w_{\mathrm{mir}} = w(z_{1,2}) = w_0 \sqrt{\frac{2}{1 + \mathrm{g}}} = w_0 \sqrt{2 \frac{\mathcal{R}}{\delta_{\mathrm{conc}}}}. \tag{3.16}$$

3.2.3 Round Trip Phase

The second condition for cavity resonance mentioned above is that the round trip phase for light within the cavity must be an integer multiple of 2π to ensure constructive interference for each reflection within the cavity.

Looking at Eq. (3.9) we see that the phase on-axis for a Gaussian beam is given by

$$\exp\left(i(kz - (n + m + 1)\psi(z))\right). \tag{3.17}$$

A beam propagating from the left mirror to the right collects a phase

$$
\begin{aligned}
&= (kz_2 - (n + m + 1)\psi(z_2)) - (kz_1 - (n + m + 1)\psi(z_1)) \\
&= kL_{\mathrm{cav}} - (n + m + 1)(\psi(z_2) - \psi(z_1)).
\end{aligned} \tag{3.18}
$$

At this point we re-express k in terms of the optical oscillation frequency f:

$$k = \frac{\omega}{c} = \frac{2\pi}{c} f, \tag{3.19}$$

so that

$$kL_{\mathrm{cav}} = \pi \frac{2L_{\mathrm{cav}}}{c} f = \pi \frac{f}{f_{\mathrm{FSR}}} = \frac{1}{2}\omega\tau, \tag{3.20}$$

where I have introduced the cavity free spectral range (FSR) $f_{\mathrm{FSR}} = \frac{c}{2L_{\mathrm{cav}}}$. Note that $\tau = 1/f_{\mathrm{FSR}}$ is the time it takes for light to traverse one full round trip through the cavity.

It can be shown with some work [102] that ϕ_G, the Gouy phase collected from one end of the cavity to the other, can be expressed as

$$\phi_G = \psi(z_2) - \psi(z_1) = \arccos(\pm\sqrt{\mathrm{g}_1 \mathrm{g}_2}) = \arccos(\mathrm{g}). \tag{3.21}$$

Here the positive sign is taken if $\mathrm{g}_1, \mathrm{g}_2 > 0$ and the negative sign is taken if $\mathrm{g}_1, \mathrm{g}_2 < 0$. The final equality follows if $\mathrm{g}_1 = \mathrm{g}_2$.

The total phase traversing from one end of the cavity to the other is given by

$$\pi \frac{f}{f_{\mathrm{FSR}}} - (n + m + 1)\phi_G. \tag{3.22}$$

The round trip phase is twice the single traversal phase. In addition, the round trip phase includes the reflection phases for each mirror, ϕ_{r_1} and ϕ_{r_2}. The condition for cavity resonance is that this round trip phase is an integer multiple of 2π:

$$\phi_{RT} = 2\pi\frac{f}{f_{\text{FSR}}} - 2(n+m+1)\phi_G - \phi_{r_1} - \phi_{r_2} = 2\pi q \tag{3.23}$$

for integer q. We will see that the Gouy and mirror phases do not adjust the frequency spacing between the cavity resonances but that they do affect the absolute position in frequency space of the family of resonances.

With the above considerations we can then solve for the cavity resonance frequencies:

$$f_q = f_{\text{FSR}}\left(q + \frac{2(n+m+1)\phi_G + \phi_{r_1} + \phi_{r_2}}{2\pi}\right) = f_{\text{FSR}}\left(q + \frac{\phi_C}{2\pi}\right), \tag{3.24}$$

where we have defined the 'extra' round trip phase

$$\begin{aligned}\phi_C =& 2(n+m+1)\phi_G + \phi_{r_1} + \phi_{r_2}, \\ =& 2\phi_G + \phi_{r_1} + \phi_{r_2}, \end{aligned} \tag{3.25}$$

where the second equality follows for the TEM_{00} mode. The resonance condition can be rewritten as

$$\omega_q\tau = 2\pi q + \phi_C. \tag{3.26}$$

For fixed values of n and m, the cavity resonance is satisfied for modes which are evenly spaced by f_{FSR}, with the lowest frequency mode occurring at $f_{\text{FSR}} \times \frac{\phi_C}{2\pi}$. Next, we see that, because of the Gouy phase, the transverse modes are spaced from each other and the fundamental modes by integer multiples of the cavity transverse mode spacing:

$$f_{\text{TMS}} = f_{\text{FSR}}\frac{\phi_G}{\pi} = f_{\text{FSR}}\frac{\arccos(\mathbf{g})}{\pi}. \tag{3.27}$$

For a near planar cavity we have $\mathbf{g} \approx 1$ which means that $\arccos(\mathbf{g}) \ll 1$ so that $f_{\text{TMS}} \ll f_{\text{FSR}}$. In this case the transverse modes cluster about the fundamental modes with higher order modes appearing at slightly *higher* frequencies. See Fig. 3.3.

For a confocal cavity we have $\mathbf{g} = 0$ meaning $\arccos(\mathbf{g}) = \frac{\pi}{2}$ so that $f_{\text{TMS}} = \frac{1}{2}f_{\text{FSR}}$. In this case all even-numbered transverse modes cluster exactly degenerately at the frequencies of the fundamental modes and the odd-numbered modes cluster at half-frequencies in-between the fundamental modes.

For a concentric cavity we have $\mathbf{g} \approx -1$ which means that $\arccos(\mathbf{g}) \approx \pi$. With this condition we have $f_{\text{TMS}} \to f_{FRS}$ while maintaining $f_{\text{TMS}} < f_{\text{FSR}}$. As concentricity is approached, transverse modes with longitudinal mode number q cluster around fundamental modes with higher longitudinal mode numbers $> q$. See Fig. 3.3 for more clarity. In this case it becomes more useful to consider not the spacing between the transverse modes and their

corresponding longitudinal mode, but, rather, the spacing between the transverse modes and the new longitudinal modes to which they are clustering. For this reason we redefine

$$f_{\text{TMS}} \to f_{\text{FSR}} - f_{\text{TMS}},$$
$$f_{\text{TMS}} = f_{\text{FSR}} \left(1 - \frac{\arccos(\mathbf{g})}{\pi} \right). \tag{3.28}$$

When looking at the spectra of a near-concentric cavity one would see longitudinal modes spaced by f_{FSR} with transverse modes clustering nearby at *lower* frequencies spaced by f_{TMS}.
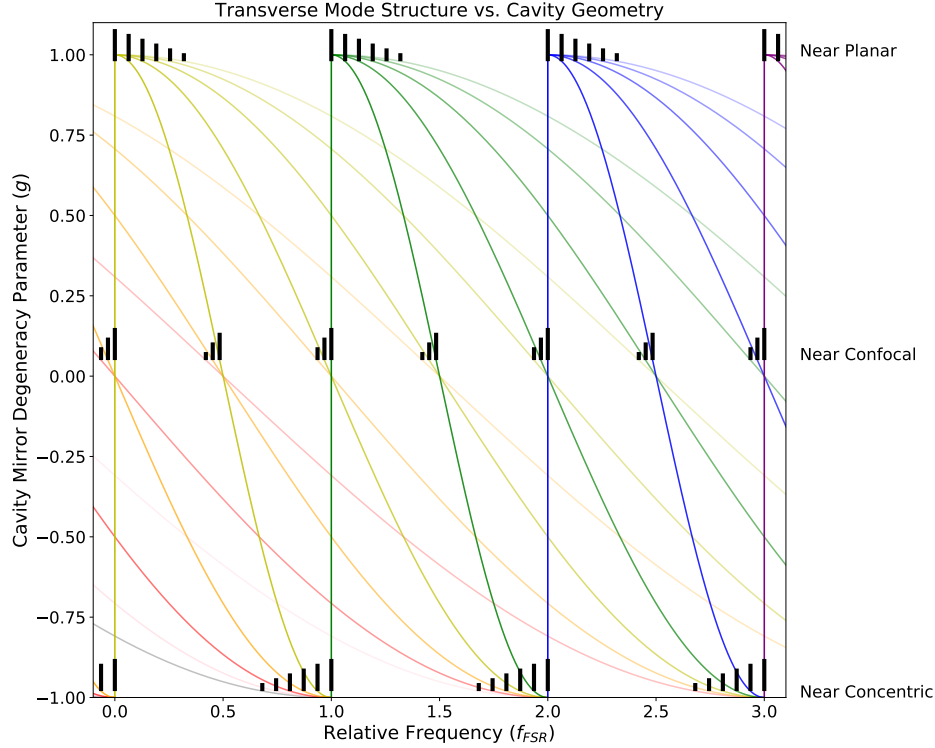
Figure 3.3: Illustration of how the transverse mode structure changes as the cavity geometry is continuously tuned from near planar (top) to near concentric (bottom) by increasing the cavity length. The colored lines represent the frequency of various cavity modes at a given value of g. Lines with the same longitudinal mode number have the same hue while transverse modes are displayed in increasingly fainter tones. The positions of modes are indicated by vertical black lines for $g = (0.98, 0.05, -0.98)$. Modes with increasing transverse mode numbers are decreasing in height. Note that for near planar cavities many transverse modes of the same longitudinal mode number are clustered whereas for a concentric cavity modes of different longitudinal mode numbers are clustered. Note also for the near confocal configuration that modes with even numbers for their transverse mode number are clustered at the positions of the TEM_{00} modes while modes with odd transverse mode number are clustered at the halfway points between the TEM_{00} modes.

3.3   Quantifying Concentricity

With the geometric and spectral conditions for cavity resonance laid out, we are now in a position to state what it means for a cavity to be near-concentric and quantify concentricity.

There is an increased interest in using near-concentric optical cavities in atomic cQED apparatuses [111–114]. Near concentric cavities are advantageous for cQED because they permit a small mode waist as needed for a large cooperativity parameter (referring back to the discussion in Sec. 1.3.4) while allowing a large NA transverse to the cavity axis. However, as a cavity is brought closer to concentricity, it becomes more sensitive to geometric misalignment, thus presenting a trade off between cooperativity and geometric tolerances. It will be useful then, for the purpose of comparing different systems, to have a method to quantify concentricity.

A number of different quantities have been used to quantify the concentricity of a particular optical cavity. Here I will present a brief glossary of the different quantities.

In Sec. 3.2.2 we encountered two quantities that already quantify the distance to concentric:

$$\delta_{\text{conc}} = 2\mathcal{R} - L_{\text{cav}}, \tag{3.29}$$

$$g = \frac{\delta_{\text{conc}}}{\mathcal{R}} - 1. \tag{3.30}$$

The distance to concentric δ_{conc} is a length that quantifies how much further the cavity mirrors could be separated while still in principle supporting an optical mode according to the cavity stability criterion $-1 < \mathsf{g}$. The mirror degeneracy parameter g is a dimensionless quantity which approaches -1 for a perfectly concentric cavity. The sum $1 + \mathsf{g}$ is a dimensionless ratio between δ_{conc} and $\mathcal{R}$. This ratio approaches 0 from above for a near-concentric cavity.

We saw in Sec. 3.2.3 that f_{TMS} depends on the cavity degeneracy parameter g.

$$\frac{f_{\text{TMS}}}{f_{\text{FSR}}} = 1 - \frac{\arccos(\mathsf{g})}{\pi}. \tag{3.31}$$

This ratio approaches 0 as the cavity approaches concentricity. The smallness of this ratio is then also a measure of concentricity. This is why near-concentric cavities are considered to be degenerate or near-degenerate.

Above, I mentioned that part of the usefulness of a near-concentric cavity is the small waist. The waist itself then can be taken as a dimension-full measure of concentricity and is useful when assessing misalignment tolerances. We can also consider the NA of the cavity

mode as a metric for concentricity that is essentially equivalent to the waist:

$$
\begin{aligned}
w_0 &= \sqrt{\frac{\lambda \mathcal{R}}{2\pi}} \left(1 - \mathbf{g}^2\right)^{\frac{1}{4}} = \sqrt{\frac{\lambda}{2\pi}} \left(L_{\text{cav}}\delta_{\text{conc}}\right)^{\frac{1}{4}}, \\
\tan(\theta) &= \frac{\lambda}{\pi w_0} = \sqrt{\frac{2\lambda}{\pi \mathcal{R}}} \left(1 - \mathbf{g}^2\right)^{-\frac{1}{4}} = \sqrt{\frac{2\lambda}{\pi \mathcal{R}}} \left(L_{\text{cav}}\delta_{\text{conc}}\right)^{-\frac{1}{4}}, \\
\text{NA} &= \sin(\theta).
\end{aligned}
\tag{3.32}
$$

Here θ is the Gaussian beam divergence angle. The waist vanishes at concentricity and the NA approaches unity. Note that when $\omega_0 \approx \lambda$ the paraxial approximation breaks down meaning some of these measures may need to be adjusted from the analytic forms presented here. This breakdown of paraxial optics as $\omega_0 \to \lambda$ also helps us to understand why no mode is possible beyond the concentric limit. See Fig. 3.4 for a comparison of these different parameters for different values of $\mathcal{R}$.

My opinion is that when building and designing an optical cavity with chosen values for $\mathcal{R}$, the distance to concentric, δ_{conc}, is the most practically useful measure of concentricity, because it relates to a quantity – length – that is directly measured and controlled as part of the cavity design. The degeneracy parameter, $\mathbf{g}$, is most useful as a dimensionless parameter for more abstract comparisons between different concentric cavities which may have different values for $\mathcal{R}$.
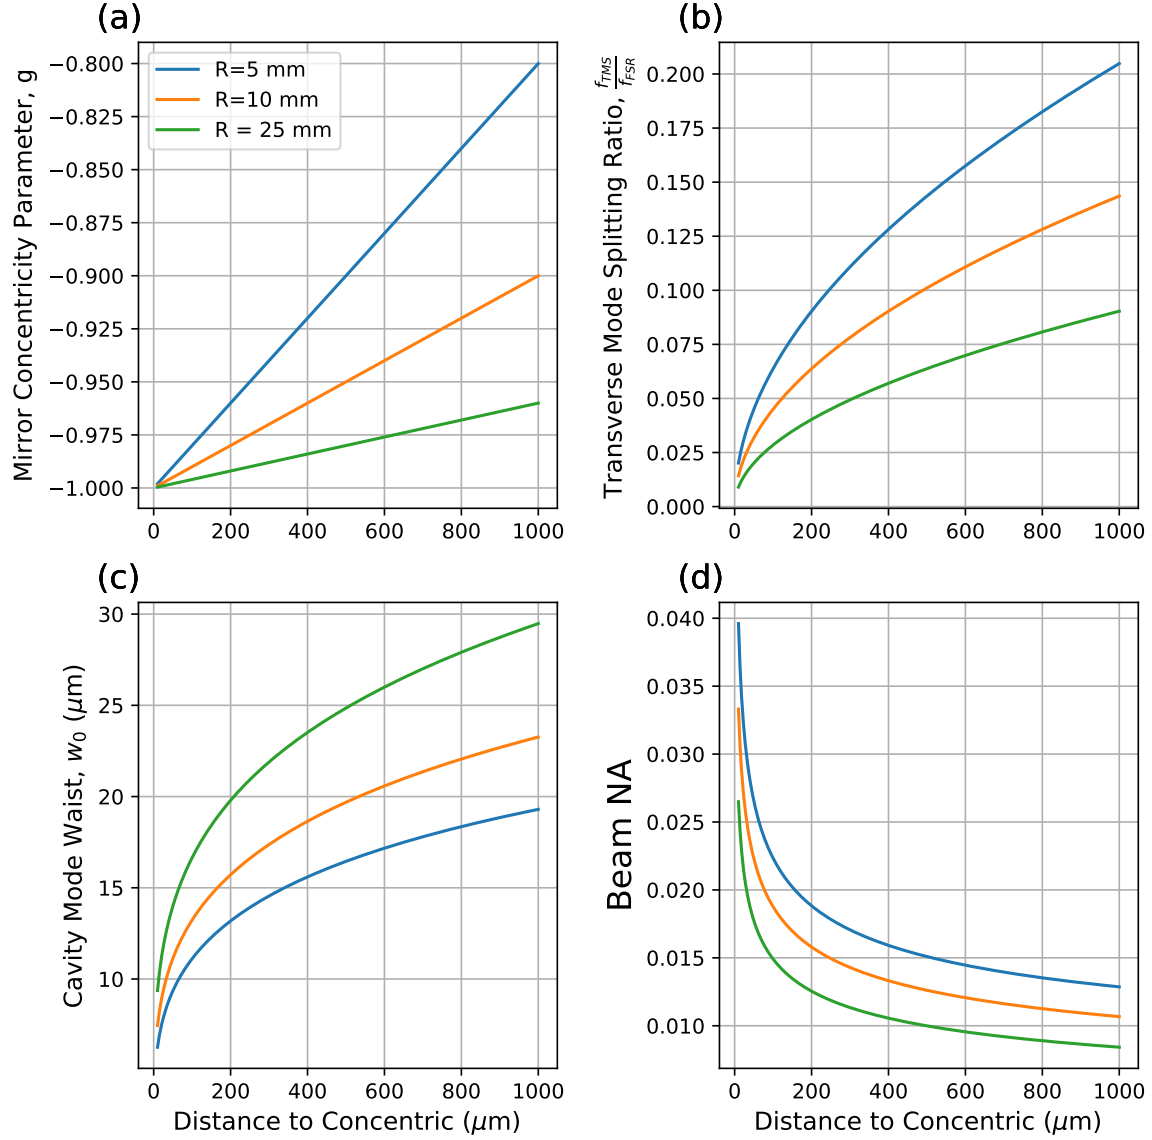
Figure 3.4: g, $\frac{f_{\text{TMS}}}{f_{\text{FSR}}}$, w_0, and NA plotted versus δ_{conc} for symmetric cavities with $\mathcal{R} = (5, 10, 25)$ mm.

3.4   Optical Cavity Dynamics and Lineshape

In this section, I will derive a number of key properties pertaining to the dynamics of a driven optical cavity in both the time and frequency domain including the impulse response function, the transfer function, the finesse, and the transmitted, reflected, and intracavity optical fields as a function of a given input field. The derivations here are similar to those in Refs. [102, 115].

We imagine an input optical field which impinges on the left side of an optical cavity described by the scalar time-dependent on-axis electric field $E_{\text{in}}(t)$:

$$E_{\text{in}}(t) = E_{\text{in}}^{(+)}(t) + E_{\text{in}}^{(-)}(t). \tag{3.33}$$

We will consider only the positive rotating part $E_{\text{in}}^{(+)}(t)$ and take the real part at the end if needed to recover the full real electric field. Note that the $e^{\mp i\omega t}$ factors are included in $E_{\text{in}}^{(\pm)}(t)$.

When a wave is incident on an optical cavity, it results in four new fields we must consider: a reflected field, a transmitted field, a right propagating intracavity field, and a left propagating intracavity field as shown in Fig. 3.5. In this section, we will determine expressions for these fields as a function of the input field and properties of the cavity and cavity mirrors. Each of these fields can be determined by taking an infinite sum of waves that have undergone a number of round trips within the cavity plus some minor additional propagation.
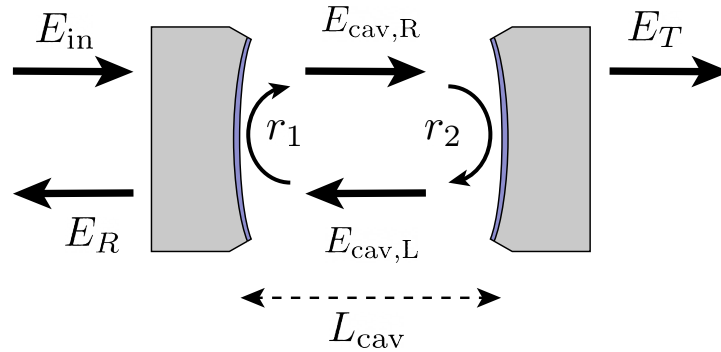
Figure 3.5: Light field E_{in} incident on one side of an optical cavity creates 4 new fields: A field reflected from the cavity E_R, a transmitted field, E_T, and right and left propagating intracavity fields $E_{\text{cav},R}$ and $E_{\text{cav},L}$. The two intracavity fields interfere to create the cavity standing wave.

The infinite sums for these four fields are as follows:

$$E_T^{(+)}(t) = t_1 t_2 \sum_{N=0}^{\infty} (r_1 r_2)^N \left(e^{-i\phi_C}\right)^N E_{\text{in}}^{(+)} \left(t - \frac{\tau}{2} - N\tau\right),$$

$$E_R^{(+)}(t) = r_1 e^{+i\phi_{r_1}} E_{\text{in}}^{(+)}(t) - t_1^2 r_2 e^{-i2\phi_G} e^{-i\phi_{r_2}} \sum_{N=0}^{\infty} (r_1 r_2)^N \left(e^{-i\phi_C}\right)^N E_{\text{in}}^{(+)}(t - \tau - N\tau),$$

$$E_{\text{cav},R}^{(+)}(z,t) = t_1 e^{-i(\psi(z)-\psi(z_1))} \frac{w_{\text{mir}}}{w(z)} \sum_{n=0}^{\infty} (r_1 r_2)^N \left(e^{-i\phi_C}\right)^N E_{\text{in}}^{(+)} \left(t - \frac{z - z_1}{c} - N\tau\right),$$

$$E_{\text{cav},L}^{(+)}(z,t) = -\, t_1 r_2 e^{-i\phi_{r_2}} e^{-i\phi_G} e^{+i(\psi(z)-\psi(z_2))}$$

$$\times\, \frac{w_{\text{mir}}}{w(z)} \sum_{n=0}^{\infty} (r_1 r_2)^N \left(e^{-i\phi_C}\right)^N E_{\text{in}}^{(+)} \left(t - \frac{\tau}{2} - \frac{z_2 - z}{c} - N\tau\right). \tag{3.34}$$

For calculation purposes, we define a round trip to begin with the beam at the left end of the cavity and propagating towards the right. The round trip is completed after the beam propagates to the right, collecting a factor of $e^{-i\phi_G}$, reflects of the right mirror collecting a factor of $-r_2 e^{-i\phi_{r_2}}$, propagates to the left now collecting a phase of $e^{-i\phi_G}$ again, and finally reflecting off of the first mirror collecting a factor of $-r_1 e^{-i\phi_{r_1}}$ for a total factor of $(-r_1)(-r_2)e^{-\left(2\phi_G+\phi_{r_1}+\phi_{r_2}\right)} = r_1 r_2 e^{-i\phi_C}$. Each round trip takes time $N\tau$ so this means that, at time t, the field which has undergone N round trips corresponds to the field that entered the cavity a time $N\tau$ in the past.

The previous paragraph explains the infinite sum which appears in the expression for each beam:

$$A\left(t,\tilde{t}\right) = \sum_{N=0}^{\infty} (r_1 r_2)^N \left(e^{-i\phi_C}\right)^N E_{\text{in}}^{(+)} \left(t - \tilde{t} - N\tau\right). \tag{3.35}$$

The prefactors to the summation and the extra temporal delay $\tilde{t}$ that appear for each term in Eqs. (3.34), correspond to parts of the optical path which occur before and after the beam undergoes round trips within the cavity. For example, for the transmitted field, the beam must first enter the cavity acquiring a factor of t_1 and, after circulating, leaves the cavity acquiring a factor of t_2. The time delay $\tilde{t} = \frac{\tau}{2}$ corresponds to the fact that, after the last *full* round trip, the beam must propagate once across the entire cavity to get to the location of the transmitted field, i.e. the right cavity mirror. For higher order cavity modes we must replace $\phi_G \to (n + m + 1)\phi_G$ and $\psi(z) \to (n + m + 1)\psi(z)$.

It's also worth explicitly pointing out that the reflected field has two terms. The first term corresponds to light which is promptly reflected off of the first cavity mirror and the second term corresponds to light which has entered the cavity, circulated within, and re-transmitted out the input mirror in the reflected direction. Note the $+$ sign for the reflection phase in the promptly reflected term. This sign appears because, in contrast to light circulating within the cavity, this reflection occurs at backside (the glass-dielectric) mirror interface rather than

the front side (the air-dielectric) mirror interface. The sign then follows from the form of the reflection scattering matrix Eq. 3.6.

Notice also that $E_{\text{in}/T/R}^{(+)}$ are defined at either the left or right cavity mirror and thus each of these fields are defined at positions along the Gaussian beam with the same waist size (at least for a geometrically symmetric cavity): w_{mir}. The intracavity field is enhanced by a factor of $w_{\text{mir}}/w(z)$. This enhancement can be significant in the case of a near-concentric cavity.

These fields can be written more compactly as:

$$E_T^{(+)}(t) = t_1 t_2 A\left(t, \frac{\tau}{2}\right),$$

$$E_R^{(+)}(t) = \left(r_1 E_{\text{in}}^{(+)}(t) - t_1^2 r_2 e^{-i\phi_C} A(t,\tau)\right) e^{+i\phi_{r_1}},$$

$$E_{\text{cav},R}^{(+)}(z,t) = t_1 e^{-i(\psi(z)-\psi(z_1))} A\left(t, \frac{z-z_1}{c}\right) \frac{w_{\text{mir}}}{w(z)},$$

$$E_{\text{cav},L}^{(+)}(z,t) = -t_1 r_2 e^{-i\phi_{r_2}} e^{-i\phi_G} e^{+i(\psi(z)-\psi(z_2))} A\left(t, \frac{\tau}{2} + \frac{z_2-z}{c}\right) \frac{w_{\text{mir}}}{w(z)}. \tag{3.36}$$

In the expression for $E_R^{(+)}(t)$, I have equivalently expressed the excess phase $-2\phi_G - \phi_{r_2} = -\phi_C + \phi_{r_1}$ for later convenience.

Now we would like to express the infinite sum $A(t, \tilde{t})$ as the convolution of the input field with a cavity convolution kernel or impulse response function. We will begin by re-expressing the reflection loss factor $r_1 r_2$:

$$r_1 r_2 = \sqrt{R_1}\sqrt{R_2}. \tag{3.37}$$

Recalling that $R_i = 1 - T_i - L_i$ and $T_i, L_i \ll 1$ we can Taylor expand to find

$$\sqrt{R_1}\sqrt{R_2} \approx 1 - \frac{1}{2}(T_1 + T_2 + L_1 + L_2) \approx e^{-\frac{1}{2}(T_1+T_2+L_1+L_2)}, \tag{3.38}$$

where we have also Taylor expanded $e^{-x} \approx 1 - x$.

We then see

$$(r_1 r_2)^N = e^{-\frac{1}{2}(T_1+T_2+L_1+L_2)N} = e^{-\frac{1}{2}(T_1+T_2+L_1+L_2)\frac{t_N}{\tau}}. \tag{3.39}$$

Here I introduced a discrete time label $t_N = N\tau$ to give the expression in terms of time rather than number of round trips. With this manipulation, and recalling $\tau = 1/f_{\text{FSR}}$, we can express

$$\frac{1}{2}(T_1 + T_2 + L_1 + L_2) f_{\text{FSR}} = \frac{\kappa}{2} \tag{3.40}$$

so that

$$(r_1 r_2)^N = e^{-\frac{\kappa}{2} t_N}. \tag{3.41}$$

Here, κ is an effective *energy* decay rate over time for a pulse within the cavity. The factor of $\frac{1}{2}$ appears because κ is an energy rather than amplitude decay rate and the factor $(r_1 r_2)^N$ appears in an amplitude expression.

We define the angular frequency FSR $\omega_{\text{FSR}} = 2\pi f_{\text{FSR}}$ and look at the ratio

$$\frac{\omega_{\text{FSR}}}{\kappa} = \frac{2\pi}{T_1 + T_2 + L_1 + L_2} = \mathcal{F}. \tag{3.42}$$

Here we have defined the important finesse parameter $\mathcal{F}$ as the ratio of the cavity's FSR to the energy decay rate for a pulse of light within the cavity. It makes intuitive sense that the decay rate should be proportional to the round trip transmission and losses.

We can write the decay time $\frac{1}{\kappa}$ in terms of the round trip time as

$$\frac{1}{\kappa} = \frac{\mathcal{F}}{2\pi}\tau = (T_1 + T_2 + L_1 + L_2)\tau. \tag{3.43}$$

For example, if there is a $T_1 + T_2 + L_1 + L_2 = 1/10,000$ chance for a pulse to transmit or scatter or absorb on a single round trip we see that $\mathcal{F} = 2\pi \times 10,000$ and that the cavity storage time is exactly $\frac{1}{\kappa} = 10,000 \times \tau$. We will see that, in addition to κ being the energy decay rate for the cavity, it is also the full-width at half-maximum (FWHM) of the cavity energy transfer function.

We can begin to rewrite the factor involving the round trip phase factor ϕ_c:

$$\left(e^{-i\phi_C}\right)^N = e^{-i\frac{\phi_C}{2\pi}\omega_{\text{FSR}}t_N}. \tag{3.44}$$

The overall expression can then be written as

$$A(t, \tilde{t}) = \sum_{N=0}^{\infty} e^{-i\frac{\phi_C}{2\pi}\omega_{\text{FSR}}t_N} e^{-\frac{\kappa}{2}t_N} E_{\text{in}}^{(+)}(t - \tilde{t} - t_N). \tag{3.45}$$

We can rewrite this as a convolution by introducing a comb of Dirac delta functions in time

$$\begin{aligned}
A(t, \tilde{t}) &= \int_{t'=-\infty}^{+\infty} E_{\text{in}}^{(+)}(t - t') \left(\sum_{N=0}^{\infty} e^{-i\frac{\phi_C}{2\pi}\omega_{\text{FSR}}t_N} e^{-\frac{\kappa}{2}t_N} \delta(t' - \tilde{t} - t_N) \right) dt' \\
&= \left(E_{\text{in}}^{(+)} * h_{\tilde{t}} \right)(t),
\end{aligned} \tag{3.46}$$

where we've defined the impulse response function

$$h_{\tilde{t}}(t) = \sum_{N=0}^{\infty} e^{-i\frac{\phi_C}{2\pi}\omega_{FSR}t_N} e^{-\frac{\kappa}{2}t_N} \delta(t - \tilde{t} - t_N). \tag{3.47}$$

Below we will calculate the Fourier transform of this impulse response function, which is the frequency transfer function for the cavity. We see that the impulse response function is

a train of impulses in time –corresponding to pulses separated by cavity round trip times τ– modulated by a decaying exponential with decay rate $\frac{\kappa}{2}$ and also by a complex exponential oscillating at angular frequency $\frac{\phi_C}{2\pi}\omega_{FSR}$. This oscillating exponential in the impulse response function will only correspond to a shift of all of the frequencies in the transfer function by an amount $\frac{\phi_C}{2\pi}\omega_{FSR}$ corresponding to the frequency offset described in Sec. 3.2.3.

We now Fourier transform the impulse response function to get the transfer function:

$$
\begin{aligned}
H_{\tilde{t}}(\omega) =&\tilde{h}_{\tilde{t}}(\omega) = \int e^{+i\omega t} h_{\tilde{t}}(t) dt \\
=& \int e^{+i\omega t} \sum_{N=0}^{\infty} e^{-i\frac{\phi_C}{2\pi}\omega_{FSR} t_N} e^{-\frac{\kappa}{2} t_N} \delta(t - \tilde{t} - t_N) dt \\
=& e^{i\omega\tilde{t}} \sum_{N=0}^{\infty} \left(e^{i(\omega\tau - \phi_C) - \frac{\kappa}{2}\tau} \right)^N .
\end{aligned}
$$
(3.48)

Since $\left| \exp\left(i(\omega\tau - \phi_c) - \frac{\kappa}{2}\tau \right) \right| < 1$, this is an easily summable geometric series:

$$
H_{\tilde{t}}(\omega) = e^{i\omega\tilde{t}} \frac{1}{1 - e^{i(\omega\tau - \phi_C)} e^{-\frac{\kappa}{2}\tau}}.
$$
(3.49)

This is the full transfer function for an optical cavity. The only approximation we have made here is that $\mathcal{F} \gg 1$.

To bring this transfer function into a more familiar Lorentzian form for the cavity line-shape, we expand this function about its resonances. We can see that this function will have the largest amplitude when the imaginary exponential in the denominator is equal to unity. This occurs when

$$
\omega\tau - \phi_C = 2\pi q,
$$
(3.50)

with integer q. This is of course the exact same expression for cavity resonance as we found above in Eq. (3.26). let us now expand the cavity transfer function for a small detuning Δ about the cavity resonance frequency ω_q. That is let $\omega = \omega_q + \Delta$. In this case we have $\omega\tau - \phi_C = 2\pi q + \Delta\tau$. If $\Delta\tau = 2\pi \frac{\Delta}{\omega_{FSR}} \ll 2\pi$ then we can expand the transfer function as

$$
H_{\tilde{t}}(\omega_q + \Delta) = e^{i(\omega_q + \Delta)\tilde{t}} \frac{1}{1 - (1 + i\Delta\tau)\left(1 - \frac{\kappa}{2}\tau\right)}.
$$
(3.51)

Expanding the denominator and dropping the doubly small term proportional to $\Delta\kappa\tau^2$:

$$
H_{\tilde{t}}(\omega_q + \Delta) = e^{i(\omega_q + \Delta)\tilde{t}} \frac{\mathcal{F}}{\pi} \frac{\frac{\kappa}{2}}{\frac{\kappa}{2} - i\Delta} = e^{i\omega\tilde{t}} \frac{\mathcal{F}}{\pi} \frac{\frac{\kappa}{2}}{\frac{\kappa}{2} - i\Delta}.
$$
(3.52)

Let us pause now and de-construct this expression for the cavity transfer function. It consists of four parts which we will explore from right to left as this is the order of importance. The far right term is the normalized cavity lineshape:

$$\chi(\Delta) = \frac{\frac{\kappa}{2}}{\frac{\kappa}{2} - i\Delta}. \tag{3.53}$$

The transfer function has a Lorentzian profile centered at $\omega = \omega_q$. The lineshape is normalized so that $\chi(0) = 1$. If we look at the squared amplitude of the lineshape (proportional to the energy transfer function) we get

$$|\chi(\Delta)|^2 = \frac{\kappa^2}{\kappa^2 + 4\Delta^2}. \tag{3.54}$$

We see that this squared amplitude falls to half of its peak value when $\Delta = \kappa/2$ so κ is the *energy* lineshape FWHM. Remember we saw earlier that $\kappa/2$, the *energy* lineshape half-width at half-maximum (HWHM), was the amplitude decay rate of the amplitude impulse response function $h_{\tilde{t}}(t)$. This is the expected relationship between the amplitude and energy decay rates and linewidths. Finally we have that

$$\int |\chi(\Delta)|^2 d\Delta = \frac{\pi}{2}\kappa. \tag{3.55}$$

The next term in the transfer function is a scaling factor $\frac{\mathcal{F}}{\pi}$. This is the term responsible for the build-up of energy within the cavity. This term captures the fact that, for resonant light, the electric field within the cavity is the sum of light which has reflected thousands or more times within the cavity, meaning that the electric field within the cavity can be many times larger than the electric field of the incident beam. The final term, $e^{i\omega\tilde{t}}$, corresponds to the extra time delay each field experiences after completing its final round trip.

Taking the inverse Fourier transform of $H_{\tilde{t}}(\omega)$ we find the single-mode cavity impulse response function, which is an exponential ringdown with decay rate $\frac{\kappa}{2}$:

$$h_{\tilde{t}}(t) = f_{\text{FSR}} e^{-i\omega_q(t-\tilde{t})} e^{-\frac{\kappa}{2}(t-\tilde{t})}. \tag{3.56}$$

We've now calculated the single-mode impulse response function and transfer function for an optical cavity. We can now, using the transfer function, easily calculate the transmitted, reflected, and intracavity fields for a monochromatic input field of the form

$$E_{\text{in}}^{(+)}(t) = E_{\text{in},0}^{(+)} e^{-i\omega t} = \left| E_{\text{in},0}^{(+)} \right| e^{-i\phi_{\text{in}}} e^{-i\omega t}, \tag{3.57}$$

with $\omega = \omega_q + \Delta$ and ϕ_{in} the inconsequential phase of the input beam. Note that

$$\left(E_{\text{in}}^{(+)} * h_{\tilde{t}} \right)(t) = E_{\text{in}}^{(+)}(t) H_{\tilde{t}}(\omega), \tag{3.58}$$

as can be confirmed by direct computation or the convolution theorem.

Using Eqs. (3.36), (3.46), (3.52) and (3.58), we can calculate the transmitted, reflected and intracavity fields as

$$E_T^{(+)}(t) = E_{\text{in}}^{(+)}(t) t_1 t_2 \frac{\mathcal{F}}{\pi} \chi(\Delta) e^{i\frac{\omega\tau}{2}},$$

$$E_R^{(+)}(t) = E_{\text{in}}^{(+)}(t) \left(r_1 - t_1^2 r_2 \frac{\mathcal{F}}{\pi} \chi(\Delta) e^{-i\phi_C} e^{i\omega\tau} \right) e^{+i\phi_{r_1}},$$

$$E_{\text{cav},R}^{(+)}(t) = E_{\text{in}}^{(+)}(t) t_1 \frac{\mathcal{F}}{\pi} \chi(\Delta) \frac{w_{\text{mir}}}{w(z)} e^{-i(\psi(z)-\psi(z_1))} e^{ik(z-z_1)},$$

$$E_{\text{cav},L}^{(+)}(t) = E_{\text{in}}^{(+)}(t) t_1 (-r_2) \frac{\mathcal{F}}{\pi} \chi(\Delta) \frac{w_{\text{mir}}}{w(z)} e^{+i(\psi(z)-\psi(z_2))} e^{ik(z_2-z)} e^{i\frac{\omega\tau}{2}} e^{-i\phi_G} e^{-i\phi_{r_2}}. \tag{3.59}$$

We can gain a little intuition by simplifying these expressions and looking at their magnitude on resonance so that $\Delta = 0$ and $\omega = \omega_q$ and approximating $r_1, r_2 \approx 1$. We recall $\omega_q \tau - \phi_C = 2\pi q$. The results are

$$\left| E_T^{(+)} \right| = \left| E_{\text{in},0}^{(+)} \right| t_1 t_2 \frac{\mathcal{F}}{\pi},$$

$$\left| E_R^{(+)} \right| = \left| E_{\text{in},0}^{(+)} \right| \left| 1 - t_1^2 \frac{\mathcal{F}}{\pi} \right|,$$

$$\left| E_{\text{cav},R}^{(+)} \right| = \left| E_{\text{in},0}^{(+)} \right| t_1 \frac{\mathcal{F}}{\pi} \frac{w_{\text{mir}}}{w(z)},$$

$$\left| E_{\text{cav},L}^{(+)} \right| = \left| E_{\text{in},0}^{(+)} \right| t_1 \frac{\mathcal{F}}{\pi} \frac{w_{\text{mir}}}{w(z)}. \tag{3.60}$$

We can calculate the real intracavity field using $E_{\text{cav}}(z,t) = E_{\text{cav},R}^{(+)} + E_{\text{cav},L}^{(+)} + c.c.$ and the equations for $E_{\text{cav},R/L}^{(+)}$ in Eq. (3.59). The calculation involves making us of the identity

$$e^{iA} + e^{iB} + c.c. = \left(e^{i\frac{A+B}{2}} + c.c. \right) \left(e^{i\frac{A-B}{2}} + c.c. \right),$$

$$2 \left(\cos(A) + \cos(B) \right) = 4 \cos\left(\frac{A+B}{2} \right) \cos\left(\frac{A-B}{2} \right). \tag{3.61}$$

It also makes use of

$$2kL_{\text{cav}} = \omega\tau = \omega_q \tau + \Delta\tau = 2\pi q + \phi_C + \Delta\tau,$$

$$\phi_C = 2\phi_G + \phi_{r_1} + \phi_{r_2}. \tag{3.62}$$

The resultant intracavity field is

$$E_{\text{cav}}(z,t) = \left| E_{\text{in},0}^{(+)} \right| 4 t_1 \frac{\mathcal{F}}{\pi} \chi(\Delta) \frac{w_{\text{mir}}}{w(z)} \cos\left(\omega t + \phi_{\text{in}} - \frac{\Delta\tau}{2} - \pi \left(q + \frac{1}{2} \right) - \frac{\phi_{r_1}}{2} \right)$$

$$\times \cos\left(kz - \psi(z) - \frac{\Delta\tau}{4} - \frac{\pi}{2}(q+1) - \frac{\phi_{r_1} - \phi_{r_2}}{4} \right). \tag{3.63}$$

Eq. (3.63) determines the optical field on axis as a function of z and t.

In Sec. 3.6 where we explore the relative z-positions of the anti-nodes of two cavity modes, we will be interested in the spatially varying part of Eq. (3.63). To simplify Eq. (3.63) we define

$$E_{\text{cav},0} = \left|E_{\text{in},0}^{(+)}\right| 4t_1 \frac{\mathcal{F}}{\pi}\chi(\Delta)\frac{w_{\text{mir}}}{w_0},$$

$$\phi_t = \phi_{\text{in}} - \frac{\Delta\tau}{2} - \pi\left(q + \frac{1}{2}\right) - \frac{\phi_{r_1}}{2},$$

$$\phi_z = -\psi(z) - \frac{\Delta\tau}{4} - \frac{\pi}{2}(q+1) - \frac{\phi_{r_1} - \phi_{r_2}}{4}, \tag{3.64}$$

which yields

$$E_{\text{cav}}(z,t) = E_{\text{cav},0}\frac{w_0}{w(z)}\cos(\omega t + \phi_t)\cos(kz + \phi_z). \tag{3.65}$$

With $\rho^2 = x^2 + y^2$, we can express the electric field magnitude everywhere within the cavity:

$$|E_{\text{cav}}(x,y,z,t)|^2 = E_{\text{cav},0}^2\left(\frac{w_0}{w(z)}\right)^2 e^{-2\frac{\rho^2}{w(z)^2}}\cos^2(\omega t + \phi_t)\cos^2(kz + \phi_z). \tag{3.66}$$

3.5 Optical Cavity Energy Flow

In this section I will derive and tabulate, for useful future reference, equations relating the intracavity peak electric field and intracavity energy to the input and output optical powers. These equations will be useful for experimentally determining, for example, intracavity photon numbers and optical trap depths. Throughout this section we assume the cavity drive field is on resonance $\Delta = 0$.

For an optical plane wave defined with electric field $E = E^{(+)} + E^{(-)}$ the local intensity can be calculated as

$$I = 2\epsilon_0 c|E^{(+)}|^2. \tag{3.67}$$

If $I = I(z)$ represents the central, peak intensity at position z along a Gaussian beam, then the power in the beam can be calculated by

$$P = \int I(z)e^{-2\frac{\rho^2}{w(z)^2}}dA = I(z)\frac{\pi w(z)^2}{2}, \tag{3.68}$$

with $\rho^2 = x^2 + y^2$. Note that $I(z) = I_0\frac{w_0^2}{w(z)^2}$, where I_0 is the peak intensity at the center of the waist of the Gaussian beam, so we see that $P = I_0\frac{\pi w_0^2}{2}$ independent of position z along the beam axis, as expected. The electric field is related to the power by $P = \epsilon_0 c\pi w(z)^2|E^{(+)}|^2$.

We can then calculate the transmitted and reflected powers on resonance using Eqs. 3.60:

$$
\begin{aligned}
P_T &= 4\eta_1\eta_2 P_{in}, \\
P_R &= (1 - 2\eta_1)^2 P_{in}.
\end{aligned}
\tag{3.69}
$$

I've defined the two cavity input and output coupling efficiencies:

$$
\eta_i = T_i\frac{\mathcal{F}}{2\pi} = \frac{T_i}{T_1 + T_2 + L_1 + L_2}.
\tag{3.70}
$$

We can account for losses by defining $\eta_L = \frac{L_1+L_2}{T_1+T_2+L_1+L_2}$ and $P_L = 4\eta_1\eta_L P_{in}$ so that $P_T + P_R + P_L = P_{in}$.

Some authors define the so-called intracavity circulating power by converting $E^{(+)}_{\text{cav},L/R}$ into powers using

$$
P_{\text{cav,R}} = P_{\text{cav,L}} = \frac{2}{\pi}\mathcal{F}\eta_1 P_{in}.
\tag{3.71}
$$

The salient point here is that the circulating power is enhanced by a factor of $\mathcal{F}$ compared to P_{in}. However, the reader is warned that, because the left and right traveling waves which compose the circulating power interfere to form a lattice, it is necessary to keep track of some factors of 2 when converting the peak electric field amplitude and the circulating power.

For the sake of calculating Rabi frequencies, it is useful to be able to convert either the input or output power to the peak amplitude of the electric field. This is easily done using Eqs. (3.60) and (3.64):

$$
\begin{aligned}
E_{\text{cav},0} &= \frac{4}{w_0}t_1\frac{\mathcal{F}}{\pi}\sqrt{\frac{P_{\text{in}}}{\pi\epsilon_0 c}}, \\
&= \frac{4}{w_0}\frac{1}{t_2}\sqrt{\frac{P_T}{\pi\epsilon_0 c}}.
\end{aligned}
\tag{3.72}
$$

Inspecting Eq. (3.66) we can see that the cavity spatial mode function is given by

$$
|\boldsymbol{f}(x,y,z)|^2 = \left(\frac{w_0}{w(z)}\right)^2 e^{-2\frac{\rho^2}{w(z)^2}}\cos^2(kz + \phi_z).
\tag{3.73}
$$

Integrating this radially and over the length of the cavity, as in Eq. (1.6), yields the cavity mode volume

$$
V_{\text{mode}} = \int |\boldsymbol{f}(x,y,z)|^2\,dV = \frac{L_{\text{cav}}}{4}\pi w_0^2.
\tag{3.74}
$$

This gives the intuitively pleasing result that the cavity mode volume is proportional to the length of the cavity and the area of the mode waist. With V_{mode} we can calculated the peak

cavity QED coupling parameter g_0 using Eq. (1.20). Finally, we can use the mode volume to calculate the total energy stored in the cavity mode[2] using Eq. (1.7):

$$U = \frac{1}{2}\epsilon_0 V_{\text{mode}} \left| E_{\text{cav, 0}} \right|^2 . \tag{3.75}$$

Using Eq. (3.72), we can express the intracavity energy U as a function of the input or output power:

$$\begin{aligned} U =& \frac{4\eta_1}{\kappa}P_{\text{in}}, \\ =& \frac{1}{\kappa\eta_2}P_T. \end{aligned} \tag{3.76}$$

We notice that this allows us to express

$$\begin{aligned} P_T =& \kappa\eta_2 U, \\ P_L =& \kappa\eta_L U, \\ P_{R,\text{leak}} =& \kappa\eta_1 U. \end{aligned} \tag{3.77}$$

Where I've introduced $P_{R,\text{leak}} = 4\eta_1^2 P_{\text{in}}$ which represents the power of the light leaking out of the cavity towards the reflection port separate from light which is promptly reflected off of the cavity. From this we get the intuitive interpretation that $\kappa\eta_i$ represents the energy loss rate into a particular cavity energy decay channel.

The discussion throughout this chapter has been purely classical. However, as per the discussion in Sec. 1.3.1, we can easily draw quantum conclusions by noticing that $U = \hbar\omega\bar{n}$, where $\bar{n}$ is the average photon number in the cavity. We see, then, that we can easily calculate intracavity photon number $\bar{n}$ if we know the transmitted or reflected power, the linewidth κ, and the relevant coupling efficiencies. This formal parallel between the classical and quantum descriptions of an optical cavity is developed further in Appendix F.

3.6 Relative Beatnote Between Two Cavity Modes

The E6 science cavity is designed to support two wavelengths of light to create a bi-chromatic optical lattice. One wavelength is 780 nm, which is near-resonant with the [87]Rb D$_2$ transition. This light is referred to as the probe light and facilitates cQED interactions. The other wavelength is 1560 nm, which is very far red-detuned from atomic resonance. This light creates a mechanical potential that serves to trap atoms optically at the anti-nodes of this lattice [45] The 1560 nm cavity mode is referred to as the cavity ODT.

For many applications one might desire that the anti-nodes of the probe beam be co-located with the anti-nodes of the cavity ODT to ensure the atoms are trapped at the points

---

[2]Note that Eq. (1.7) was expressed in terms of the amplitude of the positive rotating part of the field but Eq. (3.75) is expressed in terms of the real field amplitude, hence a relative factor of $\left(\frac{1}{2}\right)^2$.

of maximal probe coupling. More generally, one might desire arbitrary control over the relative phase of the probe and trap lattices.

In this section, I will calculate what we call the cavity beatnote. The cavity beatnote is a map, as function of distance along the cavity axis, of the local intensity of the probe at the location of the trap maxima. One can see that the advantage of using a 1560 nm cavity ODT is that, because it is precisely twice the wavelength of the probe, this beatnote can, in principle, have a very long wavelength corresponding to uniform probe coupling throughout the cavity. In fact, this trick has been used previously in cQED applications [21, 78] and a detailed discussion of some of the more subtle concepts affecting the beatnote, such as the effect of the Gouy phase, is found in Ref. [104].

We can quickly come up with a simple picture for the cavity beatnote in the center of the cavity using some naive symmetry considerations that I discuss shortly. However, there are two non-trivial details that motivate the extended discussion that appears in the rest of this section. These details are the effects of 1) the Gouy phase ϕ_G and 2) non-trivial mirror reflection phases ϕ_{r_1} and ϕ_{r_2} on the beatnote.

For the moment let us neglect these two effects by letting $\phi_C = 0$. If the longitudinal mode number for a given mode is (even) odd then an (even) odd number of half-wavelengths fit between the cavity mirrors. By symmetry this means that even (odd) modes will have a node (anti-node) at the center of the cavity. In Fig. 3.6 I show the cavity intensity profiles near the center of the cavity for all possible combinations of longitudinal parity modes for the cavity ODT and the probe. We see that, for a symmetric cavity, the trap and probe will have a favorable beatnote in any case as long as the probe drives an **odd** longitudinal mode. This tells us that it is important to have the ability to tune the cavity length so that, with a fixed probe frequency, we can drive either an even or odd longitudinal mode as desired.

The case of a symmetric cavity is reasonable to consider for a cavity made from two identical mirrors. However, as described later in Sec. 4.1.2, the E6 cavity is one-sided to optimize detection efficiency. This means there is no guarantee that $\phi_{r_1} = \phi_{r_2}$ for each wavelength. As described in Sec. 3.1, we can imagine ϕ_{r_1} and ϕ_{r_2} as shifting the field reflection surface away from the mirror's physical surface by a fraction of a wavelength. If these phases differ for the left and right mirror, then the cavity mode will no longer be centered at the geometric center of the two mirrors. While this shift of the cavity mode is not problematic for a single cavity mode, it can cause large shifts in the relative position between two cavities modes, resulting in dramatic changes to the cavity beatnote relative to the naive symmetry picture given in the previous paragraph. Below, I will calculate the beatnote across the cavity in the presence of non-trivial mirror phases and the Gouy phase and for various cavity ODT modes relative to a particular probe mode. I assume both drive fields are on resonance. The amplitude profile of light along the cavity axis is proportional to

$$\cos\left(\Phi_{cav}\right) = \cos\left(k_q z - \psi(z) - \frac{\phi_{r_1} - \phi_{r_2}}{4} - \frac{\pi}{2}(q+1)\right), \tag{3.78}$$
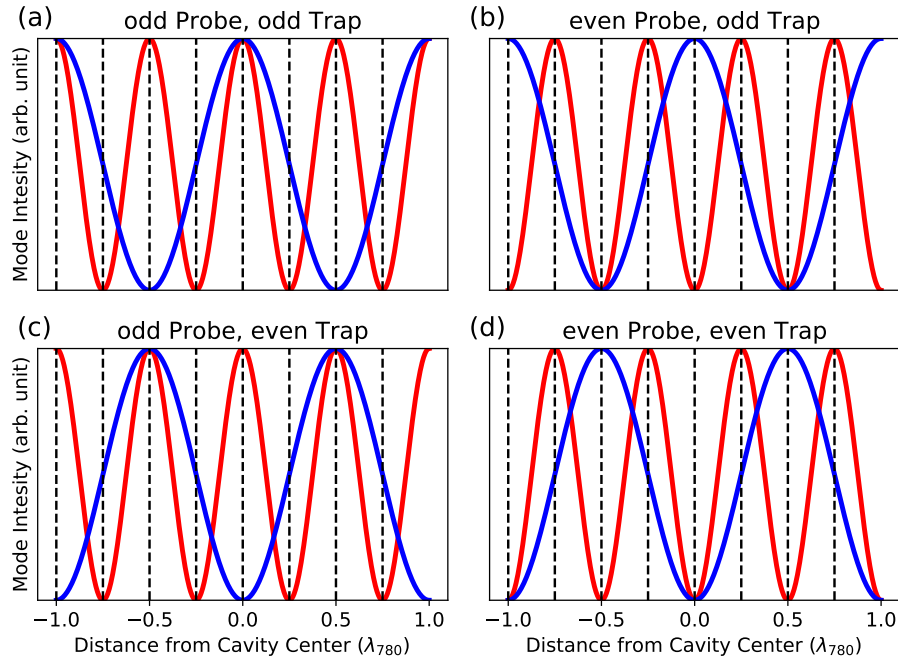
Figure 3.6: Naive cavity beatnote at the center of a symmetric cavity for 780 nm probe (red) and 1560 nm cavity ODT (blue) intracavity fields. Both configurations with an odd probe mode have anti-nodes overlapped as desired. Both configurations with even probe modes have nodes of the probe overlapped with anti-nodes of the cavity ODT which is the worst case scenario.

and we recall that

$$k_q = \frac{\omega_q}{c} = \frac{\omega_{\text{FSR}}}{c}\left(q + \frac{\phi_C}{2\pi}\right).$$ (3.79)

If we neglect the Gouy phase, we will be able to get a simple, intuitive, analytical form for the relative beatnote between the probe and the trap. This approximation is valid when the cavity length is much shorter than the Rayleigh range, which is explicitly not the case for a concentric cavity. Nonetheless, we'll pursue it for our intuition and include a numerical calculation of the beatnote including the Gouy phase at the end. The m^{th} anti-node of the cavity ODT occurs at position z^m where

$$\Phi_{cav}^T = m\pi,$$
$$k_q^T z^m = m\pi + \frac{\phi_{r_1}^T - \phi_{r_2}^T}{4} + \frac{\pi}{2}(q^T + 1).$$ (3.80)

Here the T superscript indicates the cavity ODT.

We then define

$$k_q^P = pk_q^T + \Delta k,$$
$$\Delta k = \frac{\omega_{\text{FSR}}}{c} \left(\Delta n + \frac{\phi_C^P - p\phi_C^T}{2\pi} \right) = \frac{2\pi}{L_{\text{cav}}} \left(\frac{\Delta n}{2} + \frac{\phi_C^P - p\phi_C^T}{4\pi} \right). \tag{3.81}$$

where I've defined $\Delta n = q^P - pq^T$, the number of free spectral ranges away from the perfect doubled frequency condition. For reference in the E6 cavity $q^T \approx 12,000$. The wavelength associated with Δk is the cavity length L_{cav} divided by a half-integer $\frac{\Delta n}{2}$ plus a minor correction related to the mirror phases.

