

A Cavity-Coupled Rydberg Atom Array Platform for Quantum Computing

by

Alyssa Rudelis

B.S., Stanford University (2017)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

© Alyssa Rudelis 2023. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Author

Department of Physics
May 22, 2023

Certified by

Vladan Vuletić
Lester Wolfe Professor of Physics
Thesis Supervisor

Accepted by

Lindley Winslow
Associate Department Head of Physics

A Cavity-Coupled Rydberg Atom Array Platform for Quantum Computing

by

Alyssa Rudelis

Submitted to the Department of Physics
on May 22, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

Neutral atom systems have long been the test bed for complex quantum physics. Recently, much of the focus in quantum research has shifted from fundamental science to applications in quantum computation. Although several different hardware platforms have made strides in their capabilities in this direction, each has its own impediments to scaling system size: both physically in terms of qubit number and temporally in terms of code cycles before decoherence. Specifically in neutral atom systems, the ability to non-destructively readout atomic states on timescales much faster than atomic decoherence is lacking. By pairing the geometric reconfigurability and engineered strong interactions of neutral atom Rydberg arrays with strong optical coupling to high-finesse cavities, we can build a new quantum architecture that oversteps many of the limitations of other hardware systems. In this dissertation, we lay out the case for coupling Rydberg atom arrays to cavities, discussing the connections from atomic physics to quantum computing and the fundamental physics that gives optical cavity systems an advantage over other current quantum computer implementations. We then describe the design, testing, and implementation of such a system. Our system simultaneously accommodates Rydberg excitation, reconfigurable optical tweezer arrays, selective atomic state addressing, and strong coupling to an optical cavity. We discuss in detail the risks and technical considerations of installing such a system in ultra-high vacuum, including the discovery of a new material failure mechanism for high-reflectivity mirrors. Finally, we outline concrete future steps to demonstrate proof-of-principle surface code error correction in our system, paving the way to fault-tolerant quantum computation with neutral atoms.

Thesis Supervisor: Vladan Vuletić
Title: Lester Wolfe Professor of Physics

Acknowledgments

To my mom, Tammie Rudelis, and my dad, Glen Rudelis, thank you for everything you have done to make me the person I am today. Thank you for your love and support through the high-stress academic environments and for standing by me, even when you were unfamiliar with the path I chose to pursue. Thank you for the hard work you both put in to give me and my brother Xander more opportunities than either of you could have imagined when you first started out. To Dad, this thesis is dedicated to you. If you were here now, I would tell you your favorite physics dad-joke from when I was in college: "A telegraph is just like if you have a dog stretched across the country, with its head in Boston and its tail in San Francisco. If you pull the tail, the dog barks. Quantum mechanics works exactly the same, but there's no dog."

To my partner Roberto Brenes, thank you for loving me and building me up during the highs and the lows of this experience. We started spending time together during the peak of Covid-19, and made the absolute best of everything along the way. Even when things seemed to be falling apart around us in the world (and often in each of our labs), we felt the joy in so many little things together. From watching the first robins of spring, to learning new gourmet recipes, to taking drawing classes, and even to spending a whole year a continent apart, there's no one with whom I would have rather spent the last three years, and no one with whom I'd rather face the future. Whatever comes next is still a mystery to us both, but I know it will be filled with adventure, love, and curiosity. I can't wait to embrace it all with you by my side.

Thank you, Vladan, for being such a supportive advisor during such an unprecedented time. Thank you for listening to my concerns over the years without dismissing them, and for trusting in me to lead this experiment. It's clear that you want what is best for your students, and I am truly fortunate to have worked with you and learned from you during my time at MIT. Thank you as well to Richard Fletcher and Soonwon Choi for serving on my committee and being so invested in my work.

Thank you to all of Rubidium Lab, past and present, for teaching me, working with me, and keeping the optimism and fun alive along the way. Thank you Zachary

Vendeiro and Joshua Ramette for upgrading the control system and for welcoming me into the lab right before the world shut down. I learned so much by listening to your ideas and watching you work before you each moved on to industry and theory, respectively. Thank you to Josiah Sinclair, Beili Hu, Michelle Chong, Edita Bytyqi, and Luke Stewart for building the lab into a strong team with me. After working alone for the first year of the pandemic, it was amazing to see each of you join at different points and to learn each of your unique perspectives, personalities, and ways of thinking and doing. I know it is your differences that make you a great team and that only good things are to come with you at the helm of the experiment.

Thank you to the rest of the Vuletić group for your support, mentorship, and friendship over the years. I would specifically like to thank Pablo Solano Palma, Yiheng Duan, Sergio Cantu, Yu-Ting Chen, Zeyang Li, Chi Shu, Simone Colombo, Edwin Pedrozo-Peñafiel, Wenchao Xu, Ian Counts, and Joonseok Hur for answering my many questions and showing me the ropes when I first started graduate school. I would also like to thank Eugene Knyazev, Matthew Peters, Emily Qiu, and Tamara Sumarac for being supportive friends along the way. I cannot wait to see where you end up after MIT, and I have no doubt we will cross paths in the future.

There are many others I would like to thank from all the communities I joined at MIT. Thank you to my friends and mentors in the MIT Outing Club (MITOC) who first welcomed me to Boston and built up my courage and independence in the outdoors. Thank you to the Science Policy Initiative (SPI) for showing me how to take my expertise out of academia to impact the rest of the world. Thank you to all the many graduate student organizers who leaned into democracy and spent countless hours fighting for better working conditions for graduate workers across the institute—both in the MIT Graduate Student Union (GSU) and through internal paths like the Physics Values Committee (PVC). When research was daunting, joining these groups to fight for a better culture at MIT gave me energy. To all my other friends in my cohort and across departments (you know who you are), thank you for sharing meals with me, commiserating with me, and giving me the privilege to share in even the smallest parts of your journeys at MIT.

Contents

1	Introduction	29
1.1	Atomic Physics to Quantum Computing	29
1.2	Architectures and Challenges	30
1.2.1	Scalability	31
1.2.2	Decoherence and Loss	33
1.2.3	Readout and Error Correction	35
1.3	Status of Neutral Atom Arrays	36
1.3.1	Gates	37
1.3.2	Motion and Connectivity	38
1.3.3	Readout and Error Correction	39
1.4	Cavity-Coupled Atoms	41
1.4.1	Bare Cavity Properties and Formulas	42
1.4.2	Classical Model of an Atom in a Cavity	44
1.4.3	Cavity-mediated Atomic State Readout	47
2	Vacuum Apparatus	53
2.1	Overview	53
2.2	Electrical Components	54
2.2.1	Locations, Connections, and Feedthroughs	54
2.2.2	Cavity Length	62
2.2.3	Electromagnetic Field Control	66
2.3	In-Vacuum aspheric Lenses	68
2.3.1	Relay Imaging Design	70

2.3.2	Relay Imaging Physical Tests	73
2.3.3	In-Vacuum Lens Install	79
2.4	Vacuum Chamber	86
2.4.1	Flanges and Viewports	86
2.4.2	Regular Maintenance	87
3	Cavity Apparatus	91
3.1	Cavity Geometric Properties	92
3.2	Out-of-Vacuum Alignment and Characterization	93
3.2.1	Assembly	94
3.2.2	Alignment	96
3.2.3	Characterization	107
3.3	Vacuum Installation	112
3.4	Stabilization	123
4	Cavity Mirror Materials	133
4.1	Materials and Methods	136
4.1.1	Mirror Coating Runs	136
4.1.2	Annealing and Optical Methods	136
4.1.3	Non-Optical Characterization Methods	137
4.2	Results and Discussion	138
4.2.1	Process Effects on Cavity Finesse	138
4.2.2	Process Effects on Mirror Roughness	141
4.2.3	Elemental Analysis	142
4.3	Conclusion	144
5	Atom Trapping and Addressing	147
5.1	Optical Tweezers	148
5.2	Cavity Lattice	153
5.3	Raman Addressing and Rydberg Transitions	156

6 Outlook	161
6.1 Error Correction	161
6.2 Atom Array Loading	168
A Cavity Mirror Cleaning	173

List of Figures

1-1	Bloch sphere (a) and internal energy level representations of a single atomic qubit gate in an idealized 2-level system (b) and in the ground state hyperfine levels of rubidium-87 (C).	37
1-2	Schematic of how the Rydberg blockade enables a C-Phase gate between two neutral atoms, creating entanglement between the two atoms when they are within a Rydberg blockade radius, r_B , of one another.	38
1-3	Typical arrangement of two acousto-optical deflectors (AODs) to create a two-dimensional array of optical tweezers (a). Three possible and one impossible crossed AOD optical tweezer patterns (b).	40
1-4	The rectangular Hermite-Gaussian transverse electromagnetic (TEM_{mn}) modes of a well-aligned, stable, two-mirror optical cavity in increasing mn orders from top-left to bottom-right.	43
1-5	The Lorentzian lineshape of transmission of one specific TEM_{mn} mode through an optical cavity. ν_0 is the resonant frequency of the cavity, κ is the linewidth, and ν_{FSR} is the free spectral range.	44
1-6	The landscape of stable spherical two-mirror optical cavities (pink shaded region). The unshaded regions cannot sustain a cavity mode, as any injected light is unconfined and quickly leaves the inter-mirror region after some number of reflections. The experimental cavity in this work lies directly between two limits in cavity stability, concentric and confocal. This ensures our cavity is robust against small changes in both length and mirror radius of curvature due to manufacturing and alignment tolerances.	45

1-7	Dotted lines correspond to the ratio of transmitted cavity power and input cavity powers for different light-cavity detunings. The solid lines correspond to the light scattered out of the cavity by an atom. For all plotted lines, $\Delta = \delta$ and the other parameters are chosen to be those of the installed experiment: $\eta = 2.1$, $\kappa = 127\text{kHz}$, and $\Gamma = 6.065\text{MHz}$. The blue, orange, and green dashed lines correspond to the transmission spectrum with no coupled atom, one coupled atom, and five coupled atoms, respectively. The solid red and purple lines correspond to the light scattered out of the cavity for one and five coupled atoms, respectively.	48
1-8	The ratio of transmitted cavity power and the input cavity power as a function of cavity detuning from both the atom and light, giving us a dispersive shift in the resonance of the system. From bottom to top, the curves correspond to a cavity detuning, Δ , of $\Gamma/2$, Γ , 2Γ , and 10Γ : blue, orange, green, and red, respectively.	49
1-9	The two methods of cavity-mediated atomic state readout: a) fluorescence and b) absorption.	50
1-10	Distributions of expected output photons from our experimental cavity in the absorption readout configuration for one atom (pink) and no atom (purple) coupled to the mode.	52

2-1	View of entire vacuum system, as seen from the entrance to the laboratory (a). The top portion of the chamber above the thick breadboard is the experimental chamber, with MOT coils, MOT optical access, and cavity optical access. The portion of the chamber below the thick breadboard contains the ion pump (far right), the titanium sublimation (TiSub) pump (far left), and the gate valve (bottom left). Additionally, the optical tweezers are sent in toward the experimental chamber through the bottom window. The magnified callout image shows the new in-vacuum apparatus and where it sits along the height of the experimental chamber (b).	55
2-2	a)-d) Full CAD drawing of vacuum cavity apparatus from different angles with components labeled according to Table 2.1. This image only contains the mechanical portions of the apparatus and not the electrical components that were added post-assembly.	56
2-3	Overview of all new systems required for the upgrade to 1) cavity-coupled, 2) Rydberg-excitable, 3) tweezer-trapped, and 4) individually addressed atoms. The apparatus chapters, 2, 3, and 5, will step through each of these subsystems in detail.	56
2-4	a)-b), Diagrams of the 32-pin feedthrough and plug. c) shows the labeled pinouts on the plug side of the connector. Because our cavity apparatus didn't require all 32 pins, only 20 of the pins are crimped in place.	58
2-5	One of the Macor Dsub connectors just before the second vacuum installation. The white portion at the top of the Dsub broke off completely from the remainder of the Dsub and was glued back into place using TorrSeal. Looking closely at the areas on the ceramic piece between the separate pins, scrape marks from removing shorts through the graphite coating are visible.	62
2-6	Picture of the mounted ring piezos (a) before being re-coated in chrome and gold and (b) after being recoated.	64

2-7	Picture of homemade dipole antenna for 6.8GHz transmission.	67
2-8	The new MOT coils, which were epoxied separately from the chamber for the ease of removal for future hardware upgrades. The coils are clamped into place with a custom four-jaw clamp to prevent drifting of the magnetic field over time. The clamp jaws can be further disassembled for easy installation.	69
2-9	Diagram of the two different relay lens systems, for 780nm trapping and addressing light (a) and for 480nm Rydberg addressing light (b). Both diagrams are drawn to scale to illustrate the long distance over which the light must travel in a nearly-collimated beam before reaching the atoms. This distance is set by the length of the vacuum chamber above and below the cavity apparatus to the nearest viewports. . . .	70
2-10	Schematic representation of all four inputs/outputs to the relay lens system, through both the 480nm lens and the 780nm lens, not to scale. Complete optical diagrams for these systems are documented in Ch. 5.	74
2-11	Photo of the relay lens test system in the lab during the rough, initial alignment. Here, we first collimate fiber light through an aspheric lens before switching the fiber out for the smaller SNOM tip source. Schematically, this step is illustrated in Fig. 2-12a).	75

2-12	The main four iterative steps used to align the aspheric relay lens test setup. First, we use a bare fiber as a source and a shearing interferometer as the figure of merit to set the distance between the source and the out-of-vacuum lens (a). Second, we place a pre-aligned lens-camera breadboard into the path after removing the shearing interferometer, and center the image on the camera before clamping the breadboard down at the distance set by our vacuum chamber (b). Third, we carefully mark the position of the bar fiber tip and swap it out for a SNOM tip fiber, making small adjustments until we recover a point image on the center of the camera (c). We also add a viewport and make the necessary small adjustments to recenter the point image on the camera. After the image is optimized in terms of symmetry and size, we magnify the image from the relay system to better quantify the spot size (d).	75
2-13	Minimum spot size on the physical relay lens setup at 780nm while imaging a SNOM tip.	77
2-14	FOV and DOF scans on the physical relay lens setup at 780nm while imaging a SNOM tip. The red circles overlaying the images are the diffraction limited rms radius size and are the same size as the red circles below the images in the raytracing spot diagrams.	78
2-15	USAF target imaged through our 480nm lens system while illuminated by 405nm light.	79
2-16	The 780nm (lower) lens tube before (a) and after (b) optical blackening.	80
2-17	The 780nm (lower) lenstube after a layer of Aquadag was removed via friction.	81
2-18	The 480nm (upper) lens tube with the 480nm custom aspheric lens sitting on top instead of sliding into the tube due to the thickness of the graphite layer (a) and sitting where it was designed to fit, with the exposed edge of the lens the future location of epoxy (b).	82

2-19	The physically constrained setup for aligning the high-NA in-vacuum aspheric lens to a SNOM tip fiber centered on the optical modes of the two original optical cavities. The SNOM tip is circled in magenta, directed upward. The light from the SNOM is collimated through the in-vacuum aspheric lens before being transferred to the breadboard via three, two-inch diameter mirrors.	84
2-20	An initial image of the light from the SNOM tip at the center of the cavity mode before adjusting the in-vacuum high-NA aspheric lens (a). A more aligned image of the light from the SNOM tip after adjusting the high-NA in-vacuum aspheric lens (b).	85
2-21	Calculated (line) and measured (star) reflectivities of the AR coatings on the new vacuum chamber viewports.	87
2-22	A view of the bottom of the old rubidium getter arm through the TiSub flange, looking toward the inside of the ion pump.	89
2-23	The spent and broken titanium filaments replaced during the vacuum chamber maintenance portion of the upgrade. We replaced all three filaments with fresh ones in November 2021.	90
3-1	Geometry of the experimental cavity.	93
3-2	The cavity ring mount before assembly onto the full UHV cavity support structure. The two parallel wires running from left to right in the center of the ring are the RF wires and are held in place with screwed-in washers. The small rectangular component on the bottom left of the ring is one of the UHV ceramic heaters. The other two components are glass-encapsulated 10k Ω thermistors. The four through-holes on the cavity ring are where the ring is connected to the remainder of the structure within a Viton spacer sandwich.	95

3-3	A piezo-mounted mirror immediately after being cured into place on its mount with silver epoxy. On the left of the image is a home-made clamp for our small mirrors. The clamp body is the same size as a standard Thorlabs translation stage, with through-holes for mounting directly on top of the stage. The center-left of the image shows the clamp jaw with a set screw at an angle. When tightened, this set screw holds a cavity mirror tightly into the far corner of the jaw clamp. The mirror mount is screwed into a metal holder that is clamped in place against a breadboard. The whole gluing jig holds all components in place while the mirror is centered on its mount and epoxy is applied to the edges of the mirror and surface of the mount.	97
3-4	The three beams of light hitting the tissue paper while fine-aligning the cavity output mirror in step one. We only see all three spots of light on the tissue with off-resonant light, which allows for reflections off the AR coating.	99
3-5	Schematic of full cavity mirror alignment, steps 1-4.	101
3-6	MT01 micro-needle tool from Thomas Scientific that was used to selectively block the cavity modes during the overlap procedure.	103
3-7	A view into the HEPA flow box during the cavity mode overlap process. The gold colored object clamped onto the three translation stages in the center of the image is the needle holder, extending into the region where the two cavity modes overlap.	105
3-8	The final overlap of the two crossed cavity modes before the first vacuum installation of the cavity structure. By simultaneously scanning a needle through the mode of both cavities in the Z direction, we could reliably map out the mode positions relative to one another as we made small changes to the mirror positions to maximize overlap.	106

3-9	Controls and electro-optical components used for cavity ringdown measurements. The initial ringdown method using a triggered EOM while scanning laser frequency (a). Triggering an EOM was a less reliable method to measure the ringdown time than the DPAOM method with cavity mirror piezo scanning (b).	111
3-10	View of the cavity apparatus installed in the experimental chamber from below. The highlighted rectangle is the part of the apparatus with the largest diameter, only 2mm smaller than the diameter of the chamber tube inner diameter. If the cavity apparatus strayed from parallel to the chamber tube at any point during the installation, this portion of the apparatus and electronic connections hit the sides of the chamber.	115
3-11	The ion gauge pressure reading in the experimental chamber as a function of time during the first installation of the cavity structure. We baked at a high temperature of 180°C. The spikes in current are the points at which we tried (unsuccessfully) to turn on and ramp up the voltage across the ion pump. After several rounds of hi-potting and hitting the ion pump with a rubber mallet, we were able to continue pumping down to UHV.	118
3-12	The ion gauge pressure reading in the experimental chamber as a function of time during the second and final cavity installation.	120
3-13	Block diagram of all major components of the experimental cavity stabilization scheme including passive, active, fast, and slow feedback.	122
3-14	All cavity optics paths: 1560nm locking and lattice (pink beam) and the 780nm probe (purple beam). The beams are offset for visual clarity—in reality they overlap.	123

3-15	A detailed schematic of the first stage of PDH locking electronics, where both the fast and slow servo outputs are created. The slow servo output is taken directly from the lockbox output and sent to the cavity mirror piezo. The fast servo is created from the AC Error signal from the lockbox after passing through an attenuator, homemade filters, and a homemade summing amplifier.	124
3-16	The PDH error signal for our installed experimental cavity without (a) and with (b) the fast feedback loop engaged. The sensitivity of the lock can be estimated via the slope of the central feature of the error signal without fast feedback engaged.	126
3-17	Estimated theoretical Bode plots for both the fast and slow cavity feedback loops. The slow feedback corner frequencies correspond to the first and second lockbox integrator frequencies and the piezo roll-off frequency, respectively. The fast feedback corner frequencies correspond to the cutoff frequencies of the homemade filter box shown in Fig. 3-15. This diagram is schematic and the relative gains and frequencies of the slow versus fast feedback are not plotted to scale. .	127

3-18	The RF circuit diagram for the fast feedback portion of the cavity locking scheme. The error signal resulting from the initial locking electronics in Fig. 3-15 is fed into a VCO. The resulting RF signal is mixed with an offset RF tone from an additional RF synthesizer, in this case a WindFreak. After mixing, the summed signal is filtered, attenuated, and amplified in order to create a useful gain. This final signal is sent to the cavity locking laser's fiber EOM, which creates an optical sideband at the desired frequency (i.e. at the addition of the fast feedback signal and the offset RF). In parallel, the same fast feedback signal that is sent to the mixer for the 1560nm cavity locking laser is doubled and mixed with a second offset RF tone from a second WindFreak RF synthesizer. The resulting signal is sent to the 780nm probing laser EOM to ensure that both lasers see the same fast feedback and can be detuned separately.	128
3-19	The cavity transmission spectrum without a MOT (left) and the initial Rabi splitting of the atom-cavity coupled system with a large atomic ensemble in the cavity after releasing a MOT (right).	129
3-20	Counts on the SPCM at the cavity transmission as a function of time after a loaded MOT was released by turning off the MOT light and coils both a) before and b) after adjusting the cavity lock parameters to reduce oscillations induced by MOT turn-off.	131
4-1	Cavity finesse as a function of mirror processing step. The cavity finesses decreased after the initial vacuum bake and after each subsequent recovery attempt that included a bake step. Corresponding 3D AFM images of the top surface of the mirror coatings at each step are also pictured.	135

4-2	Histogram of mirror surface heights after the various processing steps. The distribution in surface height of a single mirror becomes less Gaussian with each step, indicated by increasing skew, γ . Additionally, each processing step intended to return the mirror surface roughness to the pre-UHV-bake value actually increased the deviation in mirror roughness.	139
4-3	Expected theoretical reduction in cavity finesse with mirror surface roughness compared with measured reduction in cavity finesse with mirror surface roughness. The theory curve is calculated with no free parameters.	140
4-4	X-ray fluorescence rules out the possibility that a contaminant from the vacuum chamber caused the mirror surface roughness increase. Only Si, Ta, and artifacts of the XRF device (Rh, Pd, and Fe) were resolvable in the measurement.	143
5-1	Overview of the atom trapping and addressing systems in the new experiment platform. Not all of these systems will be integrated in the first iteration of this experiment. Currently, only the AOD tweezers are fully built and integrated with the vacuum chamber. The 6.8GHz Raman beams and 2-photon Rydberg Excitation optics have been partially built, but not yet integrated with the chamber. The SLM optical tweezers are in the design and testing process and the 10MHz Raman beams have not been started.	148
5-2	Dipole trap depth in μK as a function of trap wavelength for a rubidium-87 atom.	149
5-3	The dependence of photon scattering rate on the tweezer trap wavelength in rubidium-87. The heating rate is directly proportional to the scattering rate of the atoms in the trap, so we chose a trap wavelength of 808nm to minimize scattering.	151

5-4	Optical schematic of the optics that create the optical tweezers immediately before the vacuum chamber lower viewport. The final lens is the in-vacuum aspheric lens, AL3026, as discussed in Chapter 2. . . .	152
5-5	Detailed picture of how the many orders of light created in a crossed AOD converge/diverge and focus/defocus through the tweezer array optical path.	153
5-6	Schematic of the inter-cavity trapping lattice (pink) overlapped with the probe light standing wave (purple). Atoms feel the strongest trapping potential at the anti-nodes of the red-detuned 1560nm light. The antinodes of the trapping light line up with the antinodes of the 780nm cavity probe standing wave, allowing for the trapped atoms to be held at the point of strongest coupling along the cavity axis. Atomic coupling to the cavity mode is variable based on the position along the cavity mode.	155
5-7	Optics for creating the Raman beam pairs at 6.8GHz around 795nm with two passes through a chirped Volumetric Bragg Grating (VBG).	157
5-8	Energy diagram showing the two-photon Rydberg transition with 780nm and 480nm light. [84].	158