We can then plug these anti-node positions into the probe cavity phase,

$$\Phi_{\text{cav}}^P = \Delta k z^m + pk_q^T z^m - \frac{\phi_{r_1}^P - \phi_{r_2}^P}{4} - \frac{\pi}{2} \left(q^P + 1 \right), \tag{3.82}$$

so that the probe lattice is given by

$$\cos \left(\Delta k z^m - \frac{(\phi_{r_1}^P - p\phi_{r_1}^T) - (\phi_{r_2}^P - p\phi_{r_2}^T)}{4} - \frac{\pi}{2} (\Delta n - p + 1) + pm\pi \right). \tag{3.83}$$

We will now pick apart this phase factor to get an intuitive feel for the form of this beatnote, in particular I would like to cut through some of the complexity involving all of the different phase factors both in this final expression and in the expression for Δk. The factor $pm\pi$ tells us that, for odd p, the sign of the probe flips as we step from one cavity ODT anti-node to the next. In our case $p = 2$ so this term has no effect.

Next, we calculate the phase of the beatnote on each end-mirror by letting $z^m = \pm L_{\text{cav}}/2$ in Eq. (3.83):

$$\Phi_1^P = \frac{p\phi_{r_1}^T - \phi_{r_1}^P}{2} - \Delta n\pi + \frac{\pi}{2} (p - 1),$$
$$\Phi_2^P = \frac{\phi_{r_2}^P - p\phi_{r_2}^T}{2} + \frac{\pi}{2} (p - 1). \tag{3.84}$$

The main effect we see is that the phase of the beatnote on the end mirrors is set by the relative difference in the mirror reflection phases for the two wavelengths. Next, we notice the periodicity of the beatnote is approximately

$$\lambda_{\text{beat}} = \frac{2\pi}{\Delta k} \approx \frac{2L_{\text{cav}}}{\Delta n}. \tag{3.85}$$

In summary, then, we see that the beatnote will be a sinusoidal function with the following behavior. The phase is given by Φ_1^P at the left mirror. The beatnote has Δn half-wavelengths

within the cavity mode. The phase is given by Φ_2^P at the right end mirror. When we neglect the Gouy phase, we see that the phase of the beatnote increases linearly, meaning that the beatnote is a perfect sinusoid within the cavity.

The Gouy phase spoils the analytical treatment given above because it changes the effective wavelength of the cavity modes within the Rayleigh range, which can cause slight relative shifts in the positions of their anti-nodes giving the beatnote phase a slight nonlinearity:

$$\lambda_{eff} = \frac{2\pi}{\frac{d\Phi_{cav}}{dz}} = \frac{2\pi}{k + \frac{d\psi(z)}{dz}}. \tag{3.86}$$

At the waist of a beam we have

$$\frac{d\psi(z)}{dz} = \frac{1}{z_R}, \tag{3.87}$$

so we see that the change in the effective wavelength in the vicinity of the waist will be large when $\frac{\lambda}{z_R}$ is appreciable.

This cavity beatnote function is plotted in Fig. 3.7. We see that in some cases, the Gouy phase results in either a slight suppression or enhancement of the cavity beatnote relative to the analytic expectation.

I also plot the beatnote corresponding to the amplitude of the derivative of the probe amplitude along the cavity axis at the location of the anti-nodes of the trap (red). This latter derivative beatnote is of interest because the extrema of this beatnote correspond to the points where coupling to the probe is strongly linearly varying across the cavity ODT potential wells. These so-called linear spots have historically been used to great effect in the E3 apparatus to explore quantum optomechanics by strongly coupling the motion of an atomic cloud to the photonic cavity mode [28, 67, 68].

I'd now like to take some time to discuss the implications of these beatnote investigations on our cavity design and science outlook for this experiment. More will be said about the experimental details of the cavity mirror fabrication in Chapter 4. However, here, it suffices to say that we are using mirrors with different optical coatings for the input and output mirrors, and that we did not control the reflection phase when designing the cavity mirror coatings. In light of this, we wanted to answer the question: what can we do in the event that we find out, after the cavity is under UHV, that the mirror phases present us with the worst case cavity beatnote in the center of the cavity.

What would be a worst case scenario for us? Suppose we choose $\Delta n = 0$. The best case scenario is depicted in Fig. 3.7(a) in which the beatnote is maximal throughout the cavity. This essentially corresponds to the panels on the left side of Fig. 3.6 in which all cavity ODT anti-nodes overlap with an anti-node of the probe. One, naively, might think that Fig. 3.7(b) is the worst case scenario since the coupling is low everywhere, corresponding to Figs. 3.6(b) and (d) with an even probe mode. However, if the cavity length is adjusted by one FSR then the probe switches to an odd mode resulting in a maxima of the beatnote in the center of the symmetric cavity. This means that the scenario in Fig. 3.7(b) is actually ok as long as the length of the cavity can be tuned by at least 780 nm.
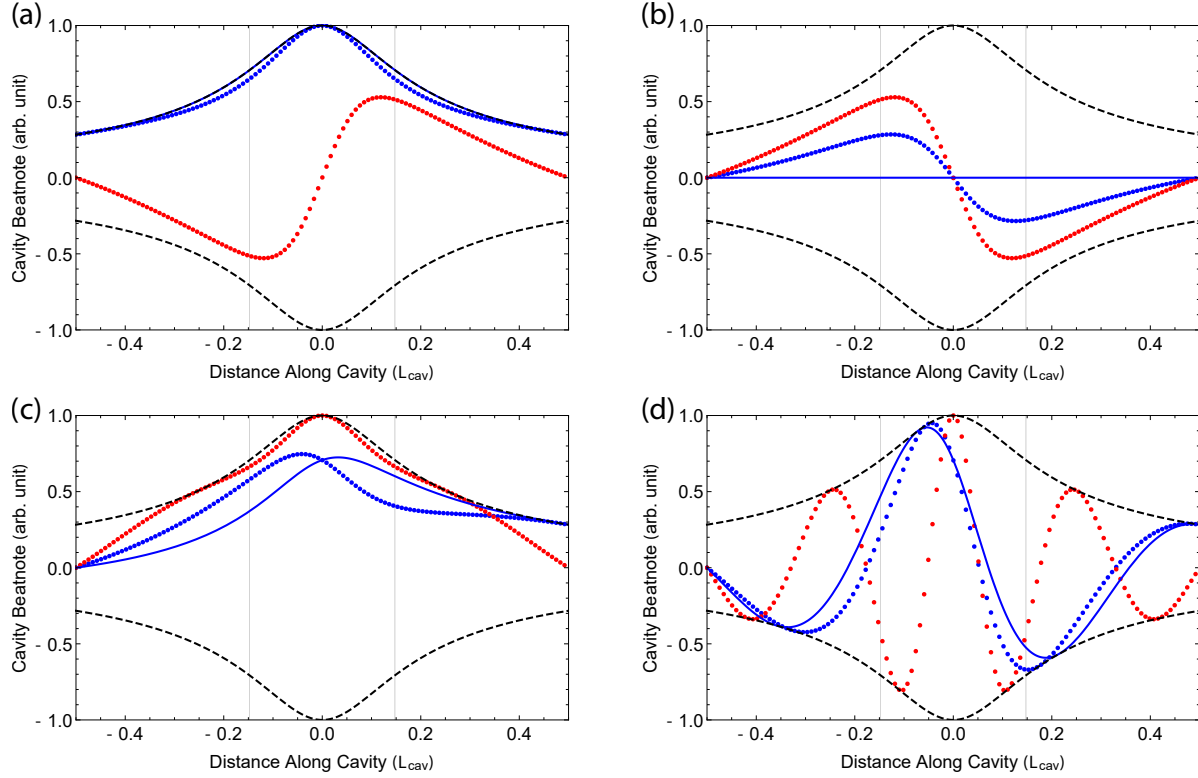
Figure 3.7: Cavity beatnotes representing the amplitude and derivative of the probe along the cavity axis at the location of the cavity ODT maxima for differing values of Δn and mirror reflection phases. Analytical calculation for beatnote neglecting Gouy phase from Eq. (3.83) (solid blue line) as well as numerical calculation including Gouy phase (blue dots). The red dots show the amplitude of the derivative of the probe intensity at the locations of the trap maxima. Maxima of the red curves are what we would refer to a 'linear spots'. The beatnote is modulated by the cavity waist profile $\frac{w_0}{w(z)}$. Vertical gray lines indicate the Rayleigh range. (a) Best case scenario $\Delta n = 0$, $\phi_{r_1}^P = \pi$, $\phi_{r_1}^T = 0$, $\phi_{r_2}^P = 0$, $\phi_{r_2}^T = \frac{\pi}{2}$. (b) Even probe mode $\Delta n = 0$, $\phi_{r_1}^P = \phi_{r_1}^T = \phi_{r_2}^P = \phi_{r_2}^T = 0$. (c) Worst case scenario $\Delta n = 0$, $\phi_{r_1}^P = \phi_{r_1}^T = \phi_{r_2}^P = \phi_{r_2}^T = 0$, $\frac{\pi}{2}$. (d) Worst case scenario with higher frequency beatnote $\Delta n = 4$, $\phi_{r_1}^P = \phi_{r_1}^T = \phi_{r_2}^P = 0$, $\phi_{r_2}^T = \frac{\pi}{2}$.

The worst case scenario is depicted in Fig. 3.7(c), in which the relative phase between left and right side of the beatnote is $\frac{\pi}{2}$. In this case, the beatnote reaches a value of $\cos\left(\frac{\pi}{4}\right) = \frac{\sqrt{2}}{2}$ in the center of the cavity and this value does not change no matter how Δn is adjusted. However, there is recourse. We can step the cavity ODT by multiple FSRs until the beatnote *does* exhibit a maxima somewhere in the vicinity of the center of the cavity, e.g. within the cavity mode Rayleigh range. This technique is shown in Fig. 3.7(d) where, for $\Delta n = 4$, a point appears within the beatnote that realizes almost maximal coupling. This strategy is possible because we can adjust the exact position that the atoms are loaded within the science cavity by adjusting the alignment of the transport ODT. The disadvantage is that the coupling has a larger spatial variation within the cavity. However, the atomic cloud will only occupy only 10s of μm along the cavity axis representing such a tiny fraction of the beatnote wavelength (proportional to $L_{\text{cav}} \approx 1\,\text{cm}$), that the atoms will all experience effectively uniform coupling.

It is worthwhile to make note of, and comparison, with the cavity beatnote utilized in E3. For details of the E3 cavity beatnote see Refs. [67, 70, 72, 73]. In E3, a similar cavity beatnote is utilized to achieve variable coupling of the atoms to probe. However, E3 does not utilize a wavelength doubled trap. E3 typically works with a ODT at approximately $850\,\text{nm}$. This choice of trapping wavelength means the beatnote has quite a large spatial frequency and varies rapidly over just a few maxima of the trap. In general it would be very difficult to place the atoms at exactly the location which has the coupling one is interested in. E3 overcomes this difficultly by utilizing the tight magnetic confinement provided by a nearby atom chip to load almost the entire atomic sample into a single well of the optical dipole trap. This confinement and control allows E3 to select maximal, minimal, or linear coupling to the probe as desired.

E6 has no atom chip to provide such tight spatial specificity. Instead, E6 relies on the long wavelength of the beatnote to ensure that atoms loaded into many neighboring wells will all experience the same coupling to the probe field. For some experiments, E6 may desire a way to realize spatially varying coupling between different subsets of atoms. It may be possible to achieve this by translating the transport ODT in-sequence to load multiple spots in the cavity with different couplings.

Chapter 4

Science Cavity

Here, I will describe, in detail, the design and construction process for the optical cavity at the heart of the E6 apparatus – the science cavity.

4.1 Design Criteria

The science cavity for E6 was designed according to three major, and several minor, design criteria. The major criteria are the following:

- High cooperativity: We saw in Sec. 1.3.4 that the cooperativity, C, is a critical figure of merit for cQED interactions. It expresses the ratio between coherent dynamics and decoherence.

- High output coupling efficiency: the output coupling efficiency, η_{out}, expresses the percentage of photons inside the cavity which leak out towards the detection setup. Each photon carries information about the atoms within the cavity. Therefore, for efficient measurement of the quantum state of the atoms within the cavity we desire to collect as many of these photons as possible

- Large transverse NA: One of the goals of E6 is to be able to locally resolve and address small spatial structures within atomic ensemble in the optical cavity. This requires a large transverse NA, for tightly focused optical fields to enter and leave the geometric volume of the cavity mode.

I'll now discuss these major criteria in greater detail.

4.1.1 High Cooperativity While Maintaining Transverse NA

In Sec. 1.3.4, the cQED cooperativity was defined as

$$C = \frac{4g^2}{\kappa\Gamma} = \frac{6}{\pi^3}\frac{\lambda^2}{w_0^2}\mathcal{F}. \tag{4.1}$$

Historically, a push was made to enhance C by increasing g by fabricating cavities with the smallest possible mode volumes motivated by Eq. 1.20 [17]. This resulted in near-planar cavities with cm-scale radii of curvature and sub-millimeter cavity lengths [15, 67, 106–109]. The finite diameter of the near-planar cavity mirrors typically dramatically limits the transverse NA. This difficulty can, in principle, be avoided by aggressively 'coning down' the mirrors to a diameter of a few tens of μm, but this is technically difficult to do while maintaining high finesse.

Alternatively, if we instead focus on the geometric interpretation of C, we see that having a small mode waist is already a sufficient condition for high cooperativity. Eq. (3.15) reveals that $w_0^2 \propto \mathcal{R}(1-\mathsf{g}^2)^{1/2}$ can be maximized by making either a near-planar cavity with $\mathsf{g} \to 1$, or a near-concentric cavity with $\mathsf{g} \to -1$. Due to the large cavity length, near-concentric cavities offer abundant transverse NA. Because of this, for this work, we choose to work with a near-concentric optical cavity.

I must mention a few cautionary points about this approach. First, for two cavity mirrors with ROCs of $\mathcal{R}$, one could make both a near-planar and near-concentric cavity with the same waist w_0 and, thus, the same cooperativity C. The difference would be that for the near-concentric cavity g_0 and κ would both be smaller with $\frac{g_0^2}{\kappa}$ staying constant between the two. The possible issue here is that one can move into a regime where $\Gamma > g_0, \kappa$. For applications in which C is truly the only figure of merit, the dominance of Γ is of no concern. However, other applications may require $g, \kappa > \Gamma$ in order that spontaneous emission can be neglected. In such cases, a near-concentric approach may be problematic. Second, if one is attempting to use the signal measured by the cavity for real-time measurement then κ will act as a low pass filter. In this case, κ must be compared to relevant dynamical time-scales within the atomic system to determine if it will be a problematic choke on the information flow rate. That said, the regime in which κ is smaller than other dynamical time-scales enables operation in the resolved sideband regime, allowing for quantum non-demolition measurements [10, 116]. Finally, it should be noted that near concentric cavities are generally much more sensitive to small geometric misalignment making the cavity assembly possibly more complicated than that of a near planar cavity. Geometric tolerances for a near-concentric cavity will be further explored in Sec. 4.4.

4.1.2 High Cooperativity While Maintaining Detection Efficiency

In Sec. 3.5, I showed that the power exiting the cavity through either mirror is related to the stored intracavity field energy by

$$P_i = \eta_i \kappa U, \tag{4.2}$$

where i indexes the cavity mirror and cavity input and output coupling efficiencies are given by $\eta_i = T_i/(T_1 + T_2 + L_1 + L_2)$). We pump our cavity from one side through mirror 1 and detect the light coming out of mirror 2; thus, we identify $\eta_{\text{in}} = \eta_1$ and $\eta_{\text{out}} = \eta_2$. In a quantum input-output formalism [117] the energy flow out of a cavity is captured by the input-output formula:

$$\hat{a}_{i,\text{out}} = \sqrt{\eta_i \kappa}\, \hat{a}_{\text{cav}} + \hat{a}_{i,\text{in}}. \tag{4.3}$$

I make the correspondence between Eqs. (4.2) and (4.3) explicit in Appendix F.

Because of the cQED interaction between the atomic sample and the optical field, the operator $\hat{a}_{\text{cav}}$ carries quantum information about the atomic system. If we want to extract maximal information about the quantum system, it is important that η_{out} does not seriously suppress this term in the output signal. Stated differently, if every photon carries quantum information, it is important that we detect as many photons as possible.

I will now address the question of how to select optical mirror coatings, i.e. setting $T_{1,2}$ and $L_{1,2}$, in light of the simultaneous requirements for large cooperativity and large detection efficiency. The first design constraint on the optical mirrors is that, unfortunately, it is not possible to make perfect lossless mirrors. This constraint means that there will necessarily be some of loss of quantum information and detection fidelity. State-of-the-art, high-reflectivity (HR) mirrors are able to achieve minimum total reflection losses at the few parts-per-million (ppm) level [118]. In Sec. 4.3, I will go into more detail about mirror losses. Here, we will simply suppose that the loss on a single mirror is limited to a minimum value L_{min} by fabrication limitations. The maximum achievable finesse and cooperativity are then given by:

$$\mathcal{F}_{\text{max}} = \frac{\pi}{L_{\text{min}}},$$
$$C_{\text{max}} = \frac{6}{\pi^3}\frac{\lambda^2}{w_0^2}\mathcal{F}_{\text{max}}. \tag{4.4}$$

A nice reference value is that $L_{\text{min}} = 10\,\text{ppm}$ corresponds to a $\mathcal{F}_{\text{max}}$ of over 3×10^5.

The question, given the non-zero value of L_{min}, is how should the values for T_1 and T_2 be chosen to set C and η_{out}? We will work under the assumption that we would like as much light as possible to go towards our detector which lies on the one of the cavity corresponding to T_2. This condition means that, to maximize $\eta_{\text{out}} = \eta_2$ we must set T_1 to be as small as possible. Such a cavity with $T_1 \ll T_2$ is called a one-sided cavity because light inside of the cavity mostly leaves through one side.

For a one sided cavity with fixed geometry and fixed mirror losses, the only parameter we have left to tune is the transmission of the second cavity mirror, $T_2 = T_{\text{out}}$. We can easily express the trade off between cooperativity and detection efficiency by noting

$$\frac{C}{C_{\text{max}}} = \frac{L_{\text{min}}}{L_{\text{min}} + T_{\text{out}}},$$
$$\eta_{\text{out}} = \frac{T_{\text{out}}}{L_{\text{min}} + T_{\text{out}}}. \tag{4.5}$$
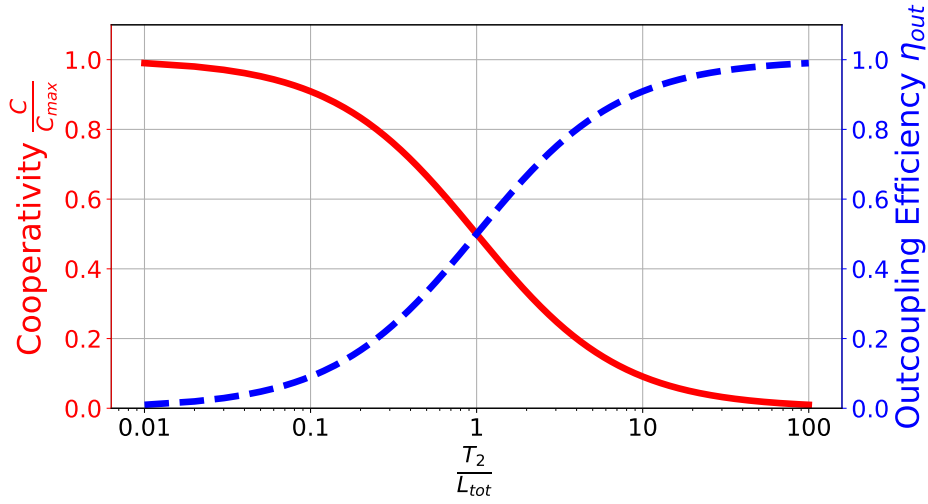
Figure 4.1: Cooperativity, C/C_{max}, and detection efficiency, η_{out} versus mirror transmission T_2 expressed as a fraction of total losses $L_{\mathrm{tot}} = L_1 + L_2$. For a critically coupled cavity with $T_2 = L_{\mathrm{tot}}$ both cooperativity and detection efficiency fall to half of their maximum value. This figure shows how the choice of mirror transmission allows us to trade cooperativity for detection efficiency.

These functional dependencies are shown in Fig. 4.1. There is a simple physical intuition behind the trade off between C and η_{out}. Cooperativity quantifies the coherence of the atom-photon interaction and is maximized if the photon remains within the cavity for as long as possible, necessitating small T_{out}. Detection efficiency requires that the cavity transmission is boosted beyond the mirror losses by having a large T_{out}.

How, then, should we choose T_{out} to balance these two figures of merit for our system? The answer depends on what science we would like to explore with the system. If we would like to explore pure coherent cQED then we should make T_{out} as small as possible to maximize C. However, if we would like to use the cavity to make a sensitive measurement of the quantum state of the atoms within the cavity, then we require a non-zero η_{out}. We explored two different thought experiments to help us come up with a figure of merit to guide our choice for mirror transmission.

4.2 Test Cases for Cooperativity vs Detection Efficiency Optimization

In this section I describe two physical test cases to demonstrate the trade-off between detection efficiency and cooperativity in cQED experiments. Consider the interaction part

of the cQED Hamiltonian when the cavity is tuned to atomic resonance:

$$\hat{H}_{\text{JC}} = g(\hat{a}^\dagger \hat{\sigma} + \hat{a}\hat{\sigma}^\dagger). \tag{4.6}$$

This gives rise to the following Heisenberg-Langevin equations of motion when cavity decay κ and spontaneous emissions Γ are included:[1]

$$\dot{\hat{a}} = -ig\hat{\sigma} - \frac{\kappa}{2}\hat{a},$$
$$\dot{\hat{\sigma}} = -ig\hat{a} - \frac{\Gamma}{2}\hat{\sigma}. \tag{4.7}$$

For large κ we can perform an adiabatic elimination of the cavity mode by formally setting $\dot{\hat{a}} = 0$:

$$\hat{a} = -i\frac{2g}{\kappa}\hat{\sigma}. \tag{4.8}$$

$$\tag{4.9}$$

This relationship between the atomic spin $\hat{\sigma}$ and the intracavity field $\hat{a}$ will be used for both test cases.

4.2.1 Continuous Readout of Atomic Excitation

For the first test case, the task is to directly monitor the light, $\hat{a}_{\text{out}}$ leaking out of the cavity and using it to extract information about the atomic spin $\hat{\sigma}$. The figure of merit will be the shot-noise-limited signal-to-noise ratio (SNR) for the photodetection of $\hat{a}$. According to Eq. (4.3), the output field is given by[2]

$$a_{\text{out}} = \sqrt{\eta_{\text{out}}\kappa}\hat{a} = -i\sqrt{\eta_{\text{out}}C}\sqrt{\Gamma}\hat{\sigma}. \tag{4.10}$$

If we detect $\hat{a}_{\text{out}}$ either directly or in a homodyne or heterodye configuration the shot-noise limited SNR will scale with the square root of the photon number, $\hat{a}_{\text{out}}^\dagger \hat{a}_{\text{out}}$. Or, in other words, the SNR will scale with $\sqrt{\eta_{\text{out}}C}$. Letting $T_{\text{in}} = 0$, $T_{\text{out}} = T$ and $L_{\text{min}} = L$, the figure of merit for this scheme is then

$$\beta_{\text{cont}} = \eta_{\text{out}}C = \frac{LT}{(T+L)^2}C_{\text{max}}. \tag{4.11}$$

This quantity is maximized when $T = L$. In this case $\eta_{\text{out}} = \frac{1}{2}$ and $C = \frac{1}{2}C_{\text{max}}$. Such a cavity where the total transmission is equal to total losses is called a critically coupled cavity. In this case detection efficiency and cooperativity are traded in equal measure to optimize the detection SNR.

---

[1]Here and below I've ignored the input fluctuation term $\hat{a}_{\text{in}}$ which arises in correspondence with the damping. Its absence won't affect any of the conclusions drawn.

[2]I've again neglected the $\hat{a}_{\text{in}}$ term.

4.2.2 Single Excitation Readout

As a second test case, we consider how to design the cavity so that it operates as the best single-photon source. We have the following setup in mind: consider that an atom is placed within the cavity in an optically excited state. After the atom/cavity system emits a photon, what is the probability that the emitted photon is received at a detector monitoring light leaking out of the output port? This probability will be the product of the probability for the photon to emit into the cavity mode (as opposed to a free-space mode) multiplied by the probability for the photon to leave the cavity towards the detector (this is exactly the detection efficiency).

Using the adiabatic elimination made above we can determine a renormalized atomic decay rate:

$$\dot{\hat{\sigma}} = -\frac{\Gamma}{2}\left(\frac{4g^2}{\kappa\Gamma}+1\right)\hat{\sigma} = \frac{\Gamma}{2}(C+1)\hat{\sigma}. \tag{4.12}$$

In this regime the spontaneous emission is enhanced by a factor of the cooperativity, C. This is exactly the physics of Purcell enhanced spontaneous emission [20, 74, 119]. One interpretation of this phenomenon is that light which is spontaneously emitted by the atom is able to reflect off of the cavity mirrors back towards the atom, thus leading to rapid stimulated emission of the atom. The decay rate into the cavity is given by $C \times \Gamma/2$ while the decay rate into free space is given, as usual, by $1 \times \Gamma/2$. For $C \gg 1$, free space emission can be neglected.

The probability for the atom to decay into the cavity is then given by:

$$\frac{C}{C+1}, \tag{4.13}$$

and the total probability for the atom to decay into the cavity and leak out towards the detector is given by

$$\beta_{\text{sp}} = \frac{C}{C+1}\eta_{\text{out}} = \left(\frac{C_{\text{max}}L}{C_{\text{max}}L+T+L}\right)\left(\frac{T}{T+L}\right). \tag{4.14}$$

I've defined the figure of merit β_{sp} which quantifies the single photon detection fidelity in this scheme.

This figure of merit, β_{sp} is maximized when

$$T = L\sqrt{C_{\text{max}}+1} \approx L\sqrt{C_{\text{max}}}. \tag{4.15}$$

We can plug this in to find that this leads to a cooperativity of

$$C = \sqrt{C_{\text{max}}+1} - 1 \approx \sqrt{C_{\text{max}}}. \tag{4.16}$$

In the large cooperativity limit, we see that this figure of merit is optimized for a much lossier cavity than in the continuous readout case.

Above I identified two figures of merit, β_{cont} and β_{sp} and a strategy to choose T to maximize each. We will keep these results in mind until Sec. 4.3.2 when I discuss our choice of optical coatings in light of actual fabrication considerations.

4.3 Fabricating Cavity Mirrors

For ideal cavity mirrors we have seen that there are essentially three important construction parameters that play a role in the resulting cQED parameters for a given cavity. These are the radius of curvature, $\mathcal{R}$, the mirror losses L, and the mirror transmission T.

Using Eq. (4.1) for the cooperativity and (3.15) for the cavity waist as a function of δ_{conc}, we can write

$$C = \frac{12}{\pi^2} \frac{1}{\sqrt{\delta_{\mathrm{conc}}}} \frac{\lambda}{\sqrt{2\mathcal{R} - \delta_{\mathrm{conc}}}} \mathcal{F} \approx \frac{12}{\pi^2} \frac{\lambda}{\sqrt{2\delta_{\mathrm{conc}}\mathcal{R}}} \mathcal{F}. \tag{4.17}$$

In this form, we see that cooperativity is increased by decreasing both $\mathcal{R}$ and δ_{conc}. We will see in Sec. 4.4 that, for a near-concentric cavity, the positioning tolerance becomes more sensitive with decreasing δ_{conc}. This motivates us to choose cavity mirrors with very small $\mathcal{R}$ because this allows us to decrease the mode waist without making major sacrifices in terms of alignment sensitivity.

The primary method for fabricating low-loss, high-reflectivity mirrors for atomic cQED experiments has been to use ground and polished cylindrical glass substrates. These substrates begin as stock rods of high quality fused silica that are ground down to an appropriate length and radius. The cylinder has two ends, one which will become the very high quality cavity mirror surface and one which will become the, typically flat, polished, back surface. The cavity mirror surface is ground down and polished to the appropriate ROC. The mirror surface is then 'super-polished' to achieve exceedingly low root-mean-square (RMS) surface roughness, σ_{RMS}, to minimize surface scattering. Polishing of glass surfaces consists of lapping the optical surface against a mold while flowing a liquid slurry with small particles to perform the polishing. Superpolishing often happens with the optic entirely immersed in the slurry and with a very specially chemically composed and calibrated slurry to ensure optimal surface quality. State of the art superpolishing can typically achieve surface roughness specifications at the $1\,\text{Å}$ level [118].

An alternative approach for building high finesse cavities for atomic cQED and other applications has been developed and is gaining popularity [120]. In this approach the end of an optical fiber is formed into a HR concave mirror by using a high power CO_2 laser to ablate a precisely shaped concavity into the fiber tip prior to the application of a HR coating. This technique has the advantages that the cavity mode is automatically fiber coupled, the mirrors are very small in size allowing for a large transverse NA (even for near-planar cavities), and the exact geometry of the mirror surface can be carefully controlled. This laser ablation technique can be used for glass substrates other than fiber tips as well.

Surface scattering dominates loss in both ground-glass and fiber cavities. The scattering losses can be calculated as a function of σ_{RMS} as [121]

$$L_{\mathrm{scatter}} = \left(\frac{4\pi\sigma_{\mathrm{RMS}}}{\lambda}\right)^2. \tag{4.18}$$

For reference, $\sigma_{\mathrm{RMS}} \approx 1\,\text{\AA}$ surface roughness with $\lambda \approx 1\,\mu\text{m}$ corresponds to scatter losses of $L_{\mathrm{scatter}} \approx 1.5\,\text{ppm}$.

The back-surface is a non critical surface as losses there do not contribute to decreases in $\mathcal{F}$ (which may be sensitive to ppm level losses), but rather only lead to a decrease in η_{out}. Reflection losses and etalons that can arise at the back-surface due to the glass-air interface are of greater concern. These losses are minimized by applying an AR coating to all back surfaces.

The next critical step in the mirror fabrication process is to apply the HR coating to the superpolished curved surfaces. The HR coating consists of a dielectric stack of dozens of alternating quarter wave ($\lambda/4$) layers of high and low index of refraction materials, often Ta_2O_5 and SiO_2 [105, 118]. The coatings are applied to the surface using an ion-beam sputtering (IBS) deposition process in which high energy ions are directed towards a target to eject target material. The substrates are geometrically within line of sight of the target so the ejected target material is deposited onto the substrate surface. Precise calibration of this process allows for precise layer thicknesses. HR coatings with transmission losses at the ppm level can routinely be achieved.

Note that the optical field penetrates some depth into the dielectric coating. This is related to the mirror reflection phase discussed in Sec. 3.1. This optical field can be absorbed by the coating leading to absorption losses. However, these absorption losses within the dielectric coating are typically at the sub-ppm level and are thus negligible.

Similarly, if the cavity mode between the mirrors is supported within an optically lossy medium, this will of course lead to additional round-trip losses. However, ultracold atomic physics experiments are performed in UHV conditions at pressures below the $10 \times 10^{-10}\,\text{torr}$ level making intracavity absorption losses non-existent. That said, it is interesting to note that absorption and subsequent spontaneous emission of photons by atoms in a cQED setting has the same effect on the cavity mode as increasing the round-trip losses.

4.3.1 Fabrication of Glass Mirror Substrates

A survey of atomic cQED apparatuses reveals that 'Kimble-style' optical cavities fabricated from super-polished glass substrates typically have a minimum mirror ROC of 1 cm to 2.5 cm [15, 21, 32, 67, 81, 106–109]. For the E6 science cavity, motivated by the discussion above, we worked with supplier PPD to have them fabricate state-of-the-art, custom, small-dimension, extra-small-ROC, superpolished, ground-glass, mirror substrates.

During this fabrication process we learned about the following technical limitations faced when trying to fabricate small ROC mirror substrates.

- The physically small size of the substrates made it difficult to handle the substrates and move them through the various steps of fabrication process. This includes difficulty during cleaning of the substrates.

- The small radius of curvature, depending on the exact diameter of the curved surface and other mirror dimensions can result in sharp features persisting at the rim of the mirror. These sharp features are especially susceptible to breakage.

- The small size of the substrates makes it difficult to hold the substrates during the super-polishing process and ensure that each substrate is subject to equal and balanced polishing pressure during the super-polishing process (multiple substrates are polished simultaneously).

- Fabrication is often limited by metrology. Surface roughness of optical substrates is often measured using an interferometric microscopy technique, or, possibly, using atomic force microscopy. Both of these imaging techniques can provide sub-nm vertical resolution of the substrate topography. However, the surface height dynamic range is limited for these measurements. For our substrates, the large curvature restricts microscope scans to small regions at the center of the substrate before the measurement is impacted by the large curvature. Scans were typically restricted to regions with sizes on the order of a few hundred μm.

Despite these difficulties, PPD has been able to repeatedly deliver substrates that meet our specifications. In early fabrication cycles, we had substrates produced with both 5 mm and 10 mm ROC substrates with the idea to use the 10 mm ROC substrates as a more reliable reference against which we could compare the more experimental 5 mm ROC substrates. We didn't notice any measurable difference in surface quality between these different batches. Note that a 5 mm ROC corresponds to a cavity length of $L_{\mathrm{cav}} \approx 1$ cm for a near-concentric cavity.

According to Eq. (4.18), a surface roughness of 1 Å would correspond to $L = 2.6$ ppm for 780 nm and $L = 0.65$ ppm for 1560 nm. During the design process, we conservatively assumed the surface roughness may be higher than 1 Å. A surface roughness of 2 Å would correspond to $L = 10$ ppm for 780 nm and $L = 2.5$ ppm for 1560 nm. We used $L = 10$ ppm as a rough estimate for losses for the 780 nm mode throughout the design process. We are less concerned about the losses for the 1560 nm mode because, as we will see in Sec. 4.3.2 below, we don't require as high of finesse for this mode.

4.3.2 Optical Coatings

We worked with FNO to produce the optical coatings for the mirror substrates. Each different optical coating involves a separate coating run in the IBS coating machine and coatings are quoted per coating run. Each fabrication round includes coating runs for the AR coatings on the backsides of the substrates as well as a number of various HR coating runs which will be described here.

As explained previously, the E6 science cavity will support an optical mode at 780 nm to serve as the cQED probe mode as well as an optical mode at 1560 nm to serve as a cavity ODT mode. This means all optical coatings must support these two wavelengths. The backside
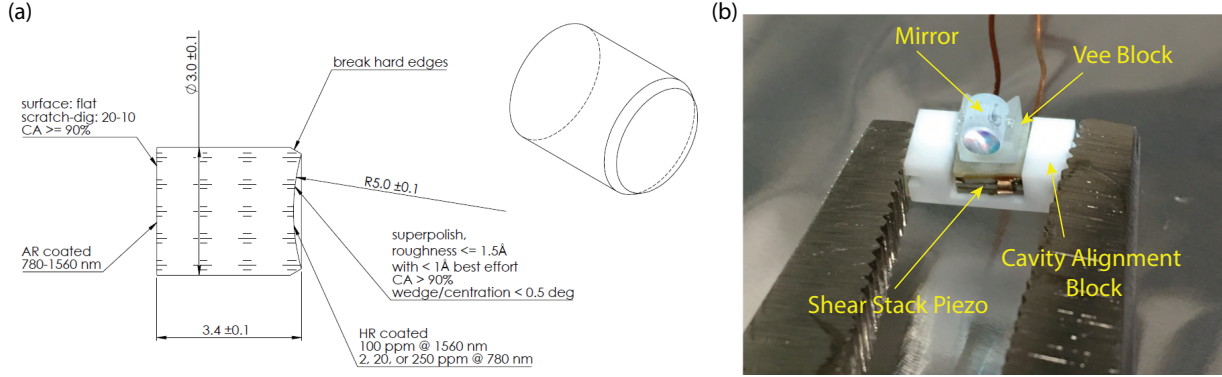
Figure 4.2: (a) Fabrication schematic for cavity mirror, dimensions in mm. 5 mm ROC, 3 mm diameter and approximately 3.4 mm overall length. Curved surface is superpolished to an RMS surface roughness $\sigma_{\mathrm{RMS}} \approx 1$ Å. There is approximately 185 μm of 'sag' of the curved surface below the 'rim' at the front of the substrate and the mirror surface is angled by about 15° to the optical axis at the edge of the mirror clear aperture (CA). (b) Photograph of cavity mirror in mirror alignment assembly. Glossy front surface visible showing dielectric HR coating. The mirror is epoxied into the glass vee block. The vee block is epoxied onto a shear stack piezo which is in turn epoxied onto the white Macor alignment block. The alignment block is held in metal vice grip for mechanical stabilization during epoxy cure.

AR coatings are dual wavelength and narrowband around 780 nm and 1560 nm. The AR coatings were measured by FNO to have reflectivity less than 100 ppm at the wavelengths of interest

State-of-the-art HR coatings can be reliably realized with transmission coefficients measured at the single ppm level. However, such coatings are most easily achieved for single wavelength and narrowband coatings. Broadband and multiwavelength coatings require more complicated patterns to be implemented in the dielectric stacks. Even for a single wavelength, the dielectric stack may be dozens of layers thick. For a complicated coating with many layers, small uncertainties in layer thickness can compound to limit coating performance.

For the E6 science cavity, we requested dual wavelength HR coatings for 780 nm and 1560 nm. Because we designed a one-sided cavity for the probe, the HR coatings for the input and output cavity mirrors were necessarily be different.

The transmission coefficients for the 1560 nm mode determine the cavity linewidth, κ, finesse, $\mathcal{F}$, and input and output efficiencies at this wavelength. The finesse and input efficiencies are important because they, through Eq. (3.72), determine the enhancement of the cavity ODT trap depth for a given input power, P_{in}. Large $\mathcal{F}$ and η_{in} are desirable because they allow us to realize a large trap depth for a small input power. If the cavity ODT light is intensity stabilized using the transmitted light then the output detection efficiency will

factor into the detection SNR and loop gain for the feedback lock circuit.

The cavity linewidth is also important for reasons having to do with stabilization of the cavity ODT. As will be described in Sec. 6.2, the science cavity length is locked to the wavelength of the cavity ODT via a PDH lock which feeds back to the science cavity piezo at low frequencies.[3] if κ_{1560} is very small, then residual frequency fluctuations of the cavity ODT laser relative to the cavity resonance will be converted, by the cavity, into amplitude fluctuations which could lead to heating of the atoms. Even more, if the cavity linewidth is small compared to the laser linewidth, it may be difficult, if not impossible, to achieve a stable or high quality laser lock.

We choose what we deemed to be moderate parameters for the 1560 nm HR coating properties. We designed to have all cavity mirrors (regardless of transmission properties for the 780 nm mode or whether the mirror is an input or output mirror) to have $T = 100$ ppm for 1560 nm. This corresponds to $\mathcal{F} \approx 25,000$ and $\kappa \approx 2\pi \times 600$ kHz and $\eta_{\text{in,out}} = 50\%$.

The choice of transmission coefficients for the 780 nm probe mode turned out to be more difficult for two reasons. First, during the design process, we had not yet finally decided for which figure of merit from Sec. 4.2, β_{cont} or β_{SP}, to optimize. Second, the optimal transmission is related to the mirror scattering losses which could not be known a-priori due to the experimental nature of the small ROC cavity mirrors we were planning. To keep our options open, we took the following approach.

The substrates were fabricated in batches of approximately 10 substrates per fabrication run. We were aware throughout the design process of the importance of having extra mirror samples in case any samples were lost due to either the mirror surface being damaged or loss of the part after an irreversible incorrect epoxying positioning step (many more details on the cavity assembly will follow in Sec. 4.5). This meant that we could divide these multiple substrates amongst a variety of different HR coating runs.

For a one-sided cavity, the input mirror should have very low transmission, T_{in}. In practice, FNO can spec transmission coefficients as small as 2 ppm. This level of transmission was low enough that to allow us to realize a one-sided cavity, but not so low that it would be impossible to couple any light into the cavity through this mirror.

Next, if we decided we wanted a cavity optimized for the continuous excitation readout described in Sec. 4.2.1, then we had to optimize β_{cont} by choosing a critically coupled cavity with $T_{\text{out}} = L_1 + L_2$. As described above, we conservatively estimated $L_1 + L_2 = 20$ ppm. Accordingly, we ordered a batch of 'critical coupling' output mirrors with $T = 20$ ppm.

Finally, if we wanted to optimize for single excitation readout, as described in Sec. 4.2.2, we had to optimize β_{sp} by choosing $T = L\sqrt{C_{\text{max}}}$. The maximum cooperativity C_{max} depends on the mirror losses and the cavity waist w_0. We did not know what size the cavity waist would be because we didn't know how near-concentric we would be able to make the cavity. Based on reviews of other near-concentric cavities and estimates of positional tolerances we targeted a 13 μm waist. Combined with total losses of 20 ppm, this corresponds to a C_{max} of

---

[3]In the future, if necessary, we can also feedback to the cavity ODT laser frequency to lock the cavity ODT wavelength to the science cavity at high frequencies.

a little more than 200. With these values, β_{sp} is maximized for T on the order of $300\,\mathrm{ppm}$. We chose to have a batch of mirrors fabricated with $T = 250\,\mathrm{ppm}$.

One of the most important science targets for E6 is the ability to perform real-time quantum feedback conditioned on the cavity photon signal. For such experiments, it will be beneficial to maximize the efficiency of detecting cavity photons. This batch of high transmission mirrors then served the dual purpose of providing a hedge for the detection efficiency against unexpectedly high losses.

This turned out to be a good strategy. Indeed, the first cavity installed into the E6 science chamber utilizes one low transmission ($T = 2\,\mathrm{ppm}$) mirror and one high transmission ($T = 250\,\mathrm{ppm}$) mirror which we used after observing unexpectedly high cavity losses during the characterization of our test cavities. More will be said about this in Sec. 4.7.2.

4.4 Near Concentric Cavity Alignment Sensitivity

In this section, I will lay out the requirements for a near-concentric cavity to be aligned, and outline the difficulties and geometric tolerances involved in meeting these requirements. What does it mean for an optical cavity to be aligned or misaligned?

The normal picture of an optical cavity is presented in Fig. 4.3(a), for example. Two cavity mirrors are aligned with their optical axes coincident such that an optical mode is supported between the two mirrors whose axis is also coincident with the optical axes of the two mirrors. In Fig. 4.3(a) only the optical axis for the cavity mode is depicted. In Fig. 3.2 the amplitude profile of the cavity mode is depicted as well. In Fig. 3.2 we see that the spatial extent of the optical mode at the mirror surface is smaller than the mirror surface ensuring that the optical energy does not 'leak out' of the cavity by passing by the outside edge of the cavity mirrors.

This final stipulation, that the optical mode does not 'fall off' of the cavity mirror will be central to the concept of cavity misalignment. Imagine the mirrors are decreased in diameter until the mirror surface is smaller than the optical mode. If we imagine the optical mode bouncing between the two mirrors as in Sec. 3.4 then we see that in addition to optical power being lost on each round trip due to mirror losses, power will also be lost due to light simply falling off of the reflective surface of the mirror. These losses are referred to as diffractive losses (because the mode diffracts out of the cavity) and contribute to round trip losses and thus reduction in the cavity finesse $\mathcal{F}$. Note that instead of decreasing the mirror diameter one could increase the cavity length closer and closer to concentricity having the effect of decreasing the mode waist at its minimum and increasing the beam waist on each cavity mirror as per Eq. (3.16). In the extreme limit that the mode is much larger than the mirror waist the finesse is ruined and it is fair to say there is no longer an optical cavity.

In addition to 'falling off' the mirror surface because the mode is too large, an optical mode can also fall off of the mirror surface due to a misalignment of the cavity mirrors. See Fig. 4.3(b). Here, the right cavity mirror is positioned as before but the left cavity mirror is slightly tilted and shifted in position. When we first began thinking about the near-
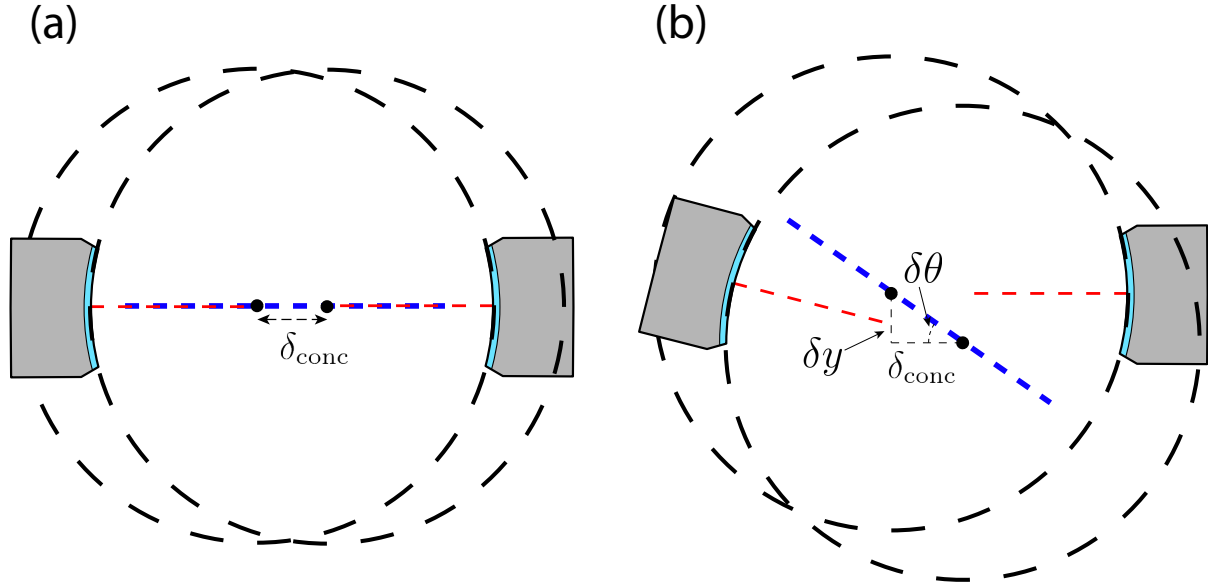
Figure 4.3: Geometric conditions for alignment of near-concentric optical cavity. Large dashed circles represent the continuation of the spherical surface for the left and right cavity mirror and the centers of these spheres are indicated by black dots. Red dashed lines indicate the optical axes of the two cavity mirrors. (a) Perfectly aligned optical cavity. The two optical axes of the mirrors are collinear resulting in an optical mode whose dominant k-vector is collinear with the optical axes of the two cavity mirrors (dashed blue line). (b) Misaligned cavity. Here we see that the optical axes of the two mirrors are misaligned. The 'mode' which would be supported by these two mirrors is indicated by the dashed blue line which passes through the centers of the two spheres. Such a mode would have the possibility to be driven if the mirrors took up a larger solid angle, but in this case the mode 'falls off' the two cavity mirrors. Distance to concentric indicated by δ_{conc}, transverse misalignment indicated by δy and angle made between cavity mode and the nominal cavity axis (horizontal) indicated by $\delta\theta$.

concentric cavity we puzzled over where the cavity mode would exist for such a misaligned cavity. For example, would it be possible to drive the original aligned cavity mode as in Fig. 4.3(a) using a cavity drive mode aligned to that original mode? Would the new mode align with the left mirror or the right mirror as depicted by the red dashed lines?

The answer to these questions came from reconsidering the discussion at the beginning of Sec. 3.2.2. Recall that for a cavity mode to be interferometrically self-sustaining and amplified in the cavity it is required that the optical mode reflects exactly back onto itself upon reflection from each cavity mirror requiring that the phase front of that mode is coincident –or fits into– the mirror surface. The only way this can occur is if the axis of the Gaussian mode passes through the center of curvature for the spherical mirror surface. In fact, this must happen for each cavity mirror meaning the optical mode must pass through the center of curvature for each mirror. Since two points in space fully determine a line, the position of the optical mode is fully determined by the spatial position of the two cavity mirrors as indicated in Fig. 4.3(b). Said more simply, we know that at the center of the beam the optical axis for the mode must be perpendicular to the mirror surface. Any line which is perpendicular to a spherical surface passes through the center of the sphere and the only way for this to hold for both surfaces is if that line passes through the center of both spheres.

The above discussion allows us to determine the spatial location of the optical mode given the positioning of the two cavity mirrors. We then use this to determine a criterion for whether the cavity is substantially misaligned or not. In fact, in Fig. 4.3(b), it is clear that the cavity is grossly misaligned because the cavity mode misses both mirrors. The criterion we used to estimate misalignment tolerances in the design for the E6 cavity is, not only that the mode axis should land on both cavity mirrors but that the $3w_0$ radius of the beam should fall on the mirrors. This condition ensures that diffractive losses are much less than $1\,\mathrm{ppm}$.

We can express this mathematically as follows. We assume the optical axes of the two mirrors are approximately horizontal. Each mirror then subtends a half-angle which is given by

$$\theta_{\mathrm{mir}} \approx \arctan\left(\frac{r_{\mathrm{mir}}}{L_{\mathrm{cav}}/2}\right), \tag{4.19}$$

where r_{mir} is the physical radius of the CA of the mirror surface. From Fig. 4.3(b) we see that the angular deflection of the optical mode is given by

$$\delta\theta = \arctan\left(\frac{\delta y}{\delta_{\mathrm{conc}}}\right), \tag{4.20}$$

where δy is the relative transverse displacement of spherical centers for the two cavity mirrors and δ_{conc} is the distance to concentric. Finally, the divergence angle for the $3w_0$ radius of the beam can be calculated from Eq. (3.32)

$$\theta_{3w_0} = \arctan\left(\frac{3\lambda}{\pi w_0}\right). \tag{4.21}$$

The condition for alignment is then

$$\delta\theta + \theta_{3w_0} < \theta_{\mathrm{mir}}. \tag{4.22}$$

We can come up with a slightly more stringent requirement for alignment which is that the $3w_0$ radius of the beam should not only fall on the front mirror surface but that it should also fall within the *back* mirror surface. While this is less critical from the perspective of the cavity mode itself, it is important for input and output coupling to ensure that the cavity mode is not aberrated by its passage through the substrate (we will discuss aberrations due to the glass substrate more in Sec. 4.6.3):

$$\theta_{\mathrm{back}} = \arctan\left(\frac{r_{\mathrm{mir}}}{\frac{L_{\mathrm{cav}}}{2} + L_{\mathrm{mir}}}\right). \tag{4.23}$$

Here L_{mir} is the length of the glass substrate. The angular tolerance is then given by

$$\delta\theta \approx \frac{\delta y}{\delta_{\mathrm{conc}}} < \theta_{\mathrm{back}} - \theta_{3w_0}. \tag{4.24}$$

For our cavity mirrors in a near concentric configuration we have that $\theta_{\mathrm{back}} \approx 170\,\mathrm{mrad}$. During our design process we imagined achieving cavity waists sizes of $w_0 \approx 15\,\mu\mathrm{m} - 20\,\mu\mathrm{m}$ which corresponds to $\theta_{3w_0} \approx 40\,\mathrm{mrad} - 50\,\mathrm{mrad}$ leading to a misalignment tolerance of $\delta\theta \lesssim 120\,\mathrm{mrad}$. With this number in hand and with some small angle approximations we can see that the positional misalignment we can tolerate is

$$\delta y \lesssim \frac{\delta_{\mathrm{conc}}}{10}. \tag{4.25}$$

So we can clearly see that as we move closer to concentric our positional tolerances become more difficult to achieve. As a rough point of comparison I will note that other alignment tasks in the lab such as fiber coupling or optical coupling of a tapered amplifier require coupling of beams to targets which are a few μm in size. This would indicate we might be able to realize distances to concentric on the order of 10's of μm using traditional alignment techniques.

The discussion above assumes purely translational misalignment of the cavity. Angular misalignment of the cavity mirrors can be analyzed similarly, however the discussion is slightly complicated because it requires a choice of rotation misalignment axis. If we assume the mirror rotates about an axis passing through the center of the mirror surface (but perpendicular to the optical axis) then a tilt of the mirror by angle $\delta\theta_{\mathrm{tilt}}$ about this axis leads to a translation of the center of the mirror sphere by $\delta y \approx \delta\theta_{\mathrm{tilt}}\frac{L_{\mathrm{cav}}}{2}$. This means the maximum tilt which can be tolerated is

$$\delta\theta_{\text{tilt}} \lesssim \frac{\delta_{\text{conc}}}{5L_{\text{cav}}}. \tag{4.26}$$

For a distance to concentric as small as $\delta_{\text{conc}} = 200\,\mu\text{m}$ this corresponds to a tilt tolerance of $\delta\theta_{\text{tilt}} \lesssim 5\,\text{mrad}$. For comparison, standard Thorlabs kinematic mirror mounts provide sensitivities of about $10\,\text{mrad}$ per rotation of one of the alignment screws and I note that it is easy to achieve alignment sensitivities at a very small fraction of a rotation of one of these micrometer screws.

In summary, depending on how close to concentric we go, we will be dealing with lateral positioning tolerances on the order of 10's of μm and angular tolerances on the order of mrad. We will keep these numbers in mind in the following sections discussing the geometric assembly and alignment of the E6 science cavity.

4.5 Science Cavity Mounting Assembly

In the previous section I described the cavity mirror design and fabrication. In this section I will describe the cavity mirror mounting hardware and in the following section I will describe how the cavity was assembled and characterized.

There were a number of constraints present when designing the optical cavity mounting assembly. First, the optical cavity would of course need to be mounted to *something* in the science chamber. We designed a three-stage Macor vibration isolation structure onto which we would mount the optical cavity. This vibration isolation system will be described in more detail in Sec. 5.2. For now, suffice it to say that the cavity sits on a flat section of the final Macor stage of the vibration isolation system, which is referred to as the science platform. See Fig. 4.4.

The optical cavity is oriented horizontally with respect to gravity in the center of the science chamber. Two re-entrant vacuum viewports (bucket windows) penetrate from the top and bottom of the Kimball physics spherical square vacuum chamber towards the optical cavity. The goal is to get these viewports as close to the optical cavity as possible so that an out-of-vacuum high-NA objective can be brought towards the atoms in the center of the cavity with as short a working distance as possible. The upper viewport is brought within a few mm of the optical cavity mirrors from above while the lower viewport comes within about a cm of the cavity from below. The lower viewport cannot come as close because it is impeded by the science platform and the rest of the cavity assembly which is used to physically hold the cavity. See Fig. 4.4. From the above discussion it is clear that the space between the re-entrant viewports in the center of the chamber is quite small so we had to be very efficient with the geometric volume for holding the mirrors.

In the next subsections I will describe how we mounted the optical cavity mirrors in a way to satisfy the geometric constraints mentioned above, to meet the positioning tolerances laid out in Sec. 4.4 and to meet additional constraints which will be introduced as they arise.
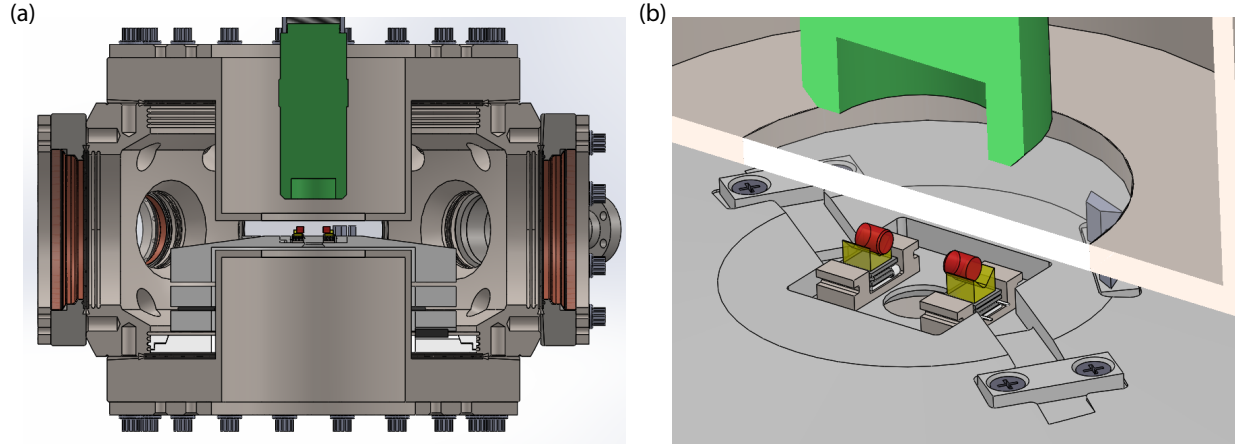
Figure 4.4: CAD visualizations of the science chamber and science platform. (a) Section view of the science chamber. Upper and lower re-entrant viewports are seen with high-NA objective (green) entering upper re-entrant viewport. The cavity mirrors (red) are mounted atop the science platform. (b) Section view zoom in on the science cavity assembly including the two mirror alignment assemblies. The circular hole in the center of the science platform allows optical access from below.

4.5.1   Vee-blocks

The mate between the glass cavity mirror substrates and the first mounting element is very important. Any strain induced within the dielectric stack mirror coatings (as may be introduced by typical optics mounts) will induce a linear birefringence which can be observed as a polarization dependent frequency splitting of the cavity modes. For very high finesse cavities this frequency splitting can result in two distinguishable peaks. While such a feature may be desirable for certain applications [122], we sought to avoid this effect. Experience from E3 taught us that an effective way to avoid birefringence in an optical cavity made from ground-glass cavity mirrors is to attach the mirrors into 'vee-blocks' using a very small amount of epoxy. The vee shape provides geometric stability, the small amount of glue ensures low strain and finally, if the vee-block is made from the same material as the cavity mirrors (fused silica in this case) strain arising from differential thermal expansion effects is suppressed.

At the suggestion of colleagues, and after some experimentation we chose to use UV-curing epoxy to attach the mirrors to the glass vee-blocks. UV-curing epoxy is epoxy which cures rapidly upon exposure to adequate fluences of ultraviolet light. The main advantage of UV-curing epoxy is that its cure time is measured in seconds, or at most up to a minute, and that it does not require heat to cure. In addition, there are UV-curing epoxies which have very low shrinkage (how much the epoxy changes size during curing), thermal expansion coefficients, and low outgassing for UHV compatibility. Finally, UV-curing epoxy is most

effective for bonding transparent components so that the ultraviolet light used for curing can penetrate through the components being bonded to ensure all epoxy is equally exposed to the UV-light. This combination of properties makes UV-curing epoxy an extremely attractive candidate for optical assemblies with very tight positioning tolerances, such as for the assembly of a near-concentric optical cavity.

We preferred UV-curing epoxies to heat curing epoxies for this assembly because we had concerns about the optics becoming misaligned due to thermal heat cycling during the long cure time for a heat cure epoxy. The long cure time would also slow down the iterative prototyping cycle for testing cavity assemblies. For cavity assembly we have chosen to use OPTOCAST 3553 UV-curing epoxy from EMIUV.[4]

The cavity mirrors were thus epoxied to the glass vee-blocks by gently placing the mirror into the vee,[5] applying the minimal amount of UV-curing epoxy possible using the tip of a small piece of high-gauge, UHV-cleaned magnet wire and then applying the UV-lamp to cure. Note that the cavity mirrors were bonded to the vee-blocks prior to cavity assembly. This means that there was no critical positioning tolerance for this bond. The design called for the cavity mirror to overhang the vee-block at the front end by about 1 mm to improve optical access to the cavity mode from below but this was not a very tight tolerance.

4.5.2 Cavity Piezos

There are a few more links in the mechanical chain connecting the cavity mirrors to the science platform. The next element in each cavity mirror assembly is a small shear-stack piezo. We required in-vacuum z-piezo positioning of the cavity mirrors for two reasons.

1. We required the ability to tune the length of the optical cavity by order $\lambda = 780$ nm to overlap the cavity resonance with the cavity probe laser which is set to a fixed detuning relative to atomic resonance. Furthermore, as discussed in Sec. 3.6, we had to be able to tuen the cavity length far enough to ensure we could overlap either an even or odd longitudinal cavity mode with the probe.

2. Though we work to vibrationally isolate the optical cavity from acoustic vibration within the lab there will always be some amount of residual vibration. High finesse cavities are especially susceptible to mechanical vibrations. The in-vacuum piezos would allow us to feedback to the length of the cavity to actively suppress acoustic vibrations of the cavity length.

---

[4]We worked with a small amount of the OPTOCAST 3553 for the science cavity assembly. However, we found NOA61 and NOA81 UV-curing epoxies (available from Thorlabs) to be very nice for other general lab purpose optics epoxying requirements such as, for example, mounting beam-splitting cubes.

[5]All small components including the mirrors themselves were typically handled using UHV-cleaned metal tweezers. It may have been more reasonable to use softer tipped tweezers such as Teflon tipped tweezers, however, in the end I saw no issues working with the metal tweezers. I was of course very careful to never let the tweezer tip touch the mirror surface.

In each cavity mirror assembly we installed a Noliac NAC2402-H2.3 shear-stack piezo to adjust the cavity length. The vee block is bonded to the top of the shear-stack piezo using MasterBond heat cure epoxy EP21TCHT-1. The piezo stacks are oriented such that the electrical leads point away from the center of the cavity towards the vacuum chamber walls where they eventually make electrical connections at electrical feedthroughs. These piezos have a 'free stroke' (maximum travel range) of $3\,\mu$m meaning that with a single piezo we are nominally able to tune the length of the cavity by a few spectral ranges for both cavity ODT and the probe mode. The travel range of the piezo is reduced when the piezo is loaded with a force or depending on details of the mounting.

Note that it is not strictly necessary for us to have a piezo for each cavity mirror, we technically only require one piezo to tune the cavity length. We put a piezo in each assembly because it (1) provides redundancy in the case of a technical failure (2) for the right driving polarities it is possible to realize a greater cavity length tuning range by driving both piezos and (3) it was easier to design two identical cavity assemblies than to have the designs differ for each mirror.

We considered including additional in-vacuum piezos or positioning stages, as in Ref. [123], to allow us to perform in-vacuum alignment of the optical cavity given the tight positioning tolerances required for a near-concentric optical cavity. However, we realized that more in-vacuum piezos wouldn't have enough throw to really help with any coarse alignment and in-vacuum translations stages would have been very bulky, increased the mechanical path length between the cavity mirrors, and introduced drift. Due to these consideration we chose to go with the monolithic design described here to minimize the overall complexity of the assembly.

4.5.3 Macor Cavity Alignment Blocks

The final pieces in the cavity mirror assemblies are the cavity alignment blocks. The purpose of the cavity alignment block is to serve as a jig which (1) supports the cavity mirror assembly including piezo, vee-block, and mirror, and (2) which is easy to grab with a macroscopic clamp to facilitate xyz, pitch and yaw positioning using macroscopic micrometer translation stages during cavity alignment. We utilized a 3D printed cavity alignment block grabber with tapped holes to grab the alignment blocks with angle tipped set screws. This grabber piece mates to a Thorlabs translation stage so that we can manipulate the pitch, yaw, and positioning of the cavity mirrors. See Fig. 4.5

The bottom of the shear-stack piezo is bonded to the center of the alignment block using the EP21TCHT-1 heat-cure epoxy. The two alignment blocks are each bonded to the science platform using the OPTOCAST 3553 UV-curing epoxy. We use the UV-curing epoxy to bond the Macor alignment block to the Macor science platform despite Macor being optically opaque because these bonds are the position critical bonds in the cavity assembly and we wanted epoxy which would cure rapidly and without heat. To avoid the epoxy 'running' under the alignment block prior to curing — which would result in uncured epoxy in the UHV chamber which could outgas — we chose a higher viscosity epoxy.
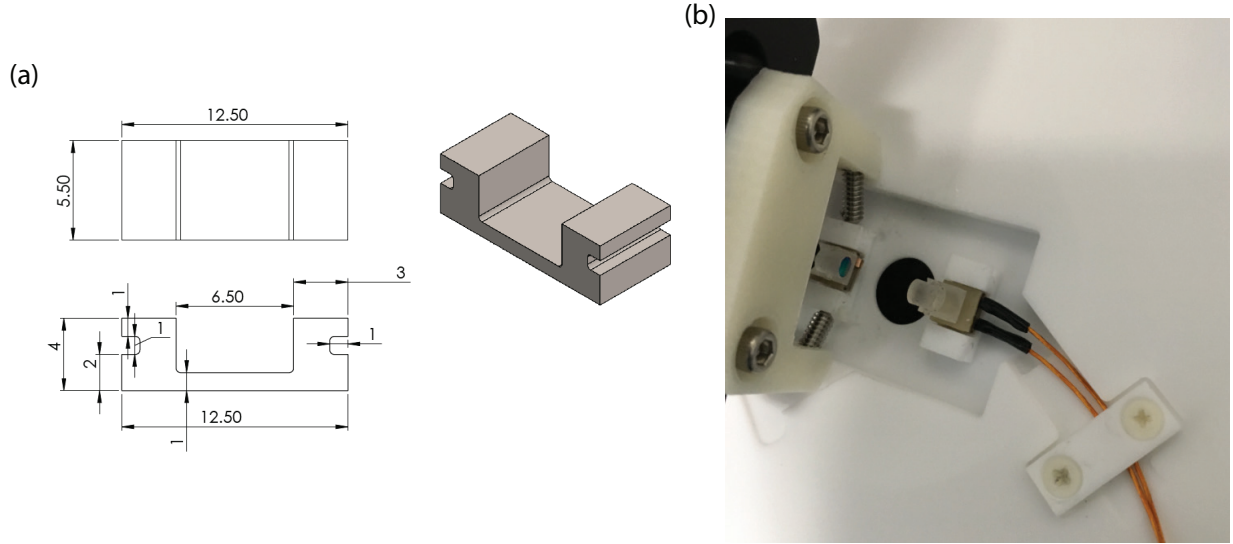
Figure 4.5: (a) CAD drawing of Macor cavity alignment block. Dimensions in mm. Central flat spot is where the shear-stack piezo is epoxied. The small grooves on the side are to make the block easier to grab with a macroscopic clamp for alignment. (b) Photograph of optical cavity near the end of assembly. The cavity mirrors, vee bocks, piezos and alignment blocks are clearly seen. The strain relief for the cavity piezo wires on the right can be seen. The cavity alignment block on the left is being grabbed by two angle tipped set screws protruding from a 3D-printed cavity grabber piece. The cavity grabber is mounted to a Thorlabs tip-tilt optics mount which is in turn mounted on a 3-axis translation stage. Here the mirror on the right is already epoxied to the science platform and the mirror on the left is being positioned in preparation for the final alignment-sensitive epoxy step to complete the cavity assembly.

In practice we made the two heat-cure bonds (alignment block to piezo and piezo to vee-block) initially. Next we made the mirror to vee-block bond. The alignment block to science platform was the last bond. This means that this was the bond that determined whether or not the cavity would be aligned appropriately. Our procedures for forming this bond are described in detail in Sec. 4.6.3.

4.6   Cavity Assembly Setup and Procedure

In this section I will describe the alignment procedure for aligning the optical cavity and the characterizations that were performed before and after installation in the UHV science chamber. The general task of the alignment of the E6 cavity mirrors was (1) to build an optical cavity which was as close to concentric as possible given the tolerance constraints outlined in Sec. 4.4 to maximize cooperativity C as per Eq. (4.1) and (2) to maintain maximal
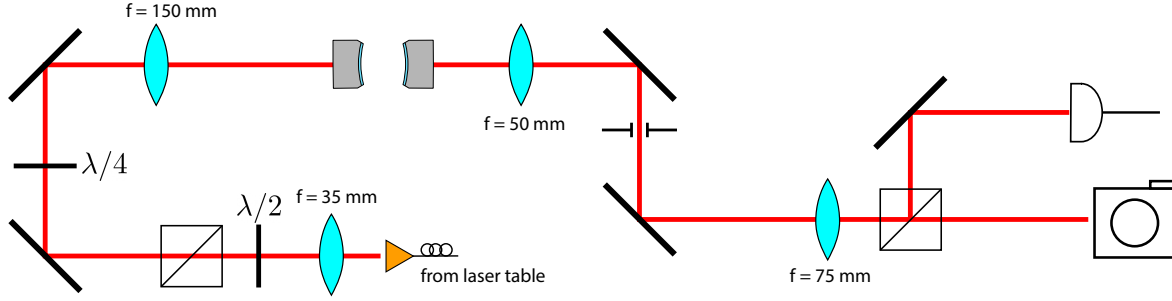
Figure 4.6: Science cavity assembly test setup. Light from an ECDL on the laser table is delivered via a fiber onto a breadboard in the flowhood assembly and mode-matched into optical cavity. The light is then analyzed using a two-lens imaging system with an APD and CMOS camera. A regular power meter can also be placed in the beampath in front of the APD.

finesse $\mathcal{F}$, again to maximize C. First I will describe the general cavity alignment procedure used to realize a stable cavity mode and next I will describe the characterizations performed to optimize the cavity geometry for maximal cooperativity.

4.6.1   The Cavity Assembly and Characterization Setup

To ensure UHV cleanliness of the cavity assembly during construction and characterization prior to installation within the science chamber we performed the cavity construction inside of a custom built flowhood. The flowhood frame was constructed out of 80-20 components which supported, about 7' above the ground, a HEPA filter salvaged from an old Budker lab. The 80-20 frame stands over a standard lab work bench on which a 3'×4' optics breadboard was placed. The flowhood walls were formed by plastic shower curtains on all four sides. The internal surfaces were wiped down with solvents to ensure cleanliness and most components and tools brought into the flowhood were likewise wiped down. While working in the flowhood, users wore gloves, clean arm coverings, a surgical face covering,[6] and a hair bouffant cap. It was, of course, discouraged for the user to bring their face into the flowhood but sometimes this was necessary to get a closer look at optics alignments. The particle count was occasionally monitored using a particle counter and we regularly realized class 100 or ISO 5 cleanroom conditions in the flowhood.

The optical layout for the science cavity test setup is shown in Fig. 4.6. $5\,\text{mW}$ to $10\,\text{mW}$ of $780\,\text{nm}$ light from the probe ECDL laser was delivered via a fiber from the laser table into the flow hood. The fiber out-coupler includes a lens on a translation stage for subsequent adjustment of the beam collimation and the ultimate position of the beam focus. The

---

[6]These face coverings, purchased in 2019, were later re-purposed for lab safety during the COVID-19 pandemic in 2020.

beam then passed through a polarization cleanup cube then hit two steering mirrors before passing through the aspheric cavity in-coupling lens with $f = 150\,\text{mm}$ and focusing down at the cavity position. This was the same cavity coupling lens which is used for the real science cavity setup on the science table. In anticipation of epoxying the cavity to the science platform, the vibration isolation stage for the science cavity was mounted below the focus of the test probe beam or 'cavity position' so that a successfully aligned cavity could be immediately epoxied down to the science platform and characterized when necessary.

We had to iterate and learn before coming up with our final optimal cavity characterization setup. Measuring the transmission coefficient for the mirrors required us to first measure the laser power directly, typically a few mW, and then measure the laser power again with a mirror in place with few ppm transmission resulting in a few nW of laser power. Such low light levels meant we were sensitive to very small amounts of stray light. We found that it was not possible to measure the mirror transmission by, for example, simply putting a power meter directly behind the cavity mirror. It would be necessary to perform spatial filtering.

In the end we used a re-imaging system to perform this spatial filtering. At the cavity output, we installed a short focal length 'objective' lens (this lens was simply a $f = 50\,\text{mm}$ singlet) to create an image of the cavity mirror and the focused input laser beam (imaged through the mirror) a few cm downstream of the cavity mirror. We then installed a second lens to relay this image onto both a complementary metal oxide semiconductor (CMOS) camera and an avalanche photodetector (APD).[7] With the camera in place we then placed an iris at the position of the first image of the cavity mirror and laser beam and irised it down so that any background light, such as that arising from diffuse scattering, was filtered out. We could then use the APD or a power meter to accurately measure the transmitted power to determine the mirror transmission value. The CMOS camera was also very important during the alignment of the cavity itself because it allowed us to observe the spatial structure of the transmitted cavity mode including the waist size of the mode, the location of the mode on each cavity mirror and it allowed us to glean some crude ideas about the surface quality of the cavity mirrors at those locations.

4.6.2 Mirror Transmission Measurement

The first step in the construction of the science cavity was characterization of the mirror transmission for any mirror which was to be tested in the setup. For this characterization a single mirror under test would be held in the position of the optical cavity. The mirror position and tilt would be adjusted so that the beam was centered on the mirror (as verified by the CMOS camera) and the reflected beam was overlapped back onto the input beam. The iris was then closed to ensure there was no diffuse background light as described above. The transmitted light level could then be measured using an optical power meter. For the lowest transmission mirrors $(2\,\text{ppm})$ it was necessary to block the probe beam and subtract

---

[7]Initially we used a regular photodetector but this did not have sufficient sensitivity to detect the pulses of light arriving as the cavity was detuned across resonance.

off a background light level. Finally the mirror could be removed by translating it and the power meter could be used to measure the input light level. The transmission could then be calculated. We typically found the mirrors to be within at least 10% of their specified transmission value. We were concerned that, due to the large mirror curvature, the transmission coefficient would vary depending on the position and angle of the beam on the mirror surface, however, fortunately, we did not find this to be the case after testing.

Each cavity mirror could be held in one of two ways. The mirror could be held directly using a pair of normally-closed tweezers or the it could be held within a mirror assembly with the cavity alignment block grabbed as demonstrated in Fig. 4.5 In either case, the xyz, pitch and yaw of each mirror could be adjusted using translation stages.

4.6.3 Aligning a Cavity

After the transmission values had been characterized for two mirrors the next task was to create an optical cavity. First the input mirror was put into the optical beampath and positioned so that the beam was centered on the mirror and the back reflection was overlapped with the input beam.

The axial position of the mirror was then adjusted following a few considerations. First, the probe beam was focused just above the central hole in the science platform. The cavity input mirror was then positioned above to one side of the hole in the science platform so that a 10 mm cavity would be centered on the hole.[8]

The task of nicely matching the axial position of the probe beam and the input and output cavity mirrors is made slightly more difficult due to what I will call the substrate lens effect. See Fig. 4.7. As described previously, the cavity mirrors are glass substrates, one side of which is curved and one side of which is flat with optical coatings on each side. The cavity mode itself has no support within the substrate itself and only exists between the cavity mirrors. However, the light which is input and output from the cavity mode does pass through the glass mirror substrates. Light which passes from the cavity mode, through the mirror surface into the substrate is unaberrated because, as necessitated by the existence of a cavity mode, the cavity mode wavefronts are perpendicular to the mirror surface. However, the flat back surface of the mirror is not matched to the diverging cavity mode at all. As the light passes through this back surface into air or vacuum the beam experiences excess divergence and spherical aberration. In a ray optic picture, this effect occurs because the substrate can be thought of as a thick lens with a focal length of $f \approx -10$ mm according to the lensmaker's equation. The effects of the substrate lens on cavity alignment are, in contrast to the case where there is no significant substrate lens effect, (1) there be a shift in the focal position of the input beam and (2) the retroreflected light from a beam which is optimally coupled to the cavity mode will appear with a divergence relative to the

---

[8]In some cases we used 10 mm ROC mirrors in which case the resultant near-concentric cavity would have been 20 mm in length.
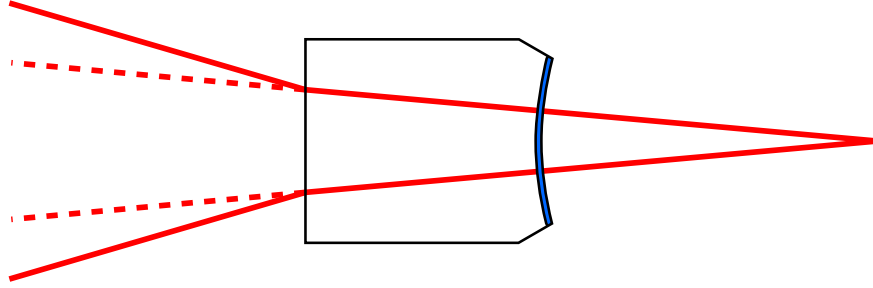
Figure 4.7: Schematic of cavity substrate lens effect. Rays emanate from the right from the center of curvature of the curved surface — approximately at the location of the cavity mode focus for a near-concentric cavity. The rays pass from vacuum into the substrate without refraction because they are already perpendicular to the surface. However, the flat back surface of the substrate is not matched to the optical wavefront and as a result the rays are deflected outwards exhibiting a larger divergence angle than they had on the other side of the substrate. In addition, this surface introduces spherical aberrations for light passing into or out of the cavity. This substrate lens effect (1) causes a relative divergence mismatch between perfectly coupled input light and the promptly retroreflected beam and (2) limits the input mode-matching efficiency achievable with a pure Gaussian input beam.

convergence of the input beam. The spherical aberration will limit the maximal input and output mode-matching efficiencies if the mode-matching is attempted with a Gaussian beam.

In practice, this was handled by positioning the mirror in the appropriate position relative to the science platform and then adjusting the position of the fiber collimation lens for the probe beam so that the retroreflected beam was roughly the same size along its beampath as the input beam. We know from the discussion in the previous paragraph that this is not the optimal condition for cavity mode matching, but it turned out to be close enough to allow us to see a first signal of a cavity mode.

One approach to mitigating the substrate lens effect is to utilize a mirror substrate whose back surface is molded into a spherical or aspherical shape so that the wavefronts of the cavity mode are matched to both the mirror surface and back surface of the substrate as was done in Ref. [124]. This process makes the substrate fabrication more complex and it would be more difficult to demonstrate very high finesse cavities with this approach.

After the input mirror was aligned, the next step involved reversibly removing this mirror from the optical path. This was done either by simply translating the mirror upwards or by removing the final stage of the translation stage (Thorlabs XR25P) and relying on the stage's 'dovetail mechanism' to position the mirror in the same spot later. The output mirror was then aligned using a second translation stage in a similar way as the first. The beam was centered on the mirror and the mirror pitch and yaw were adjusted so that the

backreflected beam was collinear with the input beam. Finally the axial position of this mirror was adjusted so that the beam size of the reflected beam matched that of the input beam. Note that for this mirror the beam does not pass through the substrate prior to retroreflection so we didn't need to worry about the substrate lens.

At this point, it was helpful to move the output coupling mirror in closer to the position of the input coupling mirror to make the resultant cavity further from concentric ensuring easier subsequent detection of and alignment of the optical cavity mode. A calipers or dimensioned spacer was used to estimate the length of the resultant cavity.

Finally the input mirror was put back into position. At this point the two cavity mirrors had been positioned such that (1) the two mirrors in fact supported a cavity mode and (2) the input beam was well coupled to low-order transverse cavity modes. It was now time to find and characterize the cavity modes by monitoring the APD and camera as the cavity-probe detuning was swept. The detuning was swept by driving shear stack piezo to modulate the cavity length by distances on the order of λ to search for cavity modes as peaks in the transmitted power as a function of piezo voltage.[9]

The very first time we formed an optical cavity with these mirrors we noticed, when looking at the image of the cavity mirrors and cavity mode on the camera, two spots on the cavity mirrors which moved around as we adjusted the mirror positions or input coupling alignment. In hindsight I believe these were the corners of high order TEM modes we were driving. We adjusted the input coupling to try to overlap these two spots and when we did so, if we were sweeping the probe detuning, we began to see low order TEM modes appear on the camera. We were also able to see cavity peaks appearing on the APD signal. Note that in early attempts it was often necessary to adjust the position of the imaging lens to ensure the output cavity mode was not aberrated in the imaging system. The necessity of this adjustment may have been related to the strong substrate lens effect.

At this time the camera and APD were monitored while adjusting the input coupling via the fiber out-coupling lens and the two input mirrors to optimize the power transmitted through the TEM_{00} mode. See Fig. 4.8 for representative images of high order modes which are observed as the cavity length is swept.

After coupling to the TEM_{00} mode is optimized, the tasks which remain are to measure the cavity finesse and transverse mode spacing and to pull the mirrors further and further apart to realize a more concentric cavity.

---

[9]For early searches, the mirrors were not mounted on piezos so we were forced to rely on modulating the ECDL current to tune the laser frequency. This proved difficult since it is difficult to tune the ECDL by a cavity FSR of 15 GHz or more without mode hops.

(a)

(b)

(c)
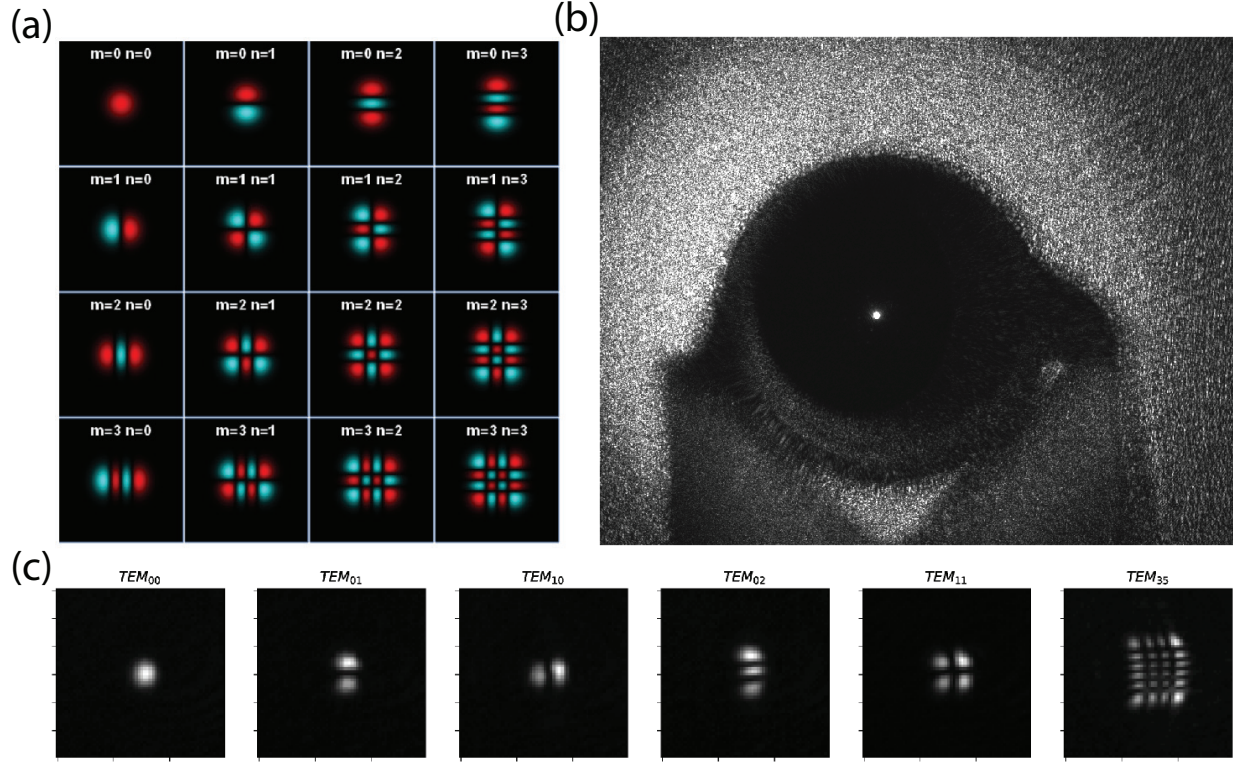
Figure 4.8: (a) Simulated amplitude profile for a number of higher order Hermite-Gaussian optical modes. Color indicates the phase of the optical field (b) Representative image of an optical cavity TEM_{00} mode seen in the cavity characterization imaging system. The circular cavity mirror can be seen as well as the glass vee-block below. The optical mode is the bright spot in the center of the mirror. (c) Images of various transverse cavity modes imaged at the cavity output.

4.7 Science Cavity Characterization

There are two important characterizations of the science cavity. The first is the measurement of the cavity finesse which is related to the cavity mirror transmission and loss coefficients and the second is the measurement of the cavity transverse mode spacing which is related to the cavity degeneracy parameter and mode waist which affects the cQED coupling parameter g.

4.7.1 Cavity Ringdown Finesse Measurement

We measure the cavity finesse by performing a fast swept cavity ring down measurement. This measurement works by observing the transmitted probe intensity in time as the detuning
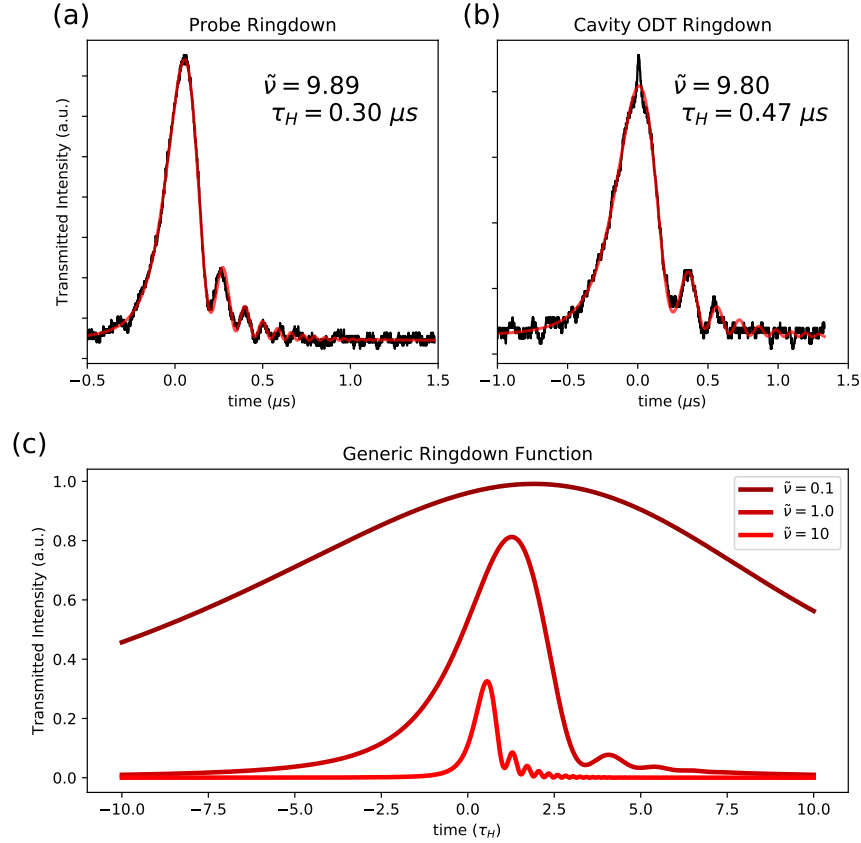
Figure 4.9: (a) Measured and fitted probe ringdown data. Assuming $L_{\mathrm{cav}} = 9.5\,\mathrm{mm}$ allows us to extract $\kappa_{780} = 2\pi \times 1.1\,\mathrm{MHz}$ and $\mathcal{F}_{780} = 14700$. (b) Measured and fitted cavity ODT ringdown data. We extract $\kappa_{1560} = 2\pi \times 0.67\,\mathrm{MHz}$ and $\mathcal{F}_{1560} = 23400$. (c) Example ringdown functions for various values of $\tilde{\nu}$. Note that as $\tilde{\nu}$ is increased the signal exhibits more pronounced ringing, the temporal width of the transmission pulse is reduced, and the peak amplitude is reduced as well.

between the probe and the cavity, Δ_{PC}, is rapidly swept across resonance.[10] If the detuning is swept such that it changes by more than κ in an amount of time which is less than $\frac{1}{\kappa}$, then the field built up within the cavity is not able to adiabatically respond to the change in the detuning. The circulating light within the cavity from the 'old' value of the detuning is still present within the cavity when light at a new detuning enters the cavity at a later time. This leads to an interference or beating effect in the transmitted optical power as a function of time as this interference evolves. See [125–127] for derivations and details.

Note that κ, and equivalently $\mathcal{F}$ could be straightforwardly measured by sweeping the detuning across resonance and extracting the width of the Lorentzian profile in frequency units. However, this would require a precise calibration of the frequency sweep. In the case of a piezo or swept laser diode current, it may be difficult to perform this calibration. One major advantage of the fast ringdown measurement, as we will see, is that no such calibration is necessary.

The lineshape for the transmitted optical field when the detuning is swept linearly across resonance at a rate $\dot{\Delta}$ is given by

$$\frac{P_{\text{trans}}(t)}{P_{\text{in}}} = \frac{2\pi}{|\tilde{\nu}|}\eta_{\text{in}}^2 \left| e^{-t' - \frac{i}{2}t'^2\tilde{\nu} + \frac{i}{2\tilde{\nu}}} \text{Erfc}\left(\frac{i - t'\tilde{\nu}}{\sqrt{2i}\sqrt{\tilde{\nu}}}\right)\right|^2, \tag{4.27}$$

with the terms defined below. In the following the cavity half-linewidth is denoted $\kappa_H = \frac{1}{2}\kappa$ and $\tau_H = \frac{1}{\kappa_H} = 2\tau = \frac{2}{\kappa}$ is the cavity field (as opposed to energy) decay time.

$$t' = \frac{t}{\tau_H},$$

$$\tilde{\nu} = \tau_H^2 \dot{\Delta} = \frac{\dot{\Delta}}{\kappa_H/\tau_H}. \tag{4.28}$$

Here t' is time scaled into units of cavity field decay time and $\tilde{\nu}$ is a measure of the detuning sweep rate, $\dot{\Delta}$, expressed in units of cavity half-linewidths κ_H per cavity field decay time, τ_H. For $\tilde{\nu} \ll 1$ the observed line shape is simply a typical Lorentzian cavity profile whose amplitude is set by η_{in}. However, for $\tilde{\nu} \gg 1$ the interference effects described above dominate giving the distinctive oscillating ringdown behavior demonstrated in Fig. 4.9(c).

Wee see that $\tilde{\nu}$ determines the oscillatory behavior of the function so for $\tilde{\nu} \gg 1$ we have that $\tilde{\nu}$ can be extracted from ringdown data independent of any other parameters. τ_H can then be independently extracted from the overall width of the transmission feature. We can then extract κ and, if f_{FSR} is known, we can extract $\mathcal{F}$.

The ringdown data and fits for both the probe and cavity ODT for the science cavity that was ultimately installed into the science chamber are shown in Fig. 4.9(a) and (b). These

---

[10]There is another kind of cavity ringdown measurement in which the cavity mode is driven with a constant intensity which is then rapidly extinguished. The transmitted power will exponentially decay with time constant κ. The interference ringdown measurement described in the main text has the advantage that it does not require a fast optical switch to extinguish the drive tone.

cavity mirrors had $T_1 = 3.5\,\mathrm{ppm}$, $T_2 = 250\,\mathrm{ppm}$ for $780\,\mathrm{nm}$ and $T_{1,2} = 100\,\mathrm{ppm}$ for $1560\,\mathrm{nm}$. For these data $f_{\mathrm{FSR}} \approx 15.8\,\mathrm{GHz}$ is estimated by assuming $L_{\mathrm{cav}} \approx 9.5\,\mathrm{mm}$ which is consistent with a crude direct measurement of the cavity length. The distance to concentric was also estimated to be about $500\,\mu\mathrm{m}$ by looking at the transverse mode spacing as will be described in the following section. In principle it would be best to directly measure f_{FSR} by placing and scanning electro-optical modulator (EOM) sidebands on the probe or cavity ODT.

The result of the ringdown measurements are that $\kappa_{780} = 2\pi \times 1.1\,\mathrm{MHz}$, $\mathcal{F}_{780} = 14700$ and $\kappa_{1560} = 2\pi \times 0.67\,\mathrm{MHz}$, $\mathcal{F}_{1560} = 23400$. This corresponds to total losses of about $170\,\mathrm{ppm}$ for $780\,\mathrm{nm}$ and about $70\,\mathrm{ppm}$ for $1560\,\mathrm{nm}$ and thus an outcoupling efficiency for $780\,\mathrm{nm}$ of about $\eta_{\mathrm{out}} = 60\%$. The input coupling efficiency for $780\,\mathrm{nm}$ is less than 1%. According Eq. (4.18) we expect $(4\times)$ larger scattering losses for $780\,\mathrm{nm}$ compared to $1560\,\mathrm{nm}$. The increase in losses for shorter wavelength light is strong evidence that the large losses we observe are due to surface roughness scattering.

Here we have relied on the out-of-vacuum measured mirror transmission coefficients to estimate the losses and detection efficiencies. In the case of perfect mode matching efficiency (and no optical path losses) it would be straightforward to measure η_{in} and η_{out} by measuring the incident, reflected, and transmitted optical power when the cavity is on-resonance and using Eq. (3.69). In the presence of non-zero mode matching losses this task becomes more complicated. For the case of a symmetric cavity it turns out that the above measurements still suffice to extract η_{in} and η_{out} if the measurement results are manipulated appropriately [105, 128]. However, for an imbalanced cavity one must repeat the measurement with the cavity flipped, i.e. with the light being sent in through the output mirror and the transmitted and reflected beams being measured appropriately.

4.7.2 Excessive Mirror Losses

The scattering losses we observe are almost $100\,\mathrm{ppm}$ for $780\,\mathrm{nm}$. This is about an order of magnitude worse than even our conservative expectation based on surface roughness characterizations from PPD. A representative surface roughness characterization is shown in Fig. 4.10(c). Unfortunately, during the cavity assembly process we were never able to reliably realize cavities that exhibited losses much lower than this level. In this section I'll give a discussion of our attempts to improve the cavity finesse and our hypotheses for why we were unable to do so.

Crunched Mirrors

During the IBS HR coating process, the mirrors were loaded into a jig which was used to hold each cavity mirror in place during the coating. Unfortunately, this jig had a geometry such that the mirror was fastened with force against the front rim around the curved superpolished surface. This rim is especially susceptible to breakage and, in fact, we observed that this breakage did occur on many cavity mirrors by inspecting them in a $(40\times)$ bright field microscope. After being loaded into these jigs, the mirrors were coated, with the
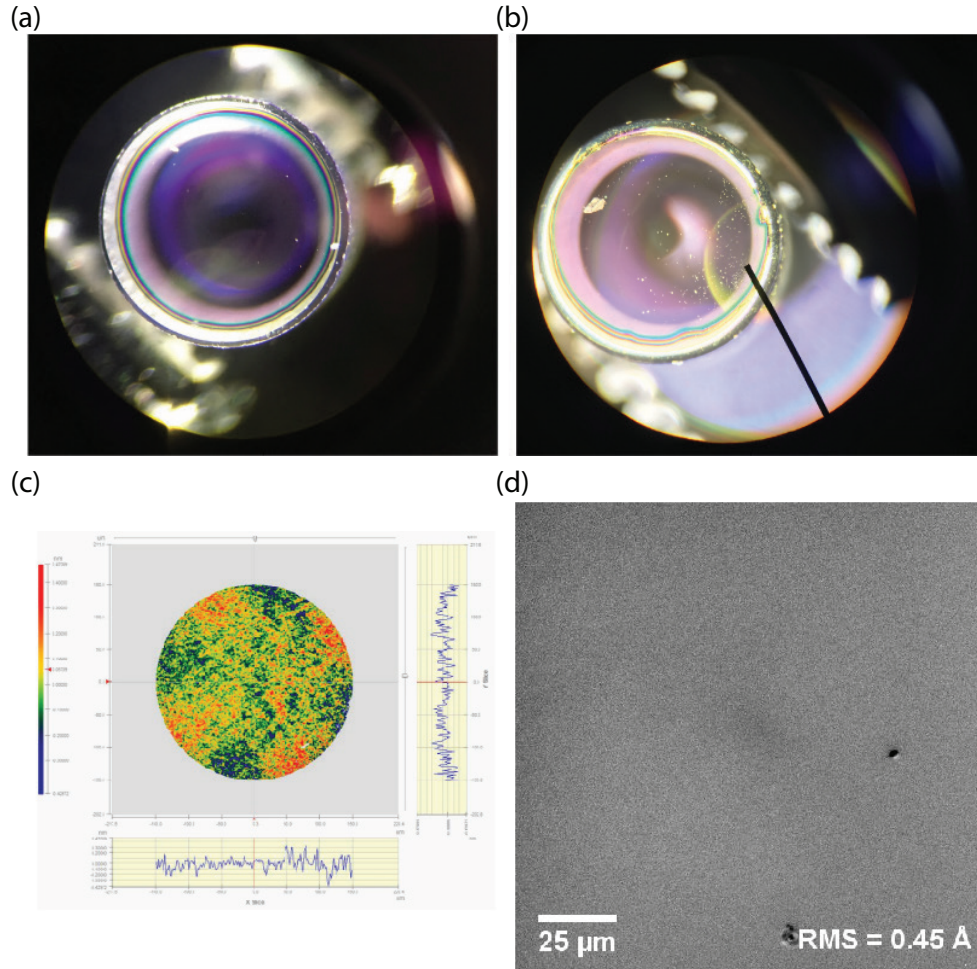
Figure 4.10: (a) (40×) magnification microscope image of 250 ppm output mirror installed into the science cavity. Note there are some visible dust features. While these features are small and difficult to see in this image they are likely at the $1\,\mu$m scale or larger which is very large compared to the $1\,\text{\AA}$ scale we are interested for ultralow losses. (b) image of a very badly damaged cavity mirror. Note that the very large glass shard at the top left of the mirror likely originated from the gash as the right side of the rim. (c) Interferometric surface roughness characterization performed by 4D Technology at the request of PPD prior to shipment of second batch of mirrors to FNO. PPD now has in-house metrology capabilities. (d) Nomarski interferometric measurement of the same mirror shown in (b) performed by FNO after we sent this mirror back to them for analysis. We were surprised to see such a clean looking surface after seeing the image in (b). We wonder if image (d) is evidence that the 'background' surface is clean and smooth with a low RMS surface roughness but that the damage created a high density of gashes or shards such as those that are seen in (d) which reduce the losses. If this is the case it would be possible, at least with a small beam, to find low loss regions of the cavity mirrors.

coating itself essentially laminating these glass shards in place so that they could never be subsequently cleared off of the surface.

The presence of large glass shards would certainly contribute to a large increase in the cavity losses. We inspected each cavity mirror we received in the microscope and tested many different cavity mirrors by using them to form cavities and measuring their finesse to estimate their losses. Confusingly, we only occasionally observed a correlation between the results of the visual inspection and the losses measured by the sensitive finesse measurement. Some mirrors that had visible large and small glass shards all across the surface sometimes had slightly lower losses and some mirrors which had very clean looking surfaces had higher losses.

We repeatedly cleaned the cavity mirrors using the procedures described below to rule out contamination on the air side of the dielectric coating which we *could* access. Assuming these cleaning procedures were effective in leaving a pristine mirror surface,[11] we have one hypothesis for the lack of correlation between on qualitative microscope inspection and the measured losses. The microscope measurement we performed did not have the spatial resolution to detect surface deviations at the few Å level which would be enough to explain the excess losses we detected. It may be possible that all of the mirrors which underwent this loading procedure exhibited some breakage of the mirror rim resulting in, at minimum, a fine dust of glass shards which is undetectable by our crude microscope inspections but nonetheless contributes to mirror scattering losses at the 100 ppm level. We had FNO perform an interferometric measurement of one of the coated mirrors to try to test this hypothesis. They generally found the surfaces to have a background RMS roughness below the 1 Å level but did see a presence of localized features such as digs, with a width of a few μm and likely similar heights. We didn't get any quantitative estimate for the spatial distribution of these features. We were not able to conclude from their data whether the features they did see on our 'clean' looking substrates were enough to explain the losses we observed via our finesse measurement.

Again, we draw no firm conclusions about the ultimate cause of the high losses. We believe that damage caused to the mirrors certainly contributed to high losses on some samples but can't be entirely certain that this is the reason for high losses on all samples.

Witness Samples

In each coating run for our science cavity mirrors we included 'witness samples'. These are larger diameter and ROC (OD 1/2", ROC 2.5 cm) stock mirror substrates from FNO which we used to serve as reliable samples for which we wouldn't have to worry about mirror size effects limiting the mirror performance. The witness samples were helpful as control test samples we could use to learn more about our processes and the small-dimension samples.

---

[11]Our cleaning procedures were shown to be effective at least on larger dimensioned witness samples. See below. It is of course still possible that our procedures had variability or issues for the small dimension samples.

We were able to make cavities using the witness samples with measured values for the finesse over of 100,000. These measurements were made with witness samples with nominal $T = 20$ ppm so we could estimate losses at the 10 ppm per mirror level. This rules out that there was some problem with our finesse measurement (such as electronic bandwidth) which limited us from measuring higher finesse than 30,000 or so.

The fact that we could measure finesses of 100,000 also gives us a baseline for the cleanliness of our flowhood environment. We had concerns that if our flowhood was not clean enough, then any mirror we brought in there would immediately become slightly 'dirty' resulting in high losses. The high finesses we measured were realized after cleaning the witness samples using the methods described below so this tells us that our cleaning procedures could be effective, at least in some cases.

One of the issues with the finesse measurement is that it only allows you to estimate the total losses L_{tot} and does not allow you to determine how those losses are split between the two mirrors in the cavity. However, if you have one mirror with confirmed low losses, such as the witness sample then, if you form a cavity using this low-loss 'reference' mirror and another 'mirror under test' and observe high losses then you know you can attribute the losses to the mirror-under-test.

10 mm ROC Substrates

Though we never saw finesse above about 26k for a cavity with two 5 mm ROC mirrors we were able to rarely see finesse as high as 60-80k for 10 mm ROC mirrors. Perhaps the 5 mm mirrors were more susceptible to breakage during the loading procedure due to the tighter radii of curvature or, perhaps there is another effect which prevented the 5 mm ROC mirrors from having lower losses. There were not significant differences between the 5 mm ROC and 10 mm ROC mirrors in the microscopy images.

Errors in the Measurement of Cavity Length

Recall that we extract the finesse, and subsequently losses, by performing a ringdown measurement to extract $\tilde{\nu}$ and t'. We then extract κ from $\tilde{\nu}$ and t', but, we can only make this conversion if we know f_{FSR} which we estimate by measuring L_{cav}. We typically measured L_{cav} with calipers. Incorrect measurements of L_{cav} would contribute to incorrect estimations of $\mathcal{F}$. However, these errors would scale linearly and are not large enough to explain the excess losses we observed.

Moving the Cavity Mode on the Mirror Surface

It is possible that the cavity modes we created were always, unluckily, landing on 'dirty' sections of the cavity mirror. Where the mode falls on the mirrors cannot be changed by changing the input coupling but it can be changed by adjusting the relative mirror positioning and tilts as explained in Sec. 4.4. We attempted adjusting the cavity alignment in this way

to see if it had any effect on the measured finesse but we saw none. Note that for our near-concentric cavity the waist on the mirror is order $100\,\mu$m.

We performed this test with different length cavities to see if the spot size affected the finesse, or if we could more easily 'get lucky' when using a smaller spot to find a small clean region of the mirror. We did tend to see that larger spot sizes correlated with slightly higher losses.

In hindsight, however, I believe that the mode positioning adjustment procedure we followed did not allow us to move the cavity mode very far along the mirror surfaces. Part of the problem was that a small motion of one of the mirrors could result in a major change to the required adjustment to the input beam positioning and pointing. Another part of the problem was that, likely due to the mirror substrate lens effect, it was also necessary to adjust the output imaging lens to accurately image the now-shifted cavity mode. At the time we did not yet understand the importance of this second effect. It is possible we may have seen a larger variation in the finesse had we adjusted the mode positioning more dramatically, adjusting both input and output optics as we did so.

Transmission Variations?

When measuring the mirror transmission we ensured that the transmission was constant for a range of input beam positions and angles so the low finesse likely cannot be explained by excessively large transmission.

Repeat Measurements

The same two mirrors did not always give the same finesse as expected. Variations from say 14k-20k might be routine for a pair of mirrors. This variation may be due to unlucky positioning of the cavity mode on the mirror surface. Occasionally we might see very large, persistent drops in the finesse. We interpreted this to be as a result of us somehow damaging the mirrors during handling.

Other Possibilities?

It is possible there is some effect of the tight mirror curvature that we are simply not accounting for which is not captured by our models for mirror losses and cavity finesse. At the edge of the mirror CA the mirror surface is angled at $\sim 15°$ to the optical axis. We've wondered if this might lead to a coating aberration that could cause issues. However, it's not clear why this would lead to losses rather than a change in transmission, and, as mentioned previously, the transmission does not vary appreciably across the mirror surface. We also only utilize a small area of the mirror near the center of the mirror.

Cleaning Procedures

There were two ways we cleaned the mirrors. The first crude way was lightly wiping the mirror surface with either a folded optics tissue or an optics cotton-tipped applicator soaked in ultra clean optical grade methanol. The second way was using First Contact cleaning solution [129].

First Contact is a polymer which you paint onto an optic similarly to painting on nail polish. Depending on the volume and thickness of the solution used the solution is left to harden in air for at least 30 minutes and as much as a few hours. The first contact is then peeled off of the optic as a single mass, pulling any dust and particulates off with it. It was a little tricky to use first contact with the very small mirrors. When using First Contact on the mirrors I would place a drop of First Contact within the mirror concavity that was up to a mm or more in thickness. Then I would place a small piece of uncoated dental floss (supplied with the First Contact) into solution, trying to be careful not to let the first contact spill over the edges of the mirror. This piece of floss is then used to pull of the first contact after hardening. It was also possible to remove the first contact by grabbing it with a tweezers but this method posed a higher risk of damaging the cavity mirrors. We also learned that, during the peel-off step, the optical surface can become charged, attracting any charged dust particles onto the surface. For this reason we installed an ionizer bar within the flow hood to eliminate any ionized particles which might contaminate the mirror surfaces.[12]

Both cleaning techniques occasionally, but not always, reduced mirror losses. Sometimes the cleaning lead to an increase in losses. This is not surprising as the cleaning procedure runs the risk of scratching or otherwise damaging the mirror surfaces.

Going Forward

Going forward, we have ordered mirrors for a third round of fabrication which did not undergo the destructive loading procedure. Measurements of the finesse of these mirrors will reveal if that procedure was the cause of the high losses. The improved procedure and our improved understanding of the cavity alignment procedure should allow us to build a new, next generation E6 cavity with a higher finesse, higher output coupling efficiency, and closer to concentric to allow us to access more interesting regimes of cQED than the first-generation cavity in the experiment now.

4.7.3 Transverse Mode Spacing Concentricity Measurement

In Sec. 3.2.3, we saw that the cavity transverse mode spacing (TMS) f_{TMS} is related to the cavity degeneracy factor $\mathbf{g}$ which is in turn related to the distance to concentric, δ_{conc}.

---

[12]We learned this tip from Dr. Osip Schwartz, then post-doc with the Mueller group. We learned a lot about mirror cleaning and characterization from Osip since he and his team were working on a high-finesse, near-concentric cavity experiment at the same time as us and had visited LIGO facilities to learn about mirror cleanliness.
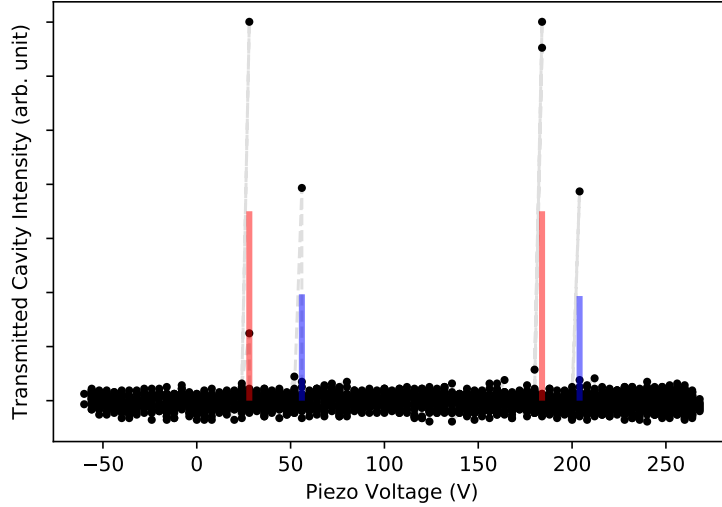
Figure 4.11: Cavity transmission versus z-piezo voltage. Two TEM_{00} modes are indicated in red and corresponding TEM_{10} modes are indicated in blue. Note that the splitting appears slightly different for the first pair compared to the second pair. This is likely due to a nonlinearity or hysteresis in the piezo response. Depending on which pair of modes is used to form the estimate of r_{TMS}, we extract $r_{\text{TMS}} = 0.179, 0.128$.

In particular we see that as the cavity is brought closer to concentric that $f_{\text{TMS}} \to 0$ which is why near-concentric cavities are described also as being near-degenerate. Additionally, δ_{conc} for given mirror ROCs $\mathcal{R}_{1,2}$ is directly related to the mode waist which is related to the cQED coupling coefficient g and the cooperativity C. For these reasons we are interested to measure f_{TMS} so that we can estimate the above parameters.

In fact we did not measure f_{TMS} directly but rather the ratio $r_{\text{TMS}} = \frac{f_{\text{TMS}}}{f_{\text{FSR}}}$ to estimate $\mathbf{g}$ according to Eq. (3.31). We performed a crude estimate of r_{TMS} by applying a sinusoidal voltage ramp to one of the cavity piezos and monitoring the piezo voltage and cavity transmission. By modulating the length of the cavity by at least $\frac{\lambda}{2}$ we observed at least two TEM_{00} modes and higher order modes in-between. The ratio r_{TMS} was then estimated from the spacing of the various peaks.

The result of this measurement are shown in Fig. 4.11. The extracted value of r_{TMS} is between 0.128 and 0.179. There appears to be a non-linearity or hysteresis in the response of the piezo which contributes an appreciable error to this measurement. Assuming the average value $r_{\text{TMS}} = 0.154$ we extract $\mathbf{g} = -0.885$. If we assume $\mathcal{R} = 5\,\text{mm}$,[13] we can extract $L_{\text{cav}} = 9.43\,\text{mm}$ and $\delta_{\text{conc}} = 570\,\mu\text{m}$. All of this leads to estimated cavity waists of $w_{0,780} = 17\,\mu\text{m}$ and $w_{0,1560} = 24\,\mu\text{m}$. Recalling the formula for the mode volume of an optical

---

[13]We have data from PPD which confirms that the deviation in the ROC from $5\,\text{mm}$ is very small.

cavity, Eq. (3.74), the formula for g_0, Eq. (1.20), and using $d_{ge} = d_0 = 2.534 \times 10^{-29}$ C m we can estimate $g_0 = 2\pi \times 3.1$ MHz.

In the future, this measurement would be improved by using an EOM sideband at a fixed cavity length to put sidebands on the probe to directly measure both f_{FSR} and f_{TMS}. Both f_{FSR} and f_{TMS} would be estimated by tuning the EOM sideband frequency until the carrier transmission peak overlaps (as measured on the APD, for example) with a TEM_{00} or TEM_{10} mode driven by the EOM sideband. f_{FSR} and f_{TMS} could then be extracted from the EOM frequency.