6-1	A schematic representation of the future of Rydberg Atom Arrays, coupled to an optical cavity. A static optical tweezer array (SLM) holds the surface code patch physical qubits outside of the cavity mode while moving tweezer traps (crossed AODs) constantly move syndrome qubits in and out of the cavity mode for error correction and the reloading of traps that lost an atom due to vacuum lifetime. On the left is a graph representation of a two-dimensional surface error correcting code, with data qubits (white) and syndrome qubits (black) connected via stabilizer measurements in the X (blue) and Z (green) bases. On the right is a simplified four-level atom in a cavity. The atom will shift the resonance of the coupled atom-cavity system if it's in state $ check\rangle$, but not if it's in state $ 0\rangle$ or state $ 1\rangle$, allowing for a determination on atomic state based on the probe transmitted power.	163
6-2	Physical internal state scheme for logical states in rubidium-87 that can be used for error diagnosis with a cavity. This scheme requires a small external magnetic field to raise the hyperfine degeneracy, allowing access to different m_F levels with Raman transitions.	164
6-3	The procedure for finding the location of an error (depicted here as the atom in the "down" state) in faster than serial times using the nondestructive, off-resonant cavity readout in a binary search scheme. This requires moving atoms out of the logical basis with Raman addressing beams to "shelve" them when the search is not in that branch of the tree, and then moving them back into the logical basis to continue the search until the error has been found.	166
6-4	The scaling of cavity error diagnosis speed with atom number for a cavity binary search method, sequential cavity search, and a parallel fluorescence readout.	167

6-5	Schematic of the cavity-mediated array loading procedure from scratch. An atomic reservoir sits outside the cavity mode, either in a large ensemble dipole trap or a MOT. A combination of 1D and 2D AOD tweezers pick up atoms from the reservoir, drag them through the cavity mode while a probe beam turns on. Based on the transmission counts at the output of the cavity on an SPCM, the tweezer is moved to the next empty spot in a stationary SLM tweezer array on the other side of the cavity mode or the tweezer is turned off, and the process continues until the 2D array is 100% filled.	170
A-1	How to fold a piece of lens tissue for cleaning our small cavity mirrors.	174
A-2	A potential path through the crossed cavity apparatus for mirror cleaning.	175
A-3	How to wipe a cavity mirror with a properly folded tissue held by a hemostat.	175

List of Tables

1.1	A summary of orders of magnitudes for key parameters of three different hardware architectures for quantum computing: superconducting, trapped ion, and atomic qubits. T_1 is the "bit-flip" coherence time, T_2 is the dephasing coherence time, T_{vacuum} is the vacuum lifetime, T_{gate} is the time required to complete a single-qubit gate (π pulse), and $T_{readout}$ is the time needed to readout one qubit state. The ratio of limiting lifetime to gate times is remarkably similar for all three architectures. Dephasing times in ions and atoms can be improved further with technical improvements to laser stability, background fields, and chip stability.	34
2.1	A summary of all new vacuum components installed as a part of the cavity structure.	57
2.2	A list of all components that connect the electrical components on the cavity apparatus to the vacuum feedthroughs. The quantities of hardware components reflect what is currently installed and working as intended. The quantities of tools and epoxies reflect how much I purchased over the course of the upgrade.	59
2.3	Feedthrough pinouts and notes on broken connections for all new electrical feedthroughs for the cavity apparatus. Some connections were damaged or removed during the course of the two vacuum installations, and are noted here.	60

2.4	A summary of the different types of electrical joints that were made at different locations and for different components in the vacuum chamber. Practicality and ease of installation dictated which joint was used in which case.	61
2.5	Key properties of the cavity piezo, both provided by the manufacturer and measured in our lab.	65
2.6	Current settings for the cavity ring mount temperature control box, Wavelength Electronics Model LFI-3751.	66
2.7	Expected performance via Zemax OpticStudio simulation of the two different relay lens systems at the designed wavelengths when perfectly aligned, on-axis. Minimum spot size is measured as the rms spot size until reaching the diffraction limit. The 780nm system is diffraction-limited. The Field of View (FOV) and Depth of Field (DOF) are both defined as the distance in XY and Z, respectively, where the spot size increases by a factor of $\sqrt{2}$	71
2.8	Simulated performance of the two different relay systems at different misalignments of the in-vacuum lens. The optical requirements take into account out-of-vacuum lens movement to minimize the spot size at the output of the full relay system. The mechanical tolerances are calculated from adding together machining tolerances of each UHV piece attached to the lens tubes. An acceptable optical requirement is defined where the spot size stays within $\sqrt{2}$ of the minimum spot size. Note that the mechanical tolerances for the 480nm lens automatically satisfy the optical requirements, where the same is not true with the 780nm lens.	73
2.9	Quantity and type of new viewports added to UHV chamber in November 2021.	87
3.1	Experimental cavity properties for both cavity mirrors.	93
3.2	Designed properties of the installed experimental cavity.	94

3.3	Ringdown times of the installed OC-OC cavity at various stages in installation. The ringdown time has fluctuated over time, potentially due to dust falling on and off the mirrors. It is also possible that in the time since installation these mirrors have experienced vacuum degradation from oxygen vacancy formation.	121
3.4	Properties of the installed and stabilized experimental cavity at the probe wavelength, 780nm.	132

Chapter 1

Introduction

1.1 Atomic Physics to Quantum Computing

The field of atomic physics has driven the development of new technologies for well over 100 years. Starting with the Bohr model of the atom in 1913, atomic physics has worked through an astonishing number of discoveries and technical feats [4]. From the unearthing of the fine structure of hydrogen in the 40s and the invention of the maser and laser in the 50s to the first coherent control of atoms in the 80s, as the experimental capabilities in the field have increased, the number of real world effects and applications has increased as well. Indeed, as our understanding of the quantum world has shifted, the field has shifted along and followed applications as they have presented themselves. For example, in the last century the invention of the laser led to novel manufacturing and surgical procedures, optical clocks have provided the most accurate and precise global positioning systems (GPS), and quantum gas microscopes have increased our understanding of condensed matter systems [60]. In the current decade the pattern of innovation has continued, this time shifting the field to the application of quantum computing. The concept of quantum computing and its relation to Turing machines, math, logic, and physics was originally introduced by Paul Benioff in 1980, but more famously, the following year Richard Feynman gave a talk titled "Simulating Physics with Computers" at MIT [13, 37]. In this talk, he predicted that a computer based on quantum mechanics could potentially open

a paradigm for discovering new physics by extending our simulation capabilities. Although it is not yet clear how long it will take to develop a universal quantum computer that is useful to academia in this exact way or useful the general public via more applied algorithms, the remarkable coherent control of atomic systems and the rate of technical improvements to those systems reveals atomic physics as a promising platform for developing quantum computing.

1.2 Architectures and Challenges

The handful of quantum computing qubit architectures currently under development—superconducting qubits, trapped ions, neutral atom arrays, nitrogen vacancy centers, and photonics—have distinct timelines, advantages, and drawbacks. Many academic institutions and large commercial developers such as Google, IBM, and Honeywell are sticking to silicon and building superconducting qubit registers and trapped ion systems [7, 8]. These two types of systems share similarities in their reliance on semiconductor manufacturing—for the qubits themselves in the case of superconducting qubits and for the trapping chips in the case of trapped ions. Because of this reliance, however, they both also share limitations in qubit geometries and are currently only able to house and control ~ 50 physical qubits at a time [19]. Neutral atom array systems, although younger compared to these other quantum computing architectures, have quickly caught up to and surpassed the qubit numbers of some of the more mature architectures. In fact, there are now a slew of start-up companies pursuing this architecture including PASQAL, QuEra, Cold Quanta, and Atom Computing, in addition to academic work on the systems [9, 73, 45]. Nitrogen-vacancy and photonic quantum information architectures are mostly aimed at addressing challenges in quantum networking rather than logic. Several reasons, including inconsistencies in nitrogen vacancy manufacturing and the difficulty of creating photonic memories, prevent these two architectures from being strong candidates for standalone quantum computing qubit architectures [25, 91]. Because of these general differences, superconducting qubit, ion trapped, and neutral atom array systems are currently the top

contenders for the future of quantum computing architectures, with photonics and vacancy centers remaining important tools in quantum networking [91, 24].

As we already alluded to, these three main computing architectures each have distinct considerations when it comes to control, reconfigurability/connectivity, loss mechanisms versus speed, readout, and finally scalability and modularity. More formally, when evaluating the suitability of physical qubits for quantum computing, there are five generally accepted criteria to meet, known as Divincenzo’s Criteria (DVC) [30]:

1. A scalable system with well-characterized qubits
2. Ability to initialize qubits to a known, simple state
3. Qubit decoherence times much longer than gate operation times
4. A universal set of quantum gates
5. Readout of individual qubit states

Superconducting, trapped ion, and neutral atom qubits all satisfy these five requirements to varying degrees. These platforms have distinct abilities in decoherence times (3), scalability (1), and readout (5) in particular. Extending DVC conceptually, all three of these properties strongly impact the potential for fault-tolerant error correction—a necessity on the path to continuous, useful quantum computation. We will explore the trade-offs in these criteria between types of physical qubits in this section, with a focus on these specific criteria 1, 3, and 5. In the process, we will reveal strengths and weaknesses across the platforms. Additionally, we will begin to motivate how cavity-coupled neutral atom systems can have an advantage moving forward with optical cavity coupling, specifically in regards to readout and scalability.

1.2.1 Scalability

Scalability of an architecture dictates if it will be able to keep up with increasing computation demands over time, especially in terms of qubit number. To be able

to successfully demonstrate a full-scale Shor’s Algorithm, for example, we require $72L^3$ quantum gates acting on $5L + 1$ logical qubits in order to successfully factor an L -bit number [12]. To create $5L + 1$ error corrected logical qubits requires many more physical qubits, with the exact number depending on the specific error correction code. The most recent proof-of-principle error corrected superconducting logical qubit required 49 physical qubits using a distance-5 surface code [5]. Overall, implementing a full-scale Shor’s Algorithm with this type of error correction would then require $\sim 10^5 - 10^6$ physical qubits to factor an $L = 1024$ -bit number (the smallest size of number typically used in the Rivest-Shamir-Adleman (RSA) encryption scheme) [51]. Currently, no quantum computing platform comes close to this scale, making scalability one of the most critical roadblocks to quantum computing progress.

Superconducting qubits and trapped ions can be size-limited by chip sizes, cryogenic dilution refrigerator volumes, and electrical connections. Superconducting qubits must be housed within a cryogenic system to keep the temperature of the chips below T_c for superconductivity in the given material, and the size of these systems has so far limited the total number of physical qubits in one system to a record of 433 [2]. Trapped ion systems can be housed in cryogenic systems to reduce vacuum requirements, but they can also operate in ultra-high vacuum (UHV) chambers at room temperature. Currently, although efforts are being made to build larger cryogenic systems for many applications, scaling up is becoming more expensive and technically difficult [3, 1]. Additionally, both superconducting qubits and ion traps rely on chip fabrication processes. In the case of superconducting qubits, this means that not every manufactured qubit will be usable and that each qubit must be characterized individually before use [39, 105]. In the case of ion trap chips, defects in manufacturing are less critical because the chips do not contain the qubits themselves, but rather are auxiliary equipment for trapping and addressing. In this case, ion traps are more limited by the geometry of the chips than anything. Indeed, until recently only one-dimensional ion trap chains were experimentally in use, and two-dimensional ion trap chips still have many technical roadblocks to scaling beyond tens of ions. Challenges include the parallelization of controls and transport and integrating sym-

pathetic cooling [79, 70]. Neutral atom systems are typically installed in a ultra-high vacuum chamber or built around a vapor cell, with most controls being in the form of laser pulses. As physical atomic qubit numbers increase, so do optical access and optical power requirements. The finite solid angle and finite available laser power restricts scalability way before the chamber volume restricts physical qubit numbers. One promising method to increase the scalability of all three of these architectures is to modularize and network smaller systems together to create a larger whole [11, 82]. In this case, quantum interconnects become the limiting technical factor. In superconducting qubits, more electrical connections into and out of a cryogenic system will make cooling more difficult. In ion and atom systems to make connected modules a reality, we require efficient photon collection and low-loss photonic interconnects.

1.2.2 Decoherence and Loss

As these quantum computing architectures scale up in size, speeds become even more critical to avoid decoherence and loss. This applies across operation types: gates, readout, and reconfigurability. Each architecture is ultimately limited by a fundamental decoherence mechanism that depends on the qubit type, usually organized into two types: T_1 (transverse or "bit-flip") and T_2 (longitudinal or "dephasing") coherence times [19, 85]. Typical times for each architecture are shown in Tab. 1.1, along with current qubit number records and other parameters. The ultimate figure of merit for each of these systems is the ratio between the time constant of the limiting decoherence mechanism and the average time it takes to perform a gate on that qubit, N_{gates} , or the possible number of gates completed before a coherence time passes. Although trapped ions and neutral atoms operate on fairly similar timescales, superconducting qubits are several orders of magnitude faster both in decoherence times and also in gate times. To compare fairly across the architectures, we have tabulated N_{gates} in Tab. 1.1.

In atomic systems, dephasing is non-negligible with T_2 time for a ground state hyperfine qubit in rubidium-87, $T_2 \sim 300\text{ms}$ [75]. Dephasing, however, can be corrected with standard error correction schemes. In a ground state atomic qubit, T_1 is infinite,

	Superconducting	Ions	Atoms
T_1	1ms	$\gtrsim 10$ s	$\gtrsim 100$ s
T_2	1ms	100ms	10s
T_{vacuum}	—	hours	1s
T_{gate}	10-100ns	10 μ s	100ns
$N_{gates} = T_{lim}/T_{gate}$	10 ⁶ -10 ⁷	10 ⁶	10 ⁷
N_{qubits}	433	30	256
$T_{readout}$	100ns	1ms	100ms

Table 1.1: A summary of orders of magnitudes for key parameters of three different hardware architectures for quantum computing: superconducting, trapped ion, and atomic qubits. T_1 is the "bit-flip" coherence time, T_2 is the dephasing coherence time, T_{vacuum} is the vacuum lifetime, T_{gate} is the time required to complete a single-qubit gate (π pulse), and $T_{readout}$ is the time needed to readout one qubit state. The ratio of limiting lifetime to gate times is remarkably similar for all three architectures. Dephasing times in ions and atoms can be improved further with technical improvements to laser stability, background fields, and chip stability.

and in clock states $T_1 \sim 100$ s, much longer than the typical vacuum lifetime of trapped atoms, which is ~ 1 -10s depending on the quality of the vacuum [54]. Vacuum loss is a source of loss that cannot be mitigated with standard error correction proposals, and so is a concrete roadblock to scaling [21, 34]. Actively reloading atoms into traps during computation is one option for combating vacuum losses. Currently, there are only a couple experimental demonstrations where atoms are reloaded during code cycles. In [90], they operate an atom array continuously by passing the atomic states back and forth between two different species of neutral atoms (cesium and rubidium) and continuing logical operations on one set of atoms while reloading the other. With a single atom species, neutral atom reloading demonstrations and proposals have relied on a finite reservoir of trapped atoms that are dynamically moved to fill in lost atoms [35]. However, both of these reloading methods depend on repeated fluorescence measurements to check for atom vacancies, which are quite slow ~ 100 ms, to diagnose the presence and location of defects. This added time allows for many more errors to be introduced during reloading. Neither trapped ions nor superconducting qubits have to contend with the vacuum lifetime issue. For trapped ions, vacuum lifetimes are typically on the order of hours to days whereas superconducting qubits

do not have to be trapped at all [106]. Trapped ion systems are limited mainly by frequency shifts and the decoherence of motional modes (used to implement Coulomb gates), through heating and dephasing [70, 97, 71]. There is a wide range in coherence times depending on the ion element and the transitions in questions, but $T_1 \sim 100\text{ms}$ and $T_2 \sim 10\text{s}$ are considered good for clock state qubits [52, 102]. Superconducting qubits have much faster T_1 and T_2 times, right around 1-100ns [42, 76, 59].

1.2.3 Readout and Error Correction

Because quantum bits are liable to accruing errors and loss through decoherence, they must be error corrected [19]. Although control, speed, and scalability are all requirements in themselves for a useful quantum computer, they are doubly critical in enabling a system to be error corrected. We will dive into more detail in Ch.6 on a specific way in which error correction can be carried out on our atomic qubit system, but here we focus on a broad understanding of the requirements for error correction. To correct errors, there are two tasks that must be performed much faster than the loss mechanisms in the system:

1. Determine if, when, and where there are qubit errors
2. Correct the found errors without introducing more errors

There are many types of algorithms that can achieve these goals, and in each has its own error threshold, meaning that all qubit manipulations must have errors less than a certain value for correction to work. In most hardware architectures, however, the combination of typical SPAM (state preparation and measurement) and gate errors are still too large for even the more lenient error correction codes to work. Widely considered one of the more attainable codes in the short term, surface error correcting codes have threshold errors $\lesssim 1\%$ [59]. In superconducting qubits, trapped ions, and neutral atom arrays, current state-of-the-art two-qubit gate errors are on the order of 0.3%, 0.2%, and 2%, respectively [59, 92, 40]. Although two-qubit gate errors do not fully explain the dynamics and errors in a quantum computing architecture,

a surface code requires four two-qubit gates on each data qubit in the surface during each code cycle, not counting the gates required for computation between each cycle. Thus, the two-qubit gate errors are certainly a good benchmark of overall control, as these errors already take the errors in each architecture higher than threshold [59]. Furthermore, since error correction relies on an error diagnosis step, a promising quantum computing architecture needs to have a fast state-readout scheme. However, as shown in Table 1.1, the readout times for both trapped ions and atoms are much longer than gate times and not significantly shorter than coherence and loss times. Ultimately, this could mean that the time taken to error correct introduces more errors than the correction code can remove. Indeed, SPAM errors are on the order of a 3% for neutral atom qubits, 0.03% for trapped ions, and 1-2% for superconducting qubits, which could be improved with faster measurement [40, 45, 6, 69]. Finally, surface code error correction is heavily geometry-dependent, making it more difficult for superconducting qubit systems to implement. A new qubit chip must be made for changing qubit geometries and thus for the express purpose of error correction. Taking into account all of these roadblocks, it is not surprise why little experimental progress has been made in implementing error correction, other than a couple recent proof-of-principle attempts [5, 16].

1.3 Status of Neutral Atom Arrays

Although neutral atom array systems have developed rapidly in recent years, there are still several capabilities that need improvement to enable fault-tolerant computing through error correction. As mentioned in the previous section, gate and SPAM fidelities, readout, and error correction all require technical improvements. This section will give more detail into how gates, readout, and error correction are currently implemented on neutral atom array systems.

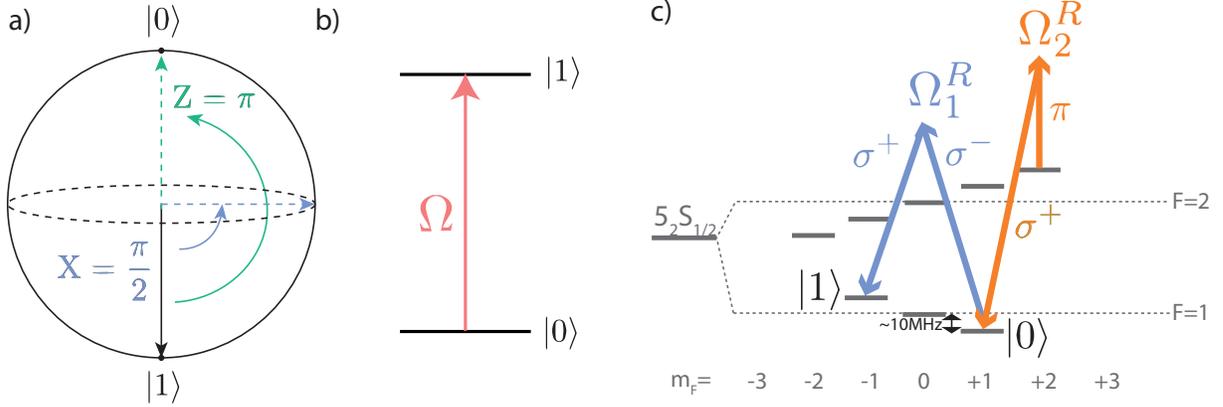


Figure 1-1: Bloch sphere (a) and internal energy level representations of a single atomic qubit gate in an idealized 2-level system (b) and in the ground state hyperfine levels of rubidium-87 (c).

1.3.1 Gates

Single qubit gates on neutral atoms are carried out via π and $\pi/2$ pulses using laser light resonant with the atomic transition(s) of interest. For example, if we have the simplest case of a 2-level atomic system as shown in Fig. 1-1b), we can use a π pulse to evolve the atom's state from $|0\rangle$ to $|1\rangle$ or use a $\pi/2$ pulse to make a superposition state, $a|0\rangle + b|1\rangle$. Similarly in a more realistic full level structure of the Zeeman shifted ground state of a Rb87 atom in Fig. 1-1c), we can use Raman π and $\pi/2$ pulses (shown in purple) to change the state of the atom from $|0\rangle$ to $|1\rangle$ or to some superposition of those logical states. In this more complex case, we can also choose to address the atom with another beam or beams to remove the atom from the logical basis entirely (shown in orange). Because atoms are not ideal two-level systems, there are many opportunities for engineering the logical basis and using other states to aid in selective excitation or readout for different atoms, often referred to as "shelving." In either the single or two-photon case, we are selecting two specific internal states of the atoms and using resonant laser beams of specific duration ($\Omega t = \pi$ or $\Omega t = \frac{\pi}{2}$) to manipulate the state of the atom on the Bloch sphere representation of the Hilbert space of those two logical states, as seen in Fig. 1-1. Using these methods, the highest fidelities of single qubit gates in neutral atoms is $\sim 0.1\%$ [36].

Higher qubit number gates can be decomposed into multiple two-qubit gates, so

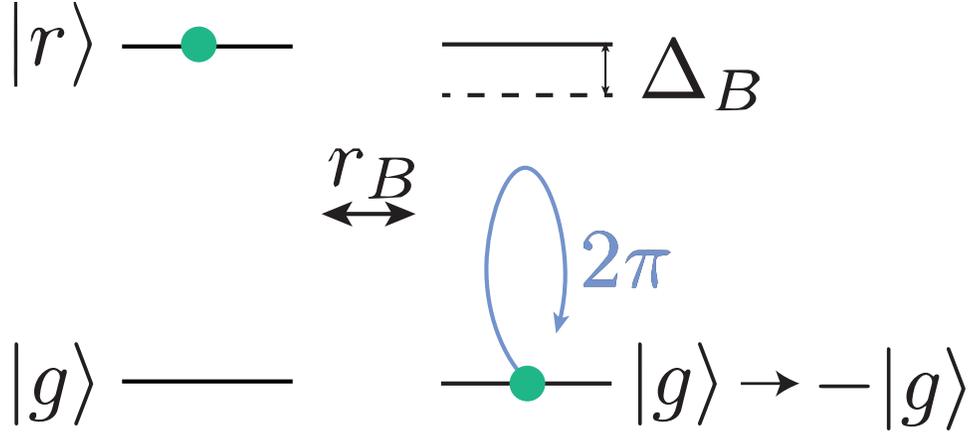


Figure 1-2: Schematic of how the Rydberg blockade enables a C-Phase gate between two neutral atoms, creating entanglement between the two atoms when they are within a Rydberg blockade radius, r_B , of one another.

here we will only discuss two-qubit gates for simplicity [74]. State of the art two-qubit gates in neutral atom array systems rely on excitation to Rydberg levels with high principle quantum number, n , and the resulting Rydberg blockade mechanism. Rydberg states are high-lying energy states of the atom, with n in the tens (depending on atomic species), up to ionization. For rubidium-87 atoms, those used in this work, there are three different Rydberg state manifolds: $nS_{1/2}$ ($n=19-65$), $nD_{3/2}$ ($n=19-57$), and $nD_{5/2}$ ($n=19-57$) [67, 65]. In these high-lying states, atoms are extremely polarizable, with polarizabilities scaling with n^7 [63]. As a result, the van der Waals forces ($\sim \frac{1}{r^6}$) between these atoms can be quite strong—strong enough to appreciably shift down internal atomic energy states, by an amount Δ_B when two atoms are within a blockade radius, r_B , of one another. Using this effect, one can construct a two-qubit entangling gate. The simplest option with the fewest operations is to construct a C-phase gate, as illustrated in Fig. 1-2.

1.3.2 Motion and Connectivity

Until recently, it was unclear how to engineer the connectivity of neutral atom array systems to enable all-to-all coupling. Because the best multi-qubit gates require atoms to be physically close to one another, most systems have been limited to nearest-

neighbor interactions [87]. However in 2016, Kim et al demonstrated a dynamic holographic atom trap with nine atoms [57]. More recently in 2022, Bluvstein et al. pioneered coherent transport of nineteen atomic qubits in tweezer arrays, allowing for near arbitrary connectivity between atoms and even demonstrating small realizations of cluster states, surface codes, and toric codes up to one code cycle. This increased connectivity was demonstrated without increasing atom loss and while staying well within coherence times [16]. The technical ability to move atoms between different locations for different operations on timescales $\sim 1\mu s$ with negligible loss has opened up much more flexibility in these systems. However, the current means of continuous tweezer control, acousto-optic deflectors (AODs), still limit these motions to particular trajectories. A two-dimensional array of optical tweezers is created in an AOD via two crossed crystals that each couple to a separate RF input. This RF input couples acoustically to the crystal, and the resulting phonons are able to deflect light passing through the crystals via momentum transfer. Because each point is deflected by the same two crystals, a fully-filled two-dimensional array cannot also arbitrarily move individual atoms to new locations without additional control optics. This complication makes calculating trajectories to perform the needed operations for a calculation much more difficult. Examples of possible deflection patterns from crossed 2D AODs are shown in Fig. 1-3. By combining static traps, created for example by spatial light modulators (SLMs), and one or two moving traps from crossed AODs, this trajectory problem could be simplified in a serial operation paradigm. Finding optimal methods for increasing connectivity with motion is a current research direction in the field [16].

1.3.3 Readout and Error Correction

The standard state readout method for neutral atom arrays is parallel fluorescence imaging [86]. This technique can achieve high fidelity, $\sim 99.8\%$, but is slow and destructive [56]. Imaging time takes $\sim 100ms$, and can cause a total loss of physical qubits via heating if too many photons are scattered within the imaging time. As discussed in Sec. 1.2.3, demonstrating a fast, nondestructive, and high-fidelity readout method for any qubit platform is critical in advancing proof-of-principle error

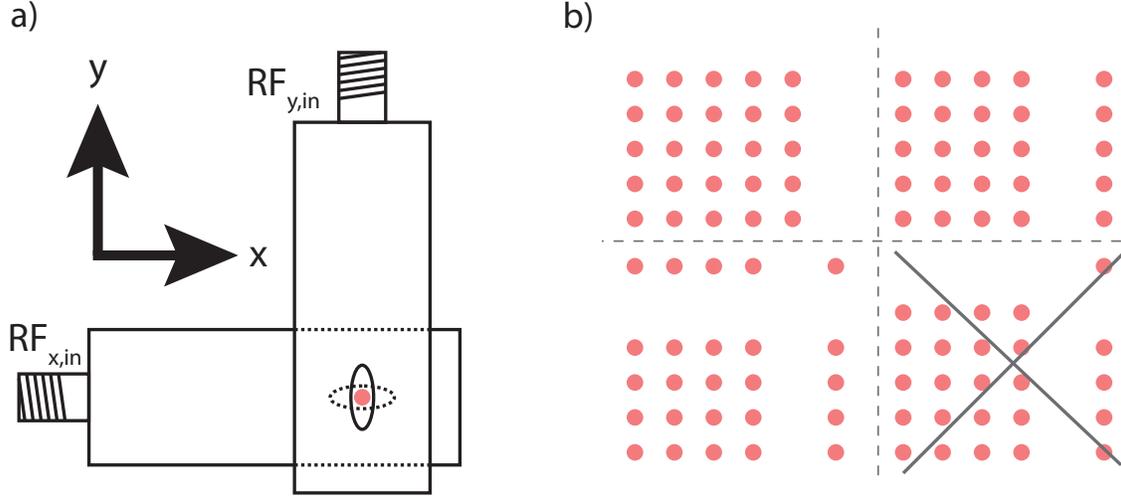


Figure 1-3: Typical arrangement of two acousto-optic deflectors (AODs) to create a two-dimensional array of optical tweezers (a). Three possible and one impossible crossed AOD optical tweezer patterns (b).

correction [85]. Focusing on reading out the state of one atom at a time (i.e. serial measurement) using time-resolved photon measurements can significantly reduce readout time to $160\mu\text{s}$, but this also results in a much lower fidelity, 97.6% [89]. This type of readout scheme that depends on single photon counting modules is fast, but the fidelity is limited by off-resonant pumping and heating. Additionally, scaling up this approach is complicated by significant crosstalk between atoms within the array, causing errors $\sim 4\%$. As such, photon counting fluorescence state readout in free space has not been demonstrated for more than a handful of atoms at a time [85, 62]. Currently, then, we are left with the most parallel, highest fidelity atomic state readout methods that take $\sim 100\text{ms}$. This long imaging time severely limits the speed of code cycles, with the next nearest time scale that of gate durations, orders of magnitude faster at $\sim 100\text{ns}$. The lack of fast and high-fidelity readout schemes in atom arrays limits proof-of-principle error correction experiments. Progress toward surface code error correction is so far limited to one code cycle on a nineteen-atom array because fast mid-circuit measurement on many atoms has not been achieved [16].

1.4 Cavity-Coupled Atoms

Now that we understand where the strengths and weaknesses of neutral atom arrays lie, we can begin to explore the capabilities of optical cavities and the benefits of coupling neutral atom arrays to them for quantum information applications. Our analysis in Ch. 1 discussed several pressing challenges to enabling scaling and error correction in neutral atom arrays: fast and non-destructive state readout, quickly loading/reloading atoms, and scaling up via modularity. Fast, non-destructive readout can be achieved by dispersively probing an atom-cavity system. As we will discuss in Ch. 6 this fast, non-destructive readout can also enable faster loading and reloading of atoms into static, fully-filled arrays. Beyond these near-term improvements to the atom array platform provided by cavity coupling, enhanced photon collection efficiencies from atoms in cavities (via the Purcell Effect) can also help create photonic interconnects between separate cavity-coupled array modules. This can lead to more modular atom array architectures that can continue scaling up even after system sizes are too large to fit within one vacuum chamber due to optical access and optical power limitations. With these improvements to the future of atomic quantum computing in mind, the remainder of this chapter will discuss operating principles behind cavities and atom-cavity systems.

The goal of this work to extend the platform of neutral atom arrays for the application of quantum computers depends uniquely on the fine control and understanding of atom-cavity interactions. Other light-atom interactions are critical, of course, from initial cooling and trapping in our magneto-optical trap (MOT) to the motion and addressing of individual atoms with AOD-based tweezers and Raman shelving beams to Rydberg state manipulation. However, the introduction of an optical cavity is what allows us to escape from previous limitations of atom array systems. To discuss the theoretical concepts behind quantum computing with cavity-coupled neutral atoms, we will first give an overview of bare cavities (i.e. cavities without coupled atoms) followed by the classical picture of an atom in a cavity. After discussing general cavity concepts, we will move on to the special case of atomic state readout via

optical cavities. We will work to understand two different cavity readout schemes, fluorescence and absorption, and some of their capabilities. The remainder of the thesis will then summarize the design, testing, final implementation, and outlook of our new experiment that takes advantage of these basic principles.

1.4.1 Bare Cavity Properties and Formulas

The resonance condition of a two-mirror optical cavity is determined by boundary conditions requiring a standing wave to match its original phase after one round-trip between the cavity mirrors. From this condition, we find cavity resonances when a half integer number of wavelengths can fit within L , as described by the allowed wavenumbers $k_q = \frac{\pi q}{L}$ where $q = (1, 2, 3, \dots)$ and L is the full length of the cavity. Converting this wavenumber condition to frequency, we find that a stationary wave, or TEM_{nm} mode is found every free spectral range (FSR) in frequency, where $\nu_{\text{FSR}} = \frac{c}{2L}$ [93]. Within one free spectral range (constant q), there are varying orders of modes that are described by the quantum indices nm which denote structure in the transverse plane of the stationary electromagnetic field. Specifically for two mirror optical cavities with concave spherical mirrors, the solution of the paraxial (weakly focused) wave equation yields rectangular Hermite-Gaussian transverse electromagnetic modes (henceforth "Gaussian modes"). Plotting the Gaussian modes for different mn values gives us the mode shapes seen in Fig. 1-4, exactly what one sees on a camera at the output of a cavity while tuning a mode-matched input probe laser's frequency across an FSR.

The width of a cavity resonance in frequency space can be found by accounting for the round-trip electric field losses in the cavity that arise from the finite diameter of mirrors, scattering losses in the mirror materials, and scattering from objects inside the cavity. The square of the resulting field in the cavity gives us a damped Lorentzian circulation intensity that is periodic in ν_{FSR} as seen in Eq. 1.1. In this equation, the finesse is $\mathcal{F} := \frac{\pi\sqrt{l}}{1-l}$ with l defined as a fractional loss in the electric field per round trip through the cavity and is conceptually the average number of round trips a photon will take through the cavity before being lost.

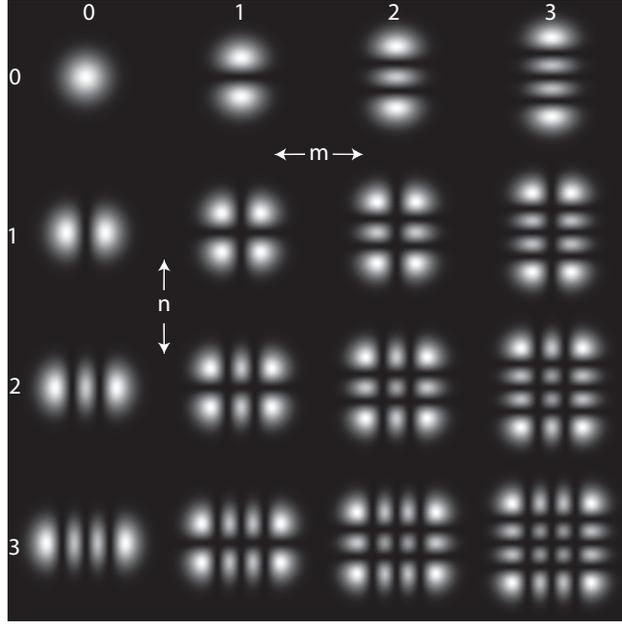


Figure 1-4: The rectangular Hermite-Gaussian transverse electromagnetic (TEM_{mn}) modes of a well-aligned, stable, two-mirror optical cavity in increasing mn orders from top-left to bottom-right.

$$I = \frac{I_{\max}}{1 + \frac{2\mathcal{F}}{\pi} \sin^2\left(\frac{\pi\nu}{\nu_{\text{FSR}}}\right)}. \quad (1.1)$$

We can extract the width of the resonant circulating intensity peaks by solving for ν in Eq. 1.1 when I drops to $\frac{I_{\max}}{2}$ and using the small angle approximation (valid for high finesse cavities). This yields a full width half maximum (FWHM) cavity linewidth of $\kappa_{\text{FWHM}} = \frac{\nu_{\text{FSR}}}{\mathcal{F}}$. The behavior of a bare cavity can be fully specified using just the free spectral range, ν_{FSR} , and the finesse, \mathcal{F} . Both of these parameters will come up frequently in this thesis, specifically in the characterization of our experimental cavity in Ch. 3, and are labeled in Fig. 1-5, which plots the typical Lorentzian lineshape of an optical cavity.

With FSR and \mathcal{F} we can fully describe any stable optical cavity. However, not any combination of spherical mirrors will produce a stable cavity. Depending on the ratio of the radius of curvature of the mirrors and the length of the cavity, the resulting standing wave can either be contained (stable) or not (unstable) [107]. In a stable cavity, the light is confined to travel between the two mirrors with low loss,

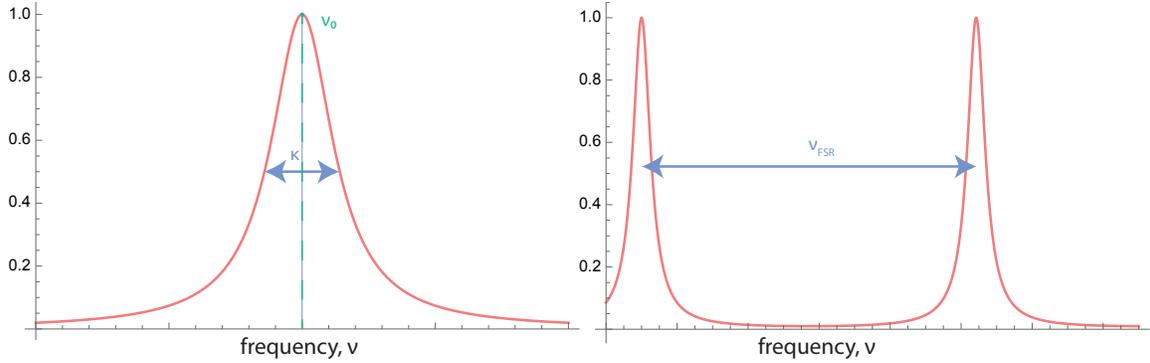


Figure 1-5: The Lorentzian lineshape of transmission of one specific TEM_{mn} mode through an optical cavity. ν_0 is the resonant frequency of the cavity, κ is the linewidth, and ν_{FSR} is the free spectral range.

allowing for interference patterns to form. Quantitatively, the stability of a cavity is parametrized by a geometric factor for each mirror, $g_i = 1 - \frac{L}{R_i}$ [107, 72]. If the product of the two g factors for a given cavity is between zero and one, the cavity is stable and confined. Plotting the stability of a given two-mirror spherical cavity based on the g factor of each mirror in Fig. 1-6, we can point out some typical types of cavities and where our designed cavity sits (more details in Ch. 3).

1.4.2 Classical Model of an Atom in a Cavity

In the limit where an atom in a cavity is not saturated by probe or cavity photons, $I \ll I_{sat} = \frac{4\pi^2\hbar c\Gamma}{3\lambda^3}$, it is both practical and accurate to describe the atom-cavity system classically [98]. This classical treatment describes the cavity modes and atomic levels as two coupled harmonic oscillators, and gives intuition for results that are often considered purely quantum mechanically. Such a treatment recovers important conclusions from the full quantum picture, such as Purcell enhancement and vacuum Rabi splitting. In this section, we will discuss the results of this classical picture most useful to other discussions in this thesis: the definition of the cooperativity parameter, η , and the interaction of a single atom within a cavity mode. For all discussions, we will only consider the simple picture of an atom at rest while ignoring light forces.

We can start by assuming an idealization of the bare cavity system we discussed in Sec. 1.4.1: a lossless spherical two-mirror cavity of length L . This idealized cavity

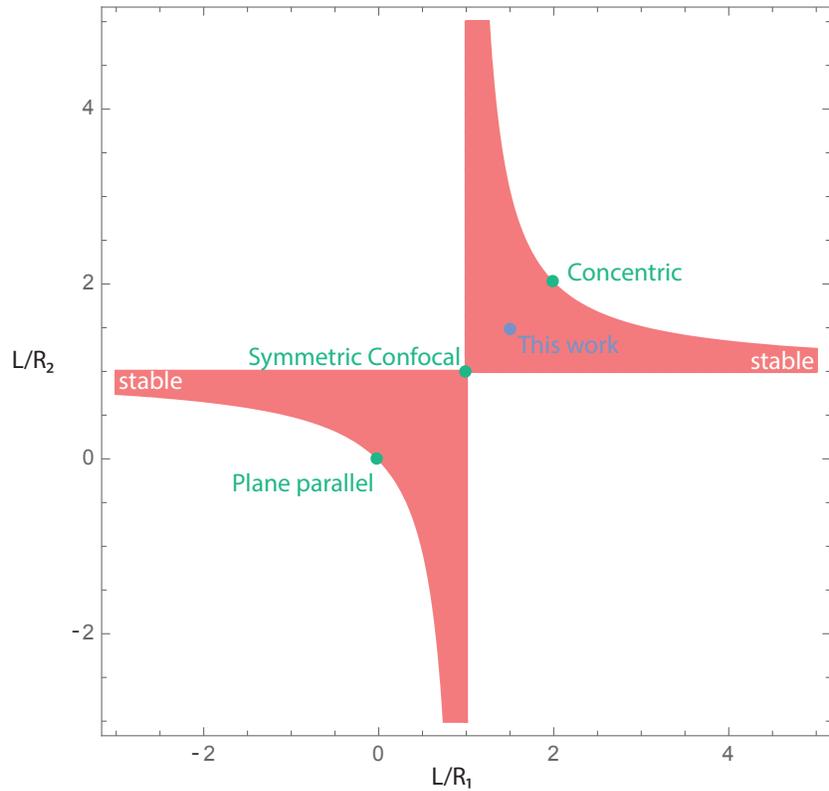


Figure 1-6: The landscape of stable spherical two-mirror optical cavities (pink shaded region). The unshaded regions cannot sustain a cavity mode, as any injected light is unconfined and quickly leaves the inter-mirror region after some number of reflections. The experimental cavity in this work lies directly between two limits in cavity stability, concentric and confocal. This ensures our cavity is robust against small changes in both length and mirror radius of curvature due to manufacturing and alignment tolerances.

has reflection (r) and transmission coefficients (iq) where $r^2 + q^2 = 1$, and we choose $q^2 \ll 1$. Now, we finally put an atom inside the cavity, close to the center of the TEM₀₀ mode and halfway between the two mirrors at $L/2$. By calculating the electromagnetic field amplitude at this position and noting how it interacts with the polarizability of the atom, we can begin to build intuition for atom-cavity coupling. The full derivation can be found in [98] or in [109], but the general procedure is first to find the steady-state traveling mode amplitude E_c by comparing the input mode amplitude E_{in} and the scattered mode amplitude E_M . By again enforcing that a field traveling in the cavity must be the same after one round trip, we can solve for the field amplitudes and convert them into a ratio of the transmitted and incident powers. Similarly, we can calculate the power scattered by the atom by integrating over the radiation field of an oscillating dipole with atomic polarizability, α . This leads us to the ratio of the scattered and incident powers. The final forms of these ratios below in Eqs. 1.2 and 1.3 are reached by introducing some useful parameters, η , Δ , δ , and Γ and by taking the rotating wave approximation (RWA), which is valid for $\Delta \ll \omega_0$. These parameters are defined as follows: $\eta = \frac{24\mathcal{F}/\pi}{k^2\omega^2} = \frac{4g^2}{\kappa\Gamma}$, where g is half the single photon Rabi frequency, $\Delta = \omega - \omega_0$, $\delta = \omega - \omega_c$, Γ is the atomic spontaneous decay rate, and the function \mathcal{L}_a and \mathcal{L}_d are the absorptive and dispersive atomic lineshapes, respectively [58]. Adding these parameters to our power ratios we find

$$\left(\frac{P_{tr}}{P_{in}}\right)_{RWA} = \frac{1}{[1 + \eta\mathcal{L}_a(\Delta)]^2 + [\frac{2\delta}{\kappa} + \eta\mathcal{L}_d(\Delta)]} \quad (1.2)$$

and

$$\left(\frac{P_{4\pi}}{P_{in}}\right)_{RWA} = \frac{2\eta\mathcal{L}_a(\Delta)}{[1 + \eta\mathcal{L}_a(\Delta)]^2 + [\frac{2\delta}{\kappa} + \eta\mathcal{L}_d(\Delta)]}. \quad (1.3)$$

In these equations the single atom cooperativity, η , is featured prominently as the strength of the atom-cavity interaction. This parameter contains all the physical intuition necessary to understand the behavior of atom-cavity systems in different regimes. We can see this in several different interpretations and manipulations of the definition. For example, when defined in terms of κ , Γ , and g , $\eta = \frac{4g^2}{\kappa\Gamma}$ is a ratio of

"good" scattering into the cavity mode to "bad" scattering out of the cavity mode. This gives us a measure of strength of the atom-cavity coupling. Using the definition involving finesse, which tells us the average number of photon round trips in the cavity, $\eta = \frac{24\mathcal{F}/\pi}{k^2\omega^2}$, η is simply the free-space single atom cooperativity, increased by a factor of finesse. In a third different algebraic manipulation of η , we see that it can be understood in terms of a ratio of the resonant scattering cross section of the atom and the cross-sectional area of the cavity mode. In this interpretation, $\eta \sim \sigma_0/A$, where $\sigma_0 = 6\pi/k_0^2$ and $A = \pi\omega_c^2$, and we again retrieve a measure of the degree to which the resonant scattering of the atom overlaps with the cavity mode—a measure of the coupling strength [98]. Conceptually, all of these pictures agree and show us that if $\eta \gg 1$, we are in the so-called "strong-coupling" regime; the regime in which the effect of back-scattering from the cavity mode onto a driven atom cannot be neglected. In this regime, we see the scattering rate of the atom is modified by the presence of the cavity and thus observe the Purcell Effect [98, 41, 81]. It is largely this effect that makes cavity-mediated atom-state readout more attractive than free-space readout.

Equations 1.2 and 1.3 together describe both the absorptive and dispersive effects of putting a single atom in a cavity mode [98, 109]. The absorptive behavior of the atom reduces the power transmitted through the cavity, and the dispersive behavior of the atom shifts the cavity's resonance. The size of the absorptive and dispersive effects depends on the single-atom cooperativity, η , the linewidth of the cavity, κ , and the detunings between the probe light and cavity, δ , and between probe light and atom, Δ . Taking the parameters of our experimental cavity, to be discussed in more detail in Ch. 3, we can begin to understand some of the expected behavior given different detunings. Fig. 1-7 illustrates the spectra of an atom-cavity system where $\delta = \Delta$ and Fig. 1-8 shows the dispersive limit where the atom and cavity are not on resonance.

1.4.3 Cavity-mediated Atomic State Readout

Because much of the motivation for future cavity-coupled atom array experiments comes directly from the ability of a cavity to quickly readout the state of a coupled

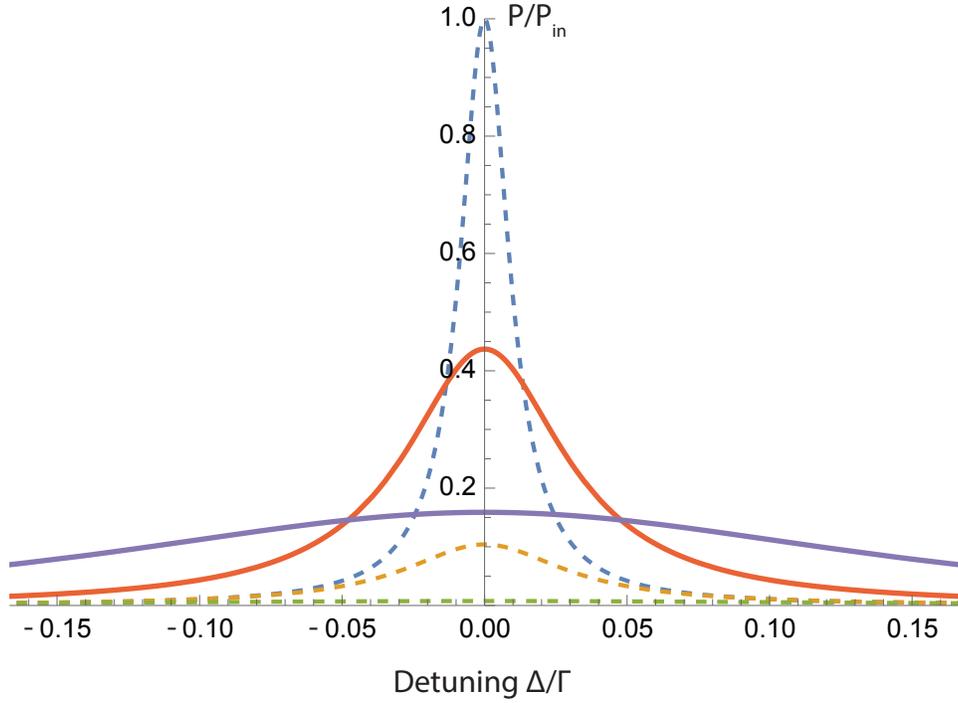


Figure 1-7: Dotted lines correspond to the ratio of transmitted cavity power and input cavity powers for different light-cavity detunings. The solid lines correspond to the light scattered out of the cavity by an atom. For all plotted lines, $\Delta = \delta$ and the other parameters are chosen to be those of the installed experiment: $\eta = 2.1$, $\kappa = 127\text{kHz}$, and $\Gamma = 6.065\text{MHz}$. The blue, orange, and green dashed lines correspond to the transmission spectrum with no coupled atom, one coupled atom, and five coupled atoms, respectively. The solid red and purple lines correspond to the light scattered out of the cavity for one and five coupled atoms, respectively.