4.7.4   Final Cavity Alignment and Epoxy Cure

Above, I've demonstrated how a TMS measurement can be used to estimate the cavity degeneracy, waist, and g_0 for the science cavity which is installed in the setup now. This TMS measurement as performed prior to epoxying down the science cavity mirrors to determine how close to concentric was a given cavity. The goal was to make the cavity as concentric as possible without the mode becoming unstable or the finesse dropping dramatically. The procedure then was to build a cavity and then, using a motorized translation stage, increase the length of the cavity to bring it closer to concentric. In this way we were able to realize cavities that were very close to concentric, I would estimate 10s of μm or less. However, these cavities were very sensitive to mechanical misalignments. A challenge was determining how close we could go to concentric while having enough mechanical tolerance to epoxy down the mirrors and vacuum bake the system while still having a well aligned optical cavity. Our plan was to target $\delta_{\text{conc}} \approx 200\,\mu$m which would have corresponded to $\mathbf{g} \approx -0.96$, $w_{0,780} \approx 13\,\mu$m, and $g_0 \approx 2\pi \times 4.0$ MHz.

In the Sec. 4.6.3 I described the procedure for aligning an optical cavity. For the final science cavity assembly these steps were performed with two mirrors: one low transmission ($T = 2$ ppm) input mirror and one high transmission ($T = 250$ ppm) output mirror, each of which was epoxied onto a cavity mirror vee block, cavity piezo, and cavity alignment block. The cavity mirrors were positioned using the translation stages so that both cavity alignment blocks were resting on Macor science platform. We found under these conditions that it was possible to form an optical cavity. We had some concern about this being possible because we had no degrees of freedom available to adjust the height of either cavity mirror. In the end, it appeared that we were able to slightly adjust the pitch of the two cavity mirrors, even with their alignment blocks in contact with the science platform.[14]

After a cavity was successfully assembled with the cavity alignment blocks resting on the science platform and we realized that it would be possible to adjust the cavity to be near-concentric we applied UV epoxy to the two corners of the output mirror alignment block closest to the hole in the science platform. We then cured the epoxy and removed the alignment block grabber by retracting the set screws. After alignment adjustments and

---

[14]This seems to have been possible by not having the alignment block be perfectly flush with the science platform.

characterizations of the cavity, the next step was to epoxy down the input coupling mirror as well. We followed a similar procedure, again applying epoxy to the two corners, or edge, of the input mirror alignment block closest to the hole in the science platform.

After the cavity demonstrated a high enough finesse and good enough concentricity, we then cured the epoxy while monitoring the output spatial modes on the camera and the cavity mode structure on the APD. After the alignment block was epoxied down, we removed the cavity alignment block grabber to determine if the cavity would still support a good mode after being unclamped. On a number of unsuccessful initial attempts, we saw the cavity mode structure drift or become badly aberrated on the camera. We also saw the mode coupling get very bad on the APD, i.e. we saw that there was no single low order mode which was coupled, but rather dozens or more modes were partially coupled. We noticed this cavity deterioration occur both during the epoxy curing process as well as during the unclamping process.

Fortunately, we found that it was possible to pretty easily pry off the input cavity alignment block from the science platform after an unsuccessful epoxy attempt to try again. If this were not the case we would have required a new science platform and cavity mirror assemblies for each attempt. After prying off the cavity alignment block we used a tool (typically a UHV-cleaned flathead screwdriver or razor blade) to scrape off excess cured epoxy from the science platform and cavity alignment block before a new attempt was made.

After some unsuccessful attempts, we decided to move a little further away from concentric in hopes that the shift during curing and unclamping would still leave a well-aligned cavity. On the final attempt we still saw a shift after curing and unclamping but we noticed that the mode alignment could be recovered by adjusting the pitch of the input mirror by applying a slight pressure with some tweezers to the back side of the cavity alignment block which had just been epoxied down. We then applied some epoxy to the backside of the alignment block, applied the pressure to recover the alignment, and then cured the epoxy with the strain in place. After the epoxy was cured we were able to release the pressure and the alignment held in place.

After the final resultant cavity was epoxied down and characterized, it sat in the flowhood until the science chamber assembly was complete a day or so later. At that point the two lower vibration isolation stages were lowered into the science chamber and finally the science platform was lowered onto the vibration isolation stack. The science platform was lowered into the vacuum chamber by stringing wire through the three vertical alignment holes around the rim of the platform and using these to lower the platform, like a puppet, into the science chamber and onto the prior vibration isolation stage then slipping out the wires.

After the science chamber was closed, it was pumped down and baked out. We were able to monitor the cavity by looking at the cavity spectrum during pump down and we saw no issues with the cavity alignment. After the bake out, we re-characterized the cavity and again saw no deterioration in finesse or alignment. The bakeout was limited to 125 °C so as not to exceed the glass transition temperature for the UV epoxy and also to ensure no thermal stresses that might have affected the cavity alignment.

I will close this section with one final note. In each case when we epoxied down a cavity

but saw the mode structure deteriorate we believed the cavity had become dramatically misaligned. This was because on the camera we could observe the cavity modes to be extremely skewed and aberrated. However, modes were present nonetheless. At one point we measured the finesse of one of these high order modes which was apparently coupled into the cavity and we saw a finesse as high as we had seen for the TEM_{00} mode. In hindsight, my suspicion is that at these times the mirrors were still geometrically able to sustain good cavity modes. The issue was that the input mirror had slightly shifted (1) causing the input mode-matching to deteriorate, thus explaining the messy mode structure we had observed and (2) causing the cavity to become misaligned relative to the imaging system thus explaining the aberrated spatial modes on the camera. We attempted to improve the input mode matching to no avail. However, I believe that if we had improved the mode imaging, we would have been able to better understand what modes we were coupling which would have allowed us to make more dramatic input mode matching adjustments to again optimize coupling to the TEM_{00} mode. I feel that had we done this procedure, we would have been able to install a more near-concentric cavity into the science chamber, achieving a value for r_{TMS} of 0.05 - 0.10. That said, it may have been the case that the substrate lens was dramatically misaligned relative to the cavity mode making it difficult or impossible to achieve nice input mode matching or the output coupling of an unaberrated mode. I believe that a next generation E6 cavity will be able to be more near-concentric than this first science cavity.

4.8   Cavity Characterization Summary

Table 4.1 contains a summary of the measured and estimated science cavity parameters. Some of these values are rough estimates. Any of these parameters related to the length of the cavity, such as f_{FSR} or f_{TMS}, could be improved with an EOM FSR measurement as described above. T_{out} and T_{in} were only measured directly long before the cavity was assembled. T_{in}, in particular, was not measured with very high precision. It would be interesting if η_{in} and η_{out} could be measured directly and in-situ. We did not perform transmission measurements for 1560 nm light, but rather rely on FNO's measurements. FNO's measurements agreed closely with ours for the 780 nm light.

Parameter	Value (780 nm)	Value (1560 nm)
$\mathcal{R}_{1,2}$	5 mm	-
L_{cav}	9.4 mm	-
δ_{conc}	600 μm	-
f_{FSR}	15.9 GHz	-
f_{TMS}	2.4 GHz	-
g	0.89	-
w_0	17 μm	24 μm
V_{mode}	0.0021 mm^3	-
g_0	$2\pi \times 3.1$ MHz	-
T_{in}	3.5 ppm	100 ppm
T_{out}	247 ppm	100 ppm
L_{tot}	177 ppm	69 ppm
$\mathcal{F}$	14,700	23,400
κ	$2\pi \times 1.1$ MHz	$2\pi \times 670$ kHz
η_{in}	0.8 %	37 %
η_{out}	58 %	37 %
Γ	$2\pi \times 6.07$ MHz	-
C	5.8	-

Table 4.1: Measured and estimated science cavity parameters.

Chapter 5

Science Chamber

In this experiment, we utilize a two-chamber design to facilitate low vacuum pressures and geometric flexibility with unencumbered optical access in the main vacuum chamber where the cavity and high-NA system are be located – the science chamber. In this chapter, I will describe, in detail, the vacuum components and all internal components inside of the vacuum chamber other than the science cavity which was described in Chapter 4.
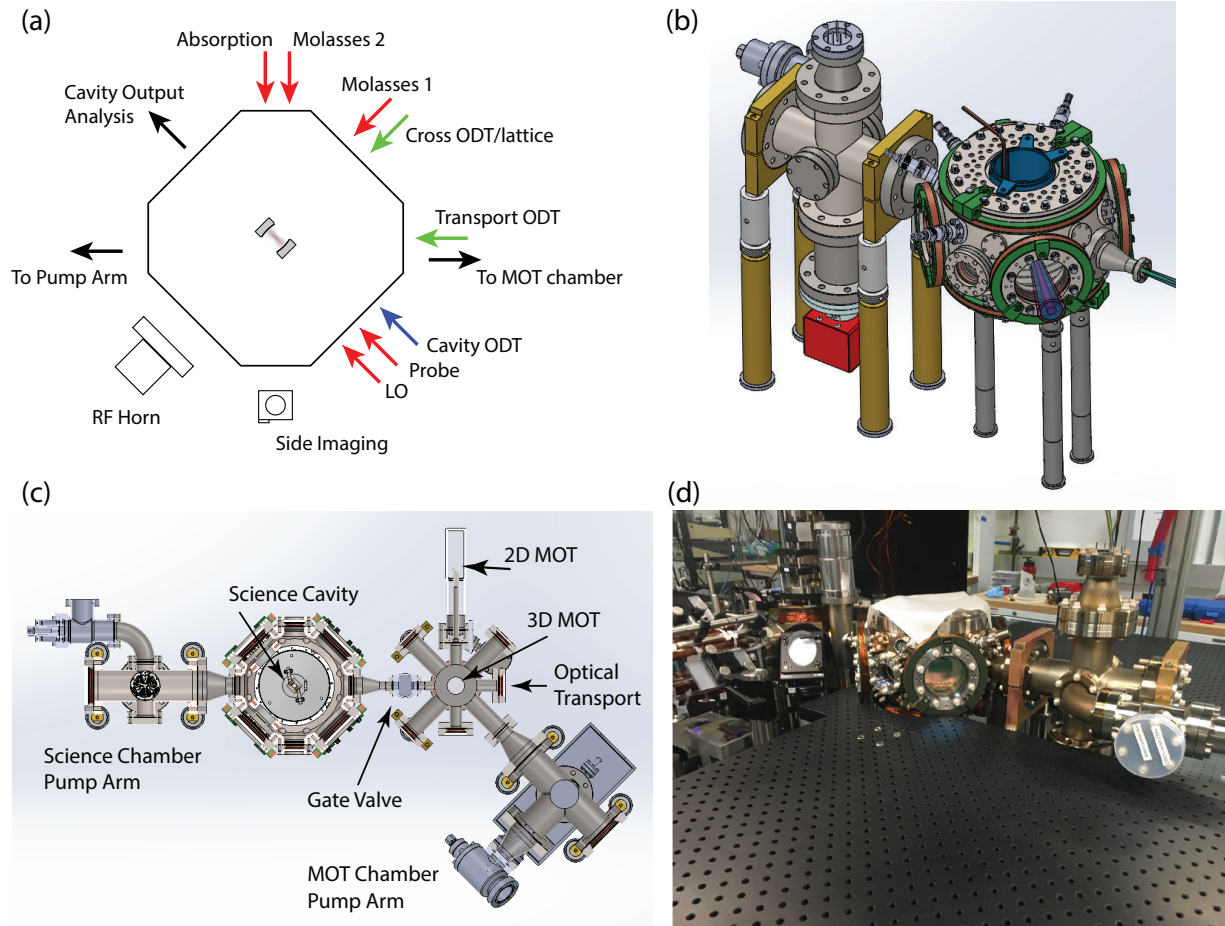
Figure 5.1: Overview of the science chamber. (a) Schematic diagram for the usage of each of the horizontal vacuum viewports on the science chamber. Note that the science cavity is at 45° to the transport axis. I will refer to the following axes throughout this text: The transport axis, the cavity axis, the vertical axis, the cross axis, and the absorption or side-imaging axis. (b) CAD image of the science chamber and science chamber pump arm. (c) CAD section view of the entire vacuum system. Atoms are collected in the 2D MOT where they are fired into the 3D MOT where they are cooled, optically trapped, and then optically transported into the science chamber. (d) Photograph of the full vacuum assembly shortly after the science chamber vacuum bakeout and before the installation of many optics.

5.1 Vacuum Components

The main science chamber is Kimball physics spherical square chamber which has two large 8" CF ports on the top and bottom, four 4.5" CF ports around the horizontal with an additional four 2.75" CF ports also around the horizontal. Additionally there are 16 small 1.33" CF ports fit between the other ports around the rims of the chamber. See Fig. 5.1(a) and (b).

The science chamber is attached to the MOT chamber through a gate valve so that the science chamber can be opened without having to replace the Rb reservoir and bake the MOT chamber. The two 8" CF ports are occupied by deep re-entrant viewports that allow us to get out-of-vacuum objectives and magnetic field coils as close as possible to the atoms in the center. The pump arm for the science chamber is a custom 6-way cross which was cannibalized from the E2 apparatus with four 4.5" CF ports and two small 2.75" ports. One of the 4.5" ports is attached to the science chamber via a conical reducer. The transport ODT in fact passes through this port and through an uncoated viewport on the opposite side. One of the other 4.5" ports has an ion gauge and the other 4.5" port hold a non-evaporable getter (NEG) pump and ion pump.

For this chamber, at the advice of our post-doc Shantanu Debnath, we chose to utilize a NEG pump rather than an ion pump. This choice was motivated by Shantanu's previous success using NEG pumps in ion trapping experiments, which, like this experiment, had a lot of possibly dirty internal components in the chamber. The NEG pump contains disks of porous carbon-like getter material. The getter material is very effective at pumping hydrogen, which is the main gas load, however, it doesn't provide any pumping for noble gases so it is necessary to have a small ion pump as well. We use a SAES NEXTorr 300 NEG/ion pump combination.

We had concerns that our vacuum pressure would be increased due to a number of worrisome materials in the chamber including the following:

- Various epoxies holding together the cavity assembly and rf coil mounts

- Graphite pencil markings on the cavity mirrors

- Large surface areas of Macor

- Polymer strain relief on cavity piezo wire connectors

- pre-baked Viton rods for vibration isolation

Despite these concerns we have achieved good vacuum pressures in our system, allowing, for example, the optical trapping of atoms in several μK deep optical traps with trapping lifetimes of 10s of s. See Fig. 5.2 for time traces from our residual gas analzyer (RGA) and temperature probes during the science chamber bakeout. Note that the bakeout began on December 20 but we were not yet recording data. The figure demonstrates repeated
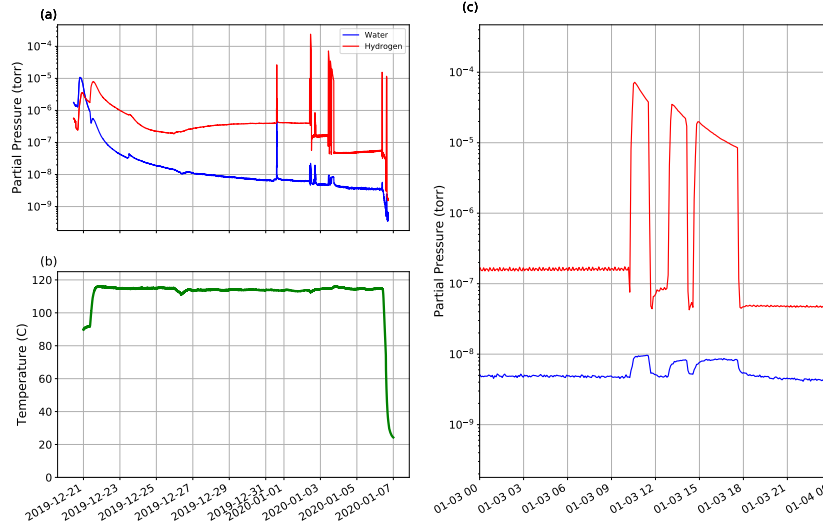
Figure 5.2: (a), (b) Pressure and temperature over the duration of the bakeout. (b) Zoom in on repeated activations of the NEG.

activations of the NEG element during which the NEG material is heated to up to over 500 C to condition the elements for subsequent pumping of hydrogen.

In Appendix C, it is explained that the ultimate vacuum pressure will be limited by outgassing Q from the various materials within the vacuum system. The purpose for baking out the chamber is to reduce the outgassing rate Q for the various in-vacuum materials [130]. A higher bakeout temperature results in a more substantial long-term reduction in outgassing rates. We had to limit our peak bakeout temperature to about 120 C due to the glass transition at 135 C for the UV-curing epoxy we used to assemble the science cavity. Such a bake out would not dramatically reduce the hydrogen gas load from the SSL chamber walls.

To address this we decided to perform an 'air bake' of the SSL vacuum fittings. There is some evidence that baking SSL vacuum components at high temperatures, at least 400 C, for example, in air (as opposed to under vacuum) can result in long-term reduction of the Hydrogen outgassing rate [131, 132]. We thus air baked all SSL vacuum fittings we could at up to 350 C, including the main chamber, before assembling the science chamber.

5.2   Vibration Isolation System

As described previously, one of the challenges for any cQED experiment is ensuring stability in the detuning between the cavity probe drive tone and the cavity resonance frequency. Such stability requires maximizing the passive absolute stability of the cavity resonance frequency by ensuring the absolute stability of the cavity length. We recall that the cavity linewidth (in cyclic units) is given by $f_{FWHM} = \frac{f_{\text{FSR}}}{\mathcal{F}}$. Following this relation, a change in the cavity length by $\frac{\lambda}{2\mathcal{F}} \approx 25\,\text{pm}$ will displace the cavity resonance frequency by one cavity linewidth. We conclude, then, that the cavity length must be stable at the pm level. In this section, I describe the vibration isolation structure we designed and installed to improve the passive stability of the cavity length.

5.2.1   Vibration Isolation and the Harmonic Oscillator Transfer Function

The simplest mathematical model for vibration isolation is as follows. Imagine one large mass that we consider as the 'base' of our system. The base may be, for example, a heavy optics table. The base is nominally motionless but it will in fact exhibit some displacement fluctuations that we are hoping to suppress. We now have a second mass, the load, which will be mounted to the base. If the load is mounted rigidly to the base then vibrations of the base will be transferred with unity gain from the base to the load. However, if the load is mounted to the base via a spring with spring constant k and viscous damping coefficient c, then position fluctuations of the base will be transferred to the load according to the following transfer function (expressed in terms of Laplace transform frequency s):

$$T(s) = \frac{x_L(s)}{x_B(s)} = 2\zeta\omega_0 \frac{s - z_0}{(s - p_-)(s - p_+)}, \tag{5.1}$$

where $\omega_0 = \sqrt{\frac{k}{m}}$ is the undamped resonance frequency, $\zeta = \frac{c}{c_c} = \frac{1}{2Q}$ where $c_c = 2\sqrt{km}$ is the critical damping value and Q is the quality factor. $z_0 = -\frac{\omega_0}{2\zeta}$ is the zero of the transfer function and $p_\pm = \omega_0\left(-\zeta \pm \sqrt{\zeta^2 - 1}\right)$ are the two poles of the transfer function.

For $Q > 1$ this transfer function exhibits a resonance at ω_0. Below ω_0 motion of the base is slow enough that the load can adiabatically follow without significant compression of the spring, resulting in a unity gain transfer function. Above ω_0 this transfer function exhibits two features. For $\omega_0 < \omega < Q\omega_0$ the transfer function behaves as if only the two poles are present, giving a rolloff in the magnitude of the gain, $|T(s)|^2$, at a rate of $-40\,\text{dB/decade}$. At the higher frequencies, $Q\omega_0 < \omega$ the zero comes into play effectively canceling out one of the poles leaving us with a rolloff of $-20\,\text{dB/decade}$. For $Q < 1$, the system behaves as a single pole lowpass filter with a cutoff frequency of $\frac{\omega_0}{Q}$, rolling off at $-20\,\text{dB/decade}$ above cutoff.

While the $-40\,\text{dB/decade}$ may sound attractive and incentivize a high Q design, in practice, the large amplification in a narrow frequency band near the low frequency resonance
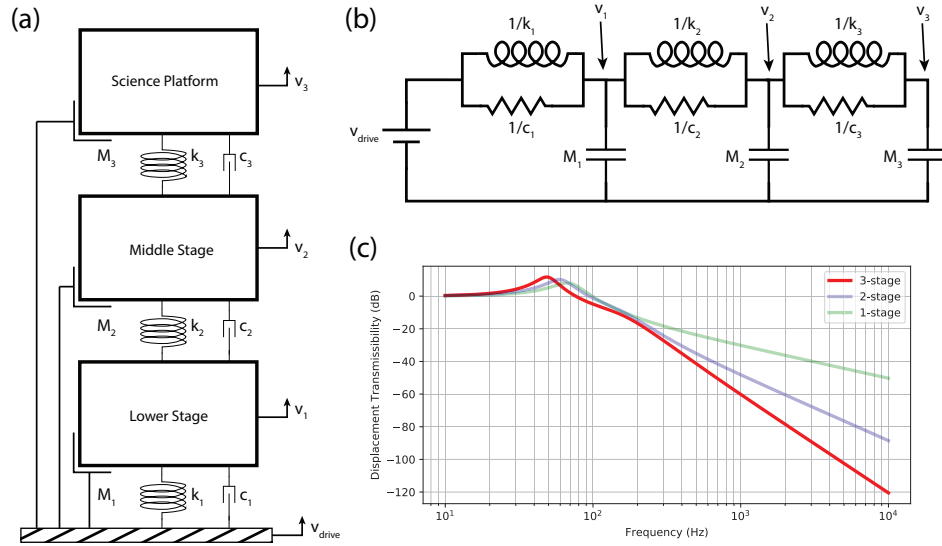
Figure 5.3: (a) Mechanical schematic of the vibration isolation system illustrating 3 masses strung together by springs and damping elements. The symbols at the left of each mass are meant to represent an inertial platform to which the position of each mass is compared. It is meant to be a visual representation of the capacitor which will subsequently appear in the circuit analog for the mechanical schematic. See Appendix D for more details. (b) Electrical circuit analog for the mechanical schematic shown in (a). The velocity of each mass is represented by the voltages at the circuit nodes at the output of each L-filter stage. The driven base is represented as constant voltage source. The output of this circuit would be the voltage or velocity appearing at the output of the third filter stage. (c) Simulated transfer function for this mechanical circuit using estimated values for the masses and spring and damping coefficients for the actual vibration isolation system.

frequency ends up being more problematic than the relatively reduced suppression at high frequencies that comes with choosing a lower Q oscillator. It seems advantageous to choose a near critically coupled oscillator with $Q \approx 1$ so that the filter cut off is not increased beyond ω_0 but so that you also do not get large amplification and instability around ω_0.

5.2.2   The E6 Vibration Isolation System

The Design

For our vibration isolation system, following the example of a number of preceding cQED experiments [32, 72, 133], we chose to have three stages of vibration isolation. In Appendix D, I describe an electrical/mechanical analogy, with which we can describe our vibration isolation system as a chain of three cascaded low pass filters [134]. The mechanical schematic,

analogous electrical schematic, and transfer function for this system are shown in Fig. 5.3.

The three masses in our vibration isolation system are made out of machined Macor. We chose Macor because we required a material that was UHV compatible, machinable, non-magnetic, non-conductive, dense, and rigid. We wanted a non-magnetic and non-conductive material because we didn't want the vibration isolation structure to become magnetized or support eddy currents that would disrupt the experiment. We preferred a dense a material so that the vibration isolation system would have low resonance frequencies. Because of the -20 dB/decade or -40 dB/decade transfer function roll offs, lower resonance frequencies lead to more noise suppression at high frequencies. We hoped for a rigid material so that the structural vibrational modes of the science platform would be as high in frequency as possible. This ensures maximal suppression of any vibration fluctuations resonant with these structural modes.

The top mass of the vibration isolation system is the science platform, described in Sec. 5.3, onto which the cavity alignment blocks, piezo strain relief, and cavity shielding mirrors are attached. The lower two masses are cylindrical annuli that wrap around the lower re-entrant view port. The three stages are not fastened together. Rather, the masses rest under gravity onto cylindrical Viton rods that fit into radial vee-slots cut into the top and/or bottom of each mass. See Fig. 5.4(a), (b). These Viton rods serve as UHV compatible spring and damping elements between the various masses.

Following Ref. [133], the clocking of the Viton rods between the different layers is offset to further suppress acoustic transfer through the structure. One way to think about this is to understand that sound waves can not travel in a straight line from the bottom to the top of the structure. Another way to understand this is to understand that the spatial eigenmodes of the vibration isolation structure are not well overlapped with a mode that uniformly drives the base.

The mating to the science chamber is accomplished via the Kimball physics groove grabber technology. The Kimball physics vacuum chamber has internal grooves machined that are concentric with the lower 8" CF vacuum port. Kimball physics supplied us with a custom SSL 'Heavy Duty Groove Grabber Assembly' that has components that form a friction fit into one of these grooves to allow us to fix hardware internally within the chamber. The central component of the assembly is a disk that, like the vibration isolation stages, has a central hole through which the lower re-entrant view port can pass and three vee-slots cut radially. The lowest stage of the vibration isolation system sits on Viton cylinders which sit in the vee-slots on this groove grabber piece.

We estimated the transfer function of the vibration isolation system. The two lower stages have a mass of about 180 g and the science platform has a mass of about 400 g. The Viton rods are approximately 4 mm in diameter and 10 mm in length. For each rod we estimated the spring constant to be $k = 0.5 \times 10^5$ N/m and damping constant to be $c = 50$ N s/m according to Ref. [135]. There are three rods between each stage so the effective spring and damping constants[1] are $k_{tot} = 1.5 \times 10^5$ N/m and $c_{tot} = 150$ N s/m.

---

[1]From Appendix D, we note that k and c are the respective impedances for the mechanical circuit
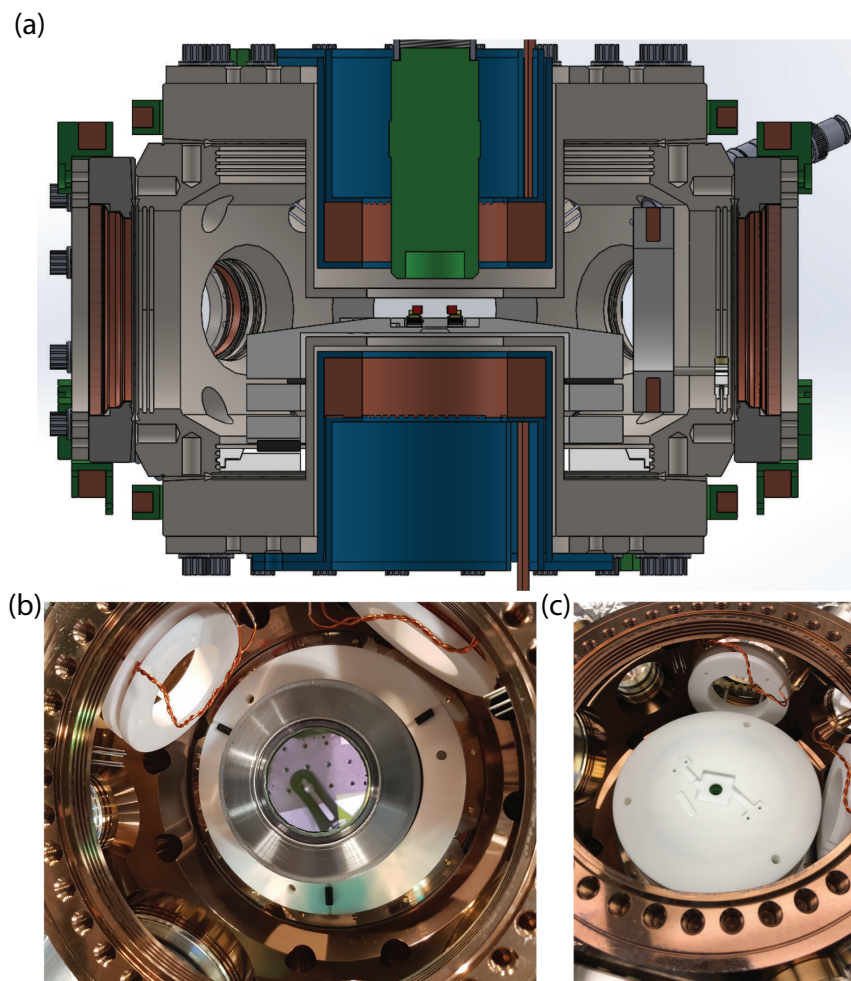
Figure 5.4: (a) Section view of CAD drawing of vibration isolation structure. Re-entrant viewports seen at center, three Macor vibration isolation stages (gray), some Viton cylinders (black), and groove grabber plate (white). Also seen are the cavity mirrors (red), the microscope objective (green), the anti-Helmholtz coils (copper), and the anti-Helmholtz coil clamps (blue). (b) Photograph of science chamber after installation of first vibration isolation stage. (3x) black Viton cylinders can be seen. (c) Photograph of the bare science platform atop the vibration isolation system during a geometry test early in the chamber assembly process.

The resultant transfer function is shown in Fig. 5.3(c). We can see from this that the dominant resonance occurs at about 50 Hz with a quality factor below 10. Above 200 Hz we see a characteristic roll off of −60 dB/decade as expected for a 3 cascaded single pole low pass filters. Of note, the transfer function exhibits 60 dB of suppression at 1 kHz and 120 dB of suppression at 10 kHz. For comparison, transfer functions are also plotted for models in which the bottom two stages are combined into a single mass (2-stage system) and in which all three stages are combined into a single mass (1-stage system). We see that, at high frequencies, the vibration isolation greatly benefits from having 3-poles.

Science Platform Structural Resonances

Are there any frequencies that we are especially interested in suppressing? Recall that this vibration isolation system is in place to ensure stability of the relative position of the two cavity mirrors. What vibrations would even cause that relative position to change? In fact, all I have discussed so far regards displacements of the entire mass of the science platform treated as a rigid body. Such rigid body motion of the science platform would not even lead to any change in the cavity length. This is one of the major advantages of the monolithic design we chose in which the cavity mirrors are mounted onto the same mass with a short mechanical path between them.

However, we know that the science platform is not a perfect rigid body, it has two zones of structural weakness that could lead to deformations under stress. Certain portions of the science platform had to be made very thin to allow for geometric clearances as described below in Sec. 5.3. One region is the circular rim close to the surface of the lower view re-entrant viewport, and the other region is the square recession where the cavity mirrors are mounted. See Fig. 5.4. The recession that holds the cavity is just over 2 mm thick and has a hole in the center of it. We had concerns that these thin pieces would contribute to a low frequency 'trampoline' or drumhead mode that would cause the cavity mirrors to deflect toward and away from each other.

A CAD vibrational normal mode analysis of the science platform was performed, using Solidworks. From this analysis, we found, indeed, that the lowest frequency deformation is exactly this drumhead mode occurring at about 2.9 kHz. This mode is shown in Fig. 5.5. Fortunately according to our estimated transfer function above, our vibration isolation system suppresses vibrations at this frequency by over 80 dB. This comparison tells us that, at least in terms of the design, we succeeded at making the science platform rigid enough so that its lowest order deformation mode is well above the resonance frequency of the vibration isolation system.

Sources of Vibration

All of the discussion so far has been quite theoretical. In the remainder of this section I'll describe some further experimental considerations and my thoughts on how this system

---

elements and that, for a mechanical system, parallel impedances add directly.
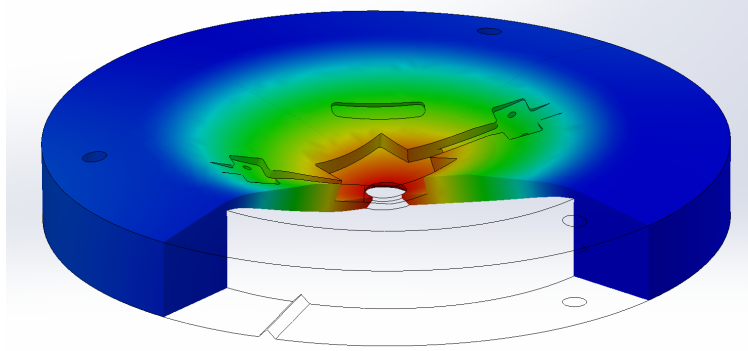
Figure 5.5: CAD simulation of the lowest frequency structural deformation of the science platform. Hotter colors indicate larger fractional strain.

could be further characterized.

One of the difficulties in designing a vibration isolation system in an optics lab is that, while we can simulate transfer functions as much as we like, it is often very difficult to specify quantitatively how much vibration we can tolerate. Further, it is even more difficult to estimate, a priori, the power spectral density (PSD) of vibrations driving the 'base' of whatever vibration isolation system we are considering. I received some advice early in the design stage: make sure to include some sort of vibration isolation system, three stages is probably good and as long as it looks like a few masses separated by springs it will probably work out.

I think this advice holds because what really matters for the vibration isolation is the resonance frequency of the system $\omega_0 = \sqrt{\frac{k}{m}}$, which only scales with the square root of k and m. This means, that, practically speaking, within the constraints of an atomic physics UHV chamber, if you put in a little effort you can get the resonance frequency below $100\,\mathrm{Hz}$. With a lot of effort, you can probably get it below $10\,\mathrm{Hz}$. But you are probably not going to get it outside of that range. Even if you were able to get a very low resonance frequency, such strong vibration isolation may already be overkill unless you have a component on board the isolation system with a very low resonance frequency that you must isolate. So, my advice to the next person designing a vibration isolation stage would be the same. Make it three stages, probably use small pieces of Viton as the springs, and try to the make the stages as heavy as you can.

The case in this experiment is also aided by the fact that there are no substantial sources of vibration near to the experiment. As long as the experimental enclosure is closed, the main sources are any building floor noise (due to people walking, air handling pumps and plumbing) and any acoustic noise (such as people talking, opening and closing drawers or handling tools). At present, we do not float the optics table on its pneumatic legs. But, in principle, this additional level of vibration reduction can be accessed if necessary.

Let me note, as a warning to future users of the E6 apparatus, that additional sources

of vibration might be added in the future to the experimental setup. In particular, we have installed magnetic field coils (described in Sec. 5.9) into the re-entrant viewports which can be used to control magnetic fields and field gradients at the location of the atoms within the cavity. Switching these coils may induce vibration through inductive forces. Further, it may be necessary to cool the coils using water cooling. The turbulent flow of that water, and also the pulsation from pumps in the water circuit, could be additional sources of vibrational noise. Note that, due to the proximity of the coils to the microscope objective (described in Chapter 7), these vibrations from the coils may also cause alignment issues for the objective as well.

Science Cavity Lock Performance

At present the cavity system is quite stable providing evidence that this vibration isolation system works well. We are able to lock the length of the science cavity to the wavelength of the cavity ODT beam using a PDH lock feeding back to a cavity piezo without too much difficulty. However, we do observe that this lock must be filtered to have a low bandwidth at the 100s of Hz or kHz level. It is possible that the instability that limits the lock bandwidth is the resonance that occurs in the vibration isolation system at around $100\,\text{Hz}$.[2] One way to test this would be to engage this lock and then stimulate the vibration isolation system using a voice coil across a range of frequencies and perform a spectral analysis of the PDH error signal. This would allow us to identify any resonances that may be causing issues.

If there is a low frequency resonance that is in fact limiting our feedback bandwidth we can use a digital finite impulse response (FIR) filter to tailor the filter response function to compensate for the presence of this resonance [136]. Hopefully at that point the appreciable suppression of high frequency signals by the vibration isolation system would then allow the feedback bandwidth to approach at least the mechanical deformation frequency of the science platform or the piezo bandwidth.

5.3 The Science Platform

The science platform is the top stage of the vibration isolation system and the Macor base to which the science cavity is epoxied as described in Chapter 4. See Figs. 5.4, 5.5, and 5.6. At the center of the science platform is a square recession with a hole in the center. The hole allows optical access to the cavity volume from below for imaging, for beams to enter from below, or for beams to pass from above. The two cavity mirrors on their respective cavity alignment blocks are epoxied into this recession. The purpose for this recession is to minimize the thickness of the science platform to minimize the distance between the lower re-entrant viewport and the science cavity. This minimum distance is set by the sum of the

---

[2]I'll point out here that this is not our only hypothesis. There are also issues with noise in the high voltage piezo driver which we heavily filter. This electronic filter combined with the PI lockbox settings may lead to the instability we observe.

thickness of the cavity mirror radius, the vee-block thickness, the shear stack piezo thickness, the Macor alignment block thickness and the science platform thickness.

There are also two angled features on either side of the cavity recession. These are piezo wire routing paths. There are two tapped holes along each of these routing paths. The piezo wires are routed along these paths and then a small Macor plate with two through holes is pressed down onto the wire and mated to the science platform using two small PEEK bolts thus locking the piezo wires in place providing strain relief. At one corner of the cavity recession is another cutout. This cutout marks the position where the transport ODT shielding mirrors are placed.

From a side view, the science platform slopes down from the center. The sloping provides additional clearance for sharply focusing optical beams, such as the probe or cavity and transport ODTs, to get to the center of the vacuum chamber without clipping on the platform. The NA accessible from the side of the chamber to the center of the cavity, limited by the science platform and upper reentrant viewport, is about 0.1.

Finally, there are three holes cut vertically in the edge of the science platform. These holes are in fact present for all three vibration isolation stages. They were used during machining to help realize the appropriate clocking of the Viton vee slots and to position routing features on the science platform. By threading wires through these holes we were able to lower in the science platform as the final step in the science chamber assembly. We needed to use this approach since there was not enough space to lower the platform in by hand, especially with the sensitive science cavity epoxied in place.

5.4 Transport ODT Shielding Mirrors

We deliver atoms into the optical cavity using our transport ODT, which is focused to a waist size of $\approx 50 \, \mu$m and has a power of up to 10 W. When the optical focus of this high power beam occurs within the science chamber, such as at the end stage of transport, the beam passes through the open space between the cavity mirrors without any clipping. Thus, even while 10 W of light is directed toward the cavity, none of it is absorbed so there is no local heating of any solid elements within the cavity setup. However, when the transport ODT is focused at intermediate locations between the starting (MOT) and final (cavity) stages of transport, the light diverging from the optical focus will enter the cavity region with a much larger beam diameter. Under such conditions, some of the 10 W of IR light used for transport will, indeed, strike the cavity mounts, and, potentially, heat them. We were concerned that this transient optical heating of the cavity setup would destabilize the cavity, or, even worse, would damage some of the components of the cavity setup.

For this reason, and inspired by Ref. [133], we installed shielding mirrors to block this light from hitting the science cavity. These mirrors clip away the portions of the transport ODT light that could otherwise strike the cavity mirrors. After hitting the shielding mirrors, the transport ODT light is deflected out of the vacuum chamber where the power can be safely dumped. See Fig. 5.6 for photographs of the shielding mirrors and more details about
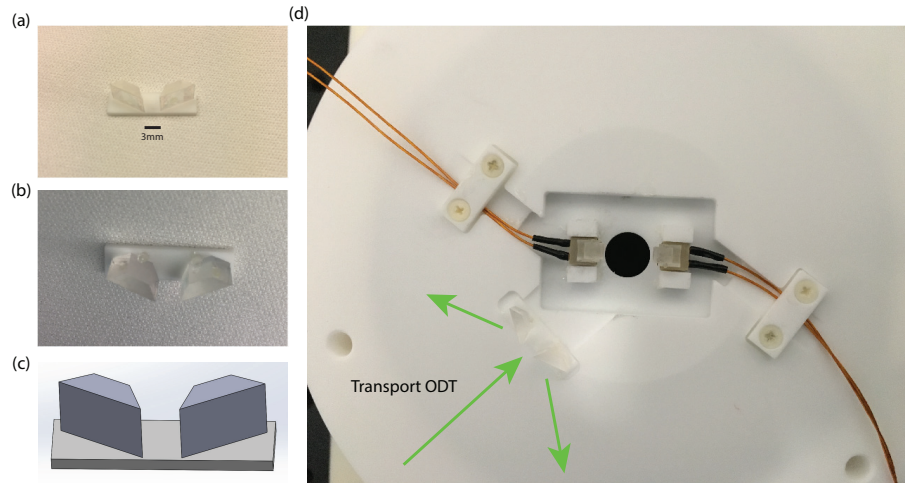
Figure 5.6: Photographs of the cavity shielding mirrors. Each mirror is constructed from a piece of a small glass prism, one side of which has been coated with a reflection coating. The shielding mirrors were epoxied using UV-curing epoxy onto a Macor shielding mirror plate. the angles were set by eye with reference to a stencil. (a) Photograph of the shielding mirrors after being epoxied. (b) Photograph of backside of shielding mirrors with drops of UV-curing epoxy visible. (c) CAD image of shielding mirrors on shielding mirror platform. (d) Image of shielding mirrors and cavity mirror assemblies on completed science platform. The shielding mirror plate is epoxied into a rectangular depression on the top of the science platform.

their construction. We do observe the cavity lineshape to move by a few linewidths over the course of an experimental sequence. This is likely due to something being heated by the cavity ODT but the motion is small and slow enough that the science cavity PDH lock (described in Sec. 6.2) has no problem staying in lock the during the sequence.

It is very important that E6 users keep careful track of transport ODT light deflected out of the chamber by the shielding mirrors. Presently, this light is safely dumped. However, the dumping paths are awkward amidst all of the other optics for the probe, cavity ODT and other systems. It is very easy, for example, to accidentally burn an IR card or piece of paper with this deflected light.

5.5   In-Vacuum rf coils

We learned from the experience of previous experiments in the group that it can be difficult to deliver rf fields generated outside of the chamber to the atoms because of geometric considerations and perhaps the magnetic properties of various vacuum components. For this reason we installed rf coils inside of the science chamber. We installed two rf coils at 90° so

that we could drive two polarizations of rf fields at the location of the atoms. See Fig. 5.4(b) for a photograph including the rf coils. We refer to the two coils as the cavity axis rf coil and the cross axis rf coil.

Each coil is wrapped around a Macor spool so that the coils have a diameter of about 2". The coils are 10 turns of 20 AWG Kapton insulated copper magnet wire. The coils are concentric with two of the 4.5" CF ports on the Kimball spherical square chamber and the center of the coils are offset by about 80 mm from the center of the chamber. The spools have small holes into which alumina rods are inserted and epoxied. The alumina rods are then clamped in place by groove grabber components which hold the rf coils in place and clamp into groove grabber slots in the 4.5" CF port.

The rods were made out of alumina so that, like the vibration isolation platform and the rf coil molds, they would be non magnetically or electrically active. Unfortunately, alumina is very brittle and the rods must sustain a moderate torque to hold the molds in place in the science chamber. On at least one occasion the alumina rods broke during assembly and had to be replaced. In hindsight, they would present such a small mass of conductive material it would have been fine had they been made of metal, especially given their proximity to the SSL chamber walls, so for any future upgrades we ordered non-magnetic SSL replacements.
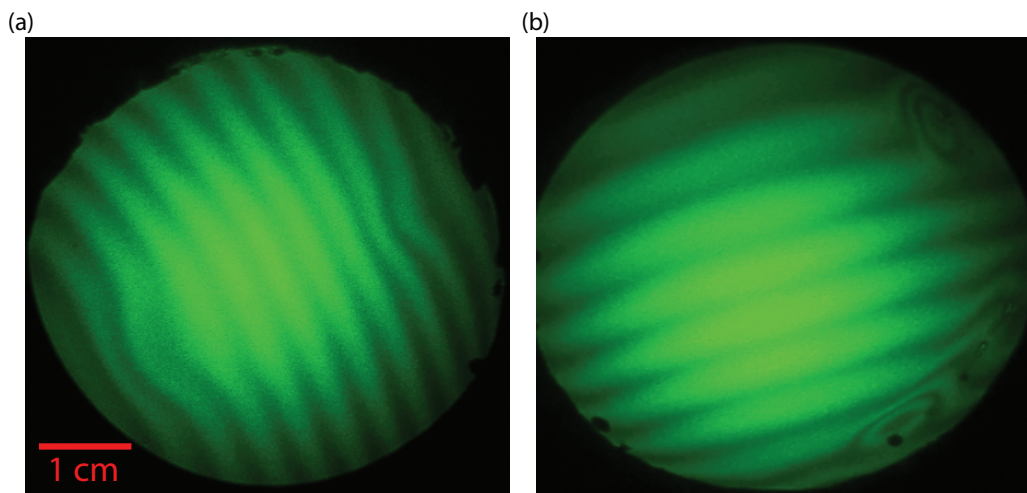
Figure 5.7: Interferometric viewport flatness measurement. (a) Upper viewport. (b) Lower viewport. Black patches and flatness deviations are seen around the rim of this viewport. This was due to some solder splatter which occurred when the viewport was being brazed into the vacuum fitting.

5.6 Viewports

All of the horizontal viewports have a broadband AR coating from below 700 nm to above 1600 nm. The one exception is the viewport on the pump arm, which we left uncoated to ensure the AR coating wouldn't block any important unforeseen wavelengths we may install in the future. For example, in the future we may want UV light to drive 2-photon transitions to Rydberg states.

Along the vertical axis we carefully considered the AR coatings for the re-entrant viewports. The two surfaces of the lower viewport both have broadband AR coatings from about 750 nm up past 1064 nm. The upper surface of upper viewport has the same AR coating as the lower viewport, but the lower (vacuum) surface of the upper viewport has an AR coating from about 770 nm up to about 900 nm but has a HR coating ($> 99\%$) at 1064 nm. This HR coating could be used to retroreflect 1064 nm light off of the upper viewport ot create a vertical lattice to strongly confine the atoms along the vertical direction.

Postdoc Johannes Zeiher performed an interferometric characterization of the surface flatness of the two re-entrant viewports. This measurement was performed because deviations in the surface flatness at even the sub-wavelength level could cause difficulties for very high-NA diffraction limited imaging. The results are shown in Fig. 5.7. The measurement was carried out by shining collimated green light directly through the viewport and monitoring the interference between the directly transmitted beam and the beam which transmits through the first surface, reflects once off of both the second and first surface and then transmits through the optic. Neighboring bright fringes correspond to thickness variations of $\lambda/2$. We see that both viewports exhibit some slight wedge across their central area but are otherwise generally quite flat. The lower viewport shows some larger surface roughness variations in the vicinity of solder which was accidentally sputtered onto the viewport when it was being brazed into the vacuum fitting.

5.7 Feedthroughs

We have four feedthroughs on the vacuum chamber, one feedthrough for each of the two piezos and one feedthrough for each of the two rf coils. We use the 1.33" ports on the Kimball sphere for the feedthroughs. All of the unused 1.33" ports are blanked off.

We have two types of feedthroughs: three of the feedthroughs are 5-pin feedthroughs and one feedthrough is a barrel BNC connector feedthrough. The BNC feedthroughs turned out to be difficult to work with because it was difficult to solder to their very thick shield conductor. We had to have the machine shop weld a small wire onto the shield conductor to which we could solder the rf coil. For the center conductor, we used a small adapter connector that makes a press fit to the feedthrough and makes a screw clamp connection with the rf coil to connect the rf coil to the feedthrough. We originally planned to use a BNC feedthrough for each rf coil, but, after one of these feedthroughs broke, we replaced it with an easier-to-work-with 5-pin feedthrough.
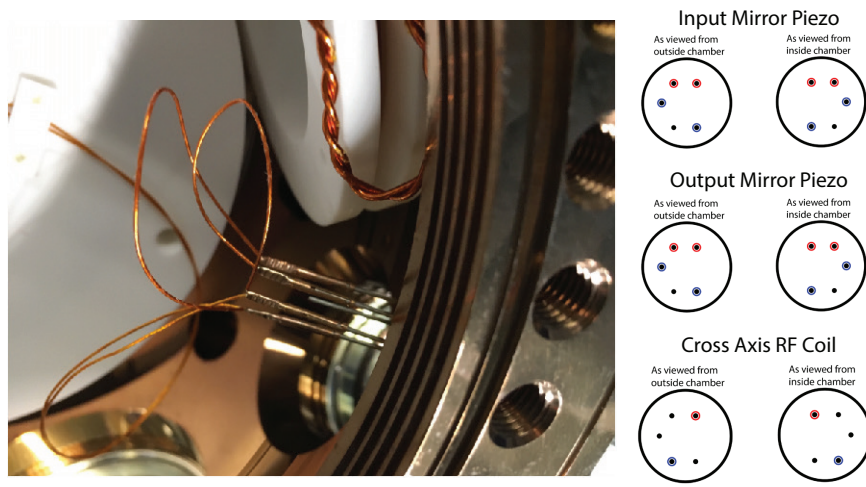
Figure 5.8: Science chamber feedthroughs. Photograph of the connection for one piezo 5-pin feedthrough. See (2x) wires coming from a cavity piezo and their jumper cables for feedthrough redundancy. Pinout diagrams for each 5-pin feedthrough. Similarly colored pins are shorted together by jumpers.

We made connections to the 5-pin feedthroughs using crimp connectors that crimped onto our wires and then made a press fit onto the pins of the 5-pin feedthroughs. We used one of these feedthroughs for each cavity mirror piezo and one for one rf coil. At each crimp connector for each of the four piezo wires we fit in a second short jumper cable that has its own crimp connector. We then attach this second jumper cable to another pin of the 5-pin feedthrough. This was done for a small bit of extra redundancy since we already had the extra pins at our disposal.

See Fig. 5.8 for pin diagrams for each 5-pin connector. Note that the piezos can be driven with positive or negative voltage. It has not yet been determined which voltage polarity corresponds to which direction of cavity mirror translation.

The feedthroughs make the assembly of the science chamber internal components difficult because they are fragile and do not have very much clearance from the vibration isolation stages during assembly or the upper re-entrant viewport. We had to take care that the wires coming from the feedthroughs did not impede any of the vacuum viewports and even removed the upper re-entrant viewport after our first attempt at installing it because we had concerns that a particular wire was blocking a viewport.

5.8   Bolt Washers

The 2.75", 4.5", and 8" viewports on the science chamber all have what we call bolt washers. When putting vacuum compression bolts into SSL UHV CF flanges we use a
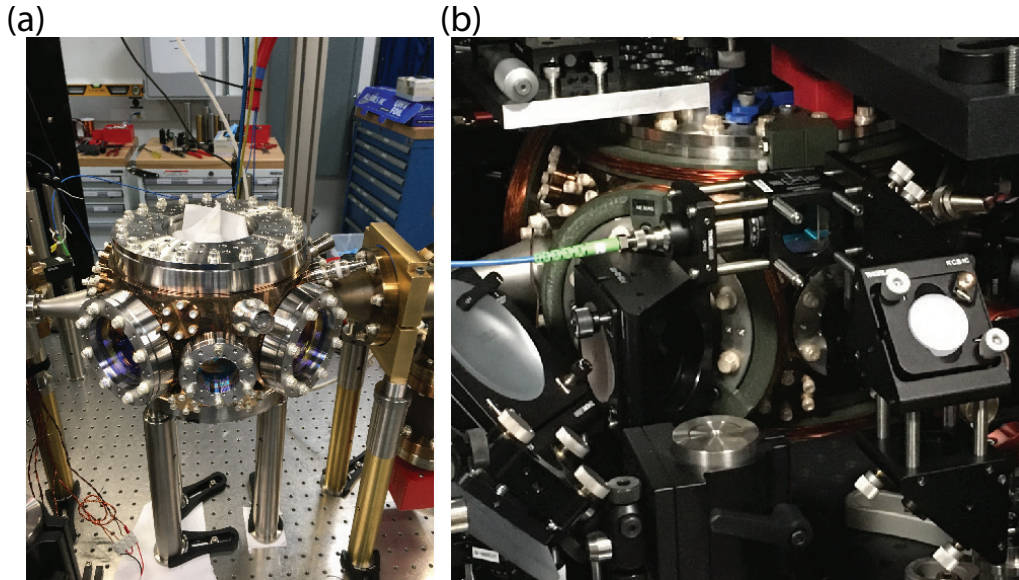
Figure 5.9: (a) photograph of science chamber shortly after construction clearly showing 2.75", 4.5", and 8" boltwashers. (b) boltwashers in use. We see a number of uses for the boltwasher in this image including mounting 3D printed coil molds to bolt washers (blue, green plastic) as well as the mounting of vertical cage system to the chamber.

washer to distribute the load of the bolt on the flange surface as a best practice. On the science chamber we have replaced the individual washers with a single SSL disk that has through holes for each vacuum bolt, but importantly, has a number of tapped holes which can be used to mount hardware directly to the vacuum chamber. Importantly, for the 2.75" and 4.5" boltwashers we took care to ensure that the bolt patterns were compatible with Thorlabs 30 mm and 60 mm cage systems so that we can install cage components directly onto the vacuum chamber. the 8" boltwashers have rings of easily accessible 1/4"-20 and 8-32 bolt holes.

So far we've used the boltwashers for the following:

- Mounting bias coils onto the science chamber. We 3D printed molds with tabs for bolts that interface with the viewport bolt washers. Coils were wrapped around the molds prior to installation.

- The 4.5" viewports closest to the MOT chamber are quite close to this chamber and it was difficult to mount a mirror onto the breadboard at those location. We were able to easily mount a cage-mounted mirror onto the chamber using these bolt washers.

- The absorption imaging light enters one of the 2.75" CF viewports via a cage system which is attached to this bolt washer. The beam is deflected vertically using this system

so that it does not need to take up space on the atom-height optics breadboard. The imaging optics on the other side of the chamber are similarly lifted using a boltwasher mounted cage system. The optical path is also shared with a molasses beam.

- The chamber feet are attached to the lower boltwasher.

- Mounting 3-axis translation stages to both the upper and lower 8" boltwashers to minimize the mechanical path length between the imaging objectives and the science chamber and science cavity. This should reduce mechanical drift relative to a design where the objectives are mounted to separate breadboards.

- Mounting irises to the chamber with machine precision centration on the viewports for optics alignment

5.9   Coils

We've mounted three pairs of bias coils onto the science chamber with axes oriented along the cavity, cross, and vertical axes. These coils are all wrapped around 3D printed molds which are fastened to the science chamber bolt washers.

The coils along the cavity and cross axes are formed from 100 turns of 20 AWG magnet wire. They have a radius of approximately 65 mm and a distance to the center of the chamber of about 130 mm. We estimate that they generate about 1.8 G/A.

The coils along the vertical axis have 25 turns of 16 AWG magnet wire. Their radius is 110 mm with a distance to the center of the chamber of 75 mm. We estimate they generate a field of 1.65 G/A.

We installed anti-Helmholtz coils inside the re-entrant viewports to allow us to generate strong magnetic field gradients at the location of the atoms. These magnetic field gradients could be used for magnetic trapping or to generate spatially dependent Larmor frequencies for MRI-like applications. These coils were custom designed and tested by Emma Deist and fabricated by Custom Coils Inc. They have an ID of 54 mm, allowing clearance for a 2" OD microscope objective, and an OD of 82 mm. They have a height of about 23 mm and are offset from the atoms by about 18 mm. There are 4 radial turns and 6 axial turns for a total of 24 turns. These coils are formed from square-profile hollow magnet wire. We estimate that they generate 1.6 G/cm/A. In initial tests, we were able to flow 2.5 mL/s of chilled water through these coils. Each coil is clamped into place within its respective re-entrant viewport using a two-piece, 3D-printed mount which is attached to the 8" bolt washer for that viewport. CAD designs for the coils can be seen in Fig 5.4(a).