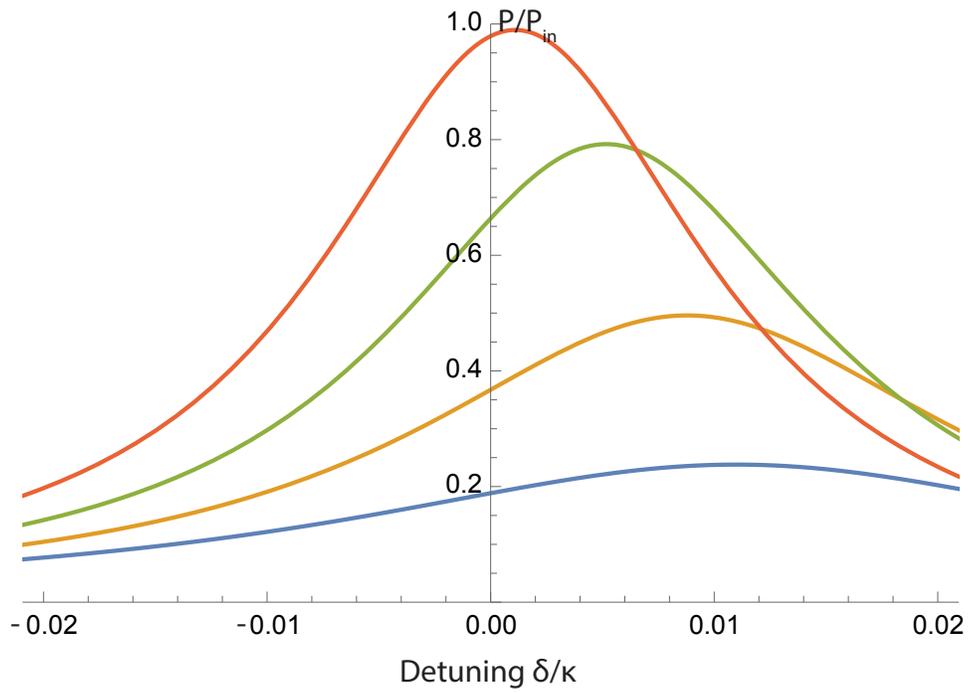


Figure 1-8: The ratio of transmitted cavity power and the input cavity power as a function of cavity detuning from both the atom and light, giving us a dispersive shift in the resonance of the system. From bottom to top, the curves correspond to a cavity detuning, Δ , of $\Gamma/2$, Γ , 2Γ , and 10Γ : blue, orange, green, and red, respectively.

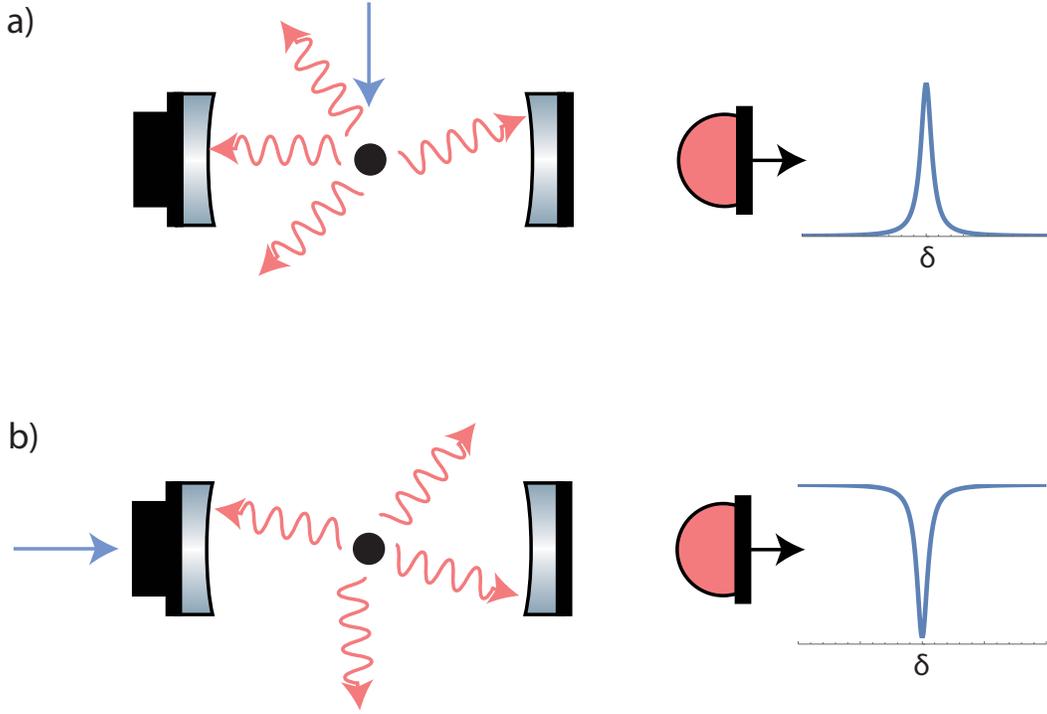


Figure 1-9: The two methods of cavity-mediated atomic state readout: a) fluorescence and b) absorption.

atom, here we will discuss two methods to do so. An atom-cavity system can be probed in two ways: through the cavity mode in absorptive measurements and via light scattering directly off the atom and into the cavity mode in fluorescence measurements. In both cases, it can be determined whether the atom's internal state is on or off-resonance with the cavity by counting the number of photons detected through the cavity mode. As illustrated above in Fig. 3-19, if the cavity system is probed through the cavity mode and an atom is on-resonance with and coupled to the cavity (Fig. 1-9b)), the number of photons in the cavity mode decreases. However, if there are no atoms present or the atoms are in a state that is not on-resonance with the cavity, the transmission counts remain high. Conversely, if an on-resonance atom is probed from the side of the cavity in fluorescence measurements (Fig. 1-9a)), the number of photons scattered into the cavity mode will increase with atom number, giving a high photon count for on-resonance and a low photon count for off-resonance. In both cases, the timescale of the readout is on the order of the ringdown time of the optical cavity, $\tau = 2\pi/\kappa$, which is usually quite fast, on the order of microseconds.

More quantitatively, both fluorescence and absorption measurements depend on the number of photons scattered into the cavity mode by the atom(s) [99]. The main difference between these two methods is how the signal-to-noise ratio (SNR) scales with atom number, implying that each method has a regime in atom number in which it is most precise. To discuss the scaling and behavior of these two different readout methods, it makes sense to consider a simple scenario with a atoms, each scattering on average m photons. In the fluorescence case, if our collection optics have a coupling efficiency of α , then we simply need to compare the collected signal of $m\alpha$ to the photon shot noise, the standard deviation of photons detected assuming a Poisson distribution $\sigma_p h = \sqrt{m\alpha a}$ [99, 32]. This yields an uncertainty in atom number of $\sigma_a = \sqrt{a/\alpha m}$, which increases with the atom number. In the absorption case, we can immediately take into account what the probe coupling is through the cavity mode in the most ideal, diffraction-limited case. This yields $\sigma_p h = \sqrt{m/4\alpha}$ and an uncertainty in atom number of $\sigma_a = \sqrt{1/4\alpha m}$, independent of atom number [99]. Thus, for the same scattering rate and one atom in the cavity, the signal to noise in atom number measurement will be less for the absorption method, making absorption readout with a cavity best suited for single-atom sensitivity.

We expect to be able to use our experimental cavity in absorption read-out in exactly this way to determine whether or not there is at least one atom coupled to the cavity mode in a few times our cavity's ringdown time. Indeed, by plotting the two Poisson distributions for no atom and one atom coupled to our experimental cavity with properties summarized in Table 3.2, we expect to be able to distinguish the atom and no atom states with an error right at 1%. This assumes an input photon number of 125, an SPCM coupling efficiency of 20%, and a measurement time of $19\mu\text{s}$ to stay below atomic saturation by a factor of three. The resulting distributions and their overlap are plotted in Fig. 1-10. By waiting longer and collecting more photons, we can further reduce the error in the atomic state determination.

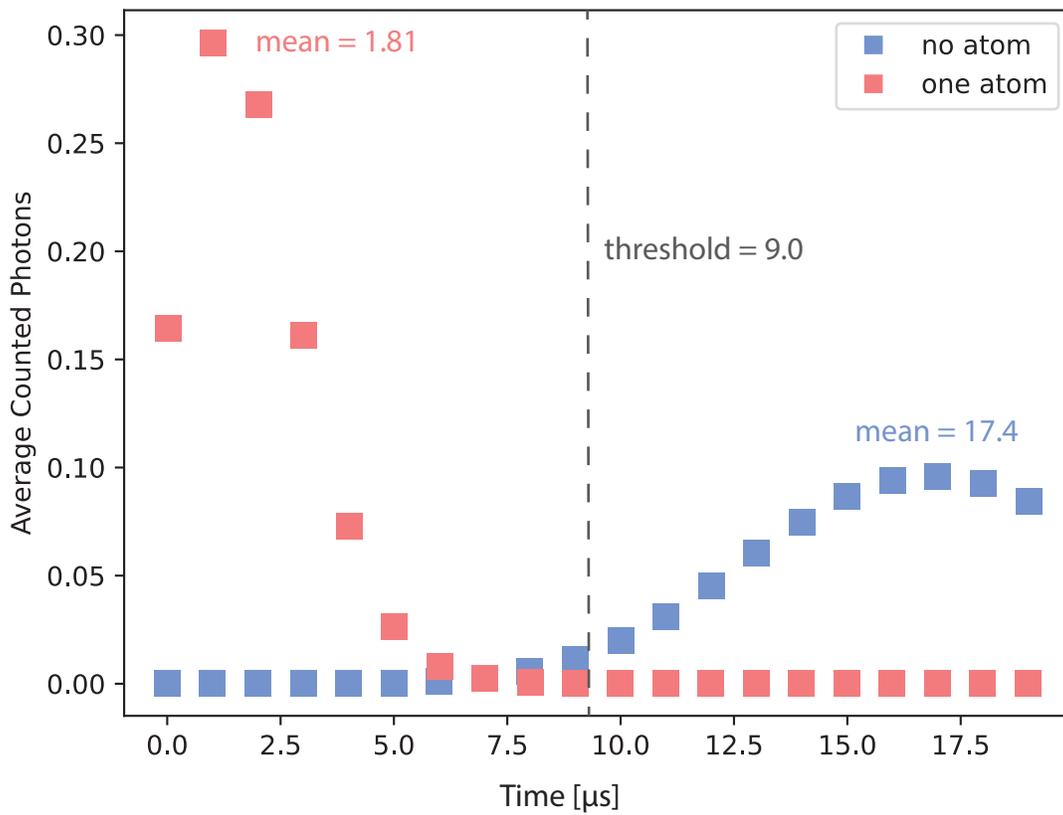


Figure 1-10: Distributions of expected output photons from our experimental cavity in the absorption readout configuration for one atom (pink) and no atom (purple) coupled to the mode.

Chapter 2

Vacuum Apparatus

2.1 Overview

The full experimental apparatus was designed to fulfill all the necessary steps for next-generation neutral atom computing: cavity coupling, optical tweezer trapping, single-site addressing, and Rydberg excitation for gates. For any of these systems to work, substantial improvements and additions had to be made to our vacuum chamber, which was previously empty (besides rubidium) for laser cooling to degeneracy experiments. The majority of the new vacuum hardware had a long design period, spanning 2015-2023. During implementation of the new system in 2022, there were two separate vacuum installation attempts, and some of the components were modified in between the attempts after damage or incorrect installation. The first, failed installation will be referenced several times during this chapter and the next.

The heart of the original vacuum apparatus design was two, two-mirror, near-concentric optical cavities with overlapping modes and distinct cooperativities, situated at an 80° angle from one another (see details in Ch. 3). The crossed cavity mount designed to hold these two cavities has dedicated spaces for stabilization electronics, atom interaction electronics, and two lens tubes for aspheric lenses to focus light of several wavelengths onto the center of the cavity modes. Additional constraints add to the complexity of the overall system, including vacuum chamber size, ultra-high vacuum (UHV) compatibility, Rydberg atom compatibility, optical access, and me-

chanical and thermal stability. The new vacuum components and their properties are summarized in Table 2.1 with labels corresponding to their location in the vacuum chamber as show in Fig. 2-2. All materials on the new structure were carefully chosen to be UHV compatible and robust at the high temperatures required for baking a chamber out to reach high vacuum. This includes the metals for the mounts, epoxies, and all electrical components. Fig. 2-1 gives an overview of the full chamber and its supports as it looks in from the entrance of the laboratory, not including the out-of-vacuum optical systems.

Many additional optical and electrical systems were built outside the chamber both to stabilize the cavity system and to trap and interact with the atoms in the cavity mode. Fig. 2-3 displays a high-level block diagram of all such systems and their relationship to the hardware in vacuum. These separate subsystems will be broken down in more detail in future chapters, namely Ch.5 for atom trapping and addressing optics. This chapter will only focus on the new hardware installed in the vacuum chamber, with the exception of the optical cavities themselves, which will be discussed in great detail in Ch. 3. Through each of these apparatus chapters, we will build a thorough understanding of the capabilities of the new system and how they were achieved, charting a clear conceptual path toward our proposals based on cavity readout in Ch. 6.

2.2 Electrical Components

In this section, we will detail each of the in-vacuum electrical components, their purpose, and connections. We start with an overview of all components before moving into the specific subsystems: cavity length control, temperature control, and electromagnetic field control.

2.2.1 Locations, Connections, and Feedthroughs

During the vacuum installs, it was critical to keep track of all the electrical connections, from their locations to their strain relief and even to their typical resistances.

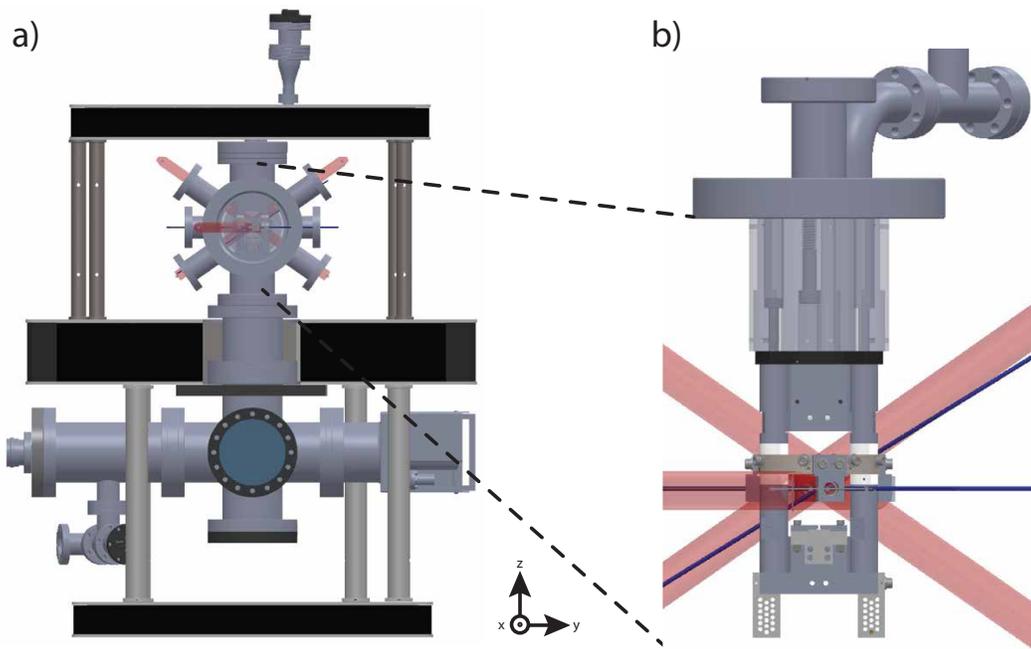


Figure 2-1: View of entire vacuum system, as seen from the entrance to the laboratory (a). The top portion of the chamber above the thick breadboard is the experimental chamber, with MOT coils, MOT optical access, and cavity optical access. The portion of the chamber below the thick breadboard contains the ion pump (far right), the titanium sublimation (TiSub) pump (far left), and the gate valve (bottom left). Additionally, the optical tweezers are sent in toward the experimental chamber through the bottom window. The magnified callout image shows the new in-vacuum apparatus and where it sits along the height of the experimental chamber (b).

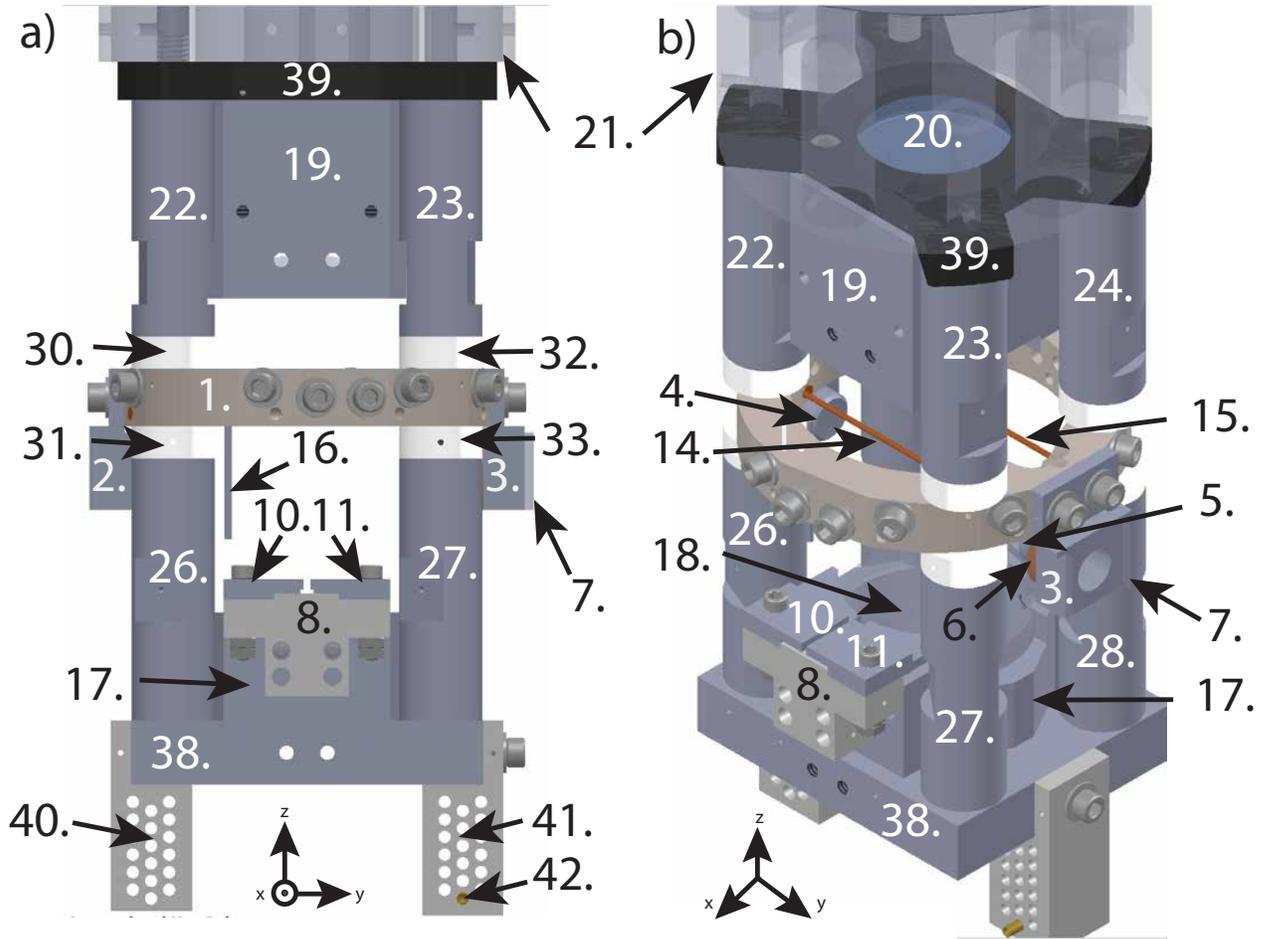


Figure 2-2: a)-d) Full CAD drawing of vacuum cavity apparatus from different angles with components labeled according to Table 2.1. This image only contains the mechanical portions of the apparatus and not the electrical components that were added post-assembly.

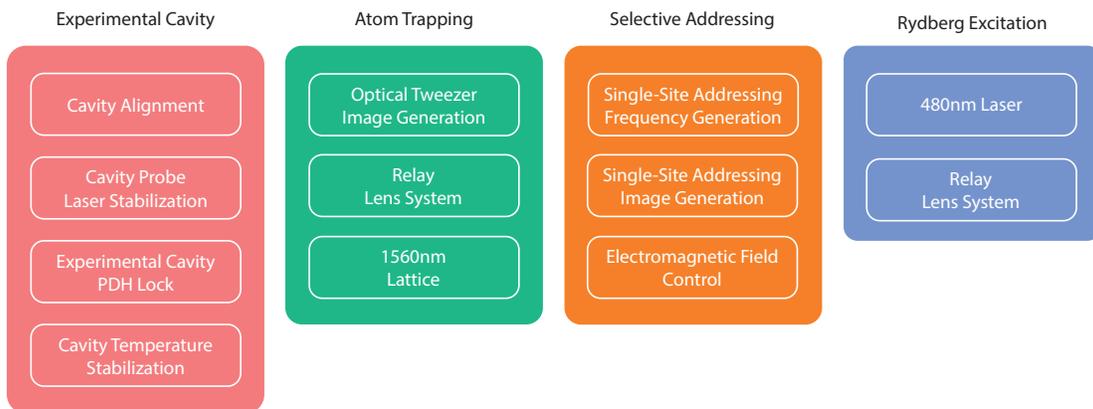


Figure 2-3: Overview of all new systems required for the upgrade to 1) cavity-coupled, 2) Rydberg-excitable, 3) tweezer-trapped, and 4) individually addressed atoms. The apparatus chapters, 2, 3, and 5, will step through each of these subsystems in detail.