5.10   Side Absorption Imaging

The horizontal science chamber axis perpendicular to the transport axis is used for absorption imaging. Onto one of the 2.75" CF port boltwashers we've installed a cage system

that couples in circularly polarized imaging and repump light from a single optical fiber. This imaging light passes through the center of the chamber where the atoms are located and then passes out the second 2.75" CF port. Another cage system is installed on the bolt washer for this other viewport onto which is mounted a $f = 75\,\text{mm}$ lens which re-images the atoms in the center of the chamber onto a CCD camera (Grasshopper USB3 GS3-U3-15S5M-C, FLIR). We've measured this imaging system to have a magnification of 0.77 by monitoring the atoms falling due to gravity.

Chapter 6

Science Cavity Laser Systems

In this experiment, the science cavity is driven by a probe laser and cavity ODT laser. To ensure coherence, reduce heating, and establish the appropriate regime of cQED, it is critical that there are minimal relative frequency fluctuations between the lasers and the science cavity. Indeed, it is also important that the relative frequency of the cavity is fixed in frequency with respect to the atomic transition absolute frequency.

To satisfy these constraints, we utilize a scheme in which we (1) use PDH locks to offset lock both the probe and cavity ODT lasers to a secondary ultra-low-expansion (ULE) cavity and (2) use another PDH lock to fix the length of the science cavity to the wavelength of the cavity ODT. By carefully choosing the laser and various offset frequencies, it is possible to satisfy all of our constraints as I will describe in detail in the rest of this chapter.

6.1   ULE Cavity

The ULE[1] cavity used for our purposes was purchased from Stable Laser Systems. The ULE cavity is housed in a vibration-isolated and temperature-stabilized vacuum enclosure to ensure excellent short- and long-term frequency stability for the cavity resonances. The ULE enclosure was based on a design from the Bloch group at the Max Planck Institute of Quantum Optics. The design was modified and assembled by Aron Lloyd and Johannes Zeiher. Photographs of the assembly are shown in Fig. 6.1. The cavity includes mirror coatings for 780 nm and 1560 nm as well as additional wavelengths to support stabilization of laser for the E8 experiment that is attempting to laser cool transition metal atoms such as titanium. The ULE cavity is 10 cm long yielding a FSR of 1.5 GHz.

We treat this cavity as an absolute frequency reference that does not drift in time and to which we can lock both lasers. The locks between the lasers and the ULE cavity reduce the

---

[1]ULE refers to a glass-like material whose thermal expansion coefficient changes sign from positive to negative at its so-called zero-crossing temperature. At this zero-crossing temperature, which can be designed to be close to room temperature, the length of the cavity is insensitive to temperature fluctuations to first order.

(a) (b)
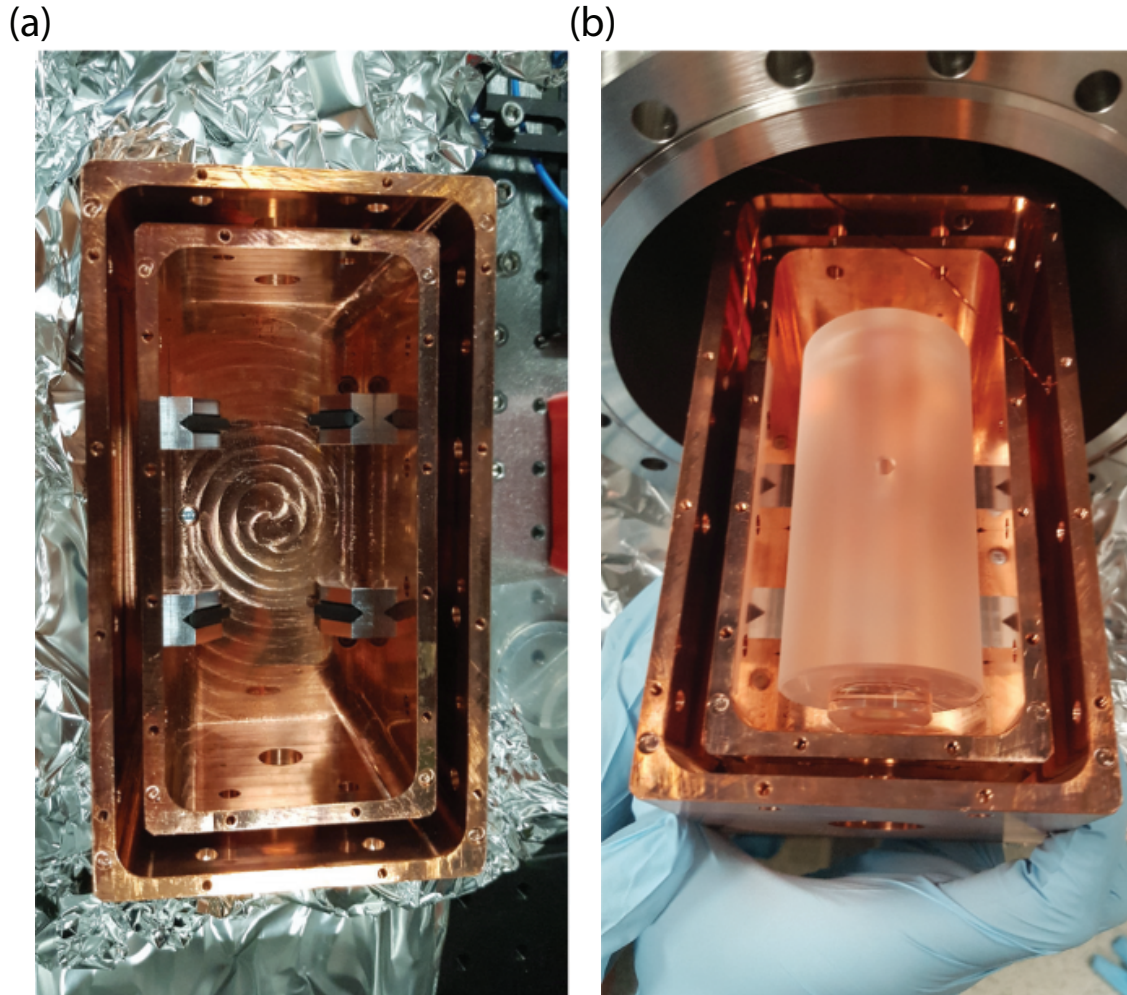
Figure 6.1: ULE cavity installation into temperature-stabilized and vibration-isolated vacuum enclosure. (a) The cavity sits on viton cylinders for vibration isolation. The cavity is housed inside of two nested copper boxes for thermal isolation from the environment. Additionally the larger box is contacted to a Peltier element and thermocouples to allow for readout and stabilization of the temperature of the cavity to the zero-crossing of the thermal expansion coefficient. Housing the cavity in vacuum reduces index of refraction and thermal fluctuations due to air currents. (b) Brushed glass cylinder is the ULE material with an overall length of approximately $10\,\mathrm{cm}$ and the transparent cylinder at bottom is one cavity mirror.

amplitude of the laser frequency fluctuation spectrum (up to the bandwidth of the feedback loop), thus reducing the laser linewidth. At low frequencies, this suppression ensures the lasers stay at the frequencies we desire for long periods of time (for example over the course of many shots of the experiment or multiple days). At high frequencies, this suppression ensures that frequency fluctuations of either laser do not result in amplitude modulation within the cavity during the course of a single experimental sequence.

6.2 Cavity Lock Scheme

6.2.1 The Frequencies Involved

To start thinking about the laser lock scheme we begin with the absolute frequency of the ^{87}Rb D$_2$ transition in the vicinity of $780\,\text{nm}$ or $\omega_A \approx 2\pi \times 384\,\text{THz}$ [66]. See Fig. 6.2. Of interest for cQED is the frequency ω_C of the science cavity TEM$_{00}$ mode that is nearest to atomic resonance. The detuning between the atomic resonance and the nearby cavity mode is given by $\Delta_{CA} = \omega_c - \omega_a$. Additionally, we must consider the frequency of the probe laser, ω_P, that drives the cavity. The detuning between the probe and the cavity is given by $\Delta_{PC} = \omega_P - \omega_C$. The detuning between the probe laser and atomic resonance is given by $\Delta_{PA} = \omega_P - \omega_A$ However, Δ_{PC} is typically on the order of κ, which, in turn, is typically much smaller than Δ_{CA}. We, then, typically consider $\Delta_{PA} \approx \Delta_{CA}$ and talk only about Δ_{CA}.

Adjusting Δ_{CA} and Δ_{PC} allows us to access different regimes of cQED physics. For example, if $\Delta_{CA} \gg g, \Gamma$, then the system is in the dispersive limit of cQED in which the
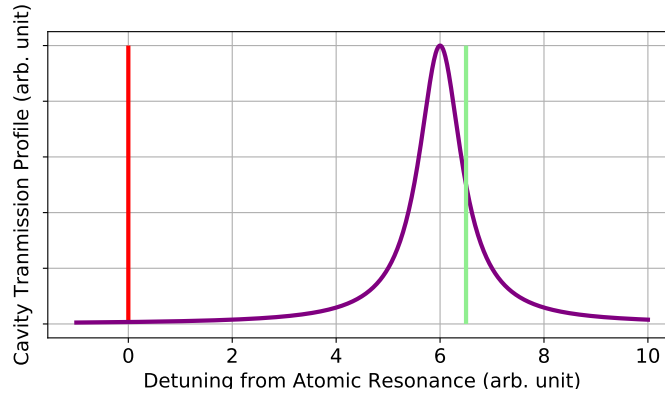
Figure 6.2: Various cavity detunings for easy reference. Arbitrary frequency scale. Atomic resonance is the red vertical line, the probe laser frequency is the green line and the cavity profile is the purple curve. In the experiment we typically have that $\Delta_{CA} \gg \kappa, \Gamma$ and that Δ_{PC} is on the order of κ. In this picture we would say the cavity is blue-detuned from atomic resonance and the probe is blue-detuned from cavity resonance.

atoms and photon fields weakly 'dress' each other, resulting in a slight frequency shift of atomic and cavity frequencies as described in Sec. 1.3.3. For $\Delta_{CA} \lesssim g, \Gamma$, the system is in the resonant cQED regime in which the atom and photon field strongly hybridize, giving rise to a vacuum Rabi splitting of the cavity transmission profile and a large excited state population. The ability to tune Δ_{PC} across the cavity resonance is of general interest for various characterizations and experiments. We would like to have arbitrary control over these two parameters Δ_{CA} and Δ_{PC}.

Recall that the FSR for the science cavity is $f_{\mathrm{FSR}} \approx 15\,\mathrm{GHz}$. In the case that the driven cavity mode is the mode which is closest to atomic resonance we neglect the presence of all of the other other modes because (1) they are relatively further detuned and (2) they are not actively pumped with energy. To date, we have always driven our cavity in this configuration meaning that we have only considered values $0 < |\Delta_{CA}| < 2\pi \times f_{\mathrm{FSR}}/2 \approx 2\pi \times 7.5\,\mathrm{GHz}$. It is an open question as to what interesting physics may arise if there are undriven modes closer to atomic resonance than the driven mode. We will be interested in positive and negative values of Δ_{PC} on the order of a few times κ, recalling that $\kappa \approx 2\pi \times 1\,\mathrm{MHz}$.

In Sec. 3.6, we concluded that to ensure a long wavelength spatial beatnote between the probe and cavity ODT, the frequency of the cavity ODT should be within one or two science cavity FSR's away from half of the probe frequency: $f_{ODT} = f_P/2 + f_{\mathrm{offset}}$. Here f_{offset} is some offset frequency less than (but possibly on the order of) f_{FSR}, which arises due to mirror reflection phase effects. See Sec. 3.6 for details.

6.2.2 Setting the Probe Frequency

I find it easiest to describe the lock-chain by describing how one performs the laser lock in practice with reference to Figs. 6.3, 6.4, and 6.6. The absolute frequency reference within the lab is the sub-Doppler saturated absorption spectroscopy used to lock our MOT cooling laser. Here we can directly observe the individual hyperfine transitions in ^{87}Rb and compare these with literature values to determine the absolute frequency of the cooling laser. We can send this cooling light onto our wavemeter (taking into account an AOM shift in the spectroscopy setup) to calibrate the wavemeter to the spectroscopy. The spectroscopy allows us to know the MOT laser frequency at the MHz level. The wavemeter has precision/accuracy at the 10s of MHz level.

The probe laser is a narrow, broadly tunable Toptica DL Pro ECDL with a linewidth $< 100\,\mathrm{kHz}$. Light is picked off from this ECDL and directed into three beam paths of interest. One path goes directly to the wavemeter, one goes to the ULE cavity, passing through a fiber EOM, and one is shifted down in frequency by a double pass AOM before being sent to the science cavity. During laser locking the probe AOM is set to a center frequency f_0 in the center of its bandwidth. See Fig. 6.4.

We first choose a desired value for Δ_{CA}. Once this value is chosen, we tune the probe, by tuning the piezo grating, laser current, and laser temperature, so that Δ_{PA} is equal to the desired value for Δ_{CA}. This ensures that when the cavity is made resonant with the probe we will realize our desired value of Δ_{CA}. We determine Δ_{PA} by monitoring the probe
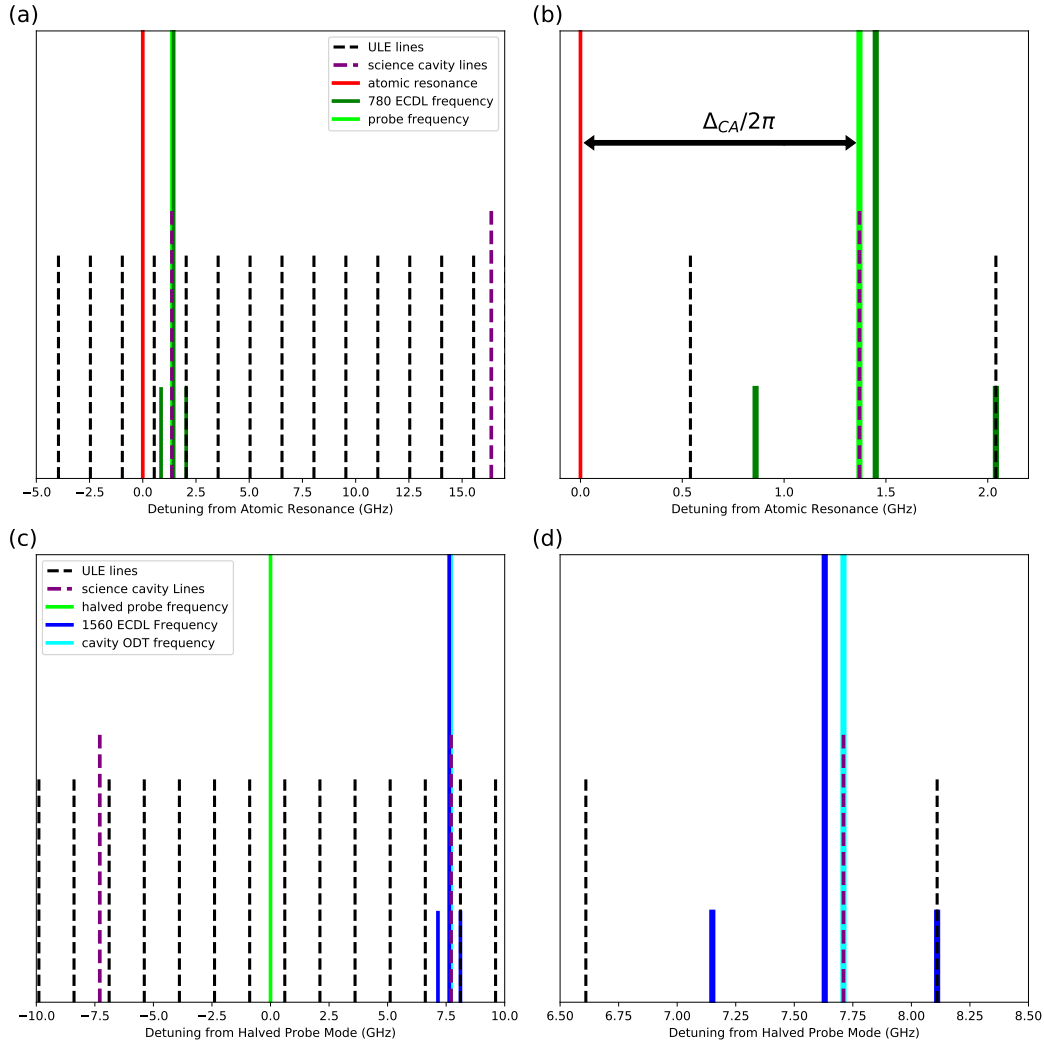
Figure 6.3: Frequency locking of the probe laser, cavity ODT laser, ULE cavity, and science cavity. The frequency spacings and offsets in these diagrams are all representative and not necessarily reflective of actual values in the experiment. See text for full details. (a), (b): The 780 ECDL (dark green) is detuned by Δ_{CA} from atomic resonance (red). An EOM sideband on this laser is locked to the nearest ULE line and finally the probe light (light green) is shifted from the 780 ECDL frequency by an AOM shift that can be adjusted to vary Δ_{PC}. The science cavity length is tuned so that one of its modes is resonant with the probe. (c), (d): The 1560 ECDL (dark blue) is tuned near to the science cavity line that is nearest to the halved probe frequency (light green). An EOM sideband is then locked to the nearest ULE line. The cavity ODT is shifted away from the 1560 ECDL carry by an AOM shift. Because of mirror reflection phase effects, the cavity ODT is detuned from the halved probe frequency by almost half a FSR. Such a configuration has implications for the properties of the cavity beatnote.

frequency on the wavemeter and taking into account the AOM shift of $-2f_0$ between the probe light and the light sent to the wavemeter. When using the wavemeter, our accuracy for setting Δ_{PA} is at the 10s of MHz level. If we require better accuracy (for example, if we want a value for Δ_{CA} less than $100\,\mathrm{MHz}$), then we will need to either perform an offset frequency lock between the probe and cooling laser or we will need to do atomic spectroscopy on ultracold atoms in the science chamber with the probe.

Our next goal is to lock the probe to this frequency using the ULE cavity. The 780 ECDL is necessarily within half a ULE FSR –approximately $750\,\mathrm{MHz}$– from the nearest ULE mode. We bridge this frequency difference with a fiber EOM that puts rf sidebands on the 780 laser before it is sent to the ULE cavity. The rf frequencies are set with a frequency synthesizer. With the 780 ECDL carrier centered on the appropriate frequency on the wavemeter, we perform small sweeps of the cavity frequency and monitor the reflection from the ULE cavity to see reflections dips from the carrier and the two sidebands resonating with the cavity. We tune the frequency of the sidebands until one sideband is resonant with the nearest ULE mode to the carrier at the same time as the carrier is at the appropriate value for Δ_{PA} as monitored on the wavemeter.

The rf tone from the synthesizer is mixed with a $20\,\mathrm{MHz}$ modulation tone before being sent into the fiber EOM. This allows us to perform a PDH lock between one of the sidebands of the 780 ECDL and ULE cavity by demodulating the light reflected off of the cavity and feeding back to the laser current and grating position. We engage this lock so that 780 ECDL frequency, and thus Δ_{PA}, are now locked. If Δ_{PA} is not at exactly the right value we can tune the rf EOM sideband frequency by small amounts while the PDH lock is engaged to realize the appropriate detuning.

6.2.3 Overlapping the Cavity ODT and Probe on the Science Cavity

The next tasks in the lock chain are to lock the cavity ODT to the ULE cavity in a manner similar to the probe lock and to lock the science cavity to the cavity ODT. The trick is to do this so that the probe and the cavity ODT can be simultaneously resonant with the science cavity!

The $1560\,\mathrm{nm}$ laser is also a narrow-linewidth Toptica DL Pro ECDL. The beam path is also split into three paths, one for the wavemeter, one for the ULE cavity (which also goes through a fiber EOM) and one that passes through a fixed frequency AOM and is sent to the science cavity as the cavity ODT. The AOM allows us to intensity stabilize the cavity ODT.

We move through the lock chain in stages of increasing frequency precision. First, by monitoring the wavemeter and tuning the ECDL, we set the cavity ODT frequency to be half the frequency of the probe to maximize the wavelength of the probe/cavity ODT beatnote. At this point, if the length of the cavity were tuned so that the cavity is resonant with the probe, the cavity ODT would be within half a FSR, $f_{\mathrm{FSR}}/2 \approx 7.5\,\mathrm{GHz}$, of the nearest science

cavity mode.

By sweeping the length of the science cavity and monitoring the transmitted 780 nm and 1560 nm light, we can observe the probe and cavity ODT TEM_{00} resonances.[2] Note that the full scan of one of our cavity piezos allows us to access 3 FSRs of the cavity ODT and about 5 FSRs of the probe. We choose to lock to a probe mode near the center of the scan range. Recall from Sec. 3.6 that the behavior of the cavity beatnote varies depending on whether this an even or an odd longitudinal probe mode. In the future, if the parity of the probe mode needs to be adjusted, one should return to this step and select a different probe mode.

After identifying a target probe mode, the 1560 ECDL should be tuned until one of the cavity ODT resonances is overlapped with the target probe mode on the science cavity. Once this condition is met, the next task is to lock the 1560 ECDL to the ULE cavity. We follow the same procedure as for the probe, adjusting the EOM sidebands until one of the sidebands is resonant with the nearest ULE cavity mode at the same time that the cavity ODT is overlapped with the probe on the science cavity.

We modulate the current to the 1560 ECDL at 20 MHz to allow us to perform the PDH lock. When the above condition is roughly met, we engage the 1560 EOM sideband PDH lock to the ULE cavity. In practice, after this lock is engaged, the probe and the cavity ODT will not be exactly overlapped on the science cavity. We make the final correction by tuning the EOM frequency with the PDH lock engaged. This has the net effect of (1) shifting the 1560 carrier relative to the EOM and thus (2) shifting the cavity ODT transmission peak on the science cavity relative to the probe. We have found that this in-loop adjustment can be done reliably for 1 MHz steps and sometimes for steps as large as 10 MHz if the 1560 lock is performing well.

Finally, with the probe and cavity ODT 'overlapped' on the science cavity, the length of the science cavity is locked to the wavelength of the cavity ODT via another PDH lock in reflection. This science cavity PDH lock feeds back to the cavity piezo. The performance of this lock is discussed in the context of the vibration isolation system in Sec. 5.2.

At this point both lasers should be locked to the ULE cavity, the science cavity should be stabilized to the cavity ODT, and both the probe and cavity ODT should be transmitting through the science cavity. This completes the E6 cavity lock chain. The detuning between the probe and the science cavity Δ_{PC} can be easily tuned around cavity resonance by adjusting the probe AOM frequency.

If necessary, this whole process can be repeated locking to a different longitudinal probe mode by changing the length of the science cavity or locking to a different longitudinal cavity ODT mode by changing the cavity ODT frequency. Finally, if Δ_{CA} is changed then this process must, in general, be repeated for the new setting.

---

[2]In the future it may be of interest to lock to higher order TEM modes.

Figure 6.4: Cavity Probe lock scheme. We generate the fiber EOM sidebands using a Valon 5009 Dual Frequency Synthesizer. We mix this with a 20 MHz tone from a Toptica PDD unit. We use this same PDD unit for detection and demodulation of the ULE PDH signal. We generate the reference signal form the probe/LO phase lock from a DDS on an Analog AD9958 evaluation board. Our VCO control boxes include voltage inputs that control both the frequency and amplitude (via a voltage controlled attenuator) of the output rf tone.

6.3 Probe Laser System

6.3.1 Probe and LO Phase Lock

We utilize optical heterodyne detection to extract the amplitude and phase of the probe light transmitting out of the cavity. The detection works by interfering the cavity output mode with a local oscillator (LO) beam, which, ideally, has no amplitude or phase fluctuations. We derive our LO beam from the same ECDL laser as the probe beam. On the laser table, the probe and LO are split and each one passes through its own double-pass AOM and is shifted down in frequency by $-2 \times f_{\text{AOM,probe/LO}}$. The difference between these two frequency shifts gives the LO detuning: $\Delta f_{\text{LO}} = -2 \times (f_{\text{AOM,LO}} - f_{\text{AOM,probe}})$.

These double-pass AOMs are set up for optimal performance close to the the center of their bandwidth tuning range near $80\,\text{MHz}$. This allows us to tune the probe and LO AOMs by as large a range as possible to allow us to tune Δ_{PC} by a large range while maintaining a fixed LO detuning.

The LO detuning Δf_{LO} sets an upper bound for the signal bandwidth that can be measured in the heterodyne detector without aliasing. The probe light coming out of the cavity is already band-limited by the linewidth κ so we sacrifice no signal as long as $2\pi \times \Delta f_{\text{LO}} \gg \kappa$. It is also beneficial to have the LO detuning be larger than κ because then LO retro-reflections, which would otherwise be coupled into the cavity, are suppressed by the cavity frequency response.

On the other hand, there are also disadvantages to making Δf_{LO} too large. The first disadvantage is that a large Δf_{LO} limits the usable frequency tuning range for the probe and LO AOMs while maintaining the required frequency splitting. Next, if we directly digitally acquire the heterodyne signal and subsequently demodulate and analyze the signal in software, then a large LO frequency requires high digital sample rates for acquisition resulting in datasets that take up a lot of memory. Finally, our heterodyne detector exhibits a $10\,\text{dB}$ increase its electronic noise floor between DC and $60\,\text{MHz}$. With all of these considerations in mind, we typically work with a LO detuning of $15\,\text{MHz}$

The two beams are transferred via independent optical fibers from the laser table onto an optics breadboard at the height of the science chamber. Our experience in E3 taught us that such fiber links add a substantial level of phase noise to the light.[3] To suppress these fluctuations, we set up a phase lock system to lock the relative phase of these two beams once they are deployed onto the experiment table.

After the probe and LO are coupled onto the breadboard, light is picked off from each beam and the picked off beams are interfered on a photodiode (Thorlabs PDA36A2). The resultant beat signal at approximately Δf_{LO} is then amplified, converted into a square wave, and then combined on a phase frequency detector (PFD)[4] (Analog Devices HMC439QS16G)

---

[3]This noise is due to the fibers experiencing mechanical strain or thermal fluctuations along their entire length ($10\,\text{m}$ each) which could arise from a number of usual sources within a standard optics lab.

[4]A PFD is a device that takes in two rf tones and outputs an error signal. For large detunings, the error signal is positive or negative depending on which tone has a higher frequency. For small detunings, the error
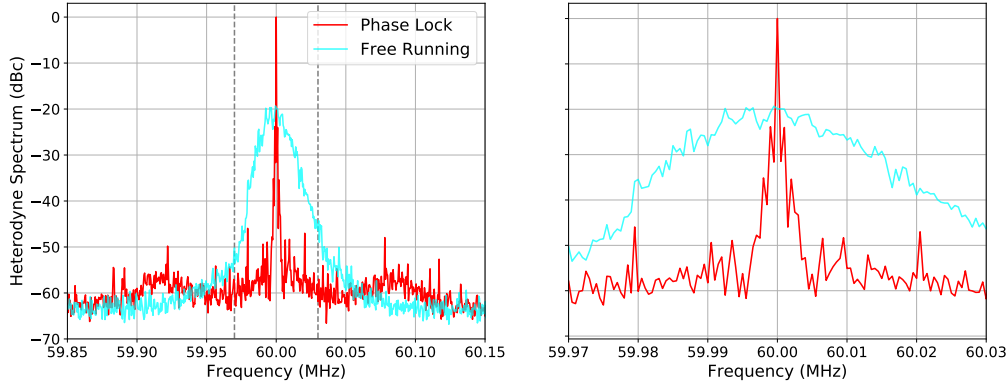
Figure 6.5: Probe/LO phase lock performance. The LO detuning was set to 60 MHz for this measurement. We see similar performance at 15 MHz.

with a phase-stable, reference rf tone generated by a direct digital synthesizer (DDS) [137]. We find that it is necessary to convert the signal to a square wave to get the desired PFD error signal when operating the PFD at frequencies below 50 MHz. Finally, the error signal is filtered through a PI lockbox whose output signal controls the frequency of the voltage-controlled oscillator (VCO) that drives the LO AOM.

We can characterize the quality of this phase lock by making another out-of-loop interferometric measurement of the phase of the two beams. Conveniently, the heterodyne detector (described in more detail in the following section) already serves as such an out-of-loop phase measurement between the probe and the LO. In Fig. 6.5, I show the PSD measured on the heterodyne detector when the phase lock is enabled and disabled. With the phase lock disabled, we see a broad frequency spectrum most likely arising due to relative fiber noise between the two beams. After the phase lock is enabled, the spectrum becomes much more narrow, now exhibiting a linewidth at the kHz level. The servo bumps, appearing at about 90 kHz on either side of the carrier, can be tuned in frequency and amplitude by adjusting the parameters of the PI lockbox in the feedback loop.

6.3.2 Cavity Heterodyne Detection

In the previous section, I described how we phase stabilize the probe and LO in preparation for heterodyne detection. In this section, I will describe the rest of the heterodyne beam path. These optical paths are shown schematically in Fig. 6.4

---

signal is positive or negative depending on which tone is advanced in phase relative to the other. The PFD plays a role similar to a mixer in a standard phase lock. However, the PFD has the advantage that it has a very large frequency capture range due to its ability to generate a meaningful error signal even for large detunings.

Early in the optical path for both the probe and LO, a small fraction of power from each beam is picked off for detection on a photodiode for intensity stabilization. After intensity stabilization, the probe passes through a neutral density (ND) filter with an attenuation of 60 dB. This reduction in optical power is necessary so that the resulting intracavity energy is at the few photon level.

The probe and the LO then take separate paths to get to the heterodyne detector. The probe must pass through the cavity and interact with the atoms, while the LO should travel directly to the detector, picking up as little phase disturbance as possible. We set up the LO beam path so that it shares as many optics as possible with the probe beam path making phase fluctuations arising from vibrating optics common mode. To this end, the LO passes through the science chamber, parallel to the probe but avoiding the science cavity. The LO shares the same input and output and lens as the cavity probe.

To realize this alignment, the probe and the LO are combined, with slightly different propagation angles, on a PBS before the science chamber. Then, when the beams pass through the input lens, the probe is coupled into the cavity but the LO is focused alongside the cavity. On the output side of the chamber, the two beams are collimated using the output lens and separated using a PBS. The LO then has a short independent path where it can be steered before, again, being combined with the probe on a PBS, this time attempting to achieve perfect mode-matching for the two beams.

The beams then reflect off of a dichroic beam-splitter that separates the probe and LO from the cavity ODT. At this point, the probe and LO have orthogonal polarizations, which will not interfere on the detector. We overcome this by passing the beams through a $\lambda/2$ waveplate so that they are diagonally and anti-diagonally polarized. Then we split the two beams on another PBS. Each beam then contributes half of its power to each output port of this final PBS. The beams at the output ports of this final PBS are then directed onto the detectors of a balanced photoreceiver (Newport 1807-FS) where the probe and LO interfere to generate our heterodyne signal.

In addition to the heterodyne detector, this path includes a flip mirror that can optionally direct the beams onto a separate beam analysis path. This path includes a short focal length lens that focuses the beams (through a beamsplitter) onto an APD and a Thorcam CMOS camera. These detectors are important diagnostics tools that allow us to easily repeat the characterization measurements described in Sec. 4.6.3 on the science cavity. In particular, when overlapping the probe and cavity ODT on the science cavity, we use the APD to observe the probe transmission spectrum and we use the camera to ensure we are locking to the TEM_{00} mode.

6.3.3 Intracavity Photon Number

We can convert the probe input power into an intracavity photon number using Eq. (3.76):

$$\bar{n} = \frac{1}{\hbar\omega}\frac{1}{\kappa}4\eta_{\mathrm{in}}P_{\mathrm{in}}. \tag{6.1}$$

This corresponds to about 18 photons per nW of input power assuming perfect input mode-matching. In practice our input mode-matching is probably 20-50%.

We can also extract the intracavity photon number by measuring the power transmitted out of the cavity.

$$\bar{n} = \frac{1}{\hbar\omega}\frac{1}{\kappa}\frac{1}{\eta_{\text{out}}}P_T, \tag{6.2}$$

corresponding to about 1 photon per transmitted pW.

6.4 Cavity ODT

1560 nm cavity ODT light is delivered onto the experiment table via optical fiber. Early in the beam path, light is picked off on a 90:10 beamsplitter for intensity stabilization. The light is then overlapped with the probe beam on a dichroic beam-splitter before being sent through the same focusing lens as the probe into the optical cavity. The light retro-reflected from the cavity exits through the other port of the 90:10 beamsplitter and is directed onto a photodiode to derive the PDH error signal that we use to lock the science cavity length to the cavity ODT. We're able to generate a PDH error signal because of the same 20 MHz
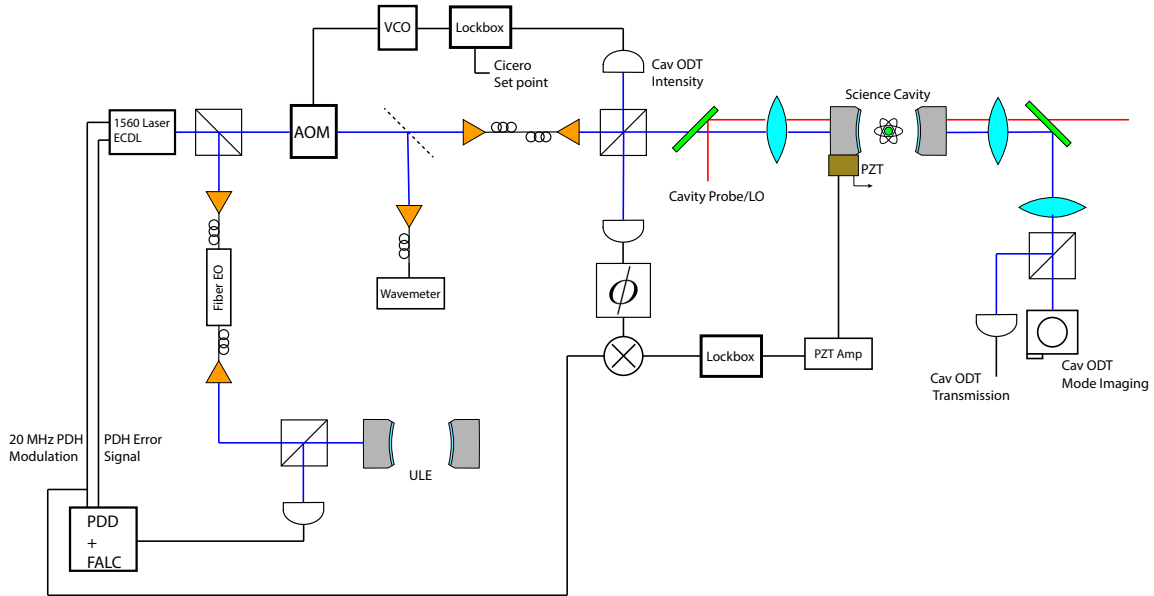
Figure 6.6: Cavity ODT optics layout

laser current modulation that was used for the 1560 ECDL lock to the ULE cavity.[5]

After passing through the the science cavity, the cavity ODT is collimated by the cavity output lens. The cavity ODT is then split from the probe/LO by a dichroic for analysis. The beam is focused through a 150 mm focal length lens and then split using a beamsplitter and directed onto an IR CMOS camera[6] and an IR APD. Here, as for the probe, we can observe the cavity ODT transmission spectrum and mode shape to ensure we are driving the TEM_{00} mode.

6.4.1 Cavity ODT Trap Depth Calculations

In this section, I will first quote some useful formulas for calculating parameters of ODTs generally, and then I will present details specific to the 1560 nm ODT that we use to form the cavity ODT.

6.4.2 Generic ODT Calculations

Presently, in E6, we form ODTs using three wavelengths of light: 1560 nm for the cavity ODT, 1064 nm for the transport ODT, and 808 nm for the microtweezers. Throughout, I will assume we are driving with π-polarized light; in other words, I will ignore vector light shifts and assume there is no tensor light shift [61, 66].[7] All three of these beams will interact appreciably with both the D_1 and D_2 lines. Our calculations will require use of the two effective transition dipole elements [66]:

$$d_{\text{eff},D_1} = 1.465 \times 10^{-29}\,\text{C m},$$
$$d_{\text{eff},D_2} = 2.069 \times 10^{-29}\,\text{C m}. \tag{6.3}$$

For an electric field with amplitude E the Rabi frequencies on these two transitions are given by

$$|\Omega_{D_i}|(\boldsymbol{r}) = \frac{|E(\boldsymbol{r})|d_{\text{eff},D_i}}{\hbar}. \tag{6.4}$$

For far-detuned light $(\Delta \gg \Gamma, \Omega)$ we can easily calculate the ground state energy shift using a second order perturbation theory [45, 61]:

$$\frac{V(\boldsymbol{r})}{\hbar} = \sum_{i=1,2} \frac{|\Omega_{D_i}(\boldsymbol{r})|^2}{4} \left(\frac{1}{\Delta_{D_i}} + \frac{1}{\Delta_{D_i}^+} \right). \tag{6.5}$$

---

[5]In principle this modulation could adversely affect the atoms. Fortunately, it is well outside the cavity bandwidth so these fluctuations are heavily suppressed inside the cavity.

[6]High quality camera sensors active at 1560 nm were more difficult to source than for 780 nm. We use a Cinology CMOS-1201 camera.

[7]The assumption of no tensor light shift is valid for the ground states of Alkali atoms, which have $J = 1/2$.

Here $\Delta_{D_i} = \omega_{\text{laser}} - \omega_{D_i}$ is regular laser detuning and $\Delta_{D_i}^+ = -\omega_{\text{laser}} - \omega_{D_i}$ is the detuning for the corresponding counter-rotating term, which gives rise to the so-called Bloch-Siegert shift. The Bloch-Siegert term becomes relevant for detunings that are large compared to the laser frequency. The Bloch-Siegert shift makes an appreciable contribution, in E6's case, for the 1064 nm and 1560 nm ODTs.

We see that the potential is attractive for ground-state atoms for red-detuned light. An atom will be attracted to the point of highest optical amplitude within such a potential: $V_0 = V(\boldsymbol{r} = \boldsymbol{0})$. If the atom has low energy (or if a gas of atoms has a low temperature $k_B T$ compared to the trap depth V_0), then we can Taylor expand around this trap minimum to extract the trap frequencies of the effective harmonic potential felt by an atom deeply confined in such an optical potential. k_B is Boltzmann's constant. If the optical field is that of a focused Gaussian beam these trap frequencies are given by:

$$\omega_{\text{axial}} = \sqrt{\frac{2V_0}{z_R^2 m}},$$

$$\omega_{\text{radial}} = \sqrt{\frac{4V_0}{w_0^2 m}}. \tag{6.6}$$

Here m is the mass of an individual atom, w_0 is the beam waist, and $z_R = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh range. If the optical field is formed by two counter-propagating beams so that it forms an optical lattice, the axial frequency is given by

$$\omega_{\text{axial}} = \sqrt{\frac{2V_0 k^2}{m}}, \tag{6.7}$$

with $k = \frac{2\pi}{\lambda}$.

A low-temperature, non-degenerate thermal cloud in such a trap would have a spatial density $\rho(x, y, z)$ given by the Boltzmann factor $\rho(x, y, z) \propto \exp\left(-\frac{1}{2}\beta m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)\right)$, resulting in a cloud with a Gaussian shape characterized by:

$$\sigma_{\text{radial}}^2 = w_0^2 \frac{k_B T}{4V_0} = \frac{w_0^2}{4\eta},$$

$$\sigma_{\text{axial, Gaussian}}^2 = z_R^2 \frac{k_B T}{2V_0} = \frac{z_R^2}{2\eta},$$

$$\sigma_{\text{axial, lattice}}^2 = \left(\frac{\lambda}{2}\right)^2 \frac{1}{2\pi^2} \frac{k_B T}{V_0} = \left(\frac{\lambda}{2}\right)^2 \frac{1}{2\pi^2 \eta}. \tag{6.8}$$

Here I've defined $\eta = \frac{V_0}{k_B T}$. Rb atoms have a large scattering cross section leading to efficient thermalization and evaporation after they are loaded into an optical trap. As a result, Rb atoms typically rapidly equilibrate, through self-evaporation, to a temperature corresponding to $\eta = 10$.
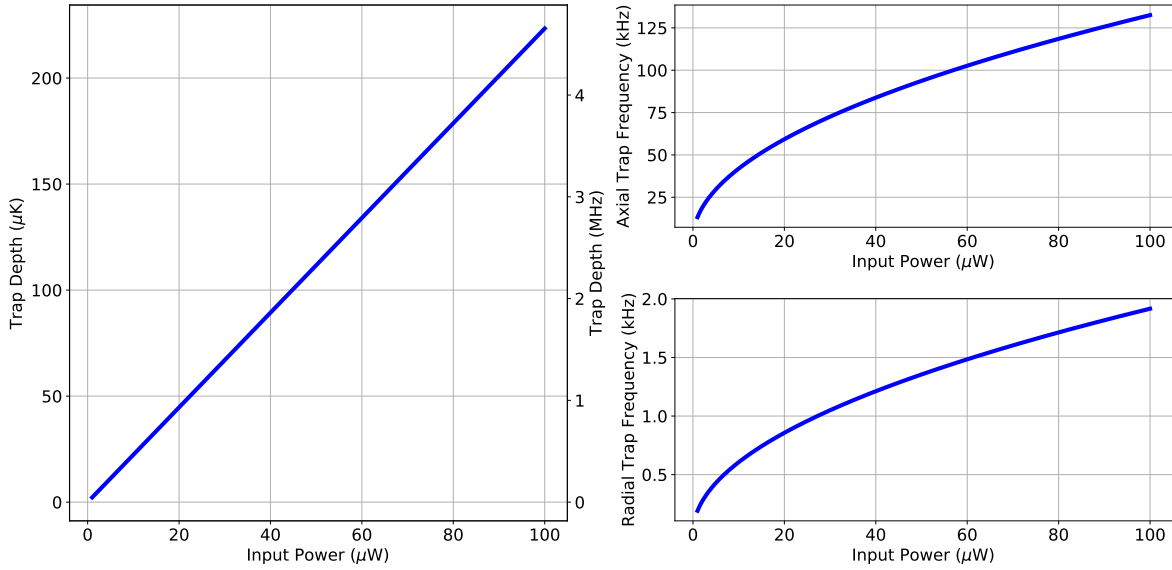
Figure 6.7: Cavity ODT trap depth and frequencies as a function of input powers assuming perfect mode-matching efficiency.

The science cavity mode for the cavity ODT has a waist of $w_0 = 24\,\mu$m and finesse of $\mathcal{F} = 23,400$. We can use this information, Eq. (3.72), and the results of this section to calculate the cavity ODT trap depth and frequencies as a function of the input power as shown in Fig. 6.7.

6.4.3   Excited State Polarizability

In the E6 experiment, all cQED probing and laser cooling to date takes place on the D_2 transition: $5\,S_{1/2} \to 5\,P_{3/2}$. In the previous section, I explicitly calculated the ground state Stark shift, or polarizability, to determine the trapping parameters for the cavity ODT. The polarizability relates the energy shift of a particular energy level to the squared amplitude of the driving electric field.

For Rb, the polarizability for the D_2 excited state, $5\,P_{3/2}$, is much larger than the ground state polarizability for driving fields in the vicinity of 1560 nm due to the presence of near resonant transitions to higher lying excited states shown in Fig. 6.8(d) [78]. In particular, there are two fine-structure transitions to $4\,D$ levels with wavelengths of about 1529 nm.

In the previous section, I showed how to calculate the polarizability 'by hand' by explicitly adding up the contributions to the Stark shift from each relevant transition. To calculate the polarizability of the excited state I used the Alkali Rydberg Calculator (ARC) software package, which already has a database of the relevant atomic transition data and a built-in method to calculate the Stark shift contributions from a large number of possible transitions.
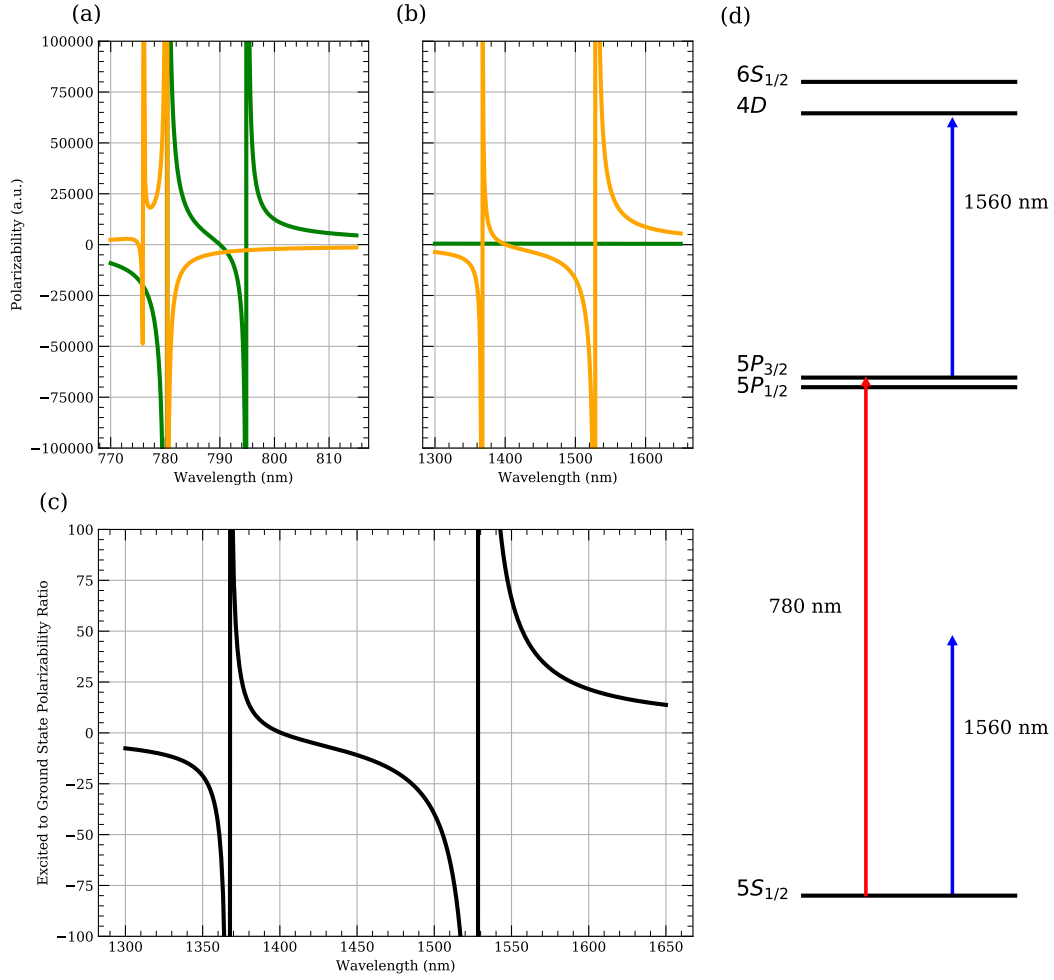
Figure 6.8: (a), (b): Scalar polarizability in atomic units for Rb $5\,S_{1/2}$ ground state (green) and $5\,P_{3/2}$ excited state (yellow). (c) Ratio of excited to ground state scalar polarizability. (d) To-scale level diagram showing the proximity of the $1560\,\mathrm{nm}$ probe light to the $5\,P_{3/2}$ doubly excited $4\,D$ manifold (which includes $4\,D_{3/2}$ and $4\,D_{5/2}$) and $6\,S_{1/2}$ transitions. Polarizabilities calculated using ARC python package [138].

The results of this calculation for the ground and excited state polarizabilities are shown in Fig. 6.8(a)-(c).

We can see that, in the vicinity of $1529\,\mathrm{nm}$, the polarizability for the $5\,\mathrm{P}_{3/2}$ state is about $(45\times)$ larger than that of the $5\,\mathrm{S}_{1/2}$ ground state. This means that a π-polarized $1560\,\mathrm{nm}$ ODT that is $100\,\mu\mathrm{K}$ deep for the ground state will be $4.5\,\mathrm{mK}$ deep for atoms in the excited state, corresponding to an excited state level shift of $\approx 100\,\mathrm{MHz}$. The vector polarizability for the excited state (not shown) is similar in magnitude to the scalar Stark shift above $1529\,\mathrm{nm}$.

This excited state shift has important implications for laser cooling and cQED. In Sec. 7.4, I will describe how we utilize an optical molasses on the D_2 line to cool and image atoms trapped in microtweezers. This molasses is typically red-detuned by 10s of MHz from the unshifted transitions. The large excited state polarizability could shift the molasses light from being red-detuned for atoms at the cavity ODT nodes to being blue-detuned at the cavity ODT anti-nodes, thus spoiling the molasses mechanism [78]. The presence of the cavity ODT also alters the cavity probe detuning by the same mechanism.

Both of these effects are important to be aware of and understand because they can cause problems with the experiment. However, this large excited state polarizability can also be seen as an advantageous feature. Indeed, in Sec. 8.3, I will show the results of an experiment in which we leverage this effect to map out the local intensity of the cavity ODT by monitoring changes in atomic fluorescence as an atom in a microtweezer is positioned at different locations within the cavity ODT.

Chapter 7

High Resolution Imaging and Addressing

One of the main science targets for E6 has been the combination a high cooperativity cQED cavity system with a high resolution imaging and addressing system to explore physics involving local degrees of freedom of a few- or many-body quantum system. We facilitated this in our system by designing the UHV science chamber with re-entrant viewports on top and bottom to allow the installation of out-of-vacuum objectives with short working distance close to the atomic sample in the center of the optical cavity as described in Chapter 5. In this chapter, I will describe the imaging and AOD based micro-tweezer systems, which are built around the high-NA objective.

7.1 The High-NA Imaging and Addressing System

Our goals with this high-NA imaging system are (1) to image individual atoms that may be separated by as little as $1\,\mu$m and (2) to generate arbitrary optical potentials in the focal plane of the objective to manipulate the atoms within the optical cavity. We target a resolution on the order of $1\,\mu$m, sufficient to resolve single sites of an IR optical lattice. This will allow us to extract local system parameters and correlation functions as has been demonstrated in various quantum gas microscope setups [35–40, 42–44]. In fact, the cavity ODT already is such a lattice with a lattice constant of $780\,$nm. It is also advantageous to have a resolution that is much smaller than the waist of the probe mode, $w_0 \approx 17\,\mu$m so we can clearly control and measure the relative positioning of atoms within the probe mode. We will typically use fluorescence imaging rather than absorption since it is very difficult to measure the presence of a single atom against intensity fluctuations in an absorption imaging beam, though it is possible [139]. We do have the ability to perform absorption imaging through our high-NA system if needed.

The addressing system is multifaceted. With high-NA access to the atomic cloud, there are several ways to realize addressing schemes for the internal and external degrees of freedom

for atoms within the cavity.

- Far-detuned light shone in from the side can create optical potentials that can be used to put mechanical forces on atoms [49, 50, 140, 141].

- Addressing light from the side could be used in combination with cavity photons to drive local multi-photon Raman transitions between internal atomic states [21].

- Addressing light from the side could be used to create atomic-state-dependent local Stark shifts [142–144].

Since the cavity field can be used to generate interactions between atoms and to measure the atoms, these addressing fields give us the ability to temporally and spatially modulate atomic interactions and measurement protocols.

We have considered two schemes to generate optical potentials in the objective focal plane. The first technique is using an AOD driven by multiple rf tones generated by a software defined radio (SDR) to create multiple beams that can be focused by the objective onto distinct spots in the focal plane [50]. In this approach, if a 2D AOD is used, then we can generate 2D arrays of optical spots whose positions and intensities are controlled by the rf signals, which are generated by the SDR. In this approach, the different focal spots will have slightly different optical frequencies due to the AOD deflection mechanism [145]. The second technique we have considered is to use a digital micromirror device (DMD). The DMD could be placed in either an object plane or Fourier plane relative to the objective focal plane to directly or holographically generate arbitrary optical potentials in the objective focal plane.

So far in the experiment we have used the high-NA system to realize high resolution fluorescence images of atoms and to generate 1D arrays of optical traps using an AOD. The rest of this chapter will be dedicated to explaining these two systems in further detail.

7.2 The Objective

The spacing of individual atoms in an optical lattice is comparable to the diffraction limited resolution for high-NA imaging systems, meaning reconstruction of the atomic density requires a large SNR in the detected image. Quantum gas microscopes overcame this challenge by leveraging simultaneous cooling and imaging techniques to allow for the collection of many photons from the single atoms. See Appendix E for details about how the diffraction limit bounds the ultimate resolution for imaging system. E6 was designed to include this capability for high resolution imaging of single atoms in conjunction with operation of the science cavity for cQED.

Many quantum gas microscopes utilize custom designed high-NA objectives to realize NAs of up to 0.8. In fact, one of the main challenges for realizing high-NA systems is the presence of the vacuum viewport, which can add appreciable aberration to the imaging system. One approach is to use an out-of-vacuum objective that is custom designed to compensate these aberrations from the viewport [41, 146]. Another approach is to suppress
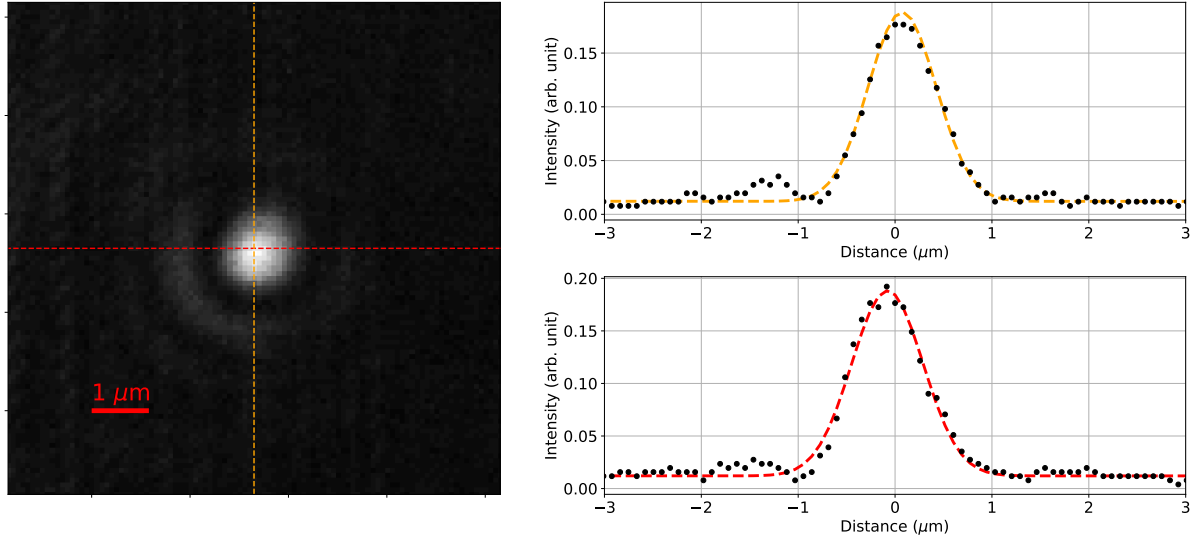
Figure 7.1: Measured point spread function (PSF) for the high-NA objective in a test setup. Here we are using the high-NA objective to measure the light coming out of a sub-micron hole in a test pattern. Line cuts with a fitted Gaussian are shown at right. The extracted Gaussian standard deviations are $\sigma = (386, 337)$ nm. The corresponding Rayleigh range spacing characterizing this PSF is then $\rho_{\mathrm{Rayleigh}} \approx 1050 \, \mathrm{nm}$.

these aberrations by placing the first high-NA optic in-vacuum so that the bundle of rays passing through the viewport has low-NA and is not heavily aberrated by the viewport [40]. However, this approach introduces the difficulty that this first optic cannot, then, be easily re-aligned.