Label	Component	Quantity	Material	Part Number
1	Cavity Ring	1	316 Stainless Steel	Custom
2-3	Cavity Mirror Mount	2	316 Stainless Steel	Custom
4-5	Cavity Mirror	2	UV Fused Silica, Ta ₂ O ₅ , SiO ₂	Custom
6	Cavity Piezo	1	Ceramic	Noliac NAC2123
7	Cavity Piezo Shield	1	316 Stainless Steel	Custom
8-9	Electrode Mount	2	Macor Ceramic	Custom
10-13	Electrode	4	316 Stainless Steel	Custom
14-15	RF Wire	2	Kapton, Copper, and Silver Plating	Accu-Glass TYP2-30'
16	RF Antenna	1	Stainless Steel, PEEK, Kapton, Copper/Gold Plating	Accu-Glass KAP-1CX-19SMA
17	780nm Lens Tube	1	316 Stainless Steel	Custom
18	780nm Aspheric Lens	1	S-LAH64	Thorlabs AL3026
19	480nm Lens Tube	1	316 Stainless Steel	Custom
20	480nm Aspheric Lens	1	S-LAH64	Custom
21	Cylindrical Flange Mount	1	316 Stainless Steel	Custom
22-29	Support Rods	8	316 Stainless Steel	Custom
30-37	Viton Spacers	8	Gold-Coated Viton	Custom
38	Base Plate	1	316 Stainless Steel	Custom
39	Top Plate	1	316 Stainless Steel	Custom
40-41	Wire Mount Organizers	2	Macor	Custom
42	Wire Socket Pins	44	Copper and Gold Plating	Accu-Glass 100190 and 100180
43	Thermistor	1	Glass, Nickel Plating, and Dumet	Digikey B57550G1103F000
44	Heater	1	Ceramic, Glass, Platinum, Nickel	Allextra 343-HEATER- 2X10-V2

Table 2.1: A summary of all new vacuum components installed as a part of the cavity structure.

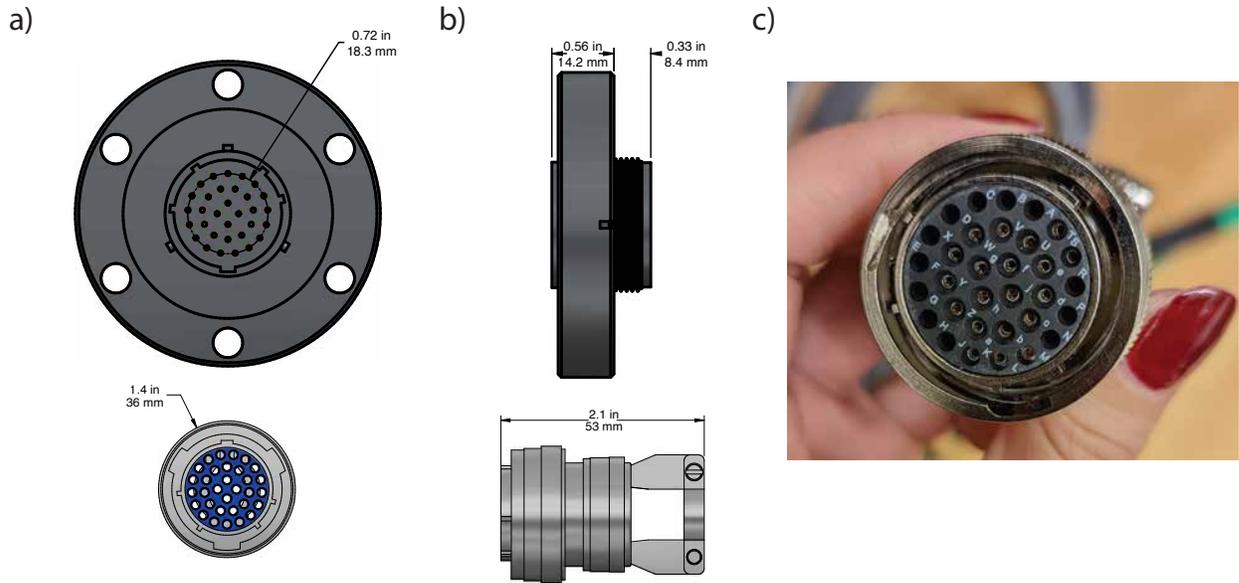


Figure 2-4: a)-b), Diagrams of the 32-pin feedthrough and plug. c) shows the labeled pinouts on the plug side of the connector. Because our cavity apparatus didn't require all 32 pins, only 20 of the pins are crimped in place.

In both installations, some of the connections were damaged. Without documentation, these issues would have been much more difficult to diagnose and fix. The most practical information for typical use of the system will be the feedthrough pinouts shown in Fig. 2-4c) and in Table 2.3.

All of the electrical connections to the cavity apparatus except for the RF antenna are connected to a single electrical feedthrough that sticks up vertically from the chamber, through the oblong slot in the upper custom breadboard. The RF antenna is connected to a single SMA connector feedthrough that sticks out below the 32-pin feedthrough at a ninety-degree angle. We purchased the 32-pin feedthrough and all its connectors from Kurt J. Lesker, and the SMA feedthrough from MDC Precision as summarized in Table 2.2. Currently, not all of the electrical components in vacuum are successfully attached to the 32-pin feedthrough, so the failures and unused connections are also summarized in Table 2.3.

We made all electrical connections with a combination of vacuum solder, silver epoxy (Epotek H21D or H27D), and crimping. To simplify assembly, all crimped connections were made between 22AWG wires. As such, all smaller gauges of wires were

Part Description	Manufacturer	Part Number	Quantity
32-pin CF Flange Feedthrough	Kurt J. Lesker	IFDRG327013	1
Air-side connector	Kurt J. Lesker	FTACIR32AC	1
Vacuum-side connector	Kurt J. Lesker	FTACIR32V	1
Crimp tool	Amphenol	Mil-C-26482	1
Air-Side Crimp Pins	Kurt J. Lesker	FTACIRCONTAC	32
Vacuum-Side Crimp Pins	Kurt J. Lesker	FTACIRCONTV	32
SMA Feedthrough	MDC Precision	9252004	1
Macor D-sub	Central Machine Shop	Custom	2
D-sub male socket pins	Accu-Glass	100170	18
D-sub female socket pins	Accu-Glass	100180	18
22AWG Kapton-Insulated Wire	Accu-Glass	100680	2
SMA Kapton-Insulated Wire	Accu-Glass	Custom	1
Silver Epoxy	Epotek	H21D	2

Table 2.2: A list of all components that connect the electrical components on the cavity apparatus to the vacuum feedthroughs. The quantities of hardware components reflect what is currently installed and working as intended. The quantities of tools and epoxies reflect how much I purchased over the course of the upgrade.

Component	Feedthrough Pins	Notes
Asymmetric/OC-OC piezo	S-T	Only cavity PZT
Old symmetric piezo	Y-b	Not connected to anything
East RF wire	K-g	
Southeast thermistor	j-L	Only working thermistor
Northeast electrode	M	
Northwest electrode	U	
Southeast electrode	V	
Southwest electrode	W	
Southwest thermistor	X-Z	Broken vacuum connection
Southeast heater	a-c	Only working heater
West RF wire	d-f	
RF antenna	SMA feedthrough	

Table 2.3: Feedthrough pinouts and notes on broken connections for all new electrical feedthroughs for the cavity apparatus. Some connections were damaged or removed during the course of the two vacuum installations, and are noted here.

soldered or epoxied to 22AWG wires before crimping into place in the Macor D-sub pieces. This allowed for easy connection between the Macor D-sub pieces and the 32-pin vacuum feedthrough, which is designed for 22AWG wires. Most initial non-crimp connections were done with silver epoxy, because it was easy to cure connections before the apparatus was fully assembled. Likewise, the heaters and thermistors were secured into place on the cavity apparatus with silver epoxy. The high thermal conductivity of this epoxy and low outgassing made it a good choice for this application. Later electrical connections made during repairs or while switching out components were soldered, because we were unable to cure the epoxy on the full assembly. For a full summary of which jointing methods were used where in the cavity apparatus, see Tab. 2.4.

Pictures of the in-vacuum Macor D-sub pieces that house sockets to connect all electrical components on the cavity apparatus to the in-vacuum side of the feedthrough are shown in Fig. 2-2. The Macor D-sub pieces are a natural point at which to repair,

Component	Joint Location/Purpose	Joint Type
Cavity Piezo	Electrode to 28AWG wire	UHV solder
	28AWG wire to 22AWG wire	UHV solder
	Piezo to mirror mount	Silver epoxy
	22AWG wire to Macor Dsub	Crimp
Heaters	Heater to cavity ring	Silver epoxy
	Heater to 22AWG wire	UHV solder
	22AWG to Macor	Crimp
Thermistor	Thermistors to cavity ring	Silver epoxy/UHV solder
	Thermistor to 22AWG wire	Silver epoxy/UHV solder
	22AWG wire to Macor	Crimp
RF Wires	22AWG wire to Macor	Crimp
Electrodes	22AWG to electrode plates	Screwed in between two nuts
	22AWG to Macor	Crimp
RF Antenna	Feedthrough	Screwed into SMA port
Ground Wire	Between cavity ring and top cylinder	Silver epoxy

Table 2.4: A summary of the different types of electrical joints that were made at different locations and for different components in the vacuum chamber. Practicality and ease of installation dictated which joint was used in which case.

remove, or add electrical components in vacuum. For future, smaller repairs or modifications (such as adding a heater or piezo connection), changing the connections here may be required. However, during assembly and installation, we realized that the Macor D-Sub design has several flaws. It was extremely easy to unintentionally short adjacent crimp pins together, either through direct contact or through contact with the Aquadag graphite coating on the Macor. After the first failed vacuum installation, it took several days of scraping off graphite from these surfaces with a razor blade and wrapping individual crimp pins in Kapton tape to remove all unwanted electrical connections. Originally, all insulating components were coated in Aquadag to prevent the collection of surface charges that could cause Rydberg state mixing, but we believe this was more trouble than what it would be worth. In future designs, commercial UHV connector options should be considered, and electrical components that are meant to insulate between connections should not be coated in graphite. If we decide to modify the current design, a useful change would be to glue in place double-sided female connector pins inside of each port in the connector. Then, connections could be plugged or unplugged easily from either side with crimped on male



Figure 2-5: One of the Macor Dsub connectors just before the second vacuum installation. The white portion at the top of the Dsub broke off completely from the remainder of the Dsub and was glued back into place using TorrSeal. Looking closely at the areas on the ceramic piece between the separate pins, scrape marks from removing shorts through the graphite coating are visible.

connector pins. Another note for future modifications or upgrades: one of the Macor Dsub pieces snapped into two pieces during the first vacuum install, so it was repaired with a small amount of TorrSeal before the second installation. Figure 2-5 shows both the triangular fault in the Macor after being epoxied back together and the scrape marks in the graphite coating after removing all the shorts between the electrical connections before the second vacuum installation.

2.2.2 Cavity Length

As described in detail later in Sec. 3.4, we control the length of our experimental cavity to hold the system at a chosen resonance frequency that can be changed with respect to the atoms and cavity probe light. The details of the actuators used to do so are documented in this section.

Piezoelectric Transducer

The cavity length is actively stabilized with one ring piezoelectric transducer (PZT or piezo) on the input mirror. The cavity piezo is part number NAC2123 from Noliac, a stacked ring piezo. The installation of the piezo included gluing it on a mirror mount, coating the mount and piezo in Faraday cage layers of chrome and gold, scraping coating off from the electrodes, and finally gluing on a cavity mirror and attaching an additional stainless steel piezo shield. The electrical connections were made with a combination of vacuum solder, silver epoxy, and crimping as shown in Tab. 2.4. Key properties of the piezo, both from the manufacturer and measured in our lab are summarized in Tab. 2.5.

Between the first and second vacuum installs, we chose to re-coat the piezos and their mounts in chrome and gold, as suggested by Boris Braverman in the initial design of the system [101]. From the time the piezos were originally coated, sometime in 2016, and the second vacuum installation in 2022, most of the original coating had worn off. Since we already replaced all cavity mirrors after the first vacuum installation, it was a natural time to replenish the gold coating. The e-beam deposition tool in MIT's Microsystems Technology Lab (MTL) that was originally used to coat the piezos by Zachary Vendeiro is no longer online, so we had to find a different coating method. We ended up using the metal evaporator in ONELab, a lab in EECS led by Vladimir Bulović. We first masked the piezo electrodes with small pieces of Kapton tape to prevent shorting, and then taped the piezo mounts flat onto a metal platform that was then slid into the evaporator chamber on a meters long translation stage, Fig. 2-6. Graduate student Roberto Brenes sputtered first a 50nm thick layer of Chrome for good adhesion to both the ceramic surface of the piezo and the stainless steel surface of the mirror mount. Second, he sputtered a 100nm thick layer of Gold for electric field shielding. We coated two piezo mounts at the same time, in preparation for potentially reinstalling two cavities into the chamber. In the end, we only installed one. Even with the masking of the electrodes before the coating process, the two electrodes still shorted on each piezo. We to manually removed thin stripes of the

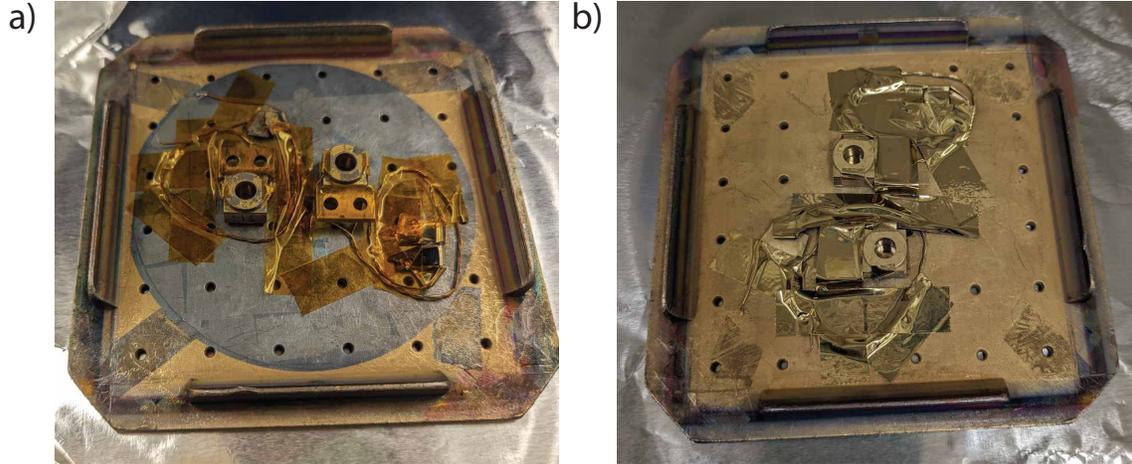


Figure 2-6: Picture of the mounted ring piezos (a) before being re-coated in chrome and gold and (b) after being recoated.

metal using a razor blade to remove the shorts. To check that the piezos were no longer shorted, we monitored the impedance across the electrodes. We applied a small voltage to the piezo in series with a resistor and measured the resulting current across the resistor. When working properly, the real part of the piezo impedance was always $\sim M\Omega$. If at any point the impedance dropped lower, we returned to scraping gold and chrome from the piezo connections until they were fully isolated. In the future, impedance measurements will remain a good diagnostic tool for the cavity piezo even while in vacuum. Note that simply measuring the resistance across the piezo with a multimeter will not yield useful results, as the piezo itself is a voltage source.

After the cavity apparatus was successfully installed in vacuum after the second installation, we characterized the roll-off frequency of the piezo. Like all real electrical systems, the response of piezos is not constant with drive frequency. To determine the high-frequency roll-off of our loaded piezo, we applied a sinusoidal signal to the piezo electrodes and measured the amplitude of the response in the cavity transmission. As we increased the frequency of this drive signal, we noted the point at which the piezo response dropped by a factor of $\sqrt{2}$, which was 135Hz. This measurement allowed us to properly design our cavity feedback system such that no frequencies above 135Hz were sent to the piezo, which would simply filter them out.

Property	Value	Notes
Outer Diameter (ϕ_{out})	12mm	from Noliac
Inner Diameter (ϕ_{in})	6mm	from Noliac
Thickness (d)	2mm	from Noliac
V_{max}	200V	from Noliac
Δd_{max}	$3\mu\text{m}$	from Noliac
Capacitance (C)	380nF	from Noliac
Max Temperature (T_{max})	200°C	from Noliac
Unloaded Resonance (f_0)	486kHz	from Noliac
Rolloff Frequency (f_{roll})	135Hz	Measured
Loaded Impedance (Z_l)	$\sim\text{M}\Omega$	Measured
Chrome Coating Thickness (d_{Cr})	50nm	Deposited
Gold Coating Thickness (d_{Au})	100nm	Deposited

Table 2.5: Key properties of the cavity piezo, both provided by the manufacturer and measured in our lab.

Temperature

Although it is not strictly necessary to temperature stabilize an optical cavity in vacuum, we chose to do so for one important reason. Because we intend to couple highly-polarizable Rydberg atoms to the cavity mode in the future, we would like to minimize the field at the piezo electrodes. One way to do this is to optimize the system and then step down the voltage at the piezo while changing the cavity ring mount temperature to compensate for length changes. If we do this, we can ensure the cavity is locked at the desired point while keeping the piezo voltage at or near zero. Practically, we control the temperature using three main pieces of equipment: in-vacuum thermistor (Digikey B57550G1103F000), in-vacuum ceramic heater (Allectra 343-HEATER-2X10-V2), and a PID temperature control box (Wavelength Electronics LFI3751). Both the thermistor and heater are at the same location on top of the cavity ring halfway between the two cavity mirrors on the east side and are epoxied in place using the H21D silver epoxy. Although the temperature control box has an auto-PID tuning function, we found it did a horrible job finding the correct gains, so we tuned

Setting	Value
T_{set}	27.15°C
I_{min}	-0.4A
I_{max}	0.0A
P	7
I	10.0
D	0

Table 2.6: Current settings for the cavity ring mount temperature control box, Wavelength Electronics Model LFI-3751.

the coefficients manually. The current settings for the box are in Table 2.6.

2.2.3 Electromagnetic Field Control

The cavity apparatus has three main categories of non-optical electromagnetic field control: electric, magnetic, and RF. In the upgrade to the cavity apparatus, the most significant additions have been in terms of electric and RF control.

Because of the eventual plans to excite atoms to highly-polarizable Rydberg states in the cavity mode, care was taken in the initial design to ensure that external electric fields could be shielded and/or actively cancelled. The largest electric field source in the chamber is the electrodes on the cavity mirror piezo. As such, the piezo is coated in a 50nm thick layer of chrome followed by a 100nm thick layer of gold, as described above in Sec. 2.2.2. In addition to this coating, a stainless steel Faraday shield is taped over the piezo. Based on measurements in Zachary Vendeiro’s thesis, it wasn’t clear whether this shielding would be enough to prevent Rydberg state mixing and decoherence [101]. Thus, there are also four separate electrodes below the cavity mode, mounted onto the lower, 780nm lens tube. A different potential can be applied to each of these metal plates to adjust either the absolute electric field or the electric field gradient at the atoms. If these four electrodes were not enough to cancel fields and gradients at the atoms, the two RF wires could also be repurposed as additional electrodes.



Figure 2-7: Picture of homemade dipole antenna for 6.8GHz transmission.

Although we did replace the main MOT coils due to a corrosion issue while baking, the geometry is almost identical to previous iterations. We have eight magnet coils, all external to the vacuum chamber. The strongest fields can be supplied by the MOT coils in the x direction, and two smaller bias coils along all three axes can provide offsets. The bias coils are the originals from previous experiments, but the new MOT coils were wound after we realized they oxidized during our multiple bake-outs in 2020 and 2021. Because of this oxidation process, we could no longer pump water through the coil wires to cool them. Our new coils are made from the same insulated hollow copper wire as before, but now each coil has 117 turns and they were designed to be removable from the chamber. That is, the coils are wound in a way that they can slide on and off the large bucket windows on either side of the chamber, as shown in Fig. 2-8. The last MOT coils were wound directly on the chamber in a way that they had to be destroyed to be removed. To hold the new coils in place, we had the MIT Central Machine Shop make us four custom clamps that holds constant the distance between the coils, as shown in Fig. 2-8. The new MOT coils are controlled with the

same electronics as the old coils, detailed in Zachary Vendeiro's thesis [101].

The RF control in the chamber is in two distinct frequency bands. The two RF wires, which are loop antennae, are intended for use in the ~ 10 MHz regime for global rotations between Zeeman split m_F levels with an external magnetic field. These antennae are considered "small-loop" tranciever antennae, with the perimeter of the loop $\sim 1/10$ the wavelength of the transmitted RF [33]. Although the efficiency of these antennae can be small, the practicality of their size and ability to transmit longer wavelengths made them a suitable choice. The higher-frequency RF antenna in the system is a dipole antenna made from a vacuum-compatible SMA cable (see Fig. 2-7) and was designed specifically for 6.8 GHz transmission ($\frac{\lambda}{2} = 2.193$ cm), allowing for global rotations between the two hyperfine ground state manifolds in rubidium [104].

2.3 In-Vacuum aspheric Lenses

In contrast to some other atom array experiments, we are trapping atoms in a stainless steel UHV chamber rather than in a small glass vapor cell. Working in a vacuum chamber allows us the flexibility to swap out vacuum hardware over time in experimental upgrades, but it also significantly reduces and complicates optical access to the atoms. While other experiments use high-numerical-aperture (NA) microscope objectives or a single high-NA lens to focus optical tweezers and single-site addressing light onto the atoms that are mere millimeters away in a vapor cell, we must span path lengths on the order of a meter between our out-of-vacuum optics and our atoms [35, 9, 10]. To address this challenge, we designed two one-to-one relay imaging systems that incorporate in-vacuum aspheric lenses for the final focusing of tweezer and addressing light onto the atoms at the cavity mode. Recently, other groups have implemented similar optical systems in order to install systems in access-limited cryogenic environments [88]. This section will discuss our relay system design, functionality, and the process of aligning and installing the in-vacuum lenses onto the cavity structure. Because we have two in-vacuum lenses designed for different purposes and wavelengths, each subsection below will generally describe the design and

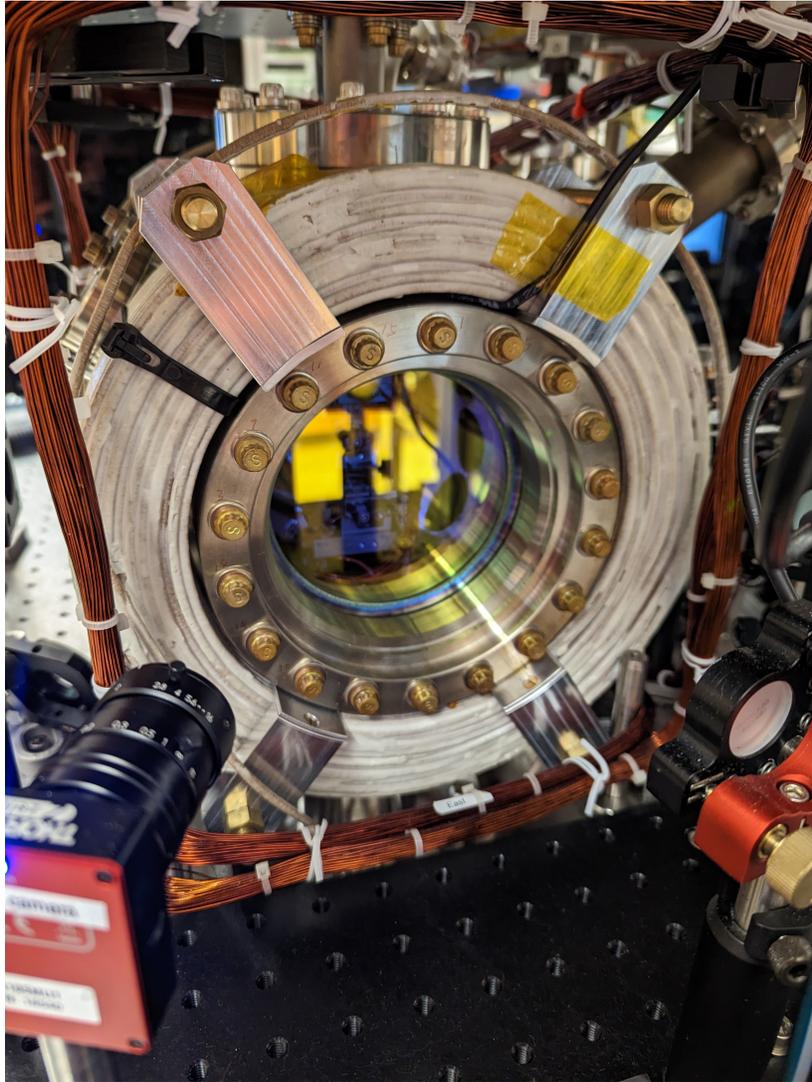


Figure 2-8: The new MOT coils, which were epoxied separately from the chamber for the ease of removal for future hardware upgrades. The coils are clamped into place with a custom four-jaw clamp to prevent drifting of the magnetic field over time. The clamp jaws can be further disassembled for easy installation.