For the first generation version of the E6 apparatus, we elected to save time and money by using an out-of-vacuum commercial objective with NA = 0.5 (Mitutoyo G Plan Apo 50x Objective). The objective has an effective focal length of 4 mm and a working distance of about 15 mm. This objective was designed to image small electronics through a thin glass layer such as the pixels in a flat screen television monitor. It is glass thickness corrected, meaning that it can produce diffraction limited images in the presence of a flat 3.5 mm thick glass viewport. The re-entrant viewports on the science chamber were designed with thickness for compatibility with this objective. The NA = 0.5 objective has a theoretical Rayleigh criterion separation distance of $\rho_{\mathrm{Rayleigh}} = 0.61 \frac{\lambda}{\mathrm{NA}} \approx 950 \, \mathrm{nm}$.[1]

---

[1]See Appendix E for more details. Note that just because two object are separated by less than ρ_{Rayleigh} doesn't necessarily mean they cannot be resolved by the imaging system. If the two objects are imaged with arbitrarily high SNR then their spacing could always be calculated by careful examination of the shape of the combined diffraction limited spots (e.g. by performing a deconvolution with the imaging system PSF). Conversely, if spots are separated by more than ρ_{Rayleigh}, but with very poor SNR, then it may be impossible to spatially resolve them.

See Fig. 7.1 for an out-of vacuum characterization of the microscope objective. This test was performed by postdoc Johannes Zeiher and undergraduate Rachel Tsuchiyama. The upper viewport was suspended with the objective inside as it would be in the experiment. 780 nm was light was shone upwards through a test pattern (Star Test TC-ST01, Technologie Manufaktur) with a sub-micron hole and the resulting PSF was imaged through the objective in combination with an $f = 300$ mm tube lens for a magnification of (75x). The PSF was fit with a Gaussian to extract an approximate effective Rayleigh criterion spacing of 1050 nm, just slightly larger than our expectation of 950 nm.

To align the objective to the vacuum chamber viewports, we used the following procedure. We first aligned a green laser guide beam passing normally upwards through the two re-entrant viewports. We centered this beam through the chamber using two temporary custom 3D printed irises that we bolted to the 8" bolt washers. We then set up another beam counter-propagating overlapped with this guide beam. After this, we installed the objective within the upper re-entrant viewport. The objective is mounted onto a tip/tilt kinematic mount, which is mounted on a 3-axis translation stage. The axial direction of this translation stage is motorized. The base of the translation stage is attached to the 8" bolt washer via an adapter piece. We translated and adjusted the tip and tilt of the objective until both the beam from below and the beam from above were perfectly retroreflected. We mechanically estimated the appropriate vertical position and this got us close enough to see first signals for imaging that we could then use to further optimize the vertical position. Further alignment of the objective would require optimization of the PSF of an imaged point source, i.e. a single atom.

7.3   Optical Microtweezers

Typical ODTs in our labs used to trap bulk atomic gases have focus sizes w_0 on the order of 10s of μm and trap depths from $10\,\mu$K to a few mK and trap atoms numbers from 10^3 to 10^7 atoms. However, in recent years, many research programs have found success working with ODTs in the microtweezer regime with waists at the sub-micron scale and in which single atoms are trapped [46, 49, 50].

Platforms based around arrays of microtweezers have exciting applications in quantum simulation and quantum information [147, 148]. The single atoms serve as well-isolated, two-level systems that can be combined into larger quantum systems using reconfigurable microtweezer arrays. Historically, Rydberg interactions have been used to induce interactions between the atoms trapped in microtweezers. In designing this apparatus, we had been curious if we could facilitate interactions between atoms using the optical cavity.

Single atoms are trapped in microtweezers by leveraging strong two-body light-assisted collisions, which occur at a high rate when atoms are trapped in a microtweezer at a very high density and are exposed to near-resonant red-detuned light. Under typical circumstances, these collisions lead to a loss of both atoms, though schemes have been demonstrated in which only one atom is lost from the trap [52]. If, initially, an even number of atoms are loaded

Figure 7.2: (33x) focused laser spots generated by imaging the AOD output using the high-NA objective. Here we see some variation in the intensity of the the different tweezer spots. In an experiment this intensity variation would need to be calibrated out and compensated for.

into the trap, then all atoms are lost after a short time in the presence of photo-association light that induces these collisions. If an odd number of atoms are initially loaded, then, after all other atoms are lost pair-wise, one single atom will remain. The photo-association light, therefore, projects the microtweezer atom number onto its parity, resulting in microtweezer occupations of zero or one atoms with approximately equal probability. Single atoms can then be prepared statistically with 50% probability. Sub-Poissionian loading of atoms, and indeed loading of single atoms, into such microtweezers was demonstrated as early as 2001 [46, 149].

Excitingly, it is also possible to image single atoms within these microtweezers using the same photo-association light and using the same objective that is used to generate the tight microtweezer beams. We use a simple scheme to image the atoms in which red-detuned light is put into a $\sigma^+ - \sigma^-$ molasses configuration [100] whereby the atoms simultaneously undergo light-assisted collisions, molasses cooling, and fluorescence scattering. The scattered light is collected using the high-NA objective and imaged onto a sensitive CCD or CMOS camera sensor while the molasses cooling ensures many photons can be scattered prior to atom loss to ensure high SNR detection of individual atoms.

7.3.1 AOD Re-imaging system

While the ability to readily prepare single trapped atoms opens the door to the exploration of interesting fundamental cQED physics, we are interested in using the cavity to engineer and monitor the dynamics of many well-controlled quantum systems, not just a single two-level system. To this end we are interested in an array of microtweezers.

We utilize an 808 nm volume holographic grating (VHG) diode laser to generate our microtweezer light. The optical path for this light is shown in Fig. 7.3. After passing through an AOM for intensity stabilization and fast switching, this light is routed via an optical fiber onto the upper breadboard above our science chamber for injection into the high-NA objective. After exiting the fiber, the 808 nm wavelength light passes through the AOD[2] (AA Optoelectronics DTSX-400-780) to be deflected into multiple beams. These

---

[2]An AOD is an acousto-optic device, identical, in principle, to an AOM. The practical difference is

beams are then relayed via a 4f imaging system, utilizing a pair of 300 mm focal length lenses, onto the back focal plane of the microscope objective. The objective then Fourier transforms the large angularly deflected beams into small spatially separated spots in the atomic plane. The AOD and back focal plane of the objective serve as Fourier planes in this imaging system. The atomic plane and an intermediate object plane in the center of the 4f relay system serve as object planes.

At the time of writing, the AOD is mounted such that it generates a 1-D array of microtweezers which is transverse to the science cavity's optical axis as shown in Fig. 8.6. However, the AOD is mounted on a rotation stage which could be rotated by 90° so that 1-D array is oriented along the cavity axis. In the future, we will install a 2D AOD so we can create lattices oriented along either or both axes.

When the AOD is driven with rf frequency f, the acousto-optic crystal sustains a traveling acoustic wave, transverse to the optical propagation axis, with an acoustic wavelength given by $\lambda_A = v_A/f$ where v_A is the speed of sound in the AOD crystal. In our case the crystal is TeO$_2$ with $v_A = 650$ m/s. The optical beam scatters off of this standing wave resulting in a duplicate copy of the original beam but offset in spatial frequency from the original beam by $\nu = \frac{1}{\lambda_A}$. This offset in spatial frequency ν corresponds to an angular deflection of $\theta = \lambda_O \nu = \lambda_O/\lambda_A$ where $\lambda_O = 808$ nm is the optical wavelength.

Our AOD is specified to have an angular deflection range of $\Delta\theta = 43$ mrad corresponding to a range in spatial frequency of $\Delta\nu = 53$ mm^{-1}. The AOD has an aperture large enough to support beams with waists as large as 4 mm. Two beams with a waist of $w_0 = 4$ mm must be separated by at least $\delta\nu = \frac{1}{w_0} = 0.25$ mm^{-1} to ensure the beams are resolved, in some loose sense, from each other. By these metrics the number of resolvable spots achievable with our AOD is approximately $\Delta\nu/\delta\nu \approx 212$. This is sufficient spatial frequency dynamic range for any experiments we are imagining at this time.

To separate two beams by $\delta\nu = 0.25$ mm^{-1} in spatial frequency space the driving tones must be separated in rf frequency by $\delta f = v_A \delta\nu \approx 160$ kHz. In such a case where we are pushing the resolution limits of the AOD these microtweezer spots will have some spatial overlap. In the regions the two microtweezers will interfere resulting in temporal intensity fluctuations at the $\delta f = 160$ kHz beat frequency between the two microtweezers. These fluctuations could result in heating of the atoms if this frequency is close to twice the tweezer trap frequencies, which is in fact quite possible given the microtweezer parameters that we choose in this work. These problems are suppressed as the microtweezers are spaced further apart. Furthermore, to date, we have only worked with up to about a dozen microtweezers, so we are nowhere near pushing the resolution limits discussed here.

Following Ref. [150] and others, we targeted microtweezer trap depths on the order of 1 mK for our ^{87}Rb atoms. As will be explained below, atoms are cooled within these microtweezers using an optical molasses which, under optimal conditions and in absence of the microtweezers, typically cools atoms to a temperature of 10s of μK. The trap depth of 1 mK

that an AOD is specially optimized to realize large angular deflections and to produce a large number of 'resolvable spots,' while an AOM is specially optimized for a high bandwidth and fast switching time.

is a couple of orders of magnitude hotter than this giving us some overhead to ensure that the laser-cooled atoms stay trapped, with long lifetimes, in the microtweezers.

For 808 nm light focused to $1\,\mu$m, this trap depth requires about $1\,$mW per tweezer. This also corresponds to theoretical radial (axial) trap frequencies of $93\,$kHz ($17\,$kHz). With up to $100\,$mW of power before the AOD, and taking into account the AOD and path efficiencies, we have been able easily generate up to 10 microtweezers with the desired trap depth. In Sec. 8.2 I will show images in which we generate and trap atoms in 8 tweezers. Note that we currently estimate the trap depth and frequencies for the microtweezers from our estimates for the beam-size and power. In the future, as in Ref. [150], we will more carefully measure trap depths using atom-loss spectroscopy on the cycling transition and we'll measure trap frequencies using parametric heating.

We drive the AOD using a PCI arbitrary waveform generator or SDR (M4i6631-x8 from Spectrum Instrumentation Corp). Undergraduate Aron Lloyd wrote a python package to allow us to easily send a number of calibrated tones to the AOD to generate numerous spots with tunable intensity by generating waveforms that are superpositions of multiple tones. With this software, we are able to generate waveforms for which the frequency of one or more of the tones is swept during the experimental cycle, allowing us to translate the atoms by moving the tweezer positions. In the future, we will be able to use this dynamic positioning ability to image atoms within our microtweezer array and reconfigure the array conditioned on the results of that imaging process [49, 50].

See Fig. 7.2 for an image of 33 spots generated in the atomic plane by the AOD. We image the microtweezers using a ($5\times$) magnification imaging system setup using an NA=0.2, diffraction-limited lens in the lower re-entrant viewport. See Fig. 7.3. The optical pattern generated in the atomic plane is imaged onto a CMOS camera mounted below the experiment on a small breadboard.
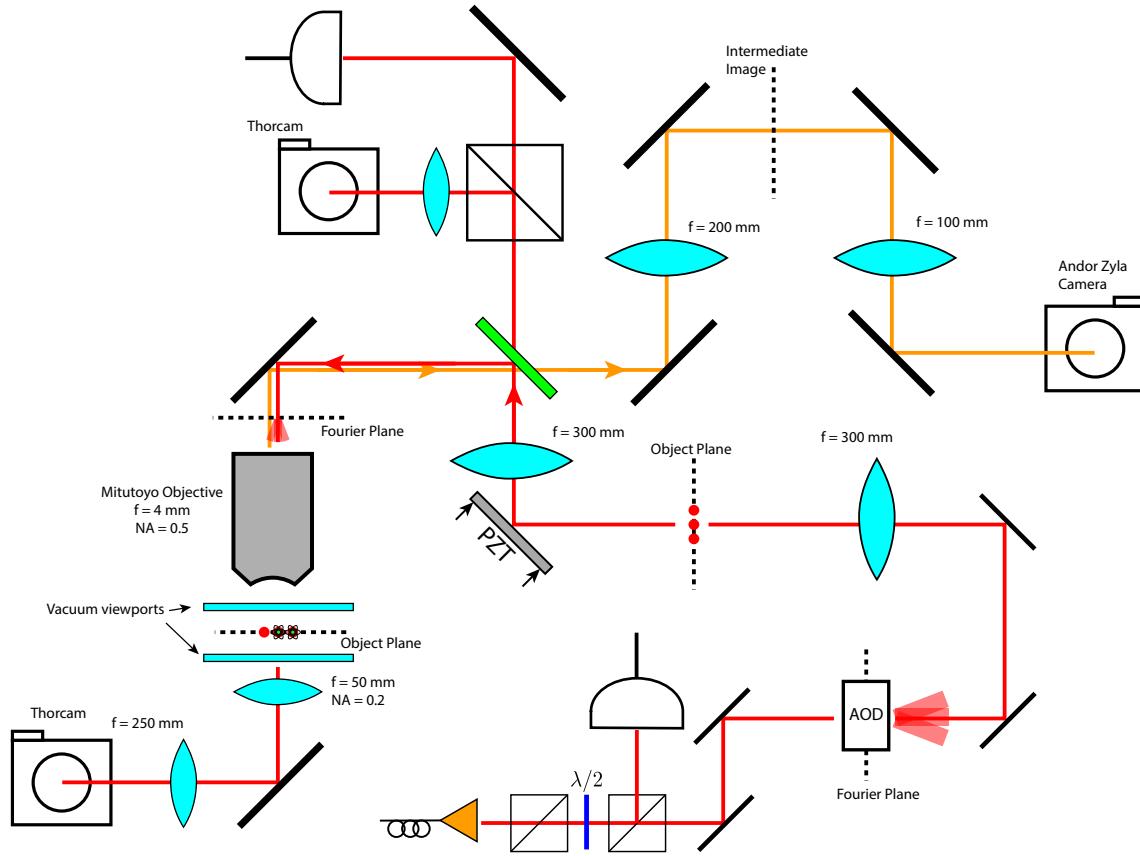
Figure 7.3: Optics scheme for the high-NA imaging and addressing system.

7.4 Fluorescence Imaging and Molasses Cooling

To achieve fluorescence imaging of the atoms, we illuminate the atoms within the microtweezers with an optical molasses in the $\sigma^+ - \sigma^-$ configuration. This molasses cools the atoms and causes them to scatter photons towards the objective. When this molasses is first turned on it also serves as the photo-association light and causes the parity projection of the atom number within the microtweezers.

7.4.1 Optical Molasses

Optical molasses laser cooling works due to an imbalance of optical forces arising from optical pumping and light shifts that affect moving atoms asymmetrically depending on the direction of the velocity vector [99, 100]. Energy dissipation and entropy reduction arises due to spontaneous emission during optical pumping. For large detunings, the scattering

rate scales with $\propto \frac{\Omega^2}{\Delta_{\text{mol}}^2}\Gamma$ and, for a perfectly balanced molasses, the final temperature scales with the light shift $\propto \frac{\Omega^2}{\Delta_{\text{mol}}}$. Here Δ_{mol} is the molasses detuning from the unshifted D_2 cycling transition and Ω is the Rabi frequency for the molasses light. Imperfections due to power imbalances, imperfect polarizations, or especially due to a dc magnetic field will limit the minimum temperature achievable by the molasses.

We operate the optical molasses in our science chamber with only two retroreflected molasses beam pairs that intersect at a 45° angle. One beam enters along the side-imaging axis (perpendicular to optical transport) in the science chamber and one enters along the cross-axis (perpendicular to the science cavity). The pair of beams entering on the cross-axis are tilted slightly out of the atomic plane of the vacuum chamber so that their wave-vector has a non-zero component along the vertical axis. We did this to ensure that we are able to cool along the vertical direction which is one of the principal components of the microtweezer harmonic trap.

We initially optimized the molasses in the science chamber by first zeroing the magnetic field by performing microwave spectroscopy of atoms in the transport ODT. We used a microwave horn to send microwaves into the chamber through a science chamber cross-axis viewport. We zeroed the magnetic field so that the ground state Larmor frequency $\omega_L = (g_F\gamma_C)B < 2\pi \times 200\,\text{kHz}$ which corresponds to a temperature of about $10\,\mu\text{K}$. $g_F = 1/2$ is the ground-state Landé g-factor for the hyperfine ground state we were using for spectroscopy and $\gamma_C = |e|/2m_e = 2\pi \times 1.4\,\text{MHz/G}$ is the classical gyromagnetic ratio.

We then optimized the polarization and alignment of the two molasses beam pairs, first separately, and then together, by monitoring the atoms from the transport ODT in TOF. In the end, this was enough to allow us to see first signals for fluorescence imaging. We first saw fluorescence images of atoms in the transport ODT, then of a density enhancements of atoms within the transport ODT in the presence of the microtweezer light, and finally of single atoms trapped in the microtweezers in absence of the transport ODT. These results will be shown in Sec. 8.2.

The detuning Δ_{mol} is very important for optical molasses within the tweezers. For example, if Δ_{mol} is very small, the molasses equilibrium temperature may exceed the tweezer trap depth leading to short atom lifetimes. When determining the detuning for the molasses light, we must recall that the effective molasses detuning for atoms within the microtweezer may differ by up to 10s of MHz from the bare detuning because of the Stark shift from the microtweezer light.

In a system with perfect polarizations, beam alignment, and no external fields, the minimal temperature achievable in an optical molasses would be, up to a prefactor, the recoil temperature: $k_B T_R = \hbar^2 k^2/2m \approx k_B \times 180\,\text{nK}$. In practice, however, residual vector shifts of the atomic energy levels can limit the minimum achievable temperature. For example, for magnetic fields slightly larger than the ideal molasses temperature, the rapid Larmor precession can spoil the optical pumping required for the sub-Doppler cooling leading to a final temperature related to $k_B T \approx \hbar\omega_L$. For large enough vector shifts the sub-Doppler cooling may be spoiled entirely so that the minimum achievable temperature is in fact only

the Doppler temperature $k_B T_D = \hbar\Gamma/2 \approx k_B \times 140\,\mu\text{K}$. For detunings $\Delta_{\text{mol}} \gg \Gamma$ the Doppler limit is proportional instead to Δ which is an even worse temperature limit [99].

Because of beyond-the-paraxial-approximation polarization effects for the tightly focused microtweezer beam with $w_0 \approx \lambda$, the polarization of the light at the focus may have a non-zero and spatially varying ellipticity that contributes to a spatially varying vector shift within the trap volume [151, 152]. The optical molasses could be spoiled by this pseudo magnetic field gradient in addition to real residual magnetic fields. However, while we're aware of this effect, preliminary calculations performed by postdoc Fang Fang indicate that, for trap depths and waists we are working with, these pseudo magnetic fields should not be limiting.

We find our molasses to operate well for tweezer depths of approximately $500\,\mu\text{K}$, with detunings from the bare transition of $\approx -20\,\text{MHz}$ ($\approx -10\,\text{MHz}$ detuned from the Stark shifted transition) and saturation parameter $s = \frac{I}{I_{\text{sat}}} = 2\frac{\Omega^2}{\Gamma^2} \approx 10$.

7.4.2 Fluorescence Imaging Microscope System

See Fig. 7.3 for a schematic layout of the fluorescence imaging system. The Mitutoyo objective we use is infinity corrected, meaning that it exhibits minimal aberration[3] when the sample is located at the focal plane of the objective within a tolerance of about 10% of the focal length or working distance. In this configuration, light from small sources in the sample plane is collimated by the objective. We place a 200 mm 'tube' lens about 400 mm away from the objective that images the light down to an intermediate image plane with a (50x) magnification.[4] We then re-image this intermediate plane onto our camera sensor using a single $f = 100\,\text{mm}$ lens in a $2f$ configuration with nominal unity magnification.

Our camera is an Andor Zyla 4.2 sCMOS camera. We chose this sensor because it had good quantum efficiency around 780 nm and very low read noise. These specifications allow us to take shot-noise limited images even at very low light levels. We place a number of optical filters in front of the detector that pass 780 nm light but block stray 808 nm tweezer light.

In the following, I provide an estimate for the photon collection efficiency of our imaging setup. Given $\Gamma \approx 2\pi \times 6\,\text{MHz}$ for resonant driving of the Rb atoms, we can conservatively estimate that in our detuned molasses an individual atom will scatter at a rate of $\approx 1\,\text{MHz}$. An objective with a given NA subtends a solid angle of $\Omega_{sr} = 2\pi(1 - \cos(\theta)) = 2\pi\left(1 - \sqrt{1 - \text{NA}^2}\right)$ where Ω_{sr} is expressed in steradians. Our NA $= 0.5$ objective then

---

[3]See Sec. E.3 in Appendix E for more details.

[4]Because we are working in an infinite conjugate setup, can we place this tube lens arbitrarily far away from the objective as is convenient? In theory, yes. In practice, no. In practice, if the tube lens is very far from the objective, then rays coming from point sources off of the optical axis will be traveling at an angle to the optical axis in the infinity space. These rays will be clipped by finite sized optics eventually leading to an effect where objects at the edge of the field-of-view appear darker, or are clipped entirely. This is called vignetting and limits the length of an infinity corrected microscope. It is not an issue for the setup described here.

collects about 6.6% of light from a single atom, assuming the light is emitted isotropically.[5] The commercial Mitutoyo objective was not optimized for 780 nm light, so it only transmits 70% of light at this wavelength. We can conservatively estimate a path efficiency throughout the rest of the optical path of at least 60% and the sensor itself has a quantum efficiency of about 50%, resulting in a total collection efficiency of about 1.4%. From this, we estimate that photons from a single atom are detected at a rate of $\approx 14\,\text{kHz} = 14\,\text{ms}^{-1}$. Given the imaging exposure time, we can then estimate how many photons we would collect from a single atom per exposure. Typical exposure times are 10s to 100s of ms.

I will show measured fluorescence images of single and multiple single atoms in Sec. 8.2.

---

[5]The atom may, in fact, emit light in a dipole pattern that could enhance or suppress this factor. The isotropic assumption, however, suffices for the crude order-of-magnitude calculation we are doing here.

Chapter 8

Experimental Results

In this chapter, I will present key experimental results we've obtained with this apparatus that represent the culmination of the multi-year design and construction process described in the preceding chapters. I will present (1) an observation of a dispersive shift of the cavity resonance frequency due to the presence of atoms trapped in the cavity ODT, (2) the observation of single atoms trapped in microtweezers, and finally (3) the observation of a variation in the single atom fluorescence as a function of the sub-wavelength positioning of the microtweezer with respect to the cavity ODT standing wave within the optical cavity. These results demonstrate that both the optical cavity and high-NA microtweezer systems are online and operational, and that we are able to see an effect that relies on their simultaneous operation. These results open the door to the further exploration of cQED with a locally addressable atomic sample.

8.1 Atomic Dispersive Shift of the Cavity Resonance

In Sec. 1.3.3, I demonstrated that, in the dispersive regime of cQED, the cavity resonance frequency is shifted by $g_C = g^2/\Delta_{CA}$ for each atom within the cavity. We have observed the shift of the cavity resonance due to the presence of atoms in the cavity probe mode using our optical heterodyne detector. The measurement results are shown in Fig. 8.1. For this experiment, about $N_A \approx 800$ atoms were transferred from the transport ODT into the cavity ODT at a trap depth between $100\,\mu\mathrm{K}$ and $200\,\mu\mathrm{K}$. The probe frequency was then swept across cavity resonance by tuning the probe VCO frequency (with the LO frequency following at a fixed detuning of Δf_{LO} because of the probe/LO phase lock described in Sec. 6.3.1). The sweep time was about $2\,\mathrm{ms}$, spanning about $15\,\mathrm{MHz}$. During this time, the heterodyne beatnote was recorded using a PCI oscilloscope (GaGe CompuScope 14200).

After the cavity sweep, the atoms were released from the cavity ODT and an absorption image was taken as shown in Fig. 8.1(a). Finally, with the cavity length still locked, the probe was again swept across cavity resonance and the heterodyne beatnote was recorded giving us an empty cavity control signal.
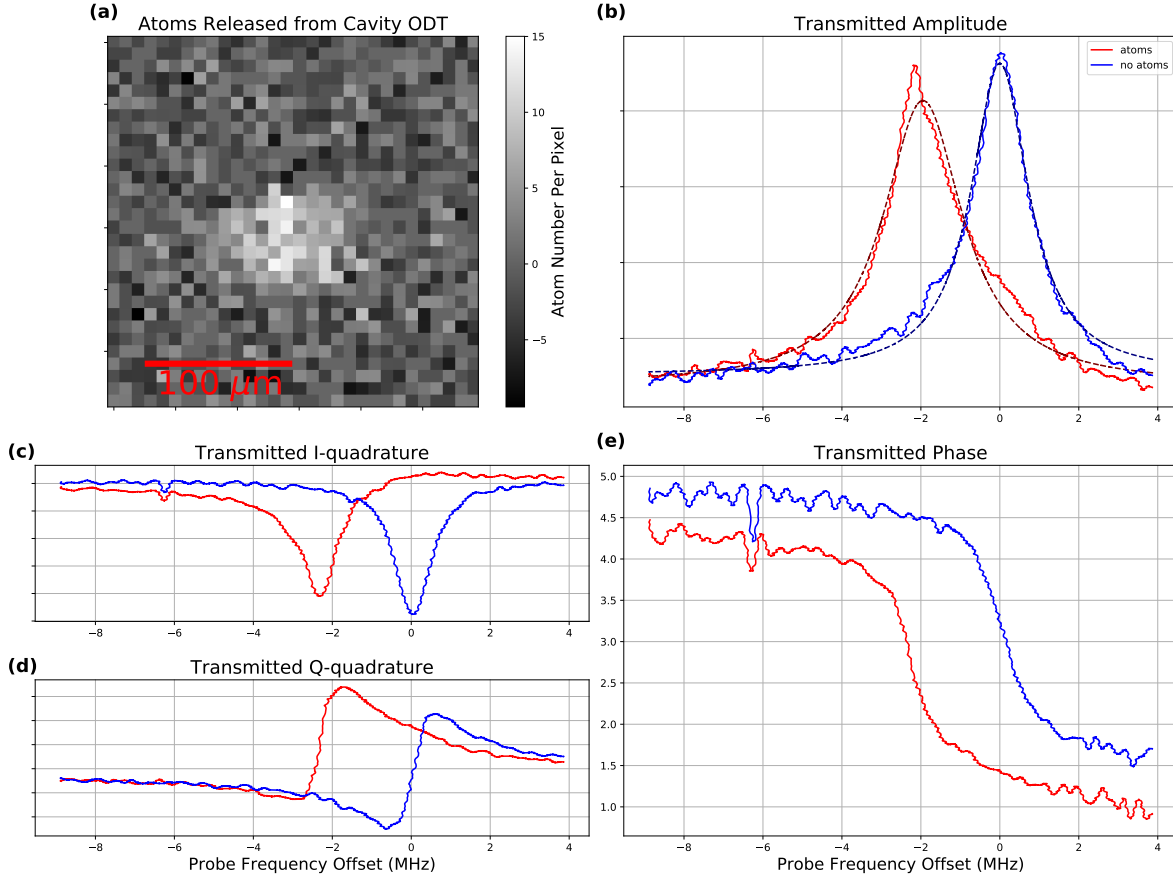
Figure 8.1: Single shot of an experimental run demonstrating cavity sweep measurements and absorption imaging of atoms trapped in the cavity ODT. (a) Absorption image of atoms after release from the cavity ODT. A Gaussian fit to the image yields an estimated atom number of $N_A = 800$. (b) Amplitude of the demodulated heterodyne signal. Note the displacement between the two distinct transmission peaks with and without atoms. Dashed lines are Lorentzian fits to the transmission peaks. (c), (d) I- and Q-quadrature for the demodulated heterodyne signal. (e) Phase of the demodulated heterodyne signal. Note the distinctive dispersive shape. Because the phase varies rapidly with detuning, it could also be used as a sensitive discriminator to estimate the cavity dispersive shift.

We carried out this experiment with the atoms unpolarized in the $F = 1$ hyperfine ground state and the probe light linearly polarized and red detuned by about $\Delta_{CA} \approx -5.6\,\text{GHz}$. The appropriate transition dipole element to use from Eqs. B.6 in this case is $d_{\text{eff}} = \left(\sqrt{2/3} \right) d_0$. This leads to an expected $g_{\text{eff}} = \left(\sqrt{2/3} \right) g_0 = 2.5\,\text{MHz}$.

We can easily extract the cavity resonance shift by looking at cavity transmission profiles with and without atoms, as shown in Fig. 8.1(b). We see that the cavity is shifted by about -2 MHz for 800 atoms. This corresponds to a single atom shift of $g_C \approx -2.5\,\text{kHz}$. From this, we can extract $g_{\text{eff}} \approx 3.7\,\text{MHz}$. Surprisingly, we measure a coupling coefficient larger than our prediction based on estimates of other experimental parameters. We attribute this discrepancy to underestimation of the atom number or a misunderstanding of the atomic state or optical polarization.

This measurement demonstrates the first interaction between atoms trapped within the cavity ODT and the cavity probe photons, giving us our first toehold to begin to explore cQED physics. In the following subsections, I will describe how this measurement is affected by the finite size of the atomic cloud and two future directions to extend this measurement.

8.1.1 Finite Cloud Size

In the previous section, I assumed that the atoms were all point-like and located at the anti-nodes of the probe so that each experienced a coupling of g_{eff}. However, recall from Eq. (1.19) that the coupling strength depends on the local magnitude of the electric field. The finite size of the atomic cloud within the spatially varying cavity mode will then lead to a spatial averaging and reduction of the total effective dispersive shift.

I will calculate the resulting dispersive shift, quantified by $g_{C,\text{eff}}$, for a Gaussian atomic cloud with N_a atoms and with variances along the principal trap axes of $\sigma_{x,y,z}^2$ and that is centered at $(x, y, z + \delta z)$ within an optical lattice with a Gaussian waist of w_0 and wavelength $\lambda = 2\pi/k$. The atomic density distribution is given by

$$\rho(x, y, z) = \frac{N_a}{(2\pi)^{3/2} \sigma_x \sigma_y \sigma_z} e^{-\frac{1}{2}\left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + \frac{z^2}{\sigma_z^2} \right)}. \tag{8.1}$$

The spatially varying coupling coefficient in the cavity is given by

$$g(x, y, z) = g_0 e^{-\frac{x^2 + y^2}{w_0^2}} \cos(kz). \tag{8.2}$$

The total dispersive shift as a function of position along the cavity axis will be given by

the convolution of the atomic and photonic distributions:

$$
\begin{aligned}
N_A g_{C,\text{eff}} &= \int \frac{g(x,y,z)^2}{\Delta_{CA}} \rho(x,y,z-\delta z) dx dy dz, \\
&= N_A g_C \frac{1}{2} \frac{1 + e^{-2k^2\sigma_z^2}\cos\left(2k^2\delta z\right)}{\sqrt{\left(\left(2\frac{\sigma_x}{w_0}\right)^2 + 1\right)\left(\left(2\frac{\sigma_y}{w_0}\right)^2 + 1\right)}}.
\end{aligned}
\tag{8.3}
$$

We see that the spatial extent of the cloud in the axial direction σ_z corresponds to a reduction in the expected contrast as a function of δz while the extent of the cloud in the radial direction, $\sigma_{x,y}$ leads to an overall reduction in the average magnitude of the shift.

For atoms trapped in the 1560 nm cavity ODT, using Eqs. (6.8), $\eta \approx 10$, and $w_{0,1560} = 24\,\mu\text{m}$, we can estimate $\sigma_x, \sigma_y = 3.8\,\mu\text{m}$ and $\sigma_z = 56\,\text{nm}$. The axial contrast is given by $\exp\left(-2k^2\sigma_z^2\right) = 0.67$. The reduction due to radial averaging, using $w_{0,780} = 17\,\mu\text{m}$, is given by

$$
\frac{1}{\left(2\frac{\sigma_{x,y}}{w_0}\right)^2 + 1} = 0.83.
\tag{8.4}
$$

8.1.2 Measure the Cavity Beatnote

One next step for this measurement would be investigating the probe/cavity ODT beatnote, which was described in Sec. 3.6. In that section, we learned that, ignoring mirror phase effects, if the cavity beatnote is maximized in the center of the cavity for one set of probe and cavity ODT modes, then the probe is driving an odd longitudinal mode. If (1) the cavity length is shifted by $\lambda_{780}/2$, shifting the probe by one FSR to an even longitudinal mode, and (2) the cavity ODT is shifted in frequency to be resonant with the shifted cavity resonance, then we would expect the beatnote to be zero in the center of the cavity and (neglecting the finite cloud size) to observe no dispersive shift.

In general, due to mirror phase effects, it may not be the case that the beatnote is maximized or minimized in the center of the cavity. To investigate the beatnote, we could load the atoms in different locations along the cavity axis by adjusting the alignment of the transport ODT (within the constraints set by the transport ODT shielding mirrors) and repeat this dispersive shift measurement of g at the trap anti-nodes along the cavity axis. We could also repeat this experiment driving the cavity on different longitudinal modes for both the trap and probe. Such measurements might allow us to extract the relative position of the two cavity lattices and estimate the relative mirror reflection phases between the two wavelengths. Such a measurement is closely related to the measurement of the relative lattice positions performed in Ref. [104], in which a scanning tip was used to induce position dependent losses in 780 and 1560 nm modes in a bichromatic high-finesse cavity.

8.1.3 Single Atom Measurement Sensitivity

Another interesting direction for this dispersive cavity shift measurement would be to improve the atom number sensitivity to single-atom precision. This would be a first step towards using the cavity to analyze single atoms trapped in microtweezers. The sensitivity of this measurement depends on how sensitively the shift in cavity resonance, g_C, can be measured as a fraction of the cavity linewidth κ. The SNR on this measurement is set by photon shot noise. Neglecting any sources of technical noise, the SNR can, then, always be improved with more detected photons by either increasing the photon number $\bar{n}$ or measurement time T. However, the presence of the near-resonant probe light will heat the atoms leading to atom loss. This atom loss will result in an increased atom number uncertainty [153].

There is, then, an optimal fluence $\bar{n}T$ when measurement noise from heating is equal to measurement shot noise and the measurement sensitivity is optimized. We know that the effect of heating noise is reduced when the atom is held in a deep trap, such as a microtweezer. We are interested to test this experiment in the lab to determine if our system is able to access a single-atom sensitive measurement regime.

Because different atomic states have different values for g, it would not be too difficult to extend this measurement to allow us to detect, possibly with high fidelity, the state of an individual atom within the cavity [154]. This would be an early step towards implementing and characterizing quantum gates or engineered Hamiltonians between different atoms held within the cavity.

It's worth noting that the discussion here focuses on attempting to use the dispersive shift of the cavity frequency to detect the presence of a single atom. Previous work has been done in the near-resonant cQED regime in which a single atom was detected via the vacuum Rabi splitting of the cavity mode [16, 107] or via enhanced fluorescence scattering into the cavity mode [154]. It would be interesting to assess and demonstrate the viability and relative advantages and disadvantages of the dispersive strategy for single atom detection.

8.2 Single Atoms Trapped in Microtweezers

Here I'll describe our results on trapping and imaging individual atoms trapped in multiple microtweezers.

To initially overlap the microtweezer beam with the transport ODT, we shone a deep microtweezer through the objective while atoms were trapped in the transport ODT and adjusted the microtweezer alignment until we noticed a density enhancement of the atoms in the transport ODT in fluorescence imaging, as shown in Fig. 8.2. In future alignments, we found it easier to coarsely overlap these two beams by monitoring the transport ODT in absorption imaging (from the side) and coupling into the microtweezer path either (1) near-resonant $F = 2 \to F' = 3$ light to blast away atoms when the beams are overlapped or
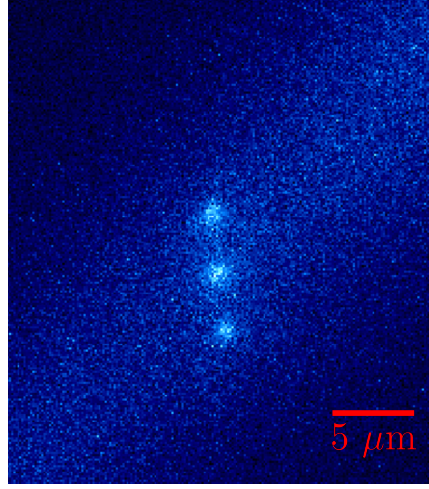
Figure 8.2: Density enhancement of multiple high power microtweezers attracting atoms within the cavity ODT. I call this a density enhancement rather than trapping is because we hypothesize that the brightness here is the exposure-time-averaged image of atoms moving in and out of the trap during the exposure (100s of ms) as a result of cooling and collisions as in [46]. However, in comparison to [46], the transport ODT out of which we load is orders of magnitude less dense than the molasses out of which those authors load their microtweezers so we would expect much longer times between atoms falling into and being kicked out of the microtweezer.

(2) repump $F = 1 \rightarrow F' = 2$ light to form a 'repump needle' that only illuminated atoms in the region near the microtweezer focus.

The central axis of the in-vacuum science cavity sets the absolute reference height for the experiment because we have no way to adjust it. We had previously overlapped the transport ODT with the cavity ODT to load atoms into the cavity ODT so we knew the transport ODT was at the height of the cavity axis. During fluorescence imaging, we adjusted the height of the objective to ensure the transport ODT was in focus, thus roughly overlapping the fluorescence imaging system focal plane with the cavity axis. Finally, to overlap the microtweezer focal position with the fluorescence imaging focal plane, we adjusted the position of the microtweezer fiber out-coupling lens to alter the collimation of the microtweezer beam entering the objective, thus translating the microtweezer focal height. We optimized this alignment to minimize the imaged size of the microtweezer density enhancement.

After overlapping the microtweezers with the transport ODT, we began efforts to show we could trap and image single atoms. A representative fluorescence image of a single atom trapped in a microtweezers is shown in Fig. 8.3(a). To realize single atom imaging, we tuned the frequency and intensity of the optical molasses and the microtweezer trap depth, bearing in mind the microtweezer light gives an AC Stark shift to the molasses transition [150]. We
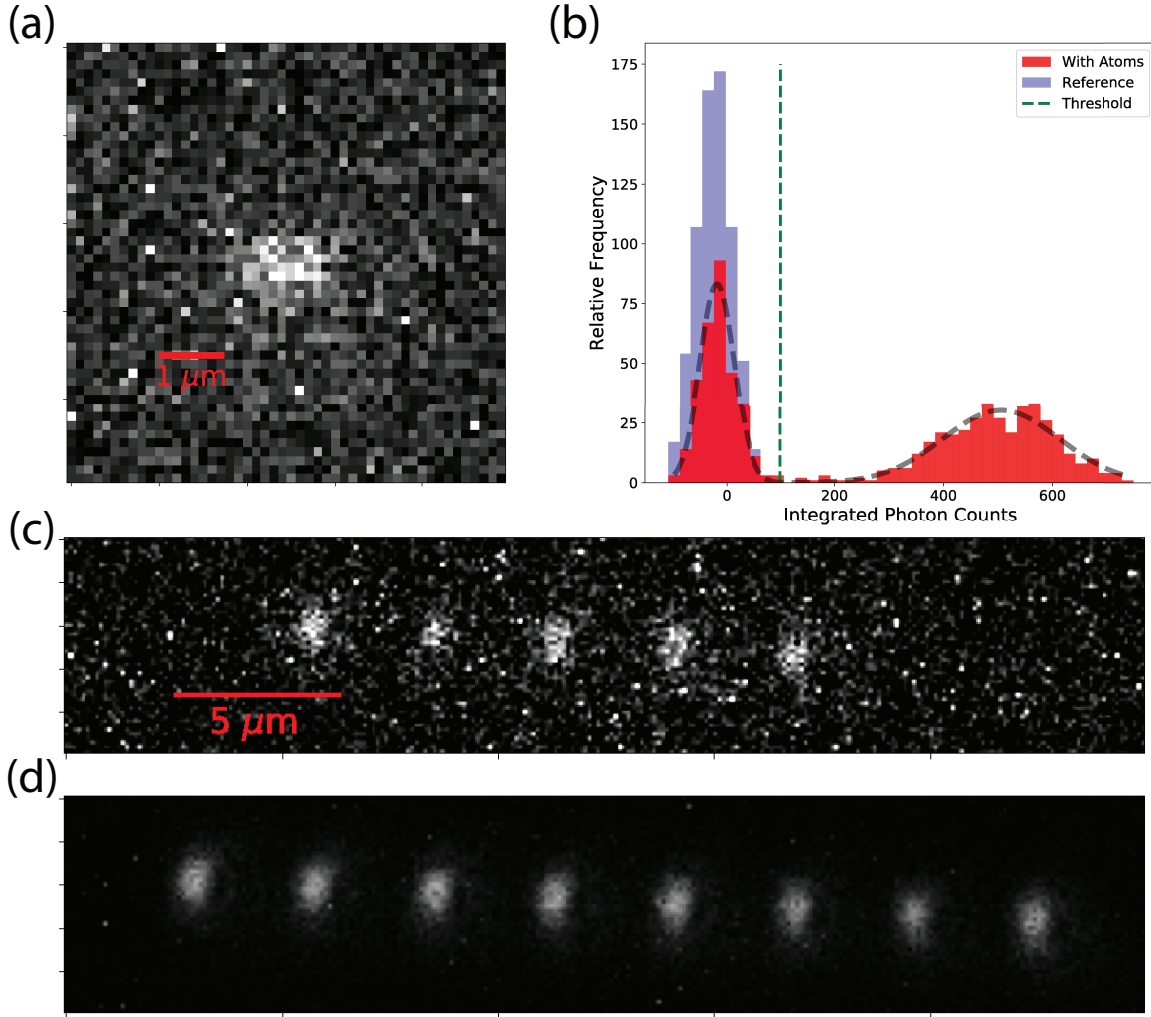
Figure 8.3: (a) Fluorescence image of a single atom trapped in a microtweezer. A Gaussian fit and our estimated object-plane pixel size yields the size of the spot to be $\sigma_{x,y} = (370, 546)\,\text{nm}$. This corresponds to a Rayleigh criterion spacing of $\rho_{\text{Rayleigh, x, y}} = (1.073, 1.58)\,\mu\text{m}$. The increased size in the y-direction may be due to aberrations from the objective that either increase the size of the tweezer (so that the atomic density has a larger distribution) and/or an increase in the size of the imaging PSF. (b) Histogram of integrated photon counts within a 20×20 grid of pixels encompassing the fluorescence region over about ~ 700 repetitions of the experiment. The bimodal structure clearly indicates shots with and without a single atom. A double Gaussian fit to the histogram yields that a threshold placed as in the figure would yield 99.99% fidelity in identifying whether or not a single atom is present in a particular frame. (c) Single image of 5 atoms statistically loaded into 8 tweezers. (d) Average of thousands of statistically loaded fluorescence images as in (c).

typically operate with a microtweezer trap depth of at least $500\,\mu\mathrm{K} \sim 10\,\mathrm{MHz}$. We find the molasses imaging to work for detunings Δ_{mol} of 10s of MHz and with an intensity saturation parameter of $s = I/I_{\mathrm{sat}} \approx 5 - 10$. An exposure time of order $100\,\mathrm{ms}$ yields good single atom images. As expected, we find, over some parameter regime, that the scattering rate (lifetime) increases (decreases) for larger saturation parameters and smaller molasses detuning.

Our experimental sequence begins with a 3D MOT and optical transport as described in Chapter 2. After optical transport, the microtweezers are loaded by turning on the microtweezer and ramping down the transport ODT. We then take a fluorescence image by turning on the molasses and triggering our camera to expose. We are able to take repeated fluorescence images in a single shot of the experiment by triggering the camera multiple times and turning on the molasses light when needed.

To prove we were seeing fluorescence from a single atom we analyzed the fluorescence counting statistics for our images. We repeated the above sequence hundreds of times and extracted the integrated photon counts within the fluorescence region for each shot. If our images, in fact, contained either 0 or 1 atoms, we would expect to see a bimodal distribution in these integrated photon counts. Indeed, as demonstrated in Fig. 8.3(b), we do very clearly see such a bimodal distribution. We can place a threshold discriminator within the histogram to allow us to classify whether an image contains an atom or not. We are free to place the threshold wherever we like but should choose it to try to reduce the number of false negatives and false positives reported by the classifier. The number of false positives (negatives) can be estimated by fitting the histogram with a double-Gaussian profile and determining the fraction of the weight of the zero (single) atom Gaussian above (below) the threshold. The false positive (negative) rate can be decreased by increasing (decreasing) the threshold value, presenting us, in general, with a trade off between these two problems. For the threshold shown here, the false positive and false negative rates are equal when threshold is chosen to be about 100 photocounts as shown in the figure and in this case the error rate is $\approx 10^{-4}$. That is, we have a 99.99% fidelity for classifying whether the image contains 1 or 0 atoms. As expected, we see that roughly half of the experimental shots contained an atom after parity projection, and half did not.

We are able to easily load multiple tweezers by driving the AOD with additional frequencies and increasing the input optical power. In Fig. 8.3(c), I show a single shot fluorescence image of 5 atoms, which have been statistically loaded into 8 microtweezers. Fig. 8.3(d) shows the average of thousands of such images showing that we are able to load all 8 microtweezers.

We characterized the lifetime of the single atom trapped in the microtweezer both only in the microtweezer and also in the presence of molasses light. We measure a dark hold lifetime of $\tau \sim 10.5\,\mathrm{s}$ and molasses lifetime of up to $\tau \sim 30\,\mathrm{s}$. The lifetime measurements are shown in Fig. 8.4. We take these data by capturing sequential $100\,\mathrm{ms}$ exposure time fluorescence images of the atoms with the molasses off (on) for the dark (molasses) lifetime measurement for the $10\,\mathrm{s}$ in the intervening time between exposures. In Fig. 8.4(b), we can see that the molasses lifetime decreases with decreasing molasses detuning. This is likely because, as Δ_{mol} is decreased, the equilibrium molasses temperature, which scales with s/Δ_{mol}, approaches the microtweezer trap depth and atom loss becomes likely.
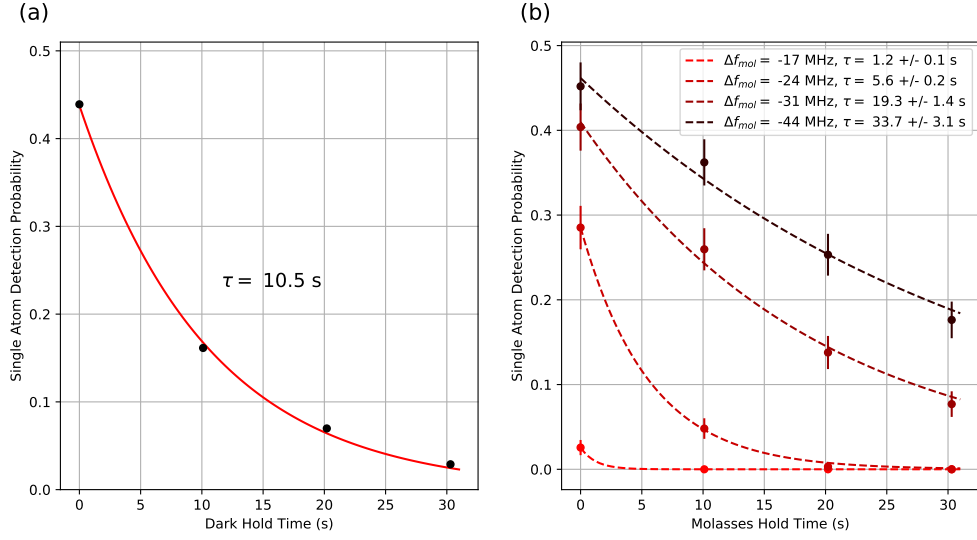
Figure 8.4: (a) Single microtweezer dark lifetime measurement. (b) Microtweezer lifetime in the presence of molasses light. Here $\Delta f_{\mathrm{mol}} = \Delta_{\mathrm{mol}}/2\pi$ indicates the detuning of the molasses light from the bare cycling transition. We roughly estimate an additional detuning due to the microtweezer AC Stark shift of $\approx -20\,\mathrm{MHz}$ for a $500\,\mu\mathrm{K}$ deep microtweezer.

8.3 Cavity ODT Contrast Measurement with Microtweezer

To demonstrate our ability to manipulate the atoms using both the high-NA system and the cavity simultaneously, we investigated the effect of the cavity ODT on the fluorescence photon count statistics. Recall in Sec. 6.4.3 that the 1560 nm cavity ODT light puts a large light shift on the $5\,\mathrm{P}_{3/2}$ excited state, which is involved in optical molasses. This light shift serves to lower the $5\,\mathrm{P}_{3/2}$ energy state so we expect the red-detuned molasses light to come closer to resonance. In the simplest picture, we can consider this light shift to act as a local adjustment to the molasses detuning, which we expect to alter the molasses scattering rate, temperature, and lifetime. Because the magnitude of the cavity ODT varies rapidly along the cavity axis, the local effective detuning has strong gradients that may lead to more complicated dynamics, or possibly heating. We will ignore such effects here and leave them as a topic for future explorations.

In this simple picture, we expect the presence of the cavity ODT to increase the scattering rate, and thus fluorescence brightness, for the atoms held in the microtweezer. This is indeed what we observe in Fig. 8.5(a), where I show the integrated photon counts as a function of the estimated cavity ODT trap depth for a range of different bare molasses detunings Δf_{mol}.

In Fig. 8.5(b) the horizontal axis for each data point is rescaled according to our estimate
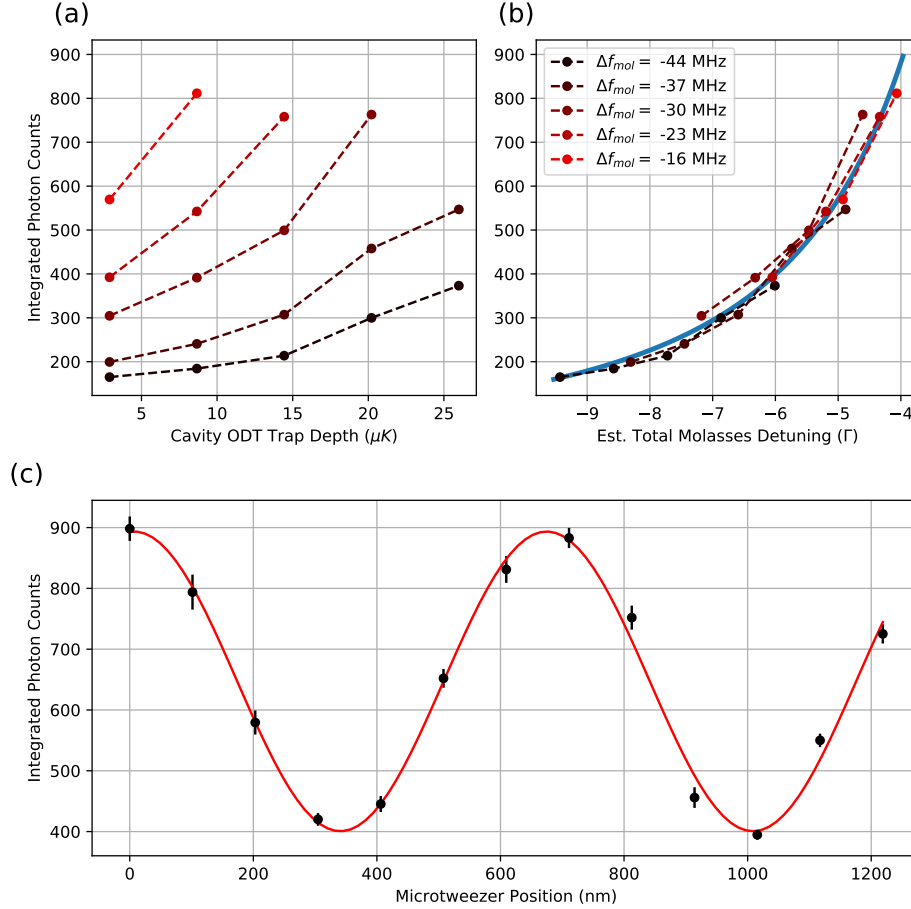
Figure 8.5: (a) Average integrated photon counts for atoms held in microtweezers in the presence of the cavity ODT as a function of cavity ODT trap depth and for various bare molasses detunings Δf_{mol}. (b) Same data as in (a) with the horizontal value for each data point adjusted to our estimated total molasses detuning for that data point according to Eq. (8.5). The blue line is a fit to a function proportional to $\Delta_{\mathrm{mol}}^{-2}$. The cavity ODT and microtweezer trap depths have been rescaled according to their respective fudge factors as described in the main text. Note that, for these data, it is not known if the atoms were held precisely in the radial center of the cavity ODT mode or if they were held axially at a node, anti-node, or somewhere in between. (c) Microtweezer and cavity ODT contrast measurement. The microtweezer is positioned at different positions along the cavity ODT axis and fluorescence images are taken at each position. We extract the average integrated photon counts and observe an expected oscillatory behavior in the single atom brightness as a function of microtweezer position. The red line is a sinusoidal fit to the data.

for the total molasses detuning for that point:

$$\Delta f_{\text{tot}} = \Delta f_{\text{mol}} + \Delta f_{\text{microtweezer}} + \Delta f_{\text{cavity ODT}}. \tag{8.5}$$

The molasses detuning Δf_{mol} from the bare $F = 2 \to F' = 3$ transition is extracted from our spectroscopy scheme and various AOM shifts shown in Fig. 2.3. We estimate $\Delta f_{\text{microtweezer}}$ by measuring the power sent into the objective and estimating the trap depth based on our estimate of the optical losses through the objective and the beam waist. We estimate the cavity ODT ground state trap depth by measuring the input power and using the results of Sec. 6.4.2. To calculate $\Delta f_{\text{cavity ODT}}$ we recall from Sec. 6.4.3 that the excited state polarizability is $\approx (45\times)$ the ground state polarizability.

Finally, I'll importantly note that for Fig. 8.5(b) I have, by hand, introduced two 'fudge' factors multiplying the cavity ODT trap depth and microtweezer trap depth: ϵ_{cav} and $\epsilon_{\text{tweezer}}$. I chose $\epsilon_{\text{cav}} = 0.37$ to make the shifted traces line up best 'by eye'. This value for ϵ_{cav} is consistent with our expectations for the cavity ODT input mode-matching, which was not taken into account otherwise in estimating the cavity ODT trap depth. Note that this factor was included in the calculation of the horizontal axis for Fig. 8.5(a) as well.