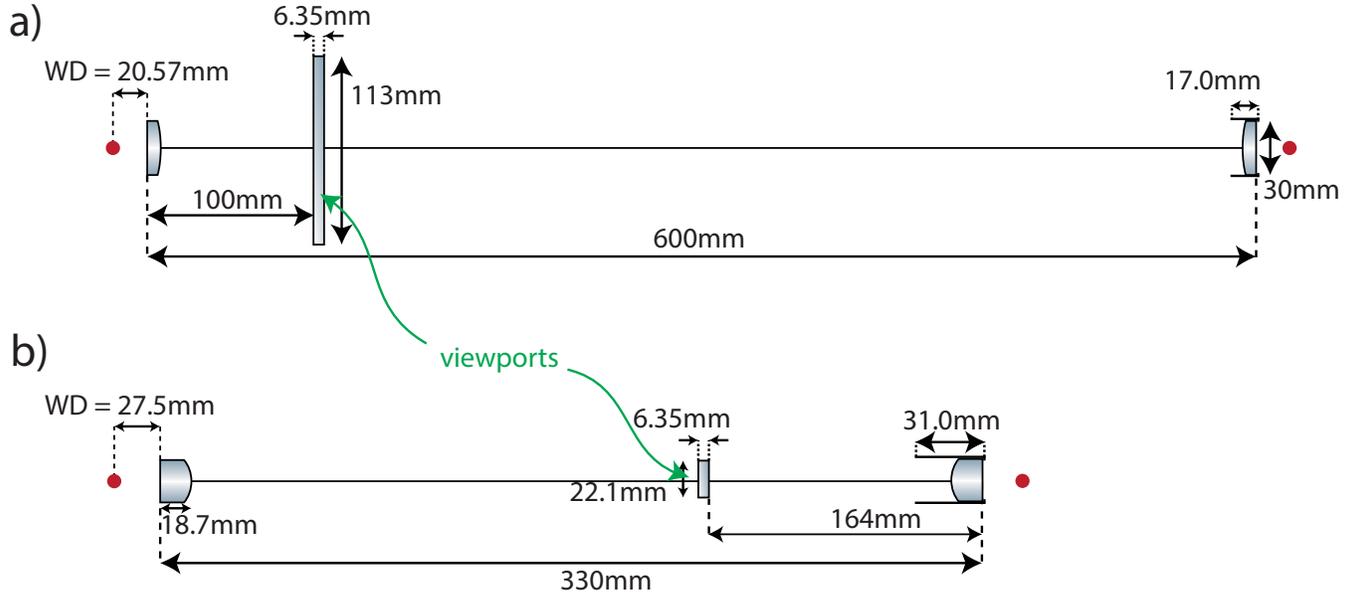


Figure 2-9: Diagram of the two different relay lens systems, for 780nm trapping and addressing light (a) and for 480nm Rydberg addressing light (b). Both diagrams are drawn to scale to illustrate the long distance over which the light must travel in a nearly-collimated beam before reaching the atoms. This distance is set by the length of the vacuum chamber above and below the cavity apparatus to the nearest viewports.

optical setups before tabulating the quantitative results for each lens system separately. The full optical schematic from lasers to AODs for the optical tweezers and addressing beams will be discussed later in Chapter 5.

2.3.1 Relay Imaging Design

The key idea behind the one-to-one relay imaging system is to first create an image of the light that we would like the atoms to see outside of the vacuum chamber before. We then approximately collimate the light with one aspheric lens before re-focusing it onto the atoms with a symmetric in-vacuum aspheric lens. For both the 480nm (upper path, lower NA) and 780nm (lower path, higher-NA) systems, the relay optics are qualitatively the same, as shown in Fig.2-9. Both systems consist of two symmetric aspheric lenses that take an image from outside the chamber and transfer it inside the chamber, passing through a viewport along the way. To confirm and fully understand the effect of this two-lens system on our trap and addressing beams, we designed and

λ	NA	Diff. Lim.	RMS Radius	z_R	FOV	DOF
780nm	0.55	355nm	355nm	$2.03\mu\text{m}$	$101\mu\text{m}$	$2.5\mu\text{m}$
480nm	0.33	364nm	500nm	$3.46\mu\text{m}$	1mm	$21\mu\text{m}$

Table 2.7: Expected performance via Zemax OpticStudio simulation of the two different relay lens systems at the designed wavelengths when perfectly aligned, on-axis. Minimum spot size is measured as the rms spot size until reaching the diffraction limit. The 780nm system is diffraction-limited. The Field of View (FOV) and Depth of Field (DOF) are both defined as the distance in XY and Z, respectively, where the spot size increases by a factor of $\sqrt{2}$.

simulated the system in a ray tracing software before ordering the lenses and checking the results with an in-lab test setup.

The curvature and size of the aspheric lenses used in the two relay paths are different due to wavelength and optical access constraints. However, all relay aspheric lenses are made from S-LAH64, superpolished to minimize wavefront aberration, and AR coated for 480nm, 780nm, 1064nm, and 1560nm by Blueridge Optics. The 480nm (upper) relay path consists of two custom thick aspheric lenses made from S-LAH64. The curvature and thickness of this lens was designed via the optimization function in Zemax raytracing software to have a large field of view at 480nm, making it insensitive to relatively large misalignment. This lower-NA (NA=0.33) system was designed to relay the Rydberg excitation light (at 480nm) to the system and focus it for higher two-photon Rabi frequencies than we could achieve with a fully out-of-vacuum optical system. The 780nm (lower) relay path is made up of two Thorlabs AL3026 aspheric lenses. Because the 780nm system is a higher-NA (NA=0.55) path, the more critical (i.e. requiring diffraction-limited performance) optical systems are relayed through it. In order to trap only one atom at a time, the waist size of each tweezer must be $\sim 1\mu\text{m}$, close to the diffraction limit of 785nm light focused with an NA=0.55, $\frac{\lambda}{2\text{NA}} = 736\text{nm}$. Thus, the 780nm system was designed to relay optical tweezer light (at 785nm) and excitation light (at 780nm or 795nm) onto the atoms. Additionally, it relays fluorescence light (at 780nm, on-resonance) out of the chamber. Details of the theoretical and technical considerations behind atom trapping and addressing optical requirements can be found in Chapter 5.

The first step in simulating these two relay systems at different wavelengths was to characterize the on-axis, well-aligned performance. Aspheric lenses are known for having good on-axis performance, so we expected to meet and exceed our requirements here. The key properties we looked at on-axis were the spot size (for addressing and trapping uses) and the resolution (for imaging uses). We use the root-mean-square (rms) radius of the spot size that results from imaging a point source through each system. The fundamental limit in each case is the diffraction limit, so a rms radius less than the diffraction limit simply tells us we are diffraction-limited. The second step in simulating these two relay systems was to measure Field of View (FOV) and Depth of Field (DOF), that is, to see how quickly imaging quality declined when the input beam was moved off-axis. XY (transverse) displacements are characterized by FOV and Z (focal) displacements are characterized by DOF. Both FOV and DOF are defined as the points at which the rms spot size has increased by a factor of $\sqrt{2}$ from the diffraction limit. We find diffraction-limited performance on axis for the 780nm path, and the 780nm path has a ten times smaller FOV and DOF than the 480nm path, as summarized in Table 2.7.

Lastly, we realized early on that it would be impossible to perfectly align the in-vacuum lens both to the atoms and to the out-of-vacuum lens. In lieu of in-vacuum alignment knobs, we needed to understand the alignment tolerance of the system and how well we could correct for imperfections by making changes only to components outside of the vacuum chamber. To capture the decline in performance as misalignment increased, we simulated spot size, resolution, FOV, and DOF at varying combinations of angular and translated misalignment of the in-vacuum lens. In each case, we used the optimize function in Zemax to move the out-of-vacuum aspheric lens to compensate for the misalignment by minimizing spot size. The results of this tolerancing method can be found in Tab. 2.8. We concluded from these simulation results that we could put up with an in-vacuum misalignment of 1.65mm in Z and 500 μ m in XY for the 480nm lens—well within our mechanical tolerances. However, we could only allow a misalignment of up to 415 μ m in Z and 500 μ m in XY for the 780nm. This result meant we would have to actively align the 780nm asphere lens to

Optical Requirements			Mechanical Tolerances	
λ	ΔX	ΔZ	ΔX	ΔZ
780nm	100 μm	415 μm	1mm	500 μm
480nm	1.65mm	500 μm	1mm	500 μm

Table 2.8: Simulated performance of the two different relay systems at different misalignments of the in-vacuum lens. The optical requirements take into account out-of-vacuum lens movement to minimize the spot size at the output of the full relay system. The mechanical tolerances are calculated from adding together machining tolerances of each UHV piece attached to the lens tubes. An acceptable optical requirement is defined where the spot size stays within $\sqrt{2}$ of the minimum spot size. Note that the mechanical tolerances for the 480nm lens automatically satisfy the optical requirements, where the same is not true with the 780nm lens.

the aligned cavity mode before installing into vacuum. There would be no other way to guarantee overlap between the focal point and the cavity mode.

Figure 2-10 shows the geometry of the relay optical paths in the chamber and which various trapping, addressing, and imaging beams are directed through the in-vacuum aspheric lenses. Although the lenses were not initially designed to combine different wavelengths through each lens, after running into technical difficulties in separating beams with similar wavelengths (e.g. 795nm and 780nm), we were able to keep all the diffraction-limited paths at the bottom of the chamber through the NA=0.55 path and put less critical paths through the top of the chamber through the NA=0.33 path. We compensate for chromatic aberration by adding a weakly focusing adjustment lens or shifting the working distance of the out-of-vacuum aspheric lens on the path. The final decision to choose 810nm light for the optical tweezers was dictated by the performance of commercial dichroic mirrors with wavelength cut-off widths on the order of tens of nanometers.

2.3.2 Relay Imaging Physical Tests

After simulating the two relay systems, we ordered and coated the aspheric lenses to build a test setup in the lab. We wanted to see 1) if the quality of the system matched the simulations and to see 2) how difficult it would be to reach the ideal

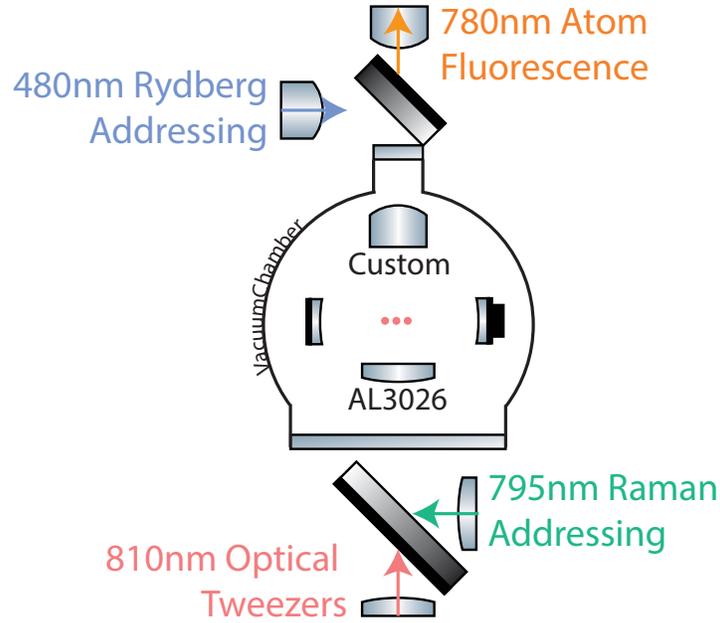


Figure 2-10: Schematic representation of all four inputs/outputs to the relay lens system, through both the 480nm lens and the 780nm lens, not to scale. Complete optical diagrams for these systems are documented in Ch. 5.

performance after aligning by hand in an imperfect world. The setup matched the dimensions in Fig. 2-9 where each lens was mounted on both a tip/tilt mount and a 3-axis translation stage, as shown in Fig. 2-11. The overall procedure for aligning the relay lens test system included setting the distance between the camera and in-vacuum aspheric lens, collimating fiber light with the out-of-vacuum lens, and finally aligning to a point source. The highest quality point source we were able to obtain was a SNOM (Scanning Near-field Optical Microscope) tip, an optical fiber that is pulled to a core diameter of $\sim 200\text{nm}$. However, it was practically impossible to align to the SNOM tip from scratch, because very small ($\sim \mu\text{m}$) displacements from the focal distance of the lens took us out of the DOF of the system. This meant we had no signal to optimize with alignment unless we were within the DOF and FOV. As quoted in Table 2.7, the DOF and FOV of the 780nm system are quite small, so the full alignment took several iterative steps in order to obtain an initial signal on the camera.

The full relay test system alignment procedure from rough to fine is illustrated in Fig. 2-12. First, we collimate the light from a bare fiber (our lab’s standard fiber,

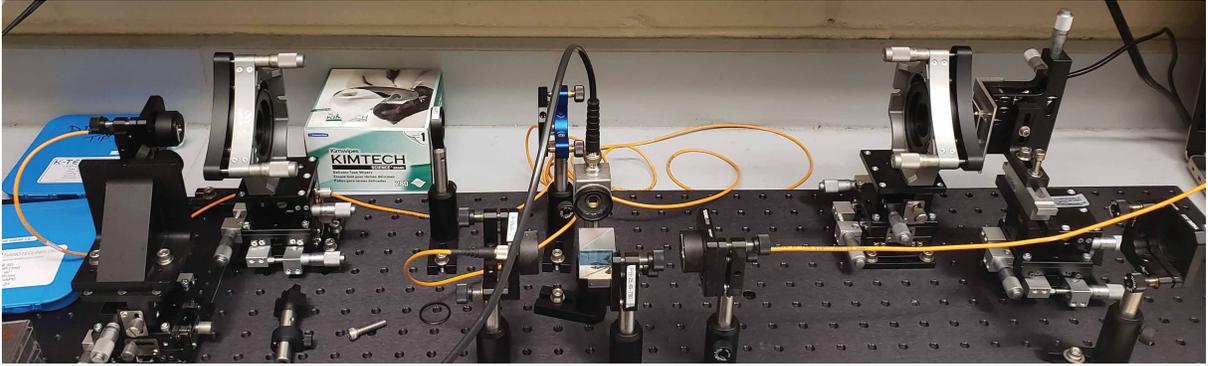


Figure 2-11: Photo of the relay lens test system in the lab during the rough, initial alignment. Here, we first collimate fiber light through an aspheric lens before switching the fiber out for the smaller SNOM tip source. Schematically, this step is illustrated in Fig. 2-12a).

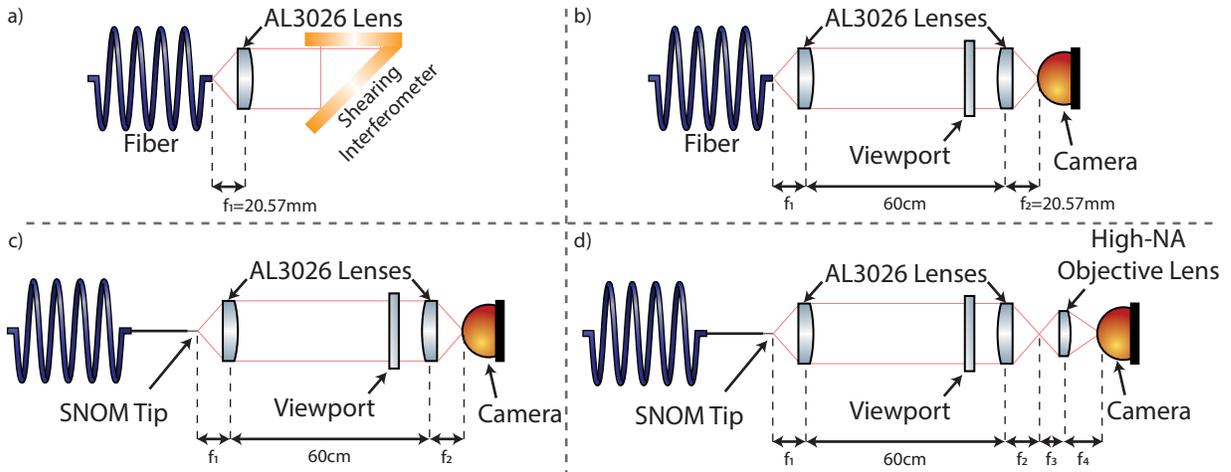


Figure 2-12: The main four iterative steps used to align the aspheric relay lens test setup. First, we use a bare fiber as a source and a shearing interferometer as the figure of merit to set the distance between the source and the out-of-vacuum lens (a). Second, we place a pre-aligned lens-camera breadboard into the path after removing the shearing interferometer, and center the image on the camera before clamping the breadboard down at the distance set by our vacuum chamber (b). Third, we carefully mark the position of the bar fiber tip and swap it out for a SNOM tip fiber, making small adjustments until we recover a point image on the center of the camera (c). We also add a viewport and make the necessary small adjustments to recenter the point image on the camera. After the image is optimized in terms of symmetry and size, we magnify the image from the relay system to better quantify the spot size (d).

OZ Optics PMJ-3A3A-850-5/125-3-1-1) using one of the relay lenses. We attempted two different collimation methods. First, we attempted using a planar mirror to recouple the collimated light from the aspheric relay lens into the bare fiber, and then maximized the recoupled power for fine collimation, shown in Fig. 2-11, where the beamsplitter sends the recoupled light to a photodiode upon reflection. However, after a handful of iterations using this method, we found that, each time, we maximized the power at this photodiode by walking the alignment to where the aspheric lens focused the light onto the retro-reflection mirror, rather than collimating it. All attempts to escape this local minimum were futile, so we moved on to a different method, pictured in Fig. 2-12, which uses a shearing interferometer to qualitatively confirm collimation with the angle of fringes on the interferometer output plate. This proved to be the most reliable method and is what we relied on for future alignment steps. We were unable to simply look at the beam far away from the setup to confirm collimation, because the beam was too large and diffuse to see easily on an IR card. After this step, we use this collimated beam to properly set the distance between the second aspheric lens and the camera. We also add a viewport in the path to ensure that all characterization steps include the refractive effect of the glass, which is present in the real apparatus. After minimizing and centering the spot on the camera by moving the camera in XYZ, we lock it in place and mark the position of the fiber source. The third alignment step is tricky, and involves replacing the initial bare fiber with a mounted SNOM tip. We screw Thorlabs optical bases into our test breadboard to mark the position of the fiber tip before mounting the SNOM. It typically took several minutes of walking the SNOM tip position with XYZ translation stages to recover a signal on the camera. The final and fourth step was to add a magnifying objective lens before the camera to allow for more precise quantification of the spot size. This lens was chosen to have a larger NA than the relay lenses.

Finally, when we had a signal from the aligned SNOM tip magnified through the relay system, we walked the relay lens positions and angles to minimize the spot size. To measure how these adjustments impacted imaging performance, we fit our spots with a Gaussian profile and noted the standard deviation of those fits. With this

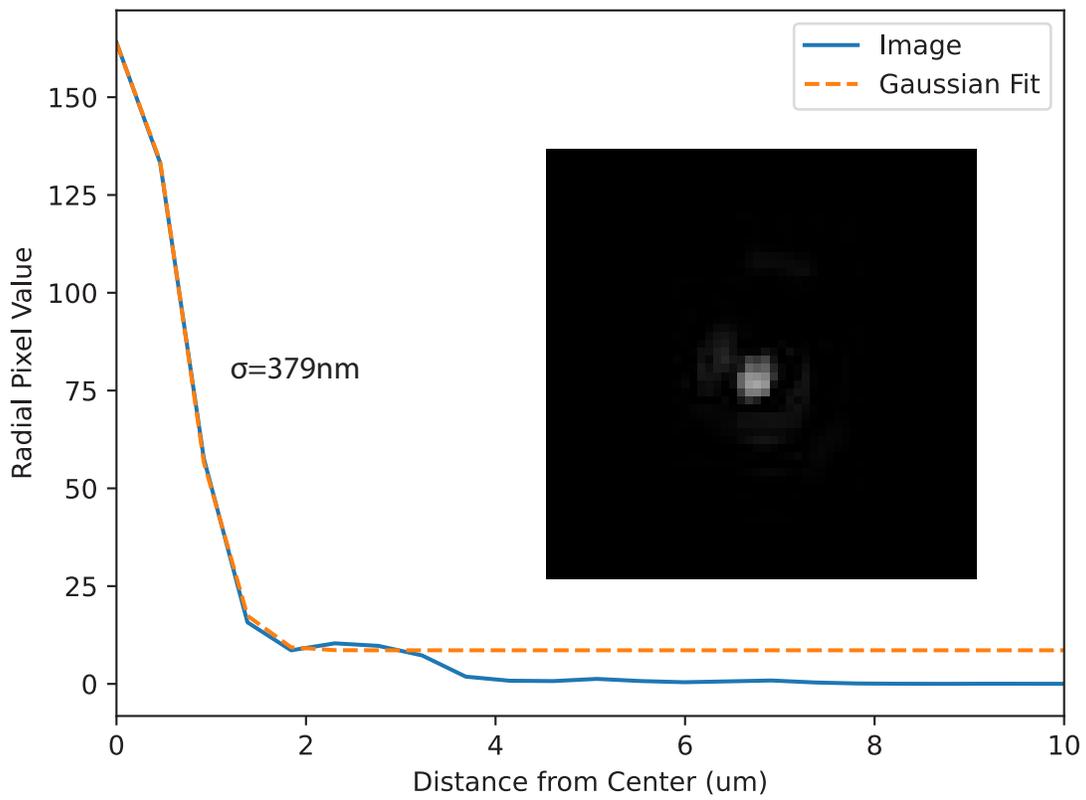


Figure 2-13: Minimum spot size on the physical relay lens setup at 780nm while imaging a SNOM tip.

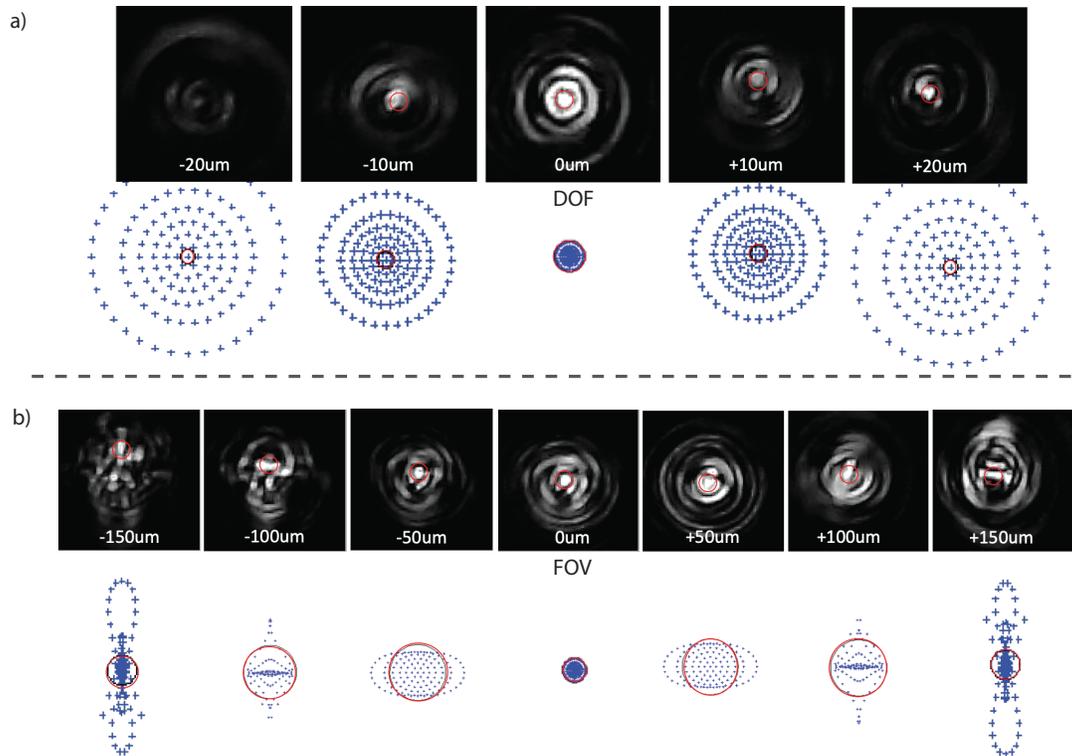


Figure 2-14: FOV and DOF scans on the physical relay lens setup at 780nm while imaging a SNOM tip. The red circles overlaying the images are the diffraction limited rms radius size and are the same size as the red circles below the images in the raytracing spot diagrams.

procedure, the smallest spot size we obtained was $\sigma = 379\text{nm}$, as shown in Fig. 2-13. To extract this number, given the low signal, we performed a radial average over the full spot and fit a Gaussian function to the resulting one-dimensional average. Our ray-tracing software quantifies spot size with rms radius, which is equivalent to one standard deviation. Our measured $\sigma = 379\text{nm}$ is consistent with our simulated rms radius, $r = 355\text{nm}$. After confirming the spot size simulations, we also did a sanity check of the FOV and DOF by scanning the SNOM tip in and out of focus as shown in Fig. 2-14. Our translation stages each had an uncertainty of $\pm 5\mu\text{m}$, but generally confirmed that our system has a $\text{DOF} \sim 10\mu\text{m}$ and a $\text{FOV} \sim 100\mu\text{m}$

Because the 480nm relay system has much less stringent requirements than the 780nm system, we did not carry out such detailed physical characterization steps. Instead, we did a rough characterization using a United States Air Force (USAF) target illuminated with 405nm light. This wavelength was chosen because of readily

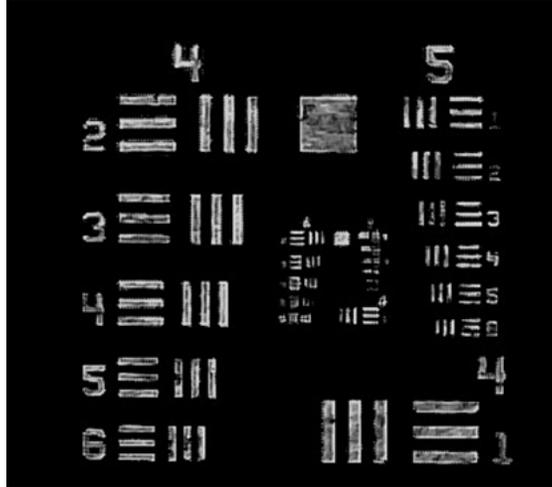


Figure 2-15: USAF target imaged through our 480nm lens system while illuminated by 405nm light.

available laser pointers in that range and also because our Rydberg laser systems had not yet been installed at the point of these tests. The resulting image, shown in Fig. 2-15 shows that the smallest resolvable bars are in group 6, element 2, and thus the resolution of the 480nm relay system is $\sim 5\mu\text{m}$. This resolution is worse than the spot sizes from our ray-tracing simulations, and is likely due to a combination of chromatic aberration, the subjective measure of a USAF target, and a lack of experience with aligning the lenses at this early point in the upgrade. However, unlike the tweezer trapping optics through the 780nm aspheric lens, the resolution of the 480nm system is not critical. Rather, it is just a method to increase the intensity of light at the atoms for larger 2-photon Rabi frequencies when exciting to Rydberg states, and as such was not a critical focus during the characterization and testing steps.

2.3.3 In-Vacuum Lens Install

The conclusions of our simulations and tests convinced us to confidently move toward installation and alignment. Installing the lenses into the cavity apparatus took several steps: cleaning, optical blackening of lens tubes, epoxying, and final alignment within the apparatus. The mechanical design of the cavity apparatus allows for the lens tubes to be moved with respect to the rest of the apparatus, allowing the lenses to be glued

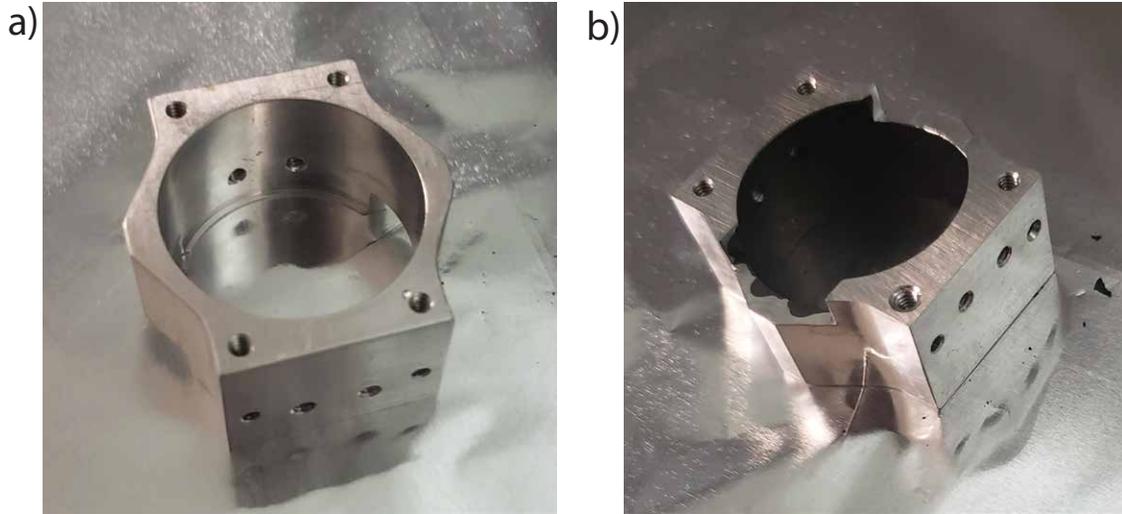


Figure 2-16: The 780nm (lower) lens tube before (a) and after (b) optical blackening.

in place without worrying about alignment.

Lens Tube Subassembly

The first step of assembling the in-vacuum lenses was to vacuum clean the custom stainless steel lens tubes and blacken the inner surface to prevent reflections that could worsen optical performance. Not many paints are UHV-compatible, so using the knowledge that graphite pencil drawings work well in UHV, we used a type of colloidal graphite ("Aquadag") to paint the surface black. Simply, Aquadag is fine ground graphite suspended in a quickly evaporating liquid. Aquadag can be purchased in two forms: a dropper bottle and a nailpolish-like brush bottle. We purchased the brush bottle and used the included brush to apply an even, single layer on the lens tubes. After playing with different methods, we found that a single layer was the easiest to make uniform. Even adding a second layer in small areas caused the Aquadag to bubble and subsequently fracture, leaving some of the stainless steel surface exposed and still reflective. The quality of the coating did not depend on whether it dried at room temperature or on a hot plate ($\sim 60^{\circ}\text{C}$), but the hotplate certainly makes the process quicker. Images of the 780nm (lower) lens tube before and after optical blackening are in Fig. 2-16. In cases where the layer was not uniform and cracked while drying, I removed the entire layer before beginning to apply a new



Figure 2-17: The 780nm (lower) lenstube after a layer of Aquadag was removed via friction.

coat. Aquadag can be removed efficiently through sonication in acetone. Otherwise, friction was the best method to manually remove it, and the results of this method are shown in Fig. 2-17.

Once the lens tubes were successfully blackened with graphite, the next task was to insert the lenses and epoxy them in place. At first, the lenses did not seem like they would fit, because the graphite layer added a non-negligible thickness to the lens tubes' inner diameters (see Fig. 2-18). However, with some patience and pressure, the lenses were able to fit into their designated positions with only minimal scraping of the dried Aquadag. Finally, once the lenses were in place, we mixed and applied a silver epoxy (Epotek H21D) to the exposed sides of the lenses and the nearby lens tube surface. We then cured the whole lens assemblies. Even without this epoxy, however, the lenses were tightly held in place by the tight diameters of the lens tubes.

In-situ Alignment

For the in-situ alignment of the lens sub-assemblies in the cavity apparatus, the procedure was effectively the same as that used during the physical lens tests in Sec. 2.3.2, but more difficult due to physical constraints and vacuum cleanliness requirements. We again used a SNOM tip as our test point source. However, with the real cavity

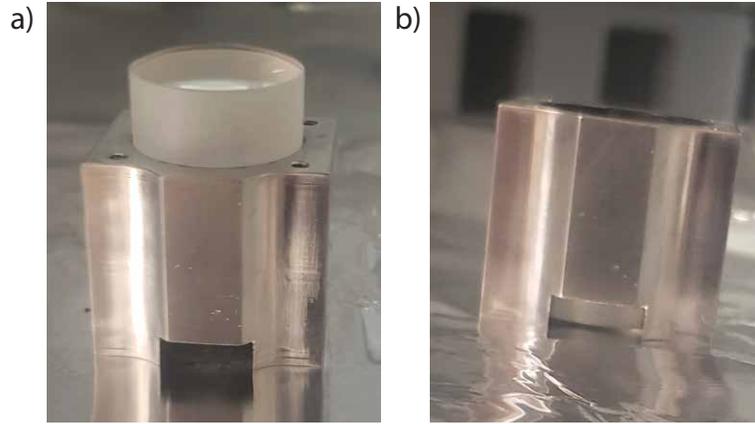


Figure 2-18: The 480nm (upper) lens tube with the 480nm custom aspheric lens sitting on top instead of sliding into the tube due to the thickness of the graphite layer (a) and sitting where it was designed to fit, with the exposed edge of the lens the future location of epoxy (b).

mount and lens tube setup, we had to align the SNOM tip to the center of our aligned cavities before making any adjustments to our aspheric lenses. The biggest difficulty in doing this with our specific apparatus was the lack of mechanical access to the center of the cavity modes. Looking at Fig. 2-2 it is clear that the solid angle surrounding the cavity modes is blocked almost completely, except for a few areas on the order of a couple square centimeters around the diameter of the cavity ring mounting structure. A picture from the actual test setup in the HEPA flow box in Fig. 2-19 shows how the electrical components and cavity mirror mounts significantly reduced the accessibility to the cavity modes. Because the SNOM tip, like any fiber, has a minimum bend radius, it was also difficult to determine how to mount the fiber such that the tip pointed upwards—a ninety degree angle from any of the access points to the cavity modes. Finally, we hammered a piece of stainless steel sheet metal into a quarter-circle with the minimum bend radius of the fiber and fed it into the middle of the cavity modes from underneath the cavity ring. However, simply mounting the SNOM tip within the cavity ring structure at the proper angle was not enough for our needs. We additionally needed fine XYZ control of the position of the tip to properly center it within the cavity modes. To do so, we screwed the sheet metal quarter-circle mount onto a Thorlabs post, connecting the SNOM tip to a manual three-axis stage

with $5\mu\text{m}$ resolution in each direction. With this SNOM setup, we were able to map out the location of the cavity modes by monitoring the cavity transmissions as a function of SNOM position, finally allowing us to center the point source.

After placing the SNOM tip at the atom position, the next challenge was to successfully direct the image of the SNOM tip light through the in-vacuum aspheric lens to a second lens and finally a camera. Because of several physical constraints, we had to set up a bulky optical path to do so. The full cavity apparatus was not only housed within a small HEPA flow box that limited overall breadboard size, but the imaging lens was also elevated off the breadboard by 40cm and directed upward to better accommodate the cavity optical paths, as shown in Fig. 2-19. This positioning of the cavity mounting structure and lens tube required a periscope to bring transfer the image back to the breadboard after passing through the imaging lens. The overall optical path length from the point source to the out-of-vacuum lens also needed to match the expected path length out of our vacuum chamber, about 60cm. After setting up this path with many mirrors, the alignment procedure was conceptually the same as in our lens physical tests described in 2.3.3 with the additional mirrors in the near-collimated beam path.

Once we iteratively walked the near-collimated beam until it was hitting all mirrors in the periscope path and was successfully parallel to the optical breadboard, we placed a pre-aligned lens-camera pair into the optical path. At first, we used a copy of the in-vacuum aspheric lens here to create a one-to-one relay imaging system. Luckily, the SNOM to in-vacuum asphere lens distance was set close enough with mechanical tolerances that we were almost immediately able to see a diffuse image of the SNOM tip on the camera, and clamped down the lens-camera pair where the beam was centered on both the lens and camera while hitting both at a ninety-degree angle without clipping. The image at this point is shown in Fig. 2-20a). From this point, it was extremely difficult to decouple the effects of the in-vacuum aspheric lens positioning and the angle of the steering mirrors along the near-collimated beam path. Eventually, we decided to walk both of these knobs to make the smallest spot size possible (a few pixels on the camera, shown in Fig. 2-20b)), while ensuring to the

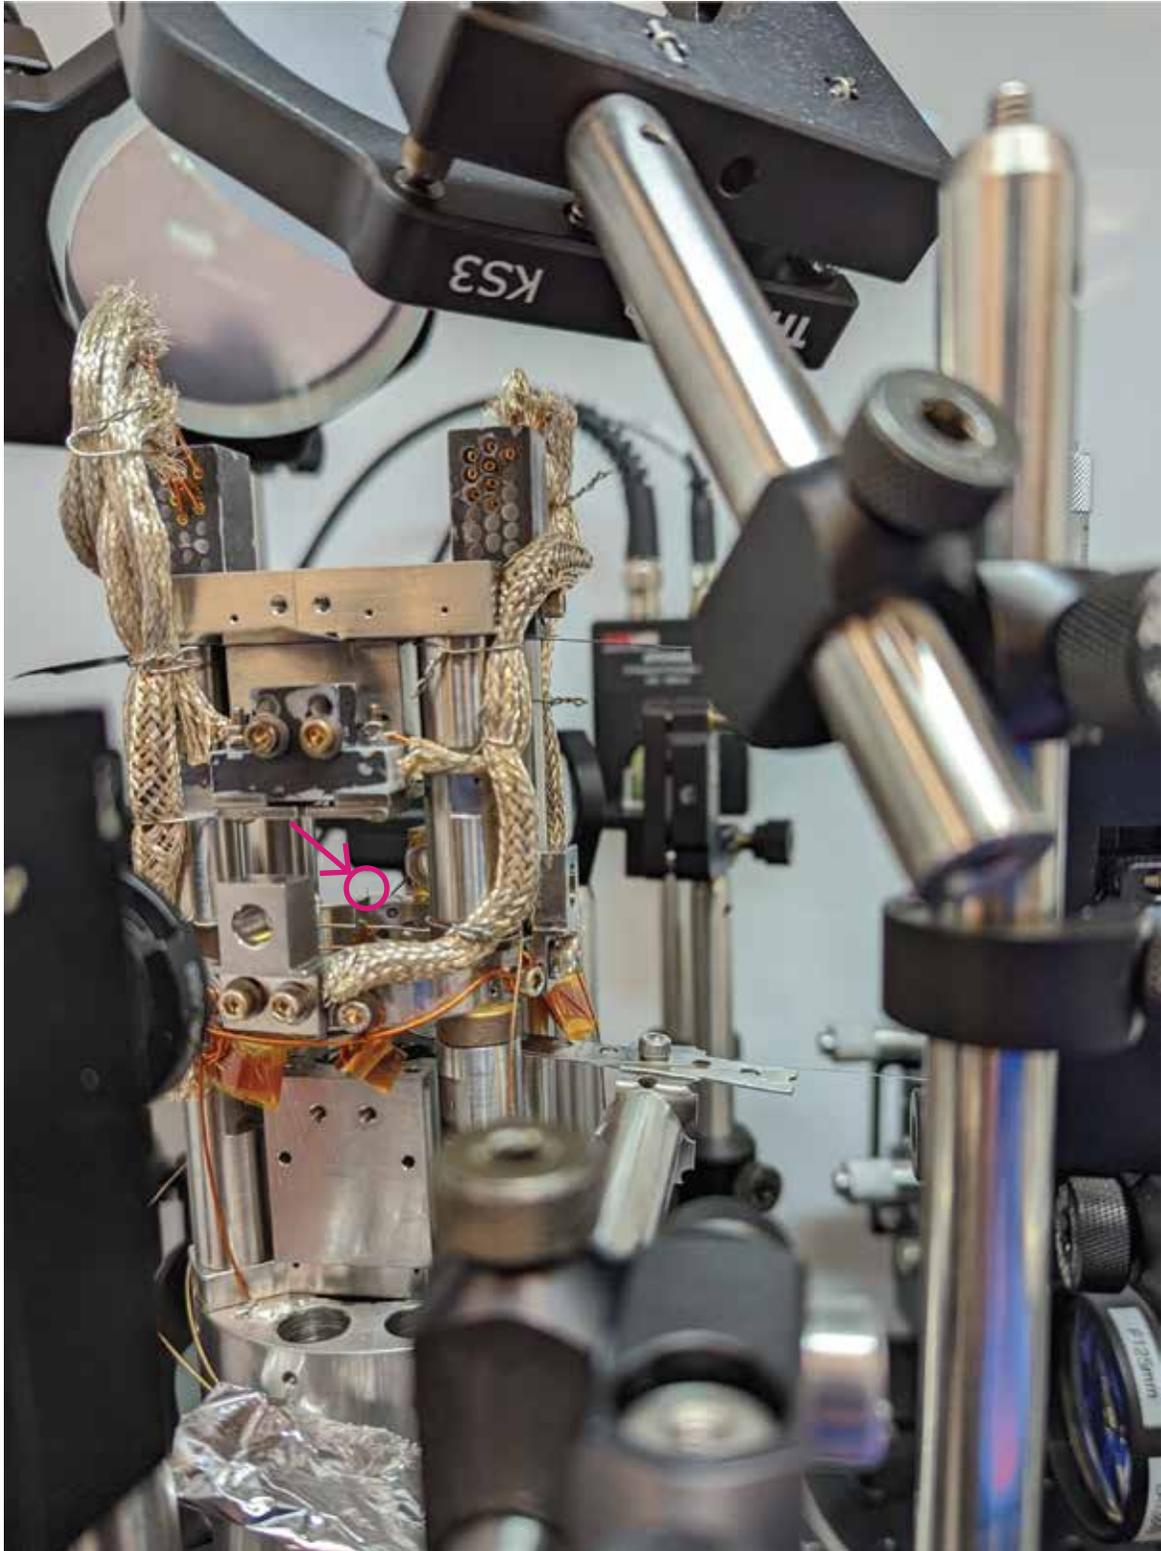


Figure 2-19: The physically constrained setup for aligning the high-NA in-vacuum aspheric lens to a SNOM tip fiber centered on the optical modes of the two original optical cavities. The SNOM tip is circled in magenta, directed upward. The light from the SNOM is collimated through the in-vacuum aspheric lens before being transferred to the breadboard via three, two-inch diameter mirrors.

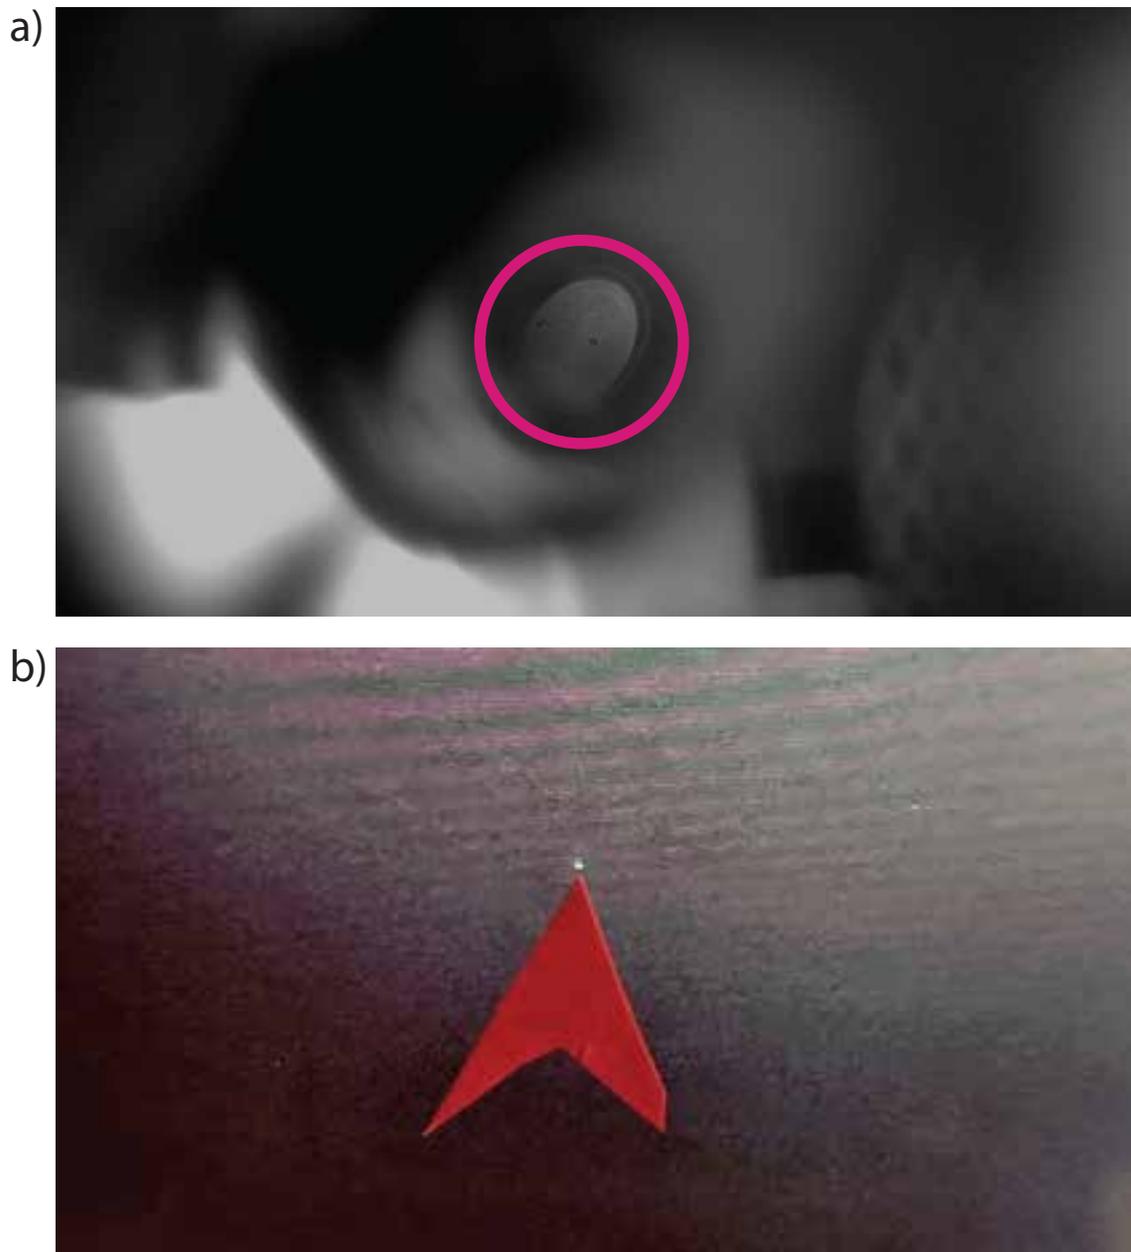


Figure 2-20: An initial image of the light from the SNOM tip at the center of the cavity mode before adjusting the in-vacuum high-NA aspheric lens (a). A more aligned image of the light from the SNOM tip after adjusting the high-NA in-vacuum aspheric lens (b).

best of our ability that the near-collimated beam hit the center of the lens-camera pair at a ninety-degree angle.