We expect the fluorescence photon counts to scale like $\propto 1/\Delta f_{\text{tot}}^2$. The blue curve in Fig. 8.5 is a fit of the shifted data to

$$N_{\text{count}} = \frac{A}{\Delta_{\text{tot}} - \Delta_0}. \tag{8.6}$$

I chose $\epsilon_{\text{microtweezer}} = 0.4$ so that the extracted value for Δ_0 was zero, as expected for the molasses scattering rate as a function of molasses detuning. This value for $\epsilon_{\text{microtweezer}}$ is consistent with a combination of clipping losses through the objective and aberrations (especially spherical) that would decrease the peak optical amplitude at the focus relative to the ideal Gaussian prediction. Neither of these had previously been taken into account in our estimations of the microtweezer trap depth.

To eliminate the need for these fudge factors in our analysis, we could much more sensitively and independently measure both $\Delta f_{\text{cavity ODT}}$ and $\Delta f_{\text{microtweezer}}$ by driving the atoms with near-resonant light on the D_2 cycling transition and performing loss spectroscopy [150].

We spatially map out this cavity-ODT-shifted effective molasses resonance by trapping an atom in a microtweezer at different positions along the cavity axis, taking fluorescence images, and analyzing the integrated photon counts as above. We control the microtweezer position by adjusting a piezo controlled mirror, shown in Fig. 7.3, in the microtweezer optical path. We are able to sensitively tune the microtweezer position by $\approx 30\,\text{nm/V}$ in the atom plane with a total scan range of $\approx 1.8\,\mu\text{m}$. In the future we will extend our scan range by controlling the microtweezer position by actuating either a second AOD or a DMD.

We think of this measurement as a type of scanning probe microscopy (SPM) in which the single atom is scanned as a sub-wavelength scanning tip whose fluorescence is imaged to map out the intensity of cavity ODT. This SPM is similar to stochastic optical reconstruction microscopy (STORM) [155] or photoactivated localization microscopy (PALM) [156] in that

the imaging of the single atom fluorescence is diffraction limited, but, because there is only a single fluorescence point source, its position can be localized to better than the diffraction limit [157]. In our case, for a single atom trapped in a microtweezer ODT with waist size $w_0 \approx 1\,\mu$m and, assuming $\eta = 10$, we estimate a cloud size of $\sigma \approx 160\,$nm, which sets the resolution for our microscopy technique. See Ref. [158] for another cold atom SPM technique.

The results of a 1D scan following the above procedure are presented in Fig. 8.5(c). These data were taken with a tweezer trap depth of $\approx 250\,\mu$K, a bare molasses detuning of $\Delta_{\mathrm{mol}} \approx -2\pi \times 40\,$MHz, a cavity ODT trap depth of $\approx 20\,\mu$K, and a molasses saturation parameter of about $s \sim 7$. Our molasses fluorescence imaging time is $100\,$ms. For each microtweezer position we estimate the average number of photon counts collected from the atom at each location along the cavity axis by post-selecting shots that contain a single atom and averaging the integrated photon counts.

The expected wavelength for this contrast curve is $1560\,$nm$/2 = 780\,$nm. A sinusoidal fit to the data yields a periodicity of $\approx 670\,$nm. The horizontal axis in Fig. 8.5(c) was estimated using the camera pixel size and a nominal magnification of ($50\times$).[1] This discrepancy in the spatial wavelength is easily explained as an overestimate of our imaging system magnification. Such an overestimate could result from the distance between the objective front principal plane and the microtweezer focus differing from the objective focal length $f = 4\,$mm by about $15\,\mu$m or from the positioning of the 'eyepiece' lens being off by a couple cm. In fact, this contrast measurement is a much more precise ruler for estimating our magnification and, going forward, we will use a magnification of ($42\times$) for this imaging system, consistent with this new calibration. The uncertainty on this new magnification will be established as we repeat this measurement and monitor the repeatability of the fitted periodicity.

8.3.1   3D Spatial Mapping

Above, I presented a measurement scheme that allows us to extract the intensity of the cavity ODT as a function of z-position along the cavity axis. By adjusting the position of the microtweezer along x (the cross axis) and y (the vertical direction) axes, we can map out the 3D structure of the cavity mode. See Ref. [159] for another approach whereby atoms passing through an optical cavity and undergoing a cQED interaction can be used to map out the spatial structure of the cavity photonic mode. In Ref. [160] a transmission electron microscope is used to spatially map out the intensity of the optical mode of a near-concentric cavity.

The scheme described in the previous paragraph requires a 3D scan of the position of the microtweezer within the cavity mode. However, since we are able to generate an array of microtweezers, we can parallelize this measurement to reduce the number of experimental cycles required. See Fig. 8.6. We can orient the AOD so that the tweezer array sweeps out along the x-axis. Then, by using the microtweezer piezo, we can translate the tweezer array

---

[1]This nominal magnification was chosen assuming the lenses had been configured as depicted in Fig. 7.3. The lenses were placed using a rulers to measure the appropriate distances with precision perhaps at the few cm level which leads to an uncertainty in the magnification at the 10% level.
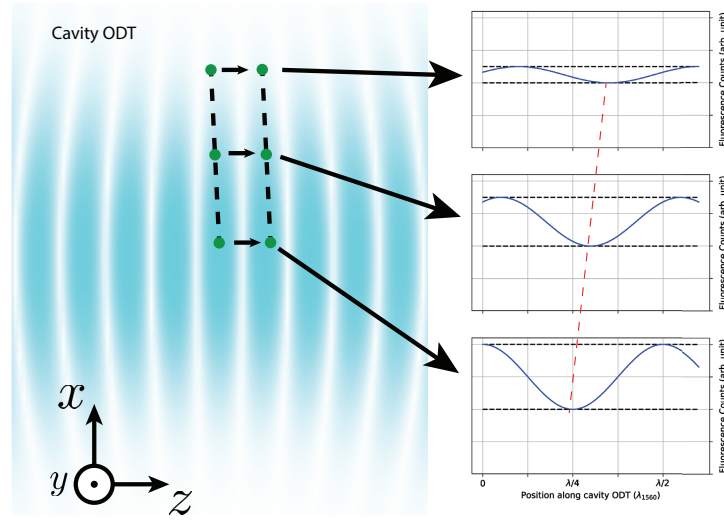
Figure 8.6: Tweezer array contrast measurement. The measurement shown in Fig. 8.5 is repeated with multiple atoms trapped in microtweezers spanning the x-direction within the cavity ODT. This measurement allows us to simultaneously extract spatial information about the probe in multiple locations within the cavity ODT.

along the z-axis mapping out the fluorescence statistics for each microtweezer in the array. The contrast of each respective z-trace is proportional to the peak intensity of the cavity ODT along that trace and the relative phase between the different traces is related to a possible angular misalignment between cavity ODT wavefronts and the microtweezer array.

If we assume the cavity mode is radially symmetric, then the above measurement suffices to map out the 3D cavity mode. If we do not make this assumption, we can map out the y-direction by translating the microscope objective to move the microtweezer focus vertically through the cavity mode. In general, this method could be used to characterize not only the TEM_{00} mode as shown here, but any higher order cavity mode as well.

If a second AOD is installed with its axis at 90° to the existing AOD, we could use this second AOD to perform the scan along z-axis instead of the microtweezer piezo. This would give us a larger scan range if desired. We could also use this second AOD to create a 2D array of microtweezers, which would allow for further parallelization of the measurement.

Chapter 9

Outlook

9.1 Summary

In the previous chapters, I have motivated the E6 apparatus by its integration of a high-NA addressing system with a cQED high-finesse cavity system. I described the design and assembly process for all of the components including the atom preparation, science cavity, and high-NA systems. In Chapter 8, I demonstrated our first results that utilized both of these systems working simultaneously. Going forward, we will leverage the unique capabilities of this apparatus to explore new physics and test new technical possibilities. I will describe some of these future opportunities in this chapter.

9.2 Near-Term

9.2.1 Continuation of the Cavity ODT Spatial Mapping Work

In the near term, we will continue the work that was described in Chapter 8. The immediate tasks are establishing a more quantitative understanding of the dependence of the single atom scattering rates and lifetimes as a function of the microtweezer depth, molasses detuning and intensity, and the depth of the cavity ODT. Once these are well characterized we will be able to invert the measured fluorescence brightness to make an accurate map of the spatial amplitude profile of the cavity ODT. We will also be able to use this information, and measured contrast curves like the one shown in Fig. 8.5(c), to quantitatively extract the contrast for our scanning probe technique and estimate our spatial sensing resolution. We expect our resolution to be about the size of the atomic distribution within the tweezer, $\sigma \approx 160\,\mathrm{nm}$.

We are already beginning to take data according to the contrast array scanning technique described in Fig. 8.6 that will allow us to map out the structure of the cavity ODT mode in more spatial dimensions. To further test this technique we could lock the cavity ODT instead to a higher order TEM mode so that we can map out a more complex spatial structure.

9.2.2 Interactions Between the Cavity Probe and Atoms in Microtweezers

The controllable interactions between the cavity and the atoms in the tweezer so far have been demonstrated with the cavity ODT. However, we are excited to use the microtweezers to modulate the cQED interaction between the atoms in the microtweezer and the cavity probe. To this end we can begin using the heterodyne detection to try to detect modifications to the cavity resonance due to atoms trapped in the microtweezers.

While the above work involved single atoms trapped in the microtweezers, I remind the reader that, prior to the burst of photoassociation light, we likely have 10s or 100s of atoms trapped in the microtweezers. By tuning Δ_{CA} it should be possible to observe a either a dispersive shift (if we work with large Δ_{CA}) or a cavity mode-splitting (if we work with small Δ_{CA}) due to the presence of these atoms with the cavity. There are many open questions about this measurement:

- Will it be possible to measure the atom number in the tweezer using the 780 nm probe light without inducing light-assisted collisions?

- Is there a parameter regime where we can realize single-atom sensitivity for this cavity read-out?

This investigation would likely yield a better understanding of probe heating mechanisms and may also provide an enhanced understanding of the photoassociation process responsible for atom number parity projection in microtweezers.

9.3 Medium-Term

9.3.1 Mapping out the Cavity Probe Using the Microtweezer Scanning Probe Microscope

Once we are able to sense the presence of a single atom using the probe, we can repeat the spatial contrast measurement we performed previously, but now using the probe rather than the cavity ODT. We can map out the relative position of the trapped atom and probe cavity mode standing wave by monitoring the position dependence of the detected cavity light.

Unfortunately, since the probe wavelength is shorter than the cavity ODT wavelength, our contrast will be $e^{-2k_{780}^2\sigma^2} \approx 0.036$, which is very poor. We may be able to decrease the cloud size to increase our resolution by either compressing the atomic thermal distribution by ramping up the microtweezer trap depth or applying more aggressive cooling techniques to cool the atom to its motional ground state within the microtweezer. Such cooling techniques might include Raman sideband cooling or some form of cavity cooling.

If we are able to realize a very large contrast with this measurement, we would be able to selectively read out whether there is an atom or not in each cell of a microtweezer array
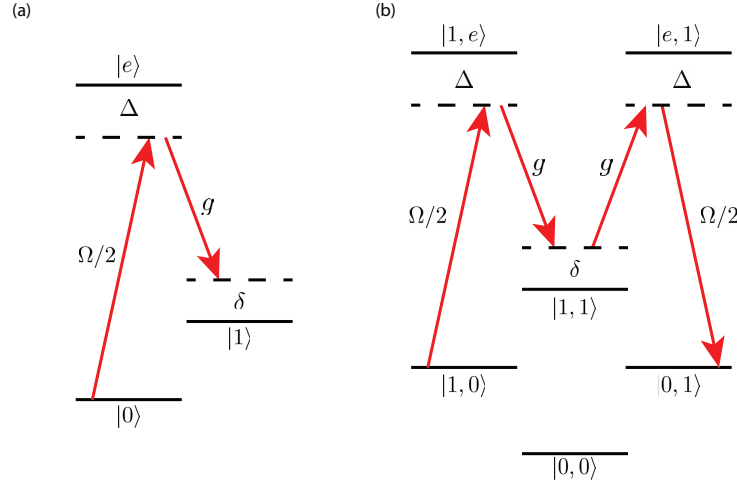
Figure 9.1: (a) Single atom level diagram. There are two grounds state $|0\rangle$ and $|1\rangle$ separated by a small energy splitting. The state $|0\rangle$ is coupled to the excited state $|e\rangle$ via an external pump photon with Rabi frequency Ω. The state $|1\rangle$ is coupled to the excited state $|e\rangle$ via a cavity photon with Rabi frequency $2g$. The frequencies of the pump and cavity photons are tuned so that the detuning of the pump photon is Δ and the two-photon Raman detuning is δ. (b) If two atoms are simultaneously coupled to the pump and cavity fields in this configuration a 4-photon flip-flip interaction will be possible that couples $|0,1\rangle$ to $|1,0\rangle$.

by placing one atom at the probe anti-node and all others at nodes and then performing our cavity readout. We could dynamically readout each atom by adjusting the microtweezer positions in real time using the AOD. The cavity is a bus to which we could enable or disable the coupling of individual atoms at will.

Furthermore, we can tune Δ_{CA} so that the cavity shift has the opposite sign for atoms in the $F = 1$ vs. $F = 2$ hyperfine ground states. This would make our single atom readout spin-sensitive giving us the ability to selectively readout the spins of individual atoms in a microtweezer array [154].

9.3.2   Locally Driven Multi-Photon Raman Interactions

It has been shown that the cavity mode, combined with external drive fields, can be used to mediate long range flip-flop and spin-mixing interactions between atoms within the optical cavity, which can be used for Hamiltonian engineering [21, 113].

To understand this system we imagine individual atoms as three-level systems that include two ground, qubit, states, $|0\rangle$ and $|1\rangle$, and one excited state $|e\rangle$. A four-photon flip-flop gate whereby the state of two atoms within the cavity can change from $|01\rangle \leftrightarrow |10\rangle$ is shown schematically in Fig. 9.1. One atom absorbs a pump photon and emits a virtual photon into the cavity changing its state in the process. The other atom then absorbs the virtual cavity

photon and emits a pump photon also changing its state. The total transition strength for this interaction is

$$g_{\text{flip-flop}} = \frac{\Omega^2 g^2}{4\Delta^2 \delta}, \tag{9.1}$$

as can be calculated using second order perturbation theory or a Schrieffer-Wolff expansion [161]. The detuning from the excited state Δ suppresses decay via spontaneous emission Γ from the excited state $|e\rangle$, and the two-photon detuning δ suppresses decay via cavity leakage κ from the single photon state $|0, 0, 1\rangle$.

This interaction has been engineered using a global pump and cavity field to realize a tunable global Heisenberg model [113]. In our apparatus, using the high-NA imaging system, we would be able to extend this scheme to realize locally tunable interactions by shining localized Raman beams onto the atoms. This scheme is analogous to gates in an ion chain in which ions are made to interact via their shared coupling to phononic modes of the ion crystal - a shared phonon bus [162, 163]. In our scheme the phonon bus is replaced by a cavity photon bus. We could also modulate this interaction by changing the overlap between the atoms and the cavity probe as described in the previous section.

In Refs. [21] and [113] this interaction was realized using the three hyperfine Zeeman sublevels in the $F = 1$ ground state manifold. We are curious if we could realize this interaction using states from the different Hyperfine levels to represent the two qubit states. For example, $|0\rangle = |F = 1, m_F = +1\rangle$ and $|1\rangle = |F = 2, m_F = +2\rangle$.

9.3.3 Cavity Optodynamics with Atomic Ensembles

In E3, we were able to use the cavity to realize cavity spin optodynamics, the linear coupling of the collective spin of an ensemble of atoms within the cavity to the cavity probe field [56, 69, 164]. We had hoped to realize the cavity readout and coupling of two spatially-separated spin ensembles with distinguishable Larmor frequencies. In particular, we had hoped to initialize one spin ensemble in the spin down state and the other in the spin up state. We were curious if, by coupling these two ensembles using the cavity, we could realize interesting spin-squeezed or cat states.

However, we were unable to carry out this measurement. In E3, we were able to load ensembles of atoms into neighboring wells of the cavity ODT. To realize distinct Larmor frequencies for the two ensembles we generated a large magnetic field gradient using the E3 atom chip wires. Unfortunately, we found that when we applied a magnetic field gradient large enough to distinguish the spin ensembles spectroscopically we also introduced large curvatures into the magnetic field that led to decoherence of the atomic spin ensembles due to their motion within the E3 cavity ODT.

In E6 we may be able to get around this issue by (1) using the transport ODT to load atoms into the cavity ODT at much larger spatial separations (10s or 100s of μm is easily imaginable) and (2) using our powerful anti-Helmholtz coils to generate the magnetic field gradients necessary to distinguish the atomic ensembles. Alternatively, we may be able to

create fictitious magnetic fields optically by locally shining in laser beams using the objective to apply vector Stark shifts to the atoms.

9.3.4 Spin Squeezing

Optical cavities have a successful history of being used to squeeze the collective spin of ensembles of atoms with the cavity mode [22–24, 165]. See Refs. [76, 166, 167] for detailed explanations of spin squeezing.

In short, spin squeezing is a technique whereby the variance of the spin projection of a large spin system $\langle \hat{S}_z^2 \rangle$ is massively reduced. Spin squeezing can be implemented by subjecting a spin ensemble to the Hamiltonian

$$\hat{H}_{SS}/\hbar = \chi \hat{S}_z^2 \tag{9.2}$$

where χ is the spin squeezing strength. Spin squeezed states can be used for quantum-enhanced metrology beyond the standard quantum limit [168].

If the spin ensemble is, in fact, the collective spin of N independent spins, then spin squeezing results in a massive entanglement between all of these different spins. Historically, spin squeezed states have been measured by measuring statistics of the collective spin ensemble such as $\langle \hat{S}_z^2 \rangle$. Such collective statistics do not reveal much information about the underlying entanglement structure of the spin squeezed state. One science target for E6 could be to realize spin squeezing of an ensemble of atoms trapped in microtweezers and to, then, use a single spin readout technique to extract statistics for the individual spin states and correlations between different spins. In principle, this could allow us to perform a full spin tomography of the spin squeezed state.

9.4 Long-Term

9.4.1 Few-to-Many Body Quantum Dynamics

Much of the physics considered so far involves either a few atoms trapped into microtweezers or a small number of atomic ensembles within the cavity. However, it would be interesting to use the apparatus to explore many-body quantum dynamics. One approach to do this would be to load a large number of atoms within the microtweezer array. This would require a 2D AOD and additional laser power to allow us to generate many tweezers. Another approach would be to create a 2D or 3D optical lattice within the cavity and to load a BEC into this lattice so that we could explore low temperature Hubbard physics in this system. We would have access to the local degrees of freedom via the high-NA objective just like any other quantum gas microscope, but, in addition we would have access to the measurements and interactions described above that are made possible via the cQED interaction with the science cavity. Additionally, we might be able to controllably introduce dissipation into the system by transferring excitations within the many-body system into photonic excitations

that, then, leak out of the cavity at rate κ. We could explore (1) how this dissipation affects the many-body dynamics and (2) what can be learned about the many-body system by monitoring this light that leaks out of the cavity using our heterodyne detector.

9.4.2 Quantum Feedback

Throughout this text I have given examples of how both the science cavity and the high-NA system can be used to both measure and control and quantum atomic sample at the heart of this experiment. One of the ways we could most tightly intertwine these three systems — the science cavity, the high-NA system, and the atoms — would be to leverage the unique precision measurement capabilities of the cavity to monitor the atomic dynamics in real time and process this data to feedback to arbitrary optical fields that we use to drive the atoms using the high-NA system. Though we haven't yet proposed detailed experimental schemes, we are confident that the technical capability to carry out such ideas opens the door to exciting and unexplored regimes of many-body quantum feedback.

Note that here I am specifically talking about measurement-based feedback in which a 'classical' detector is included in the measurement loop. Measurement-based-feedback is in contrast to 'coherent' quantum feedback in which the dynamics of one quantum system (such as the optical cavity) alter the dynamics of another quantum system (such as the spin or motion of atoms within the cavity) [169–172]. Such measurement-based quantum feedback is related to quantum error correction protocols, which are important for quantum information processing applications. Measurement-based quantum feedback has been realized in cQED systems to realize a type of cavity cooling or cold damping [173] and for the capture and release of conditional states [174]. Interesting measurement-based feedback protocols have also been implemented in superconducting circuit quantum electrodynamics experiments that we could try to replicate in our atom-based system [175]. There are a number of other theoretical proposals for quantum feedback in a cQED system [176–178]. Finally, there are a number of very interesting proposals that consider the dynamics of many body quantum systems under the influences of back action and quantum feedback that may be possible to realize in a system like E6 [179–183].

9.4.3 The Next Generation Apparatus

There are years of research that can be carried out in the existing apparatus along the lines described throughout this dissertation and in this chapter. Likewise, the 5 grams of ^{87}Rb that were installed in the 2D MOT reservoir should also last for many years before needing replacement. This means that the UHV vacuum chamber can, in principle, stay under vacuum for a very long time.

However, a time will likely come in the next few years when the new E6 team finds themselves wishing for improvements in the science cavity parameters. There are two major gains that can be made in terms of the science cavity specifications. First, had the mirrors not been damaged during assembly, the losses would have been an order of magnitude lower

and we would have been able to realize a cavity with simultaneously higher $\mathcal{F}$ (and thus C) and detection efficiency η_{out}. Second, in hindsight, we realized that during the cavity characterization process we likely assembled well-aligned cavities that were much closer to concentric than the one that is installed in the science cavity now with $\delta_{\mathrm{conc}} \approx 600\,\mu\mathrm{m}$. In the future, we should be able to assemble and install a second generation science cavity that has better coating properties and that is closer to concentric, allowing us to realize much higher cooperativities.

I will make one point here regarding the concentricity of the cavity. The gains in cooperativity for a more near-concentric cavity come as a result of the decreasing cavity waist w_0. There is an important downside to decreasing the cavity waist which is that a smaller waist limits the physical size of the quantum system that can be loaded into the focus of the cavity. When designing the next generation cavity, this effect should be taken into account, because, in addition to targeting improved cavity specifications, the next generation experiment will also likely target large system sizes.

However, an advantage to very near-concentric cavities is that the transverse mode spectrum becomes near degenerate as in Ref. [32]. In a regime where f_{TMS} is small it would be possible for the atoms to scatter photons between different transverse modes allowing for a richer atom-cavity dynamics. Considerations for such possibilities should also be made when designing the next generation science cavity.

In addition to an improved cavity, it will also likely be beneficial for the next generation of this experiment to include an improved high-NA objective. Major improvements would include increased transmission efficiency in the IR for near-resonant and tweezer beams and also a higher NA to improve the imaging resolution. When this new objective is tested I would recommend also purchasing and testing a new, high-optical-quality, re-entrant viewport.

9.5   Conclusion

I have detailed the design and construction of a new ultracold atomic physics apparatus that was motivated by the ambitious goal of combining an atomic cavity QED system with a high-NA imaging and addressing system. We have successfully built this system and demonstrated the tip of the iceberg for the experimental results to come. It is difficult, here at the unveiling of this new apparatus, to predict which scientific directions future researchers will take this project, but I can say that I believe these directions will be interesting and I am very excited to see what is to come.

Bibliography

[1] J. C. Maxwell, "On physical lines of force", *Philosophical Magazine*, **90**, 11 (1861) (cit. on p. 1).

[2] M. Planck, "Über eine Verbesserung der Wien'schen Spectralgleichung", *Verhandlungen der Deutschen Physikalischen Gesellschaft*, **2**, 202 (1900) (cit. on p. 1).

[3] M. Planck, "Zur Theorie des Gesetzes der Energieverteilung im Normalspectrum", *Verhandlungen der Deutschen Physikalischen Gesellschaft*, **2**, 237 (1900) (cit. on p. 1).

[4] J. S. Bell, "On the Einstein Podolsky Rosen Paradox", *Physics*, **1**, 195 (1964) (cit. on p. 1).

[5] R. P. Feynman, "Simulating physics with computers", *International Journal of Theoretical Physics*, **21**, 467 (1982) (cit. on p. 1).

[6] D. Sullivan, "Time and Frequency Measurement at NIST: The First 100 Years", in IEEE International Frequency Control Symposium (2001) (cit. on p. 2).

[7] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard, "Trapping of Neutral Sodium Atoms with Radiation Pressure", *Physical Review Letters*, **59**, 2631 (1987) (cit. on pp. 2, 26).

[8] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, "Observation of bose-einstein condensation in a dilute atomic vapor", *Science*, **269**, 198 (1995) (cit. on pp. 2, 23).

[9] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, "Bose-Einstein Condensation in a Gas of Sodium Atoms", *Physical Review Letters*, **75**, 3969 (1995) (cit. on pp. 2, 23).

[10] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification", *Reviews of Modern Physics*, **82**, 1155 (2010) (cit. on pp. 2, 71).

[11] I. M. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation", *Reviews of Modern Physics*, **86**, 153 (2014) (cit. on p. 2).

[12] J. Preskill, "Quantum Computing in the NISQ era and beyond", *arXiv:1801.00862*, (2018) (cit. on p. 2).

[13] C. L. Degen, F. Reinhard, and P. Cappellaro, "Quantum sensing", *Review of Modern Physics*, **89**, 035002 (2017) (cit. on p. 2).

[14] E. T. Jaynes and F. W. Cummings, "Comparison of Quantum and Semiclassical Radiation Theories with Application to the Beam Maser", *Proc. IEEE*, **51**, 89 (1963) (cit. on p. 2).

[15] M. G. Raizen, R. J. Thompson, R. J. Brecha, H. J. Kimble, and H. J. Carmichael, "Normal-mode splitting and linewidth averaging for two-state atoms in an optical cavity", *Physical Review Letters*, **63**, 240 (1989) (cit. on pp. 3, 71, 77).

[16] R. J. Thompson, G. Rempe, and H. J. Kimble, "Observation of Normal-Mode Splitting for an Atom in an Optical Cavity", *Physical Review Letters*, **68**, 1132 (1992) (cit. on pp. 3, 17, 160).

[17] H. J. Kimble, "Strong Interactions of Single Atoms and Photons in Cavity QED", *Physica Scripta*, **T76**, 127 (1998) (cit. on pp. 3, 19, 46, 71).

[18] A. Boca, R. Miller, K. M. Birnbaum, A. D. Boozer, J. Mckeever, and H. J. Kimble, "Observation of the Vacuum Rabi Spectrum for One Trapped Atom", *Physical Review Letters*, **93**, 233603 (2004) (cit. on pp. 3, 17).

[19] H. J. Kimble, "The quantum internet", *Nature*, **453**, 1023 (2008) (cit. on p. 3).

[20] M. Brekenfeld, D. Niemietz, J. D. Christesen, and G. Rempe, "A quantum network node with crossed optical fibre cavities", *Nature Physics*, **16**, 647 (2020) (cit. on pp. 3, 75).

[21] E. J. Davis, G. Bentsen, L. Homeier, T. Li, and M. H. Schleier-Smith, "Photon-Mediated Spin-Exchange Dynamics of Spin-1 Atoms", *Physical Review Letters*, **122**, 10405 (2019) (cit. on pp. 3, 22, 46, 64, 77, 146, 171, 172).

[22] Z. Chen, J. G. Bohnet, S. R. Sankar, J. Dai, and J. K. Thompson, "Conditional spin squeezing of a large ensemble via the vacuum rabi splitting", *Physical Review Letters*, **106**, 133601 (2011) (cit. on pp. 3, 18, 22, 173).

[23] N. J. Engelsen, R. Krishnakumar, O. Hosten, and M. A. Kasevich, "Bell Correlations in Spin-Squeezed States of 500 000 Atoms", *Physical Review Letters*, **118**, 140401 (2017) (cit. on pp. 3, 173).

[24] I. D. Leroux, M. H. Schleier-Smith, H. Zhang, and V. Vuletić, "Unitary cavity spin squeezing by quantum erasure", *Physical Review A*, **85**, 013803 (2012) (cit. on pp. 3, 173).

[25] M. H. Schleier-Smith, I. D. Leroux, and V. Vuletić, "States of an ensemble of two-level atoms with reduced quantum uncertainty", *Physical Review Letters*, **104**, 073604 (2010) (cit. on pp. 3, 18).

[26] R. McConnell, H. Zhang, J. Hu, S. Cuk, and V. Vuletić, "Entanglement with negative Wigner function of almost 3,000 atoms heralded by one photon", *Nature*, **519**, 439 (2015) (cit. on p. 3).

[27] D. W. C. Brooks, T. Botter, S. Schreppler, T. P. Purdy, N. Brahms, and D. M. Stamper-Kurn, "Non-classical light generated by quantum-noise-driven cavity optomechanics", *Nature*, **488**, 476 (2012) (cit. on p. 3).

[28] S. Schreppler, N. Spethmann, N. Brahms, T. Botter, M. Barrios, and D. M. Stamper-Kurn, "Optically measuring force near the standard quantum limit", *Science*, **344**, 1486 (2014) (cit. on pp. 3, 67).

[29] S. Reick, K. Mølmer, W. Alt, M. Eckstein, T. Kampschulte, L. Kong, R. Reimann, A. Thobe, A. Widera, and D. Meschede, "Analyzing quantum jumps of one and two atoms strongly coupled to an optical cavity", *Journal of the Optical Society of America B*, **27**, A152 (2010) (cit. on p. 3).

[30] V. Vuletić, H. W. Chan, and A. T. Black, "Three-dimensional cavity Doppler cooling and cavity sideband cooling by coherent scattering", *Physical Review A*, **64**, 033405 (2001) (cit. on p. 3).

[31] R. Reimann, W. Alt, T. Kampschulte, T. Macha, L. Ratschbacher, N. Thau, S. Yoon, and D. Meschede, "Cavity-Modified Collective Rayleigh Scattering of Two Atoms", *Physical Review Letters*, **114**, 023601 (2015) (cit. on p. 3).

[32] A. J. Kollar, "Observation of Supermode-Density-Wave-Polariton Condensation in a Multimode Cavity QED-BEC System", PhD thesis (Stanford University, 2016) (cit. on pp. 3, 77, 114, 175).

[33] J. Léonard, A. Morales, P. Zupancic, T. Esslinger, and T. Donner, "Supersolid formation in a quantum gas breaking a continuous translational symmetry", *Nature*, **543**, 87 (2017) (cit. on p. 3).

[34] L. W. Clark, N. Schine, C. Baum, N. Jia, and J. Simon, "Observation of Laughlin states made of light", *Nature*, **582**, 41 (2020) (cit. on p. 3).

[35] C. A. Regal, C. Ticknor, J. L. Bohn, and D. S. Jin, "Creation of ultracold molecules from a Fermi gas of atoms", *Nature*, **424**, 47 (2003) (cit. on pp. 3, 145).

[36] C. J. Kennedy, W. C. Burton, W. C. Chung, and W. Ketterle, "Observation of Bose-Einstein condensation in a strong synthetic magnetic field", *Nature Physics*, **11**, 859 (2015) (cit. on pp. 3, 145).

[37] G. Jotzu, M. Messer, R. Desbuquois, M. Lebrat, T. Uehlinger, D. Greif, and T. Esslinger, "Experimental realization of the topological Haldane model with ultracold fermions", *Nature*, **515**, 237 (2014) (cit. on pp. 3, 145).

[38] T. H. Barter, T. H. Leung, M. Okano, M. Block, N. Y. Yao, and D. M. Stamper-Kurn, "Spatial coherence of a strongly interacting Bose gas in the trimerized kagome lattice", *Physical Review A*, **101**, 011601 (2020) (cit. on pp. 3, 145).

[39] M. Aidelsburger, "Artificial gauge fields and topology with ultracold atoms in optical lattices", *Journal of Physics B: Atomic Molecular and Optical Physics*, **51**, 193001 (2018) (cit. on pp. 3, 145).

[40] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner, "Probing the superfluid-to-Mott insulator transition at the single-atom level.", *Science*, **329**, 547 (2010) (cit. on pp. 3, 22, 145, 147).

[41] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, "Single-atom-resolved fluorescence imaging of an atomic Mott insulator.", *Nature*, **467**, 68 (2010) (cit. on pp. 3, 22, 146).

[42] J. Simon, W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner, "Quantum simulation of antiferromagnetic spin chains in an optical lattice.", *Nature*, **472**, 307 (2011) (cit. on pp. 3, 145).

[43] J.-y. Choi, S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross, "Exploring the many-body localization transition in two dimensions", *Science*, **352**, 1547 (2016) (cit. on pp. 3, 145).

[44] P. B. Patel, Z. Yan, B. Mukherjee, R. J. Fletcher, J. Struck, and M. W. Zwierlein, "Universal sound diffusion in a strongly interacting fermi gas", *Science*, **370**, 1222 (2020) (cit. on pp. 3, 145).

[45] R. Grimm and M. A. P. H. I. A. S. Weidemuller, "Optical Dipole Traps For Neutral Atoms", *Advances in Atomic, Molecular and Optical Physics*, **42**, 95 (2000) (cit. on pp. 3, 63, 140).

[46] N. Schlosser, G. Reymond, I.Protsenko, and P. Grangier, "Sub-poissonian loading of single atoms in a microscopic dipole trap", *Nature*, **411**, 1024 (2001) (cit. on pp. 3, 22, 148, 149, 161).

[47] J. P. Covey, I. S. Madjarov, A. Cooper, and M. Endres, "2000-Times Repeated Imaging of Strontium Atoms in Clock-Magic Tweezer Arrays", *Physical Review Letters*, **122**, 173201 (2019) (cit. on p. 4).

[48] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, "High-Fidelity Control and Entanglement of Rydberg-Atom Qubits", *Physical Review Letters*, **121**, 123603 (2018) (cit. on p. 4).

[49] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays", *Science*, **354**, 1021 (2016) (cit. on pp. 4, 22, 23, 146, 148, 151).

[50] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletić, M. Greiner, and M. D. Lukin, "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays", *Science*, **354**, 1024 (2016) (cit. on pp. 4, 22, 23, 146, 148, 151).

[51] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletić, and M. D. Lukin, "Coherence and Raman sideband cooling of a single atom in an optical tweezer", *Physical Review Letters*, **110**, 133001 (2013) (cit. on pp. 4, 22).

[52] P. Sompet, A. V. Carpentier, Y. H. Fung, M. Mcgovern, and M. F. Andersen, "Dynamics of two atoms undergoing light-assisted collisions in an optical microtrap", *Physical Review A*, **88**, 051401 (2013) (cit. on pp. 4, 148).

[53] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, "Rapid Production of Uniformly Filled Arrays of Neutral Atoms", *Physical Review Letters*, **115**, 073003 (2015) (cit. on pp. 4, 22).

[54] B. J. Lester, A. M. Kaufman, and C. A. Regal, "Raman cooling imaging: Detecting single atoms near their ground state of motion", *Physical Review A*, **90**, 011804 (2014) (cit. on p. 4).

[55] V. Lienhard, P. Scholl, S. Weber, D. Barredo, S. De Léséleuc, R. Bai, N. Lang, M. Fleischhauer, H. P. Büchler, T. Lahaye, and A. Browaeys, "Realization of a Density-Dependent Peierls Phase in a Synthetic, Spin-Orbit Coupled Rydberg System", *Physical Review X*, **10**, 021031 (2020) (cit. on p. 4).

[56] J. Kohler, J. A. Gerber, E. Dowd, and D. M. Stamper-Kurn, "Negative-Mass Instability of the Spin and Motion of an Atomic Gas Driven by Optical Cavity Backaction", *Physical Review Letters*, **120**, 013601 (2018) (cit. on pp. 4, 9, 18, 22, 172).

[57] N. Spethmann, J. Kohler, S. Schreppler, L. Buchmann, and D. M. Stamper-Kurn, "Cavity-mediated coupling of mechanical oscillators limited by quantum back-action", *Nature Physics*, **12**, 27 (2015) (cit. on p. 4).

[58] J. A. Gerber, S. Berweger, B. T. O'Callahan, and M. B. Raschke, "Phase-Resolved Surface Plasmon Interferometry of Graphene", *Physical Review Letters*, **113**, 055502 (2014) (cit. on p. 6).

[59] A. Keshet and W. Ketterle, "A Distributed GUI-based Computer Control System for Atomic Physics Experiments", *arXiv:1208.2607*, (2012) (cit. on p. 7).

[60] J. Kohler, J. A. Gerber, E. Deist, and D. M. Stamper-Kurn, "Simultaneous retrodiction of multimode optomechanical systems using matched filters", *Physical Review A*, **101**, 023804 (2020) (cit. on p. 9).

[61] D. A. Steck, *Quantum and Atom Optics* (2007) (cit. on pp. 12, 13, 140, 192).

[62] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, 1995) (cit. on p. 12).

[63] C. C. Gerry and P. L. Night, *Introductory Quantum Optics* (Cambridge University Press, 2005) (cit. on p. 12).

[64] R. Loudon, *The Quantum Theory of Light*, Third (Oxford University press, 2000) (cit. on p. 12).

[65] R. Littlejohn, *Physics 221AB Quantum Mechanics Lecture Notes*, 2019 (cit. on pp. 12, 13, 192).

[66] D. A. Steck, *Rubidium 87 D Line Data*, tech. rep. (2019) (cit. on pp. 16, 19, 26, 32, 130, 140, 192, 193, 197).

[67] T. P. Purdy, D. W. C. Brooks, T. Botter, N. Brahms, Z. Y. Ma, and D. M. Stamper-Kurn, "Tunable cavity optomechanics with ultracold atoms", *Physical Review Letters*, **105**, 133602 (2010) (cit. on pp. 18, 46, 67, 69, 71, 77).

[68] T. Botter, D. W. Brooks, S. Schreppler, N. Brahms, and D. M. Stamper-Kurn, "Optical readout of the quantum collective motion of an array of atomic ensembles", *Physical Review Letters*, **110**, 153001 (2013) (cit. on pp. 18, 67).

[69] J. Kohler, N. Spethmann, S. Schreppler, and D. M. Stamper-Kurn, "Cavity-assisted measurement and coherent control of collective atomic spin oscillators", *Physical Review Letters*, **118**, 063604 (2017) (cit. on pp. 18, 172).

[70] N. Brahms, T. P. Purdy, D. W. C. Brooks, T. Botter, and D. M. Stamper-Kurn, "Cavity-aided magnetic resonance microscopy of atomic transport in optical lattices", *Nature Physics*, **7**, 604 (2011) (cit. on pp. 18, 69).

[71] M. H. Schleier-Smith, I. D. Leroux, H. Zhang, M. A. Van Camp, and V. Vuletić, "Optomechanical cavity cooling of an atomic ensemble", *Physical Review Letters*, **107**, 143005 (2011) (cit. on pp. 18, 20).

[72] T. P. Purdy, "Cavity QED with UltraCold Atoms on an Atom Chip", PhD thesis (University of California, Berkeley, 2009) (cit. on pp. 19, 69, 114).

[73] T. C. M. Botter, "Cavity Optomechanics in the Quantum Regime", PhD thesis (University of California, Berkeley, 2013) (cit. on pp. 19, 69).

[74] E. M. Purcell, "Spontaneous Emission Probabilities at Radio Frequencies", *Physical Review A*, **69**, 681 (1946) (cit. on pp. 19, 75).

[75] H. Tanji-Suzuki, I. D. Leroux, M. H. Schleier-Smith, M. Cetina, A. T. Grier, J. Simon, and V. Vuletić, "Interaction between Atomic Ensembles and Optical Resonators. Classical Description", *arXiv:1104.3594*, (2011) (cit. on pp. 19, 20).

[76] M. H. Schleier-Smith, "Cavity-Enabled Spin Squeezing for a Quantum-Enhanced Atomic Clock by", PhD thesis (Massachusetts Institute of Technology, 2011) (cit. on pp. 19, 173).

[77] J. M. Raimond, M. Brune, and S. Haroche, "Colloquium: Manipulating quantum entanglement with atoms and photons in a cavity", *Reviews of Modern Physics*, **73**, 565 (2001) (cit. on p. 20).

[78] J. Lee, G. Vrijsen, I. Teper, O. Hosten, and M. A. Kasevich, "A Many-Atom Cavity QED System with Homogeneous Atom-Cavity Coupling", *Optics Letters*, **39**, 4005 (2014) (cit. on pp. 22, 64, 142, 144).

[79] G. E. Marti, R. Olf, and D. M. Stamper-Kurn, "Collective excitation interferometry with a toroidal Bose-Einstein condensate", *Physical Review A*, **91**, 013602 (2015) (cit. on p. 22).

[80] A. J. Kollar, A. T. Papageorge, V. D. Vaidya, Y. Guo, J. Keeling, and B. L. Lev, "Supermode-density-wave-polariton condensation with a Bose-Einstein condnesate in a multimode cavity", *Nature Communications*, **8**, 14386 (2017) (cit. on pp. 22, 46).

[81] H. W. Chan, A. T. Black, and V. Vuletić, "Observation of collective-emission-induced cooling of atoms in an optical cavity", *Physical Review Letters*, **90**, 063003 (2003) (cit. on pp. 22, 77).

[82] C. K. Thomas, T. H. Barter, T.-H. Leung, M. Okano, G.-B. Jo, J. Guzman, I. Kimchi, A. Vishwanath, and D. M. Stamper-Kurn, "Mean-field scaling of the superfluid to Mott insulator transition in a 2D optical superlattice", *Physical Review Letters*, **119**, 100402 (2017) (cit. on p. 22).

[83] C. Georges, J. Vargas, H. Keßler, J. Klinder, and A. Hemmerich, "Bloch oscillations of a Bose-Einstein condensate in a cavity-induced optical lattice", *Physical Review A*, **96**, 063615 (2017) (cit. on p. 22).

[84] M.-S. Chang, Q. Qin, W. Zhang, L. You, and M. S. Chapman, "Coherent spinor dynamics in a spin-1 Bose condensate", *Nature Physics*, **1**, 111 (2005) (cit. on p. 22).

[85] J. Hu, A. Urvoy, Z. Vendeiro, V. Crépel, W. Chen, and V. Vuletić, "Creation of a Bose-condensed gas of rubidium 87 by laser cooling", *Science*, **358**, 1078 (2017) (cit. on p. 23).

[86] T. Mukai and M. Yamashita, "Efficient rapid production of a Bose-Einstein condensate by overcoming serious three-body loss", *Physical Review A*, **70**, 013615 (2004) (cit. on p. 23).

[87] M. Horikoshi and K. Nakagawa, "Atom chip based fast production of Bose-Einstein condensate", *Applied Physics B: Lasers and Optics*, **82**, 363 (2006) (cit. on p. 23).

[88] D. M. Farkas, K. M. Hudek, E. A. Salim, S. R. Segal, M. B. Squires, and D. Z. Anderson, "A compact, transportable, microchip-based system for high repetition rate production of Bose-Einstein condensates", *Applied Physics Letters*, **96**, 093102 (2010) (cit. on p. 23).

[89] Y. J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, "Rapid production of R 87 b Bose-Einstein condensates in a combined magnetic and optical potential", *Physical Review A*, **79**, 063631 (2009) (cit. on pp. 23, 37).

[90] M. D. Barrett, J. A. Sauer, and M. S. Chapman, "All-optical formation of an atomic bose-einstein condensate", *Physical Review Letters*, **87**, 010404 (2001) (cit. on p. 23).

[91] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, "Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms", *Physical Review Letters*, **123**, 170503 (2019) (cit. on p. 23).

[92] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. Ovchinnikov, and T. Pfau, "Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap", *Physical Review A*, **66**, 023410 (2002) (cit. on p. 25).

[93] T. Uehlinger, "A 2D Magneto-Optical Trap as a High-Flux Source of Cold Potassium Atoms", PhD thesis (2008) (cit. on pp. 25, 26).

[94] J. R. Kellogg, D. Schlippert, J. M. Kohel, R. J. Thompson, D. C. Aveline, and N. Yu, "A compact high-flux cold atom beam source", *arXiv:1107.5602*, (2011) (cit. on p. 25).

[95] B. Holtkemeier, "2D MOT as a source of a cold atom target", , Diploma Thesis (University of Heidelberg) (2011) (cit. on p. 25).

[96] H. Busche, "Efficient loading of a magneto-optical trap for experiments with dense ultracold Rydberg gases", , Diploma Thesis (University of Heidelberg) (2011) (cit. on p. 25).

[97] K. Dieckmann, R. Spreeuw, M. Weidemüller, and J. Walraven, "Two-dimensional magneto-optical trap as a source of slow atoms", *Physical Review A*, **58**, 3891 (1998) (cit. on pp. 25, 26).

[98] S. Chaudhuri, S. Roy, and C. Unnikrishnan, "Realization of an intense cold Rb atomic beam based on a two-dimensional magneto-optical trap: Experiments and comparison with simulations", *Physical Review A*, **74**, 023406 (2006) (cit. on pp. 25, 26).

[99] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer-Verlag, New York, 1999) (cit. on pp. 26, 152, 154).

[100] J. Dalibard and C. Cohen-Tannoudji, "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models", *Journal of the Optical Society of America B*, **6**, 2023 (1989) (cit. on pp. 27, 149, 152).

[101] J. Léonard, M. Lee, A. Morales, T. M. Karg, T. Esslinger, and T. Donner, "Optical transport and manipulation of an ultracold atomic cloud using focus-tunable lenses", *New Journal of Physics*, **16**, 093028 (2014) (cit. on pp. 35, 36).

[102] A. E. Siegman, *Lasers*, edited by A. Kelly (University Science Books, Sausalito, California, 1986) (cit. on pp. 39–41, 47, 54).

[103] A. Yariv and P. Yeh, *Photonics - Optical Electronics in Modern Communications*, 6th ed. (Oxford University Press, New York, New York, 2007) (cit. on p. 39).

[104] S. Garcia, F. Ferri, J. Reichel, and R. Long, "Overlapping two standing-waves in a microcavity for a multi-atom photon interface", *arXiv:2003.0273*, (2020) (cit. on pp. 41, 64, 159).

[105] C. J. Hood, H. J. Kimble, and J. Ye, "Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity", *Physical Review A*, **64**, 033804 (2001) (cit. on pp. 41, 77, 98).

[106] M. Hijlkema, B. Weber, H. P. Specht, S. C. Webster, A. Kuhn, and G. Rempe, "A single-photon server with just one atom", *Nature Physics*, **3**, 253 (2007) (cit. on pp. 46, 71, 77).

[107] M. Khudaverdyan, W. Alt, T. Kampschulte, S. Reick, A. Thobe, A. Widera, and D. Meschede, "Quantum jumps and spin dynamics of interacting atoms in a strongly coupled atom-cavity system", *Physical Review Letters*, **103**, 123006 (2009) (cit. on pp. 46, 71, 77, 160).

[108] S. Gupta, K. L. Moore, K. W. Murch, and D. M. Stamper-Kurn, "Cavity Nonlinear Optics at Low Photon Numbers from Collective Atomic Motion", *Physical Review Letters*, **99**, 213601 (2007) (cit. on pp. 46, 71, 77).

[109] T. Bourdel, T. Donner, S. Ritter, A. Öttl, M. Köhl, and T. Esslinger, "Cavity QED detection of interfering matter waves", *Physical Review A*, **73**, 043602 (2006) (cit. on pp. 46, 71, 77).

[110] V. D. Vaidya, Y. Guo, R. M. Kroeze, K. E. Ballantine, A. J. Kollár, J. Keeling, and B. L. Lev, "Tunable-range, photon-mediated atomic interactions in multimode cavity QED", *Physical Review X*, **8**, 11002 (2017) (cit. on p. 46).

[111] C. H. Nguyen, A. N. Utama, N. Lewty, K. Durak, G. Maslennikov, S. Straupe, M. Steiner, and C. Kurtsiefer, "Single atoms coupled to a near-concentric cavity", *Physical Review A*, **96**, 031802 (2017) (cit. on pp. 46, 51).

[112] T. E. Northup, B. Casabone, K. Friebe, K. Schüppert, F. R. Ong, M. Lee, D. Fioretto, K. Ott, S. Garcia, J. Reichel, and R. Blatt, "An ion-cavity interface for quantum networks", in Proc. SPIE 9615, Quantum Communications and Quantum Imaging XIII, Vol. (2015), p. 961506 (cit. on pp. 46, 51).

[113] E. J. Davis, A. Periwal, E. S. Cooper, G. Bentsen, S. J. Evered, K. Van Kirk, and M. H. Schleier-Smith, "Protecting Spin Coherence in a Tunable Heisenberg Model", *Physical Review Letters*, **125**, 60402 (2020) (cit. on pp. 51, 171, 172).

[114] O. Schwartz, J. Axelrod, D. R. Tuthill, P. Haslinger, C. Ophus, R. Glaeser, and H. Müller, "Near-concentric Fabry-Pérot cavity for continuous-wave laser control of electron waves", *Optics Express*, **25**, 14453 (2017) (cit. on p. 51).

[115] Y. Xiao, D. N. Maywar, and G. P. Agrawal, "Optical pulse propagation in dynamic Fabry–Perot resonators", *Journal of the Optical Society of America B*, **28**, 1685 (2011) (cit. on p. 54).

[116] A. Schliesser, O. Arcizet, R. Rivière, G. Anetsberger, and T. J. Kippenberg, "Resolved-sideband cooling and position measurement of a micromechanical oscillator close to the Heisenberg uncertainty limit", *Nature Physics*, **5**, 509 (2009) (cit. on p. 71).

[117] C. Gardiner and M. Collett, "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation", *Physical Review A*, **31**, 3761 (1985) (cit. on pp. 72, 209).

[118] G. Rempe, R. Lalezari, R. J. Thompson, and H. J. Kimble, "Measurement of ultralow losses in an optical interferometer", *Optics Letters*, **17**, 363 (1992) (cit. on pp. 72, 76, 77).

[119] J. Gallego, W. Alt, T. Macha, D. Pandey, and D. Meschede, "Strong Purcell Effect on a Neutral Atom Trapped in an Open Fiber Cavity", *Physical Review Letters*, **121**, 173603 (2018) (cit. on p. 75).

[120] D. Hunger, T. Steinmetz, Y. Colombe, C. Deutsch, T. W. Hänsch, and J. Reichel, "A fiber Fabry-Perot cavity with high finesse", *New Journal of Physics*, **12**, 065038 (2010) (cit. on p. 76).

[121] J. M. Bennett, "Recent developments in surface roughness characterization", *Measurement Science and Technology*, **3**, 1119 (1992) (cit. on p. 76).

[122] M. Uphoff, M. Brekenfeld, G. Rempe, and S. Ritter, "Frequency splitting of polarization eigenmodes in microscopic Fabry-Perot cavities", *New Journal of Physics*, **17**, 013053 (2015) (cit. on p. 86).

[123] E. Davis, "Engineering and Imaging Nonlocal Spin Dynamics in an Optical Cavity", PhD thesis (Stanford University, 2020) (cit. on p. 88).

[124] K. Durak, "A Diffraction Limited Fabry-Perot Cavity with a Strongly Focused Mode", PhD thesis (National University of Singapore, 2014) (cit. on p. 93).

[125] M. J. Martin, "Quantum Metrology and Many-Body Physics : Pushing the Frontier of the Optical Lattice Clock", *Thesis*, 308 (2013) (cit. on p. 97).

[126] M. J. Lawrence, B. Willke, M. E. Husman, E. K. Gustafson, and R. L. Byer, "Dynamic response of a Fabry – Perot interferometer", *J. Opt. Soc. Am. B*, **16**, 523 (1999) (cit. on p. 97).

[127] M. Rakhmanov and A. Arodzero, "Dynamics of Fabry-Perot resonators with suspended mirrors. I. Nonlinear coupled oscillators", *arXiv:physics/9809038*, (1998) (cit. on p. 97).

[128] C. J. Hood, "Real-time Measurement and Trapping of Single Atoms by Single Photons", PhD thesis (California Institute of Technology, 2000) (cit. on p. 98).

[129] M. Phelps, *First Contact Application and Removal Procedure*, tech. rep. (2012), (cit. on p. 103).

[130] J. F. O'Hanlon, *A User's Guide to Vacuum Technology* (2003) (cit. on pp. 112, 194, 196, 197).

[131] M. Bernardini, S. Braccini, R. De Salvo, D. Di Virgilio, A. Gaddi, A. Gennai, G. Genuini, A. Giazotto, G. Losurdo, H. B. Pan, A. Pasqualetti, D. Passuello, P. Popolizio, F. Raffaelli, G. Torelli, Z. Zhang, C. Bradaschia, R. Del Fabbro, I. Ferrante, F. Fidecaro, P. La Penna, S. Mancini, R. Poggiani, P. Narducci, A. Solina, and R. Valentini, "Air bake-out to reduce hydrogen outgassing from stainless steel", *Journal of Vacuum Science & Technology A*, **16**, 188 (1998) (cit. on p. 112).

[132] S. Avdiaj and B. Erjavec, "Outgassing of Hydrogen from a Stainless Steel Vacuum Chamber", *Materials and Technology*, **46**, 161 (2012) (cit. on p. 112).

[133] J. Leonard, "A Supersolid of Matter and Light", PhD thesis (ETH Zurich, 2017) (cit. on pp. 114, 115, 120).

[134] C. W. de Silva, *Vibration and Shock Handbook* (CRC Press, Boca Raton, Fl, 2005) (cit. on pp. 114, 198).

[135] A. I. Oliva, M. Aguilar, and V. Sosa, "Low- and high-frequency vibration isolation for scanning probe microscopy", *Measurement Science and Technology*, **9**, 383 (1998) (cit. on p. 115).

[136] A. Ryou and J. Simon, "Active cancellation of acoustical resonances with an FPGA FIR filter", *Review of Scientific Instruments*, **88**, 013101 (2017) (cit. on p. 119).

[137] M. Müller, "Realization and characterization of a phase locked laser system for coherent spectroscopy of fiber-coupled cesium atoms", , Diploma Thesis (Johannes Gutenberg (2010) (cit. on p. 137).

[138] N. Šibalić, J. D. Pritchard, C. S. Adams, and K. J. Weatherill, "ARC: An open-source library for calculating properties of alkali Rydberg atoms", *Computer Physics Communications*, **220**, 319 (2017) (cit. on p. 143).

[139] E. W. Streed, A. Jechow, B. G. Norton, and D. Kielpinski, "Absorption imaging of a single atom", *Nature Communications*, **3**, 933 (2012) (cit. on p. 145).

[140] E. Guardado-sanchez, A. Morningstar, B. M. Spar, P. T. Brown, D. A. Huse, and W. S. Bakr, "Subdiffusion and Heat Transport in a Tilted Two-Dimensional Fermi-Hubbard System", *Physical Review X*, **10**, 11042 (2020) (cit. on p. 146).

[141] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith, and Z. Hadzibabic, "Bose-einstein condensation of atoms in a uniform potential", *Physical Review Letters*, **110**, 200406 (2013) (cit. on p. 146).

[142] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schau, T. Fukuhara, I. Bloch, and S. Kuhr, "Single-spin addressing in an atomic Mott insulator", *Nature*, **471**, 319 (2011) (cit. on p. 146).

[143] T. Fukuhara, A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr, "Quantum dynamics of a mobile spin impurity", *Nature Physics*, **9**, 235 (2013) (cit. on p. 146).

[144] A. Rubio-Abadal, J. Y. Choi, J. Zeiher, S. Hollerith, J. Rui, I. Bloch, and C. Gross, "Many-Body Delocalization in the Presence of a Quantum Bath", *Physical Review X*, **9**, 41014 (2019) (cit. on p. 146).

[145] M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman, "Seconds-scale coherence on an optical clock transition in a tweezer array", *Science*, **366**, 93 (2019) (cit. on p. 146).

[146] C. Robens, S. Brakhane, W. Alt, F. Kleissler, D. Meschede, G. Moon, G. Ramola, and A. Alberti, "High numerical aperture (NA = 0 . 92) objective lens for imaging and addressing of cold atoms", *Optics Letters*, **42**, 1043 (2017) (cit. on p. 146).

[147] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, "High-fidelity entanglement and detection of alkaline-earth Rydberg atoms", *Nature Physics*, **16**, 857 (2020) (cit. on p. 148).

[148] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys, "Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms", *Science*, **365**, 775 (2019) (cit. on p. 148).

[149] N. Schlosser, G. Reymond, and P. Grangier, "Collisional Blockade in Microscopic Optical Dipole Traps", *Physical Review Letters*, **89**, 023005 (2002) (cit. on p. 149).

[150] A. Kaufman, "Laser cooling atoms to indistinguishability : Atomic Hong-Ou-Mandel interference and entanglement through spin exchange by", PhD thesis (University of Colorado, 2015) (cit. on pp. 150, 151, 161, 166).

[151] B. Albrecht, Y. Meng, C. Clausen, A. Dareau, P. Schneeweiss, and A. Rauschenbeutel, "Fictitious magnetic-field gradients in optical microtraps as an experimental tool for interrogating and manipulating cold atoms", *Physical Review A*, **94**, 061401 (2016) (cit. on p. 154).

[152] S. Garcia, J. Reichel, and R. Long, "Improving the lifetime in microtraps by using elliptically polarized dipole light", *Phys. Rev. A*, **97**, 023406 (2018) (cit. on p. 154).

[153] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, "Energy distribution and cooling of a single atom in an optical tweezer", *Physical Review A*, **78**, 033425 (2008) (cit. on p. 160).

[154] J. Bochmann, M. Mücke, C. Guhl, S. Ritter, G. Rempe, and D. L. Moehring, "Lossless state detection of single neutral atoms", *Physical Review Letters*, **104**, 203601 (2010) (cit. on pp. 160, 171).

[155] M. J. Rust, M. Bates, and X. Zhuang, "Stochastic optical recontstruction microscopy (STORM) provides sub-diffraction-limit image resolution", *Nature Methods*, **3**, 793 (2006) (cit. on p. 166).

[156] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution", *Science*, **313**, 1642 (2006) (cit. on p. 166).

[157] M. Bates, B. Huang, and X. Zhuang, "Super-resolution microscopy by nanoscale localization of photo-switchable fluorescent probes", *Chemical Biology*, **12**, 505 (2008) (cit. on p. 167).

[158] M. Gierling, P. Schneeweiss, G. Visanescu, P. Federsel, M. Häffner, D. P. Kern, T. E. Judd, A. Günther, and J. Fortágh, "Cold-atom scanning probe microscopy", *Nature nanotechnology*, **6**, 446 (2011) (cit. on p. 167).

[159] M. Lee, J. Kim, W. Seo, H.-g. Hong, Y. Song, R. R. Dasari, and K. An, "Single atoms localized by a nanohole array", *Nature Communications*, **5**, 3441 (2014) (cit. on p. 167).

[160] O. Schwartz, J. J. Axelrod, S. L. Campbell, C. Turnbaugh, R. M. Glaeser, and H. Müller, "Laser phase plate for transmission elecron microscopy", *Nature Methods*, **16**, 1016 (2019) (cit. on p. 167).

[161] S. Bravyi, D. P. DiVincenzo, and D. Loss, "Schrieffer-Wolff transformation for quantum many-body systems", *Annals of Physics*, **326**, 2793 (2011) (cit. on p. 172).

[162] D. Leibfried and C. Monroe, "Quantum dynamics of single trapped ions", *Rev. Mod. Phys.*, **75**, 281 (2003) (cit. on p. 172).

[163] H. Häffner, C. F. Roos, and R. Blatt, "Quantum computing with trapped ions", *Physics Reports*, **469**, 155 (2008) (cit. on p. 172).

[164] N. Brahms and D. M. Stamper-Kurn, "Spin optodynamics analog of cavity optomechanics", *Physical Review A*, **82**, 041804 (2010) (cit. on p. 172).

[165] M. H. Schleier-Smith, I. D. Leroux, and V. Vuletić, "Squeezing the collective spin of a dilute atomic ensemble by cavity feedback", *Physical Review A*, **81**, 021804 (2010) (cit. on p. 173).

[166] M. Kitagawa and M. Ueda, "Squeezed spin states", *Physical Review A*, **47**, 5138 (1993) (cit. on p. 173).

[167] C. Gross, "Spin squeezing, entanglement and quantum metrology with Bose-Einstein condensates", *arXiv:1203.5359*, (2012) (cit. on p. 173).

[168] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen, "Spin squeezing and reduced quantum noise in spectroscopy", *Physical Review A*, **46**, R6797 (1992) (cit. on p. 173).

[169] K. Jacobs, "Twenty open problems in quantum control", *arXiv:1304.0819*, (2013) (cit. on p. 174).

[170] K. Jacobs, *Quantum Measurement Theory and its Applications* (Cambridge University Press, 2014) (cit. on p. 174).

[171] K. Jacobs, X. Wang, and H. M. Wiseman, "Coherent feedback that beats all measurement-based feedback protocols", *New Journal of Physics*, **16**, 073036 (2014) (cit. on p. 174).

[172] H. M. Wiseman and G. J. Milburn, *Quantum measurement and control* (Cambridge University Press, 2010) (cit. on p. 174).

[173] M. Koch, C. Sames, A. Kubanek, M. Apel, M. Balbach, A. Ourjoumtsev, P. W. Pinkse, and G. Rempe, "Feedback cooling of a single neutral atom", *Physical Review Letters*, **105**, 173003 (2010) (cit. on p. 174).

[174] W. P. Smith, J. E. Reiner, L. A. Orozco, S. Kuhr, and H. M. Wiseman, "Capture and release of a conditional state of a cavity QED system by quantum feedback", *Physical Review Letters*, **89**, 1336011 (2002) (cit. on p. 174).

[175] Z. K. Minev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutierrez-Jauregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, "To catch and reverse a quantum jump mid-flight", *Nature*, **570**, 200 (2019) (cit. on p. 174).

[176] G. Crowder, H. Carmichael, and S. Hughes, "Quantum trajectory theory of few-photon cavity-QED systems with a time-delayed coherent feedback", *Physical Review A*, **101**, 023807 (2020) (cit. on p. 174).

[177] J. E. Reiner, W. P. Smith, L. A. Orozco, H. M. Wiseman, and J. Gambetta, "Quantum feedback in a weakly driven cavity QED system", *Physical Review A*, **70**, 023819 (2004) (cit. on p. 174).

[178] W. P. Smith and L. A. Orozco, "Quantum feedback in a non-resonant cavity QED system", *Journal of Optics B: Quantum and Semiclassical Optics*, **6**, 127 (2004) (cit. on p. 174).

[179] G. Mazzucchi, W. Kozlowski, S. F. Caballero-Benitez, and I. B. Mekhov, "Collective dynamics of multimode bosonic systems induced by weak quantum measurement", *New Journal of Physics*, **18**, 073017 (2016) (cit. on p. 174).

[180] G. Mazzucchi, W. Kozlowski, S. F. Caballero-Benitez, T. J. Elliott, and I. B. Mekhov, "Quantum measurement-induced dynamics of many-body ultracold bosonic and fermionic systems in optical lattices", *Physical Review A*, **93**, 023632 (2016) (cit. on p. 174).

[181] G. Mazzucchi, S. F. Caballero-Benitez, and I. B. Mekhov, "Quantum measurement-induced antiferromagnetic order and density modulations in ultracold Fermi gases in optical lattices", *Scientific Reports*, **6**, 31196 (2016) (cit. on p. 174).

[182] T. J. Elliott and I. B. Mekhov, "Engineering many-body dynamics with quantum light potentials and measurements", *Physical Review A*, **94**, 013614 (2016) (cit. on p. 174).

[183] D. A. Ivanov, T. Y. Ivanova, and I. B. Mekhov, "Incoherent quantum feedback control of collective light scattering by Bose-Einstein condensates", *arXiv:1601:02230*, (2016) (cit. on p. 174).

[184] J. Kohler, "Optodynamical Measurement and Coupling of Atomic Motion and Spin", PhD thesis (University of California, Berkeley, 2018) (cit. on p. 192).

[185] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, "Simplified system for creating a Bose-Einstein condensate", *Journal of Low Temperature Physics*, **132**, 309 (2003) (cit. on p. 197).

[186] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics* (John Wiley & Sons, Inc., New York, NY, 1991) (cit. on pp. 201, 204).

[187] J. W. Goodman, *Introduction to Fourier Optics*, 3rd (Roberts & Co. Publishers, Englewood, CO, 2005) (cit. on p. 201).

Appendix A

Fourier Transform Conventions

Convention choices abound for Fourier transforms. Here I will state the conventions taken in this document.

For temporal functions, I take the Fourier transform and its inverse to be defined as

$$\tilde{g}(\omega) = \int g(t)e^{+i\omega t}dt \qquad g(t) = \frac{1}{2\pi}\int \tilde{g}(\omega)e^{-i\omega t}d\omega. \tag{A.1}$$

Under this convention, the Fourier transform $\tilde{g}(\omega)$ of a function like $g(t) = e^{-i\omega_0 t}$, with positive ω_0, has support for positive ω, so we say it is a positive frequency function.

For spatial functions in d-dimensions, I take the Fourier transform and its inverse to be defined as

$$\tilde{g}(\boldsymbol{k}) = \int g(\boldsymbol{x})e^{-i\boldsymbol{k}\cdot\boldsymbol{x}}d\boldsymbol{x} \qquad g(\boldsymbol{x}) = \frac{1}{(2\pi)^d}\int \tilde{g}(\boldsymbol{k})e^{+i\boldsymbol{k}\cdot\boldsymbol{x}}d\boldsymbol{k}. \tag{A.2}$$

Under this convention, the Fourier transform $\tilde{g}(\boldsymbol{k})$ of a function like $g(\boldsymbol{x}) = e^{+i\boldsymbol{k}_0\cdot\boldsymbol{x}}$, with positive $\boldsymbol{k}_0$ (meaning all components of $\boldsymbol{k}_0$ are positive), has support for positive $\boldsymbol{k}$, so we say it is a positive frequency function. The opposite sign convention for temporal and spatial functions is chosen so that a right-propagating wave, with positive k_0 and ω_0, is represented by a function like $e^{i(k_0 z - \omega_0 t)}$.

The two conventions above deal with angular temporal or spatial frequencies ω or $\boldsymbol{k}$. Sometimes it is useful to work with cyclic frequencies f or $\boldsymbol{\nu}$. In these cases I use the following conventions

$$\tilde{g}(f) = \int g(t)e^{+i2\pi ft}dt \qquad g(t) = \int \tilde{g}(f)e^{-i2\pi ft}df, \tag{A.3}$$

$$\tilde{g}(\boldsymbol{\nu}) = \int g(\boldsymbol{x})e^{-i2\pi\boldsymbol{\nu}\cdot\boldsymbol{x}}d\boldsymbol{x} \qquad g(\boldsymbol{x}) = \int \tilde{g}(\boldsymbol{\nu})e^{+i2\pi\boldsymbol{\nu}\cdot\boldsymbol{x}}d\boldsymbol{\nu}. \tag{A.4}$$

These conventions have the same conditions for positive/negative frequency functions as above. The main advantage of working with cyclic frequencies is that there are no prefactors

of $1/2\pi$ to keep track of in the definition of the Fourier or inverse Fourier transform. When working with signal processing (for example processing the measured heterodyne signal or thinking about imaging systems as spatial signal processors), I prefer to work with the cyclic frequency Fourier transforms.

Appendix B

Multilevel Atomic Transition Dipole Elements

In this appendix, I describe the calculation of transition dipole elements for multilevel atoms in more detail than is given in Sec. 1.3.2. A full treatment of this multilevel atom would include a transition dipole element $\boldsymbol{d}_{ij}$ between each pair of ground and excited states $|g_i\rangle$, $|e_j\rangle$. The $\boldsymbol{d}_{ij}$ between different hyperfine sublevels are related by geometric factors having to do with the symmetries of the states in question and can be calculated using the Wigner-Eckart theorem and related identities. An alternative, and more thorough, treatment of the extension of cQED to the multilevel and multi-atom case can be found in Ref. [184]. General information about dipole transitions and angular momentum rules can be found in [61, 65, 66].

Here I will specify a few of these $\boldsymbol{d}_{ij}$ which will be important for this work. I follow Refs. [61, 66]. First, we have the fine-structure reduced dipole element

$$d_{JJ'} = \langle J||\hat{\boldsymbol{d}}||J'\rangle = \sqrt{\frac{3\pi\epsilon_0\hbar c^3}{\omega_0^3}}\sqrt{\frac{2J'+1}{2J+1}}\sqrt{\Gamma}, \tag{B.1}$$

where ϵ_0 is the permittivity of free space, c is the speed of light, ω_0 is the transition frequency for the $J \to J'$ atomic transition, and Γ is the decay rate out of the J' excited state. This parameter is of central importance because it can be experimentally extracted by measuring the excited state decay rate Γ.

We can calculate corresponding hyperfine reduced dipole transition elements by

$$d_{FF'} = \langle F||\hat{\boldsymbol{d}}||F'\rangle = d_{JJ'}\,(-1)^{F'+J+1+I}\,\sqrt{(2F'+1)(2J+1)}\begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix}, \tag{B.2}$$

where $I = 3/2$ is the nuclear spin for ^{87}Rb and the quantity in brackets is a Wigner 6-j symbol. From $d_{FF'}$, we can calculate the transition dipole elements between any two hyperfine sublevels driven by light with polarization $q = 0, \mp 1$, corresponding to π, and $\sigma^{\pm}$

polarization:

$$d^{FF'}_{m_F m'_F q} = d_{FF'} (-1)^{F'-1+m_f} \sqrt{2F+1} \begin{pmatrix} F' & 1 & F \\ m'_F & q & -m_F, \end{pmatrix},$$ (B.3)

where the symbol in parenthesis is a Wigner 3-j symbol.

In general, one must consider the effect of coupling between any hyperfine ground state to any hyperfine excited state. However, there are a few special combinations of these elements which are particularly useful. The D_2 transition supports a cycling transition: $|F = 2, m_F = +2\rangle \to |F = 3, m_F = +3\rangle$. This transition has the largest transition dipole element given by

$$d^{2,3}_{2,3,-1} = d_0 = \sqrt{\frac{2J+1}{2J'+1}} d_{JJ'} = \frac{1}{\sqrt{2}} d_{JJ'}.$$ (B.4)

I have notated this particular transition element as d_0. This will be the transition strength to which all other transition strengths are compared, and, as described in Sec. 1.3.2, which will be used to calculate g_0.

The other useful combination of transition dipole elements is the effective far-detuned transition dipole element

$$d_{\text{eff}} = \frac{1}{\sqrt{3}} d_{JJ'}.$$ (B.5)

When the driving field is detuned much further than the hyperfine splittings this element can be used to calculate the scalar Stark shift due to the driving field. For Alkali atoms, there is no tensor Stark shift of the ground state (because $J = 1/2$) and, if the driving field is π-polarized, there is no vector Stark shift. This element will be used when making calculations about linearly polarized ODTs.

The numerical values for these elements are [66]:

$$\begin{aligned} d_{JJ'} &= 3.584 \times 10^{-29} \, \text{C m} = 4.228 \, ea_0, \\ d_0 &= 2.534 \times 10^{-29} \, \text{C m} = 2.989 \, ea_0, \\ d_{\text{eff}} &= 2.069 \times 10^{-29} \, \text{C m} = 2.441 \, ea_0. \end{aligned}$$ (B.6)

Here a_0 is the Bohr radius and e is the electron charge.

Appendix C

Vacuum Pressure Calculation

In this appendix, I will present a general formalism for calculating vacuum pressures in UHV chambers using the 2D and 3D MOT chambers as an example case. The formalism presented here was used throughout the design of both the atom preparation and science chambers to ensure we would achieve sufficiently low vacuum pressures.

In a sealed vacuum chamber, the ultimate pressure, at a fixed temperature T, is set by a balance of particles being extracted from the system by a vacuum pump (pump speed S expressed in L/s) and particles flowing into the system [130]. We will consider particle flow into the gaseous (fluid) system via two mechanisms: first, due to sublimation at a fixed temperature from a solid sample, which results in a fixed vacuum pressure in that chamber section (expressed in torr) and second, due to outgassing from surfaces in the chamber including the SSL walls. Outgassing for a particular material is often quantified with units of $\mathrm{torr} \cdot \mathrm{L} \cdot \mathrm{s}^{-1} \cdot \mathrm{cm}^{-2}$ and the total outgassing rate Q can be calculated by multiplying this material constant by the material surface area to give a flow rate in $\mathrm{torr} \cdot \mathrm{L} \cdot \mathrm{s}^{-1}$. Finally, pressure differentials can be supported within different sections of the vacuum chamber due to apertures with finite gas conductance C (expressed in L/s).

An electrical circuit analogy can be made to aid in calculations of vacuum chamber pressures. In this analogy, different chambers (identified as distinct volumes separated by small apertures) can be identified with conductive nodes in a circuit where the chamber pressure P is analogous to the voltage at that node in the circuit. Conductances are represented as resistors between nodes (chambers), vacuum pumps are represented as resistors to ground, sublimating sources are represented as voltage sources (with their negative terminal grounded), and outgassing sources are represented as current sources from ground. Part of the intuition for these identifications is that ground represents a source or sink for particles flowing into or out of the gaseous state (the system) within the chamber.

See Fig. C.1 for a geometric and electrical schematic diagram of the 3D and 2D MOT chambers. The two chambers are separated by a conductance C_{diff} of the differential pumping tube. The 3D MOT chamber is attached to a vacuum pump with speed S through a conductance C_{pump}, representing the vacuum path between the 3D MOT chamber and the pump. The 2D MOT chamber includes a sublimating source and the 3D MOT chamber
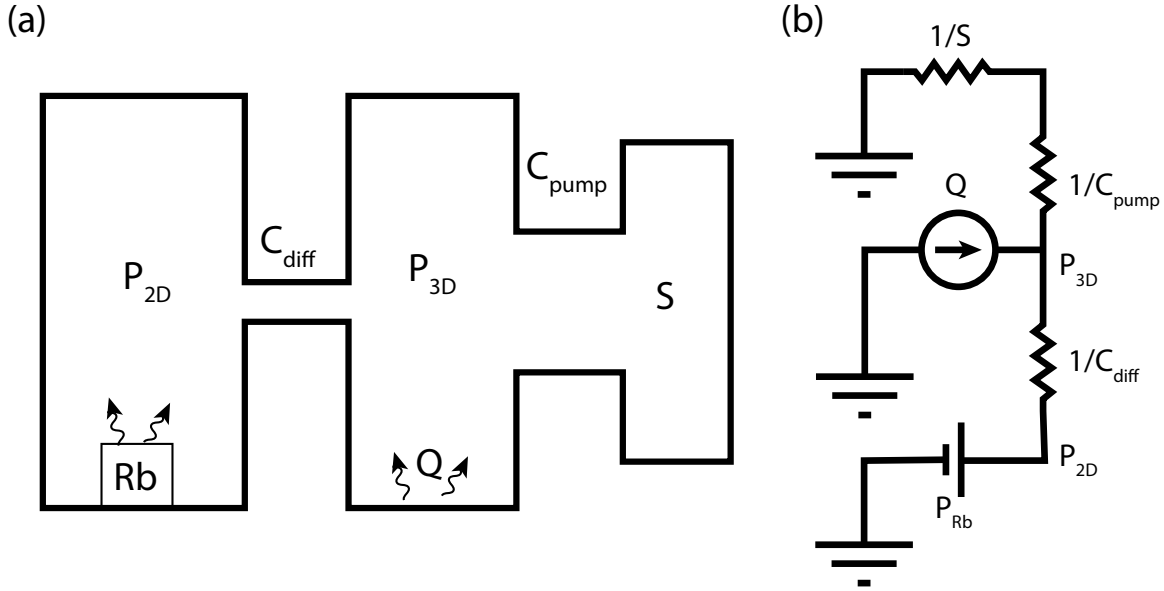
(a)

(b)

Figure C.1: Circuit analog for differential pumping between 2D and 3D MOT chambers. (a) The 2D MOT chamber has a gas load due the sublimation of a mass of Rb and the 3D MOT chamber has a gas load Q due to outgassing of the SSL walls. The two chambers are attached via the differential pumping tube conductance C_{diff} and the 3D MOT chamber is pumped out by a pump with pump speed S through a conductance C_{pump}. (b) Circuit analog for the schematic representation of the chamber in (a). The pump is represented by a resistor to ground with value $\frac{1}{S}$, the outgassing load is represented as a current source from ground with value Q, the sublimating Rb is represented as a voltage supply from ground with value P_{Rb}, and the conductances are represented as resistors between nodes whose voltages represent chamber pressures.

includes an outgassing source. The corresponding electrical circuit appears in Fig. C.1(b).

Given reasonable values for the various components, we may roughly estimate the vacuum pressure within the 3D MOT chamber. We desire for the gas pressure within the 3D MOT chamber to be well below the 1×10^{-10} torr level, ideally approaching the 1×10^{-11} torr level. In this calculation, I determine whether the differential pumping provided by the 2D MOT differential pumping tube is sufficient to isolate the 3D MOT chamber from the sublimating Rb mass.

To determine the gas load presented by the Rb mass to the 3D MOT chamber, we analyze the circuit in Fig. C.1(b) while setting $Q = 0$. Under those conditions the first bit of intuition we identify is that the pump and the pump conductance can be combined into

a single effective pump speed or resistance by adding the resistors in series:

$$\frac{1}{S_{\text{eff}}} = \frac{1}{S} + \frac{1}{C_{\text{pump}}},$$

$$S_{\text{eff}} = S || C_{\text{pump}} = \frac{SC_{\text{pump}}}{S + C_{\text{pump}}}. \tag{C.1}$$

Note that for $C_{\text{pump}} \ll S$, we can neglect C_{pump} and it is as if the pump S is attached directly to the chamber. However, for $C_{\text{pump}} \gg S$, S_{eff} is dominated by C_{pump}. This tells us that it is a waste to use a pump with a large pump speed if the effective pump speed is just going to be limited by conductance anyways.

We see that the remaining circuit is a voltage divider consisting of the effective pump speed and the differential pumping conductance giving

$$P_{3D} = \frac{\frac{1}{S_{\text{eff}}}}{\frac{1}{S_{\text{eff}}} + \frac{1}{C_{\text{diff}}}} P_{2D}$$

$$= \frac{C_{\text{diff}}}{C_{\text{diff}} + S_{\text{eff}}} P_{\text{Rb}}. \tag{C.2}$$

The differential pumping tube is designed to have a low conductance so that $C_{\text{diff}} \ll S_{\text{eff}}$ so we can approximate

$$P_{3D} \approx \frac{C_{\text{diff}}}{S_{\text{eff}}} P_{\text{Rb}}. \tag{C.3}$$

We see that ratio of pressures between the two chambers is given by the ratio of the differential pumping conductance to the effective pump speed on the 3D MOT chamber.

We can obtain an estimate of the Rb gas load in the 3D MOT chamber by estimating these parameters. First, the formula for conductance of a high-aspect-ratio tube is given by [130]:

$$C = \frac{4}{3} \sqrt{\frac{2\pi k_B T}{m}} \frac{r^3}{L} = \frac{2\pi}{3} \bar{v} \frac{r^3}{L}, \tag{C.4}$$

where k_B is the Boltzmann constant, $\bar{v}$ is the mean thermal velocity, and r and L are the radius and length of the differential pumping tube. The Rb conductance of the differential pumping tube is dominated by the first section with a small radius ($R = 750\,\mu$m, $L = 31\,$mm, see Fig. 2.2):

$$C_{\text{diff}} = 7.6 \times 10^{-3}\,\text{L/s}. \tag{C.5}$$

We will assume that the effective pump speed S_{eff} is limited by the vacuum conductance of the tube from the central MOT chamber over to the pump arm ($R = 17\,$mm, $L = 83\,$mm):

$$S_{\text{eff}} = 33\,\text{L/s}. \tag{C.6}$$

For a room-temperature Rb vapor pressure $P_{\text{Rb}} = 3 \times 10^{-7}$ torr [66], we find

$$P_{\text{3D}} \approx 6.9 \times 10^{-11} \text{ torr.} \tag{C.7}$$

This is within an order of magnitude of what we were hoping.

Eq. C.4 assumes operation in the molecular flow regime and that particles that collide with the wall scatter off at a random angle [130]. Rb, in fact, has the property that it typically sticks to the walls of the chamber upon collision [185]. In terms of the circuit analogy this can be thought of in at least two ways: as a reduction of the conductance of the differential pumping tube, or as a pump for Rb that is distributed throughout the circuit - a distributed resistance to ground. In any case, this should be convincing that we've given an overestimate for the Rb gas load to the 3D MOT chamber and that we have sufficient differential pumping. Indeed, experimentally, we measure a vacuum pressure at the 1×10^{-11} torr level in our 3D MOT chamber and observe vacuum-limited atomic trap lifetimes of 10s of s.

Appendix D

An Electrical Circuit Analogy for Mechanical Systems

In Sec. 5.2 while describing the E6 vibration isolation system, I indicate that the three stages of vibration isolation can be thought of as three cascaded low pass filters that suppress mechanical fluctuations. In this appendix, I will show that the analogy between a mechanical circuit and an electrical circuit can be made to be exact. I'll give a quick introduction to this analogy here, a more thorough treatment can be found in Ref. [134]. The reason the analogy can be made between electronics and mechanical systems is because both systems can be described by second order differential equations and because both systems can be decomposed into lumped elements that represent either nodes or links between nodes. In an electrical circuit, nodes are conductive wires or terminals, and circuit components act as links (or edges) between those nodes. Examples of circuit components are resistors, capacitors, inductors, and current or voltage supplies. To every edge corresponds a voltage difference and a current. The relationship between voltage **across** a node and the current **through** a node is constrained by the type of that node.

For example, in the Laplace domain, the voltage and current through a capacitor[1] are related by $V = \frac{I}{sC} = \frac{Q}{C}$.[2] The constitutive relations for each edge, the boundary conditions determined by voltage and current sources, a current conservation law at nodes, and a voltage conservation law around loops (Kirchoff's rules), provide enough analytic equations to solve for the voltage and current at each edge and thus 'solve' the circuit.

An exact analogy can be made with mechanical systems. Just like equations of motion for the voltage in a circuit can be expressed as a function of the current itself and the time derivative and integral of the current, the forces in a mechanical system can be expressed as a function of the velocity of masses within a circuit, and the integral (position) and derivative (acceleration) of the velocity of those masses. It can be seen that force (due to Newton's

---

[1]In this section, an upper case C refers to a capacitance while a lower case c refers to a mechanical damping coefficient. Likewise, upper case V refers to a voltage while lower case v refers to a velocity.

[2]Here I recall that current is the time derivative of charge $I = \dot{Q}$ and the relationship between a function and its derivative in the Laplace domain: $\dot{Q} \rightarrow sQ$.

third law) will satisfy a node conservation law similar to current, and that velocity will satisfy a loop conservation law similar to voltage. We will thus identify force as a **through** variable and velocity as an **across** variable.

Typical circuit components will be springs and dampers connecting one mass to another, and that have constitutive equations such as $v = F\frac{s}{k}$ and $v = \frac{F}{c}$. If a mass is driven with a certain velocity, it is as if that mass is driven by a constant voltage source in the electrical analogy. In an electrical circuit schematic, nodes represent conductive pieces of metal to which other components (edges) can be connected. In a mechanical circuit schematic, nodes represent discrete pieces of mass, each of which can have a particular velocity relative to some inertial reference frame. The mechanical circuit has the curious feature that every mass is subject to Newton's law: $F = ma \rightarrow v = \frac{F}{sm}$. This acts as a constitutive relation that relates the force on every mass to the derivative of the velocity of the mass relative to some ideal stationary inertial reference mass. This ideal reference mass is the ground node in the mechanical circuit. Newton's law tells us that every mass in the system needs to be connected to this 'ground' reference node via a capacitor with a value of m.

Here are the $v - F$ relationships for inertia, dampers, and springs and their analogous electric $I - V$ relationships for capacitors, resistors, and inductors:

$$v = \frac{F}{sm} \leftrightarrow \frac{I}{sC} = V,$$
$$v = \frac{F}{c} \leftrightarrow IR = V,$$
$$v = F\frac{s}{k} \leftrightarrow ILs = V. \tag{D.1}$$

We thus see that Newton's law enforces that masses are connected to ground via a capacitor with a value of m. Springs are like inductors[3] between masses where the inductance is given by $\frac{1}{k}$. Dampers are like resistors between masses whose value is given by $\frac{1}{c}$. Note that above I have spoken about **through** and **across** variables. We can also consider the idea of **effort** and **flow** variables. We see that the analogy mixes up these different notions. While current and force are both **through** variables, we have that current is a **flow** variable while force is an **effort** variable and vice-versa for voltage and velocity. Impedance is generally defined as

$$Z = \frac{\textbf{effort variable}}{\textbf{flow variable}}. \tag{D.2}$$

This means that according to the node and loop conservation rules impedance adds in series in electrical circuits while impedance adds in parallel in mechanical circuits. What we are rather interested in is the ratio of the **across** variable to the **through** variable that is given by mobility in the mechanical case (mobility is the inverse of mechanical impedance) and normal impedance in the electrical case[4].

---

[3]Isn't it convenient that their schematic symbols are so similar?

[4]Unfortunately I haven't come across a standard term that generally refers to the ratio of **across** to **through** variable for all electrical analogies.

In the circuit analogy when one mass is attached to another via a spring it looks like an electrical low pass filter between the two corresponding circuit nodes. This is because the spring acts like an inductor between the two nodes but the second node is shunted to ground via its Newton's-law capacitor. If we build a chain of masses, each attached to the previous by a spring, this is equivalent to a series of low pass filters, each of which suppresses high frequency noise. The E6 vibration isolation system, shown schematically in Fig. 5.3 is just such a chain of mass connected by springs.

Appendix E

Diffraction Limited Optical Systems

In this appendix, I will lay out some basic principles in Fourier optics and use them to develop an understanding of the so-called diffraction limit. The diffraction limit tells us that the resolution of a far-field imaging system can never be better than approximately λ, the wavelength of light used for imaging. The diffraction limit also tells us that if, instead, we are using an imaging system to create small focused spots, that the size of those spots can only be about as small as λ, but no smaller.

For me there are two meanings to the diffraction limit, (1) a fundamental limit having to do with the electromagnetic wave equation and (2) a technical limit having to do with a particular optical system involving lenses. First, I will first describe this fundamental diffraction limit. Next, I will describe this technical diffraction limit having to do with (even ideal) imaging systems. Finally, I will describe how in a real imaging system aberrations can and must be quantified to determine if the imaging system realizes diffraction limited performance. These ideas are treated more thoroughly in Refs. [186, 187].

E.1 The Fundamental Diffraction Limit

The scalar wave equation for the electric field is given by

$$\left(c^2 \nabla^2 - \frac{\partial}{\partial t} \right) E = 0. \tag{E.1}$$

This wave equation is satisfied by the complex plane wave solution

$$U^{\mathrm{PW}}(\boldsymbol{r}, t) = e^{i(\boldsymbol{k} \cdot \boldsymbol{r} - \omega t)}, \tag{E.2}$$

$$\tag{E.3}$$

with $\frac{\omega^2}{|\boldsymbol{k}|^2} = c^2$ with $|\boldsymbol{k}| = k = \frac{2\pi}{\lambda}$. Going forward, for convenience and in keeping with Ref. [186], I will use cyclic spatial frequencies rather than angular: $\boldsymbol{\nu} = \frac{1}{2\pi} \boldsymbol{k}$ with $\nu = |\boldsymbol{\nu}| = \frac{1}{\lambda}$.

When considering an optical beam, we typically single out a particular direction and call that the dominant propagation axis and align our coordinate z with that axis. Notice that

the value of $\nu_z = \sqrt{\nu^2 - \nu_x^2 - \nu_y^2}$ is constrained by the values of ν_x and ν_y through $|\boldsymbol{\nu}| = \frac{1}{\lambda}$. We see, then, that the value of the plane wave for all space and time is fully determined by the values of ν_x and ν_y, or, equivalently, is fully determined by the value of the wave on the $z = 0$ plane at $t = 0$:

$$U_{\nu_x,\nu_y}^{\text{PW}}(x, y, 0, 0) = e^{i(2\pi\nu_x x + 2\pi\nu_y y)}. \tag{E.4}$$

Recalling that any superposition of solutions of the wave equations is also a wave equation, and motivated by knowledge of the Fourier transform we write

$$E(x, y, z, t) = \int \tilde{E}(\nu_x, \nu_y) U_{\nu_x,\nu_y}^{\text{PW}}(x, y, z, t) d\nu_x d\nu_y,$$

$$E(x, y, 0, 0) = \int \tilde{E}(\nu_x, \nu_y) e^{i(2\pi\nu_x x + 2\pi\nu_y y)} d\nu_x d\nu_y. \tag{E.5}$$

Here we have done nothing interesting beyond summing up a number of plane waves. What is more interesting, is interpreting this expression as the inverse Fourier transform of a wavevector space function, $\tilde{E}(\nu_x, \nu_y)$, and writing down the corresponding Fourier transform expression:

$$\tilde{E}(\nu_x, \nu_y) = \int E(x, y, 0, 0) e^{-i(2\pi\nu_x x + 2\pi\nu_y y)} dx dy. \tag{E.6}$$

From Eqs. E.5 and E.6, and recalling that any sufficiently well-behaved function $\tilde{E}(\nu_x, \nu_y)$ has a Fourier transform, we can draw the profound conclusion that any solution $E(x, y, z, t)$ to the wave equation can be expanded as a sum of plane waves of the form $e^{i(2\pi\nu_x x + 2\pi\nu_y y)}$. The coefficients, $\tilde{E}(\nu_x, \nu_y)$, can be found by taking the Fourier transform of $E(x, y, 0, 0)$. This establishes the basic correspondence between plane waves and a Fourier decomposition of an optical field. This correspondence forms the basis of Fourier optics.

However, there is a slight wrinkle in this story coming from the constraint that $\nu = \frac{1}{\lambda}$. If we require that each of ν_x, ν_y, ν_z is real, as is the case for a propagating plane wave, then we see that we must have $\nu_x, \nu_y, \nu_z \leq \nu$.

What happens then if we specify the electric field at the $z = 0$ plane, $E(x, y, 0, 0)$, to include transverse spatial frequencies ν_x, ν_y greater than $\frac{1}{\lambda}$? Such a situation can easily arise, for example, if a normal plane wave is made to impinge upon an opaque mask with a hole in it whose diameter is much less than λ. The short answer is that such a situation poses no problem for the machinery described above. Eq. (E.5) can still be used to determine $E(x, y, z, t)$ through the Fourier transform Eq. (E.6), even if $E(x, y, 0, 0)$ contains high spatial frequencies.

The strange feature is that $\nu^2 - \nu_x^2 - \nu_y^2$ will be negative, meaning that ν_z will be imaginary. The term $e^{i2\pi\nu_z z}$ then represents exponential decay at a rate $2\pi\nu_z$ which increases as the amount by which ν_x and ν_y exceed ν is increased. Such an optical field, with complex ν_z, is referred to as an evanescent field or near-field because it only has appreciable support in a small region of space near to where the field was generated. Such fields are most commonly

encountered in the case of total internal reflection and are of central importance in near-field super-resolution microscopy techniques.

We can get a nice interpretation for these evanescent fields by defining the transfer function for free space as a function of transverse spatial frequency:

$$\frac{U^{\text{PW}}_{\nu_x,\nu_y}(x,y,z,0)}{U^{\text{PW}}_{\nu_x,\nu_y}(x,y,0,0)} = T_z(\nu_x,\nu_y) = e^{i2\pi\nu_z z} = e^{i2\pi\sqrt{\nu^2-\nu_x^2-\nu_y^2}}. \tag{E.7}$$

We see that the transfer function that describes propagation from the plane $z = 0$ to another plane is given by $T_z(\nu_x, \nu_y)$. For plane waves with $\nu_x, \nu_y < \nu$, we see that the plane wave simply collects a phase factor given by $2\pi\nu_z z$. However, for evanescent waves, we see that the field is attenuated by a factor $e^{-2\pi|\nu_z|z}$.

The first, fundamental, version of the diffraction limit can be succinctly stated as 'for transverse spatial frequencies greater than ν, the transfer function of free space is exponentially suppressed with distance z away from the source plane'. This has the implication that if, on a source plane, the electromagnetic field has structure with spatial frequencies greater than $\frac{1}{\lambda}$ (that is features smaller than λ), then information about those features will not be propagated into the far field. I emphasize that this definition is not a statement about, for example, the resolution of any particular imaging system but is rather a statement about the propagation of optical fields *in free space* under the wave equation. It should be clear at this point that it will be impossible for any linear, far-field imaging system to accurately image a sample that has spatial features much smaller than λ, the wavelength of light used for imaging.

E.2 The Particular Diffraction Limit for Ideal Imaging Systems

The second version of the diffraction limit does have to do with the details of specific imaging systems. This version of the diffraction limit will give us the conclusion that the resolution of any optical system depends on the appropriately defined NA of that imaging system. There are a number of optical schemes we could imagine but they will all give us the same result: that the length scale of the minimum resolvable feature is proportional to λ/NA.

Here are a few optics schemes that could be analyzed to explore the physics of the diffraction limit:

- Consider a lens of diameter D and focal length f. Imagine a plane wave impinges on the lens. What is the size of the resultant focused spot distance f away from the lens?

- Imagine a $4f$ imaging system with a pinhole of diameter D placed in the Fourier plane. What will be the spot size of a re-imaged point source? Equivalently, what is the spatial frequency cutoff due to the spatial filter?
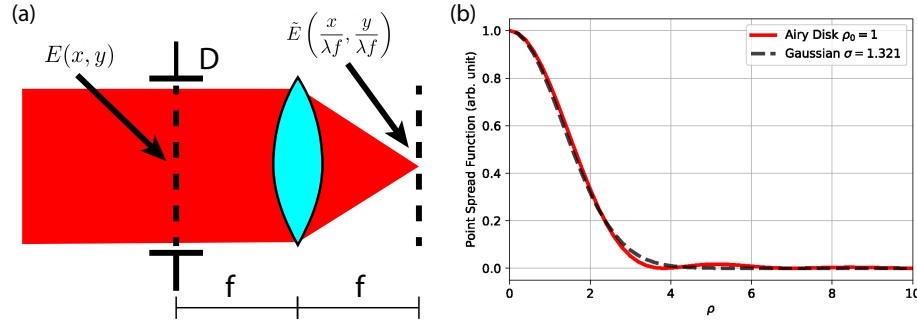
Figure E.1: (a) Fourier imaging schematic. A beam with profile $E(x, y)$ impinges from the left passing through an iris of diameter D is focused by a lens with focal length f. In the paraxial limit, the spatial profile on the output plane is proportional to a scaled version of the Fourier transform of the input profile at the iris: $\tilde{E}\left(\frac{x}{\lambda f}, \frac{y}{\lambda f}\right)$. (b) Airy disk point spread function with scale factor $\rho_0 = 1$ (red) and Gaussian fit to this Airy disk (black). The best fit yields $\sigma = 1.321\rho_0$.

- Consider a single lens imaging system satisfying the image formation equation $\frac{1}{d_i} = \frac{1}{d_o} + \frac{1}{f}$. Imagine a point source is imaged by this system. What will be the size of the resultant spot in the image plane?

Other possible schemes can be considered. Here I'll consider the first bullet point above which is also depicted in Fig. E.1(a).

One of the main principles in Fourier optics is that, in the paraxial limit, an ideal lens implements an optical Fourier transform [186]. Specifically — ignoring effects due to finite aperture sizes — if a field $E(x, y)$ with Fourier transform $\tilde{E}(\nu_x, \nu_y)$ appears at an input plane, and a lens with focal length f is placed f away from that input plane, then, f away on the other side of the lens, the field will be given by[1] $\frac{1}{(\lambda f)^2}\tilde{E}\left(\frac{x}{\lambda f}, \frac{y}{\lambda f}\right)$. That is, the Fourier transform of $E(x, y)$ is mapped into the spatial output plane.

Consider again Fig. E.1(a). If a plane wave impinges on a circular lens with diameter D, then the beam will be clipped as if it passed through an aperture at the back focal plane of the lens.[2] The resultant spot will then be given by the Fourier transform of the circular aperture. The Fourier transform of a circle is an Airy disk. Defining $\rho_\nu = \sqrt{\nu_x^2 + \nu_y^2}$, we can write down the Fourier transform of a circular aperture:

$$\tilde{E}(\nu_x, \nu_y) = \frac{D}{2\rho_\nu} J_1\left(\pi \rho_\nu D\right), \tag{E.8}$$

---

[1]The pre-factor is unimportant conceptually but ensures conservation of energy.

[2]Note that a lens of diameter D is not strictly equivalent to an ideal lens plus an iris in the back focal plane. However, the difference is negligible especially for small angle beams and the conceptual conclusions are unchanged.

where $J_1(x)$ is a first order Bessel function. The spot resulting from the plane wave will then be given by (with $\rho = \sqrt{x^2 + y^2}$):

$$E_{\text{out}}(\rho) = \frac{1}{\lambda f} \frac{D}{2\rho} J_1\left(\frac{\pi \rho D}{\lambda f}\right) = \frac{1}{2\pi \rho_0^2} \frac{J_1\left(\rho/\rho_0\right)}{\rho/\rho_0}. \tag{E.9}$$

We see that the output field is given by an Airy disk characterized by a scale factor $\rho_0 = \frac{\lambda f}{\pi D}$. The spatial pattern that is generated by an imaging system when it is illuminated with a plane wave is referred to as the PSF. The Airy disk PSF is shown in Fig. E.1(b).

For the simple imaging system considered here, the NA is defined by $\text{NA} = \sin(\theta) = \sin\left(\arctan\left(\frac{D}{2f}\right)\right)$. For small angles, this gives $\text{NA} \approx \frac{D}{2f}$ so the scale factor can be written as $\rho_0 = \frac{\lambda}{2\pi \text{NA}} = \frac{1}{k\text{NA}}$. We see clearly that the width of the spot is inversely proportional to NA. The Rayleigh criterion defines two spots as being 'resolvable' if they are spaced such that the second spot is further from the center of the first spot than the radius of the first zero of the Airy disk.[3] The first zero of the first order Bessel function occurs at approximately $z_1 = 3.8317$. The distance required for the Rayleigh criterion is then

$$\rho_{\text{Rayleigh}} = \rho_0 z_1 \approx 0.61 \frac{\lambda}{\text{NA}} \approx \frac{\lambda}{2\text{NA}}. \tag{E.10}$$

This is the diffraction limit for an imaging system. Here it has been derived by determining the minimum size of a focused spot but a similar argument works for the reverse problem of determining when two point sources can be resolved by an imaging system.

Note our derivation of the Rayleigh criterion relied at least in two places on making the paraxial approximation. The final result, expressed in terms of NA, apparently holds even beyond the paraxial approximation.[4] For a lens (with equal index of refraction media on the input and output sides) the maximum achievable NA is 1 meaning the ultimate resolution limit is of order λ, in accordance with what was found in the first section on the fundamental diffraction limit.

In practice, we occasionally characterize imaging systems by fitting their PSFs with a Gaussian function and extracting the Gaussian standard deviation σ. It is useful, then, to be able to convert between Gaussian σ and the size of the corresponding Airy disk. In Fig. E.1(b) I show the result of fitting a 2D Gaussian to an Airy disk with $\rho_0 = 1$. The best fit Gaussian has $\sigma = 1.321$. So from this we can conclude

$$\sigma = 1.321 \rho_0 = 0.345 \rho_{\text{Rayleigh}} \approx 0.21 \frac{\lambda}{\text{NA}}. \tag{E.11}$$

---

[3]Whether two spots are resolvable is not a matter only of the size of the PSF compared to the spacing but also of the SNR with which the spots are imaged.

[4]I've personally had a surprisingly difficult time finding a convincing derivation of this result!

E.3 Diffraction Limited Optical Systems and Aberrations

The third, and final, concept I would like to discuss regarding diffraction limits is that of a diffraction limited optical system. In simple terms, a diffraction limited system is one that can resolve or generate spots approaching the size of the diffraction limit, described in the previous section, of approximately $\frac{\lambda}{2\text{NA}}$. However, any real imaging system will exhibit optical aberrations that spoil the resolving power of the system. These aberrations will change the PSF from being a diffraction limited Airy disk to being smeared and stretched, sometimes asymmetrically, to a size larger than $\frac{\lambda}{2\text{NA}}$.

I'll describe now two ways in which aberrations can be quantified in a ray tracing picture. The picture I'll describe here coincides with how ray tracing software, such as Zemax, can quantify aberration. Every (cylindrically symmetric) imaging system has a so-called field stop that limits the radius of the field of view of an imaging system and an aperture stop that limits the angles of light emanating from the object or image plane that will pass through the imaging system. The field stop can be imaged by the imaging system into the object or image space to define the entrance and exit pupils. The imaging system is analyzed by launching a bundle of rays from the entrance pupil and tracing them through the system until they intersect the exit pupil. The so-called chief ray is the ray that passes through the center of the field stop (and therefore the center of the entrance and exit pupils as well).

After these various pupils have been determined, a ray tracing analysis proceeds by launching a bundle of rays[5] from the entrance pupil, tracing them through the imaging system undergoing refraction according to Snell's law at all surfaces, until they intersect the exit pupil.

For concreteness, we will consider the type of system depicted in Fig. E.1(a). The first analysis we can consider is a spot diagram. In the ideal case all parallel rays from the entrance pupil will be focused down to a single point in the exit pupil. However, optical aberrations will cause those rays to fall at slightly different points within the exit pupil. The distribution of those intersection points is a measure of the optical aberrations within the system. In particular, if the spread of the intersection points —the RMS deviation of the points from the optical axis, for example— is much larger than $\lambda/(2\text{NA})$, then aberrations will dominate the size and shape of the PSF and diffraction effects can neglected. Here NA is defined according to the angle corresponding to the ray bundle spread in the image space. Conversely, if the spot diagram of intersection points is much smaller than the Airy disk, then diffraction will be the leading effect in determining the PSF and we say the system is diffraction limited.

---

[5]Depending on the type of imaging system the entering bundle of rays may be a cone of rays emanating from a single point within the entrance pupil or a bundle parallel rays filling the entrance pupil and likewise for the exit pupil. For example, an imaging system takes a cone of rays from a point in the entrance pupil and focus it down as a cone of rays towards a point in the exit pupil. A laser focusing system takes a bundle of parallel rays in the entrance pupil and focuses it down to a cone of rays towards a point in the exit pupil.

A second analysis is an optical path difference (OPD) analysis. The optical path length (OPL) for each ray is calculated by adding up the length of each segment of a ray's path scaled by the index of refraction of the medium for the segment. The OPD for each ray is calculated by subtracting the OPL for the chief ray from the OPL for the particular ray under consideration. The OPD can then be plotted against position or angle of that ray in the entrance pupil. For an ideal imaging system the OPD will be constant for all rays while an aberrated imaging system will lead to deviations in the OPD measured in units of λ. The quality of the imaging system can be quantified by the maximal or RMS value of the OPD across all rays. If the OPD is much less than λ, then the system will exhibit ideal behavior and the PSF will be limited by diffraction. However, if the OPD exceeds λ then deleterious interference effects will lead to an increase in the size of the PSF and the system will no longer be diffraction limited.

Here I have given two definitions of a diffraction limited optical system: one based on examining the spot diagram and one based on the OPD of a bundle of rays. I'll assure the reader, without proof, that these two definitions generally agree as to if a particular imaging system is diffraction limited or not.

Note that the two definitions above for whether a system is diffraction limited depend importantly on the properties of the bundle of rays that is used to analyze the system. For example, in Fig. E.1(a) if the diameter of the input beam is reduced then both the spread of points in the spot diagram and the OPD will be reduced. However, the NA of the system will also be reduced meaning the size of the Airy disk will be increased! If we continue decreasing the size of the input bundle of rays we are certain that, at some point, the decreasing size of the spot diagram will become smaller than the increasing size of the Airy disk. At this point, the system will become diffraction limited according to the definitions above.

What I describe in the previous paragraph is a generic feature of optical systems. Any optical system can typically be made to be diffraction limited by artificially reducing the NA of the system by installing an iris.[6] Of course, we must remember that, in the end, we do not care about whether a particular imaging system is diffraction limited, we only care about its absolute resolution. The challenge then is to have a system that is diffraction limited at a large NA. Designing imaging systems, such as microscope objectives, that meet this challenge is the job of qualified optical engineers.

Finally I'd like to clarify one final point about lens selection. Often, when purchasing a lens, the NA will be specified. However, to the surprise of naive shoppers, this NA very often does not correspond to the arctangent of the radius of the lens over the focal length of the lens as we normally expect based on ideal lens calculations. How then is the NA of these lenses specified? The answer is that the specified NA is the maximal NA of the lens that can be utilized while still maintaining diffraction limited performance as defined above. Diffraction limited performance may be defined as the RMS spot size being smaller than a

---

[6]As the NA is decreased the lens exhibits more and more 'ideal' or paraxial behavior. This can be leveraged as a technique in software to determine nominal characteristics about the imaging system such as its effective focal length, the location of principle planes and to find the location of the field and aperture stops and entrance and exit pupils.

certain fraction of the Airy disk, or as the RMS OPD being less than some fraction of λ. It may be possible to use the lens with a larger bundle of rays, but the NA indicates that the lens performance will then become aberration limited.

Appendix F

Input-Output Formalism

The quantum input-output theory [117] is a formalism in which an optical cavity is coupled to optical fields entering and exiting the cavity. This coupling results in the previously closed cavity system becoming an open quantum system that now admits dissipation.

The quantum input-output formula is

$$\hat{a}_{\text{out}} = \sqrt{\eta\kappa}\,\hat{a}_{\text{cav}} + \hat{a}_{\text{in}}, \tag{F.1}$$

where $\hat{a}_{\text{in,out}}$ are the quantum input and output fields, $\hat{a}_{\text{cav}}$ is the cavity photon annihilation operator, κ is the cavity energy decay rate and η is the cavity coupling efficiency.

In this appendix, I will demonstrate a formal parallel between the quantum input-output formalism and energy flow within a classical optical cavity. Relevant formulas and ideas for the classical case appear in Sec. 3.5. Recall that we can write the electric field of an electromagnetic wave as

$$E = E^{(+)} + E^{(-)}. \tag{F.2}$$

In this formulation, the energy in a cavity with waist w_0 and length L_{cav} is given by

$$U = 2\epsilon_0 V_{\text{mode}} E_{\text{cav}}^{(-)} E_{\text{cav}}^{(+)} = \pi w_0^2 \epsilon_0 \frac{L_{\text{cav}}}{2} E_{\text{cav}}^{(-)} E_{\text{cav}}^{(+)}. \tag{F.3}$$

We can define a 'classical' photon number n and amplitude a by

$$\begin{aligned} n =& a^* a = \frac{U}{\hbar\omega}, \\ a =& \sqrt{\frac{\pi w_0^2 \epsilon_0}{\hbar\omega}} \sqrt{\frac{L_{\text{cav}}}{2}} E_{\text{cav}}^{(+)}. \end{aligned} \tag{F.4}$$

I like to think of the photons in the cavity as 'stationary' photons because their energy density is localized within the volume of the cavity and there is no energy flow.

In contrast, we also deal with 'flying' photons. These are photons that are flying through space and that do have a non-zero energy flow. For these photons we focus on the energy

flow rate rather than the fixed energy. The power in a Gaussian mode with waist w_0 is given by

$$P = \pi w_0^2 \epsilon_0 c E^{(-)} E^{(+)}. \tag{F.5}$$

We can define a 'classical' flying photon number (or photon flux) and amplitude by

$$n = a_{\text{fly}}^* a_{\text{fly}} = \frac{P}{\hbar \omega},$$

$$a_{\text{fly}} = \sqrt{\frac{\pi w_0^2 \epsilon_0}{\hbar \omega}} \sqrt{c} E^{(+)}. \tag{F.6}$$

Note here that a_{fly} has dimensions of $\text{s}^{-1/2}$ in contrast to a, which is dimensionless.

We now revisit Eqs. (3.60). In particular we will focus on the relationship between the reflected field and the intracavity field. Not worrying about phases on various optical fields for the moment, we have:

$$E_R^{(+)} = E_{\text{in}}^{(+)} - t_1^2 \frac{\mathcal{F}}{\pi} E_{\text{in}}^{(+)},$$

$$E_{\text{cav}}^{(+)} = E_{\text{cav, R}}^{(+)} + E_{\text{cav, L}}^{(+)} = 2 t_1 \frac{\mathcal{F}}{\pi} E_{\text{in}}^{(+)}. \tag{F.7}$$

If we identify $E_R^{(+)} = E_{\text{out}}^{(+)}$, we can rewrite this as

$$E_{\text{out}}^{(+)} - E_{\text{in}}^{(+)} = -\frac{1}{2} t_1 E_{\text{cav}}^{(+)}. \tag{F.8}$$

Plugging in Eqs. (F.4) for $E_{\text{cav}}^{(+)}$ and (F.6) for $E_{\text{in,out}}^{(+)}$ we see that we can write this as

$$a_{\text{out}} - a_{\text{in}} = -\frac{1}{2} t_1 \sqrt{\frac{2c}{L_{\text{cav}}}} a_{\text{cav}} = -\sqrt{\eta \kappa} a_{\text{cav}}, \tag{F.9}$$

which is exactly the quantum input-output relation Eq. (F.1) up to a phase factor and missing the hats. Note that $a_{\text{out/in}}$ are 'flying photons' and a_{cav} are 'stationary photons'. We see that the quantum input output relation is simply a quantum statement that the power coming out of some port for an optical cavity (related to $\hat{a}_{\text{out}}$) is the sum of the power entering the cavity and promptly reflecting off of that port (related to $\hat{a}_{\text{in}}$) plus the light that is leaking out of the cavity that has a power proportional to $\eta \kappa U$ where $\eta \kappa$ is the contribution of that particular port to the full energy loss rate κ.

Classically, it is possible for a_{in} to be zero such as for the transmission port in a cavity that is only driven on one side. Quantum mechanically, however, we can never strictly set $\hat{a}_{\text{in}} = 0$ because, depending on the quantum optics formalism used to handle the problem, even a vacuum field reflecting off of the cavity can have non-trivial implications for the statistics of the output field because of the bosonic commutation relations:

$$\left[a_{\text{in,out}}, a_{\text{in, out}}^\dagger \right] = 1. \tag{F.10}$$