After achieving the smallest spot size we could with the one-to-one imaging system, we swapped out the camera-lens pair for a pre-aligned pair with a much longer focal length lens ($f_2=300\text{mm}$), giving us a magnification of about 14.6. At this magnification, we fine-tuned the in-vacuum aspheric lens alignment in XYZ with two goals: 1) make the spot as symmetric and aberration-free as possible and 2) minimize the spot size. We were ultimately unable to rig up a micrometer system that could reliably position the lens tube, so this final alignment was fully manual and largely stochastic with the help of stainless steel shims. Fig. 2-19 shows the placement of these untrimmed shims during alignment—near the top of the structure between the rectangular plate and the lens tube on all four sides. The lens tube is held into place with four screws through the rectangular plate, and tightening these screws in different orders to different torques noticeably effects lens alignment. Finally, we were able to image a spot with an rms radius of $1\mu\text{m}$, without the addition of any shims in the Z direction. Not having to shim the lens tube was pure coincidence and depended on where exactly we mounted the cavity mounts in Z.

2.4 Vacuum Chamber

2.4.1 Flanges and Viewports

Because our cavity-coupled array system requires more and different wavelengths than previous experiments in our chamber, it became necessary to replace the viewports with new ones with anti-reflective (AR) coatings. Additionally, so much new hardware in the vacuum chamber significantly reduces our optical access to the atoms. To remedy this, we swapped out the two large, standard viewports on the east and west sides of the chamber with custom "bucket" or "re-entrant" windows. In total, these changes meant we replaced nine viewports on the chamber, each with new AR coatings for 480nm, 780nm, 1064nm, and 1560nm. When we tested the performance

MDC - MBAR @ <.5 % R @ 480nm- <.25 % R @ 780nm & 1064nm < 4 % R @ 1560nm

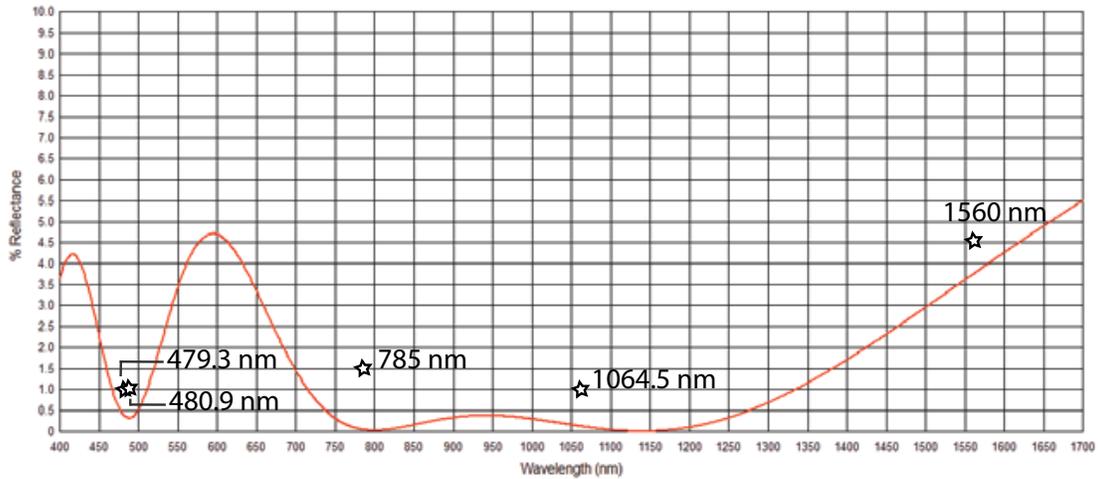


Figure 2-21: Calculated (line) and measured (star) reflectivities of the AR coatings on the new vacuum chamber viewports.

Viewport CF Size [in]	Quantity	Notes
6.75	2	- Custom re-entrant - Front and back of experimental chamber
2.75	7	- Around perimeter of experimental chamber
4.625	2	- Bottom and front of lower chamber

Table 2.9: Quantity and type of new viewports added to UHV chamber in November 2021.

of these AR coatings, we found the results plotted as starts in Fig. 2-21. The exact dimensions and quantities of these new viewports are summarized in Table 2.9.

2.4.2 Regular Maintenance

While we already had the vacuum chamber open and exposed to air to replace the viewports, we decided to complete all possible maintenance on the chamber. We replaced all three titanium filaments in the titanium sublimation (TiSub) pump (see Fig. 2-23 for a view inside the TiSub before filament replacement), replaced the ion gauge, removed the old electrical feedthroughs for the rubidium getters, and also removed the old, spent rubidium getters themselves. The rubidium getters and their support structure were directly in the way of the future location of the new cavity

apparatus, and were not straightforward to remove. A getter support arm was screwed into place inside the vacuum chamber on a custom piece mounted between two flanges in the lower chamber structure. This piece was extremely difficult to access without completely disassembling the chamber (see the bottom portion of the arm in Fig. 2-22). To reduce the amount of disassembly, we simply cut the old getter structure into pieces using a large vacuum-clean wire cutter and then pulled the pieces out of chamber through the already open ion gauge port. Before removal, the rubidium getters had already been sitting empty for years. Our rubidium source is now a vapor cell attached to the upper north viewport flange with a tee and valve. The amount of rubidium in the chamber is dependent both on the degree the valve is open and the temperature of the vapor cell. In the previous laser cooling to degeneracy experiment, the vapor cell was not heated and the amount of rubidium atoms was only controlled by the degree to which the valve was open. After installing the new cavity apparatus, however, we began heating the vapor cell to $\sim 40^\circ$ to increase the atom number and allow for a larger MOT.

In future chamber openings, a few more maintenance tasks should be strongly considered. Because of the large number of vacuum bakes in 2021 and 2022 for the hardware upgrade in this thesis, we opened and closed the gate valve of the chamber many times. The gate valve on our chamber has a finite lifetime depending on the number of cycles, and should certainly be replaced when convenient to maintain the integrity of the vacuum seal over time. Only after our vacuum modifications were complete for this upgrade did we realize that our gate valve digs further and further into a gasket like structure with each closing, and requires periodic replacement. We baked our chamber a total of five different times during 2022. Although we are not sure how many baking cycles the chamber has gone through during its ~ 20 year lifetime, we surely decreased the lifetime significantly from its recommended maximum thirty opening and closing cycles. Additionally, during the first vacuum installation of the cavity structure (details in Ch. 3), the ion pump had issues with shorting across the electrodes. We were able to recover the pump in a dramatic fashion for the time being, but it likely needs to be removed and fully cleaned by

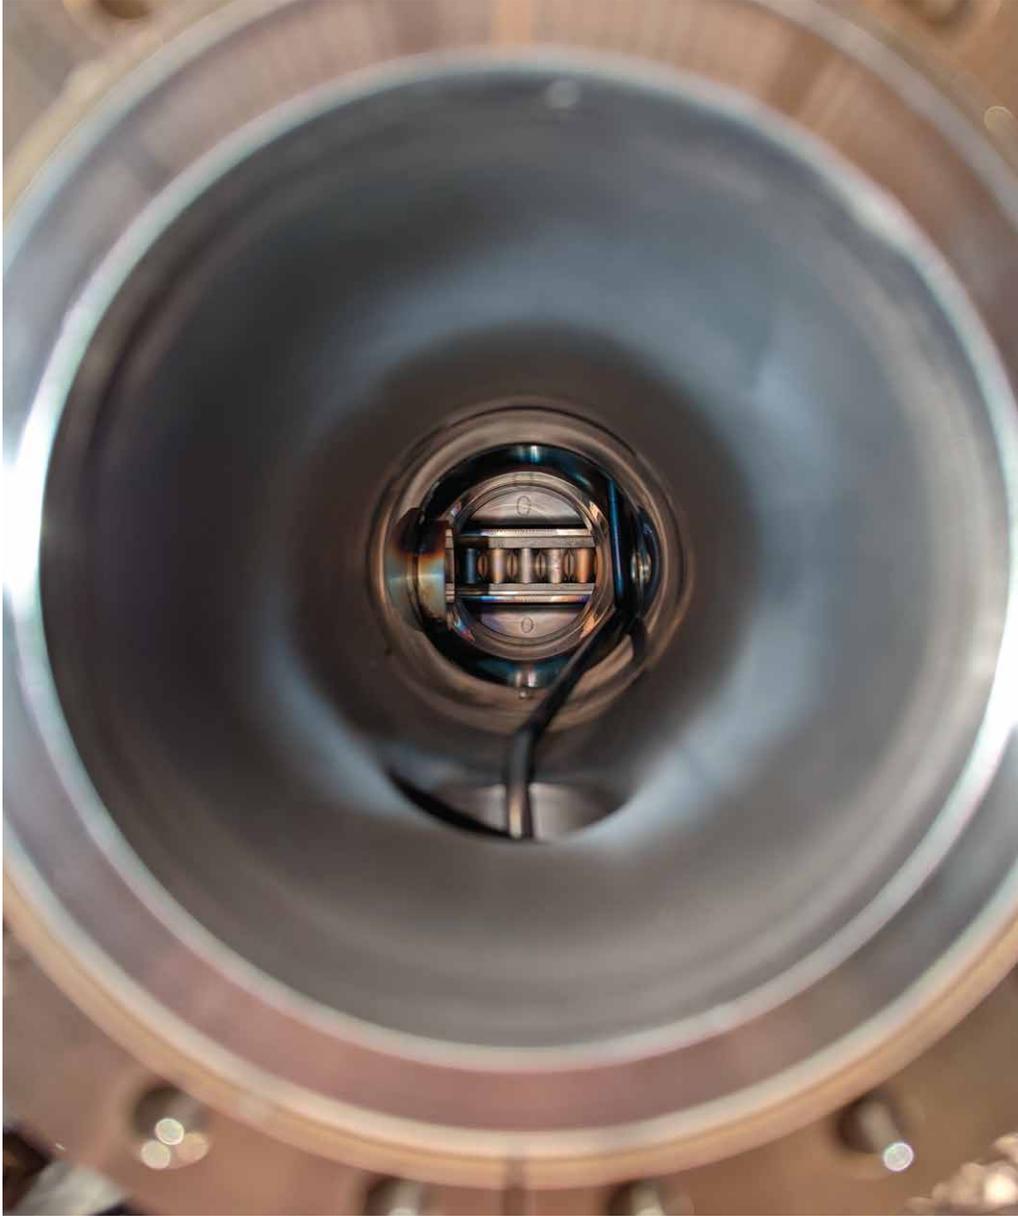


Figure 2-22: A view of the bottom of the old rubidium getter arm through the TiSub flange, looking toward the inside of the ion pump.



Figure 2-23: The spent and broken titanium filaments replaced during the vacuum chamber maintenance portion of the upgrade. We replaced all three filaments with fresh ones in November 2021.

the manufacturer to prevent future failures. During the next substantial vacuum upgrade, we certainly recommend first removing the ion pump and either replacing it with a new pump immediately, or capping off the ion pump flange and pumping out the system with a turbo pump while the original pump undergoes maintenance. This should be completed before moving forward with any more sensitive vacuum installs (especially if another optical cavity is involved).

Chapter 3

Cavity Apparatus

Based on our understanding of atoms and cavities from Ch. 1, this chapter will discuss both the intended cavities and the final implemented cavity through a narrative description of the design, alignment, and installation stages of the hardware upgrade. Initially, the goal was to install two, two-mirror crossed cavities with the same geometry but different optical properties. Specifically, we planned to install one high-finesse symmetric cavity (made up of two high-reflectivity mirrors) and one lower-finesse asymmetric cavity (made up of one high-reflectivity and one output-coupling mirror). This would have allowed us to choose which cavity we couple atoms to, depending on the application. For example, the lower-finesse asymmetric cavity would be best-suited for fast readout of atomic states whereas the higher-finesse symmetric cavity would be best-suited for cavity-mediated entanglement. As discussed in detail later in chapter 4, however, during the first bake, all three of the high-reflectivity (HR) cavity mirrors degraded irreversibly. This degradation forced us to pursue a secondary path: creating a single symmetric cavity with two of our output coupling mirrors. This gave us one cavity with slightly worse properties than both of the intended cavities, but allowed us to move forward quickly with enough functionality to work toward our proof-of-principle experiments involving atom arrays in an optical cavity. With the exception of a discussion on how to overlap two, crossed-mode cavities in Section 3.2.2, this chapter will focus on the installation, characterization, and capabilities of the system that was successfully installed in vacuum in fall 2022, with

one cavity made from two output-coupling (OC) mirrors. For the original crossed cavity design, intent, and testing from 2015-2019, see Ch. 8 of Zachary Vendeiro’s thesis [101].

3.1 Cavity Geometric Properties

All of our attempted optical cavities were designed to be mounted on a ring mount ($\phi = L = 4.39\text{cm}$) as two-mirror, near-concentric cavities. Depending on the desired optical properties of a cavity, we assembled it with either two HR, two OC, or one HR and one OC mirror. Each of these mirrors are geometrically the same on the macroscopic scale, and the HR and OC mirrors simply differ by the thickness of the dielectric stacks on the mirrors’ reflective surfaces. The designed mirror properties are summarized in Table 3.1 for each type of mirror. The path length between the two mirrors, L , was designed to be much greater than 1cm in order to ensure all dielectric surfaces will be far from atoms in high-lying Rydberg states. Both mirror types (OC and HR) are from the same coating run from Advanced Thin Films (ATF) in Boulder, CO and have the same dimensions and radius of curvature. The mirror coating run for the output coupling mirrors specified transmission of 50ppm at 780nm and transmission of 300ppm at 1560nm with an anti-reflective (AR) coating on the planar side of the mirror. The coating run for the high-reflectivity mirrors specified transmission of 2ppm and lowest possible losses at 780nm. The geometric cavity properties are summarized in Table 3.2 and are the same for all combinations of mirror coatings. For more background on the definitions of these cavity parameters, see Sec. 1.4.1 in the introduction.

Combining all the geometric and material properties together, we get a cavity system represented schematically in Fig. 3-1. A cavity with this geometry is considered a stable resonator because the product of the g factor, $g = 1 - \frac{L}{R} = -0.756$, of each mirror $0 \leq g_1 g_2 = 0.572 \leq 1$. In other words, rays of light, when well-matched to the cavity mode, continue to bounce between the mirrors rather than being deflected out of the mode after a certain number of bounces due to beam divergence. Thus,

Mirror Type	Both	OC only	HR Only
Diameter (mm)	7.75		
Thickness (mm)	4		
Radius of Curvature (cm)	2.5		
T/L@780 (ppm)		50/10	5/7
T/L@1560 (ppm)		345/0	1.7/1.8
Substrate	UV Fused Silica		
Dielectric Stack	SiO ₂ //Ta ₂ O ₅		

Table 3.1: Experimental cavity properties for both cavity mirrors.

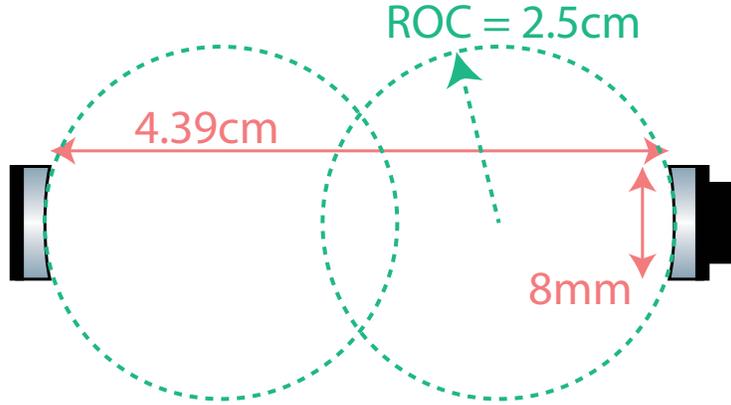


Figure 3-1: Geometry of the experimental cavity.

light only leaves the cavity via scattering and absorption in the mirrors, transmission through the mirrors, and scattering via atoms within the cavity mode. A cavity of these dimensions is not either special case of concentric or confocal, but it is a combination of the two (see Fig. 1-6), making it easier to align than a confocal cavity while still retaining a relatively small waist ($45\mu\text{m}$).

3.2 Out-of-Vacuum Alignment and Characterization

Before installing the cavities in vacuum, it was necessary to assemble, align, and test the components outside of the vacuum chamber. For this purpose, we used a large HEPA flow box to keep the setup in a relatively clean (dust-free) environment that could protect the finesse of the cavity mirrors over the course of assembly and

Length (cm)	4.39
Mode Waist @780nm (μm)	45
Beam Width at Mirror@780nm (μm)	129
Rayleigh Range @780nm (mm)	8.16
Mode Waist @1560nm (μm)	64
Beam Width at Mirror@1560nm (μm)	182
Rayleigh Range @1560nm (mm)	8.25
g-Factor	-0.755
Cavity Type	near-concentric
FSR (GHz)	3.42
Adjacent TEM Mode Spacing (GHz)	4.89GHz

Table 3.2: Designed properties of the installed experimental cavity.

alignment. Before the apparatus was assembled, all components and necessary tools were vacuum cleaned. Any surfaces that could come into contact with the vacuum components but were unable to be cleaned directly due to material or size constraints were wrapped in clean vacuum foil to further prevent contamination. All assembly and alignment procedures detailed in this section were completed wearing powder-free nitrile or latex gloves. This prevented skin oil and other contamination from getting into the system. We were careful to maintain the cleanliness of the apparatus while making any alignment adjustments outside of the vacuum chamber. This meant frequently changing gloves and re-cleaning tools or components when dirty. This environment was used for all assembly, alignment, and characterization steps before vacuum installation. The specific procedures for each assembly, alignment, and characterization will be detailed in the remainder of this section.

3.2.1 Assembly

Before aligning the cavity mirrors and looking for transmitted modes, the cavity structure had to be assembled and the mirrors and their piezos had to be epoxied into place. The assembly of the mechanical structure was based on a CAD design from

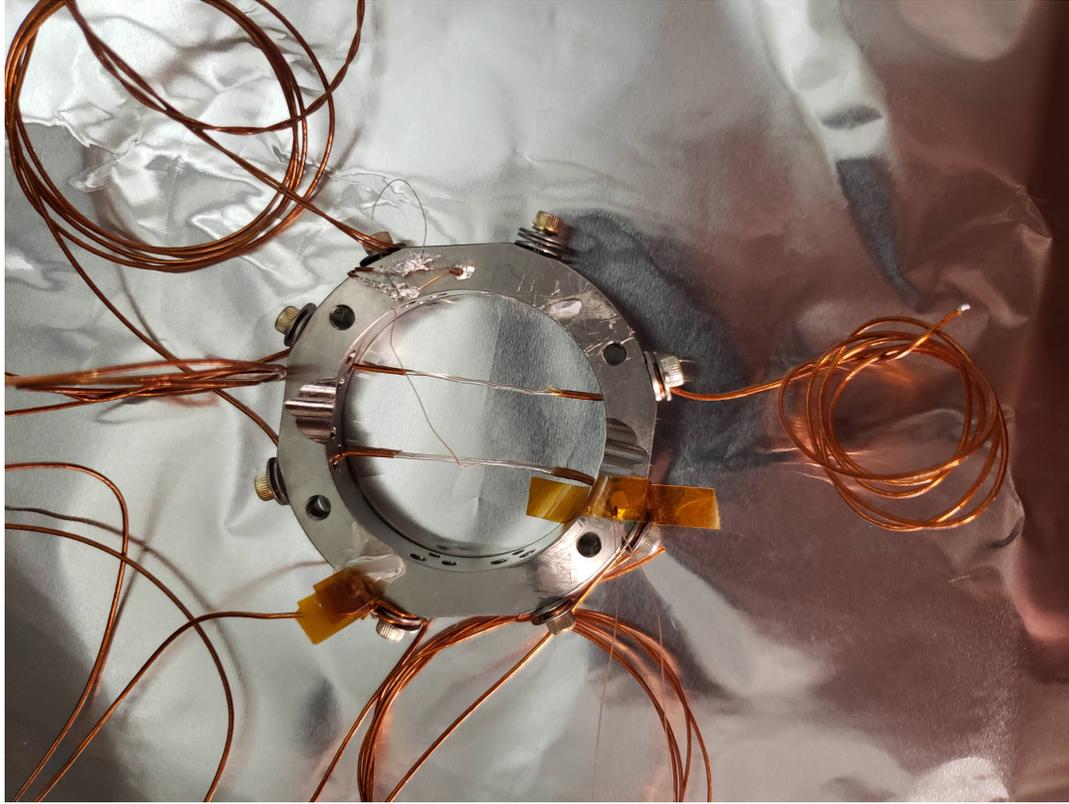


Figure 3-2: The cavity ring mount before assembly onto the full UHV cavity support structure. The two parallel wires running from left to right in the center of the ring are the RF wires and are held in place with screwed-in washers. The small rectangular component on the bottom left of the ring is one of the UHV ceramic heaters. The other two components are glass-encapsulated $10\text{k}\Omega$ thermistors. The four through-holes on the cavity ring are where the ring is connected to the remainder of the structure within a Viton spacer sandwich.

previous students, with a handful of modifications in the electronics. Special care was taken to only use vented screws to prevent virtual leaks in the future. Additionally, several heaters and thermistors along with a grounding wire were epoxied onto the metal cavity ring mount and the RF wires were pulled taught along the length of one cavity axis on the structure before it was screwed into place between the damping posts. A detailed view of the cavity ring mount immediately after adding these electronic components is shown in Fig. 3-2. After the entire mounting support structure was fully assembled as shown in Fig. 2-2, the last remaining assembly step was to epoxy the mirrors into place on their mounts along with their piezos.

We epoxied the cavity mirrors into place with the help of a small custom mirror

clamp that could be easily attached to a translation stage and a holder that held the stainless steel mirror mount during the epoxying process. The idea behind this clamp was to be able to easily hold on to and adjust the position of the mirror held above its mount to properly center it on the mount's hole. The procedure for doing this assumed a tighter mechanical tolerance than the actual machined tolerance of the clamp, however, so the centering itself was done by eye rather than by referencing off of the machined surfaces of the clamp. Even so, the custom mirror-sized clamp was helpful in holding the small, slippery mirrors while centering them—a task that could have easily been catastrophic with tweezers. The full mirror-gluing jig setup is shown in Fig. 3-3. The mirror in its clamp was first centered about a millimeter above the surface of the mirror mount surface. Next, the clamp was loosened so that the mirror fell softly onto the mount surface in its centered location before being re-clamped at that height. With the clamp tightened again, we added one or two dots of epoxy along the vertex between the mirror and its mount or piezo. We then carefully loosened the clamp and cured the mirror in place on a hotplate covered with a foil lid. Finally, we removed the clamping structure entirely and applied a third dot of epoxy to the mirror before completing one more curing round.

Both cavity piezos were glued into place on their mounts sometime around 2015, so the details of that procedure have been lost. However, the manufacturer, Noliac, has some recommendations for specific clamping forces during the epoxy process for their ring piezos that should be followed in future installations. Because each of our cavities is stabilized with only one piezo, the two mirror mounts are geometrically different, accounting for the difference in thicknesses.

3.2.2 Alignment

Once the mirrors were epoxied onto their mounts, they could easily be connected to the cavity ring mount structure. All cavity mirror mounts have two through-holes above the location of the mirror, and 4-40 screws can attach this flat part of the mirror mounts to corresponding tapped holes in the cavity ring mount. However, no alignment pins or references were designed or machined into these locations, meaning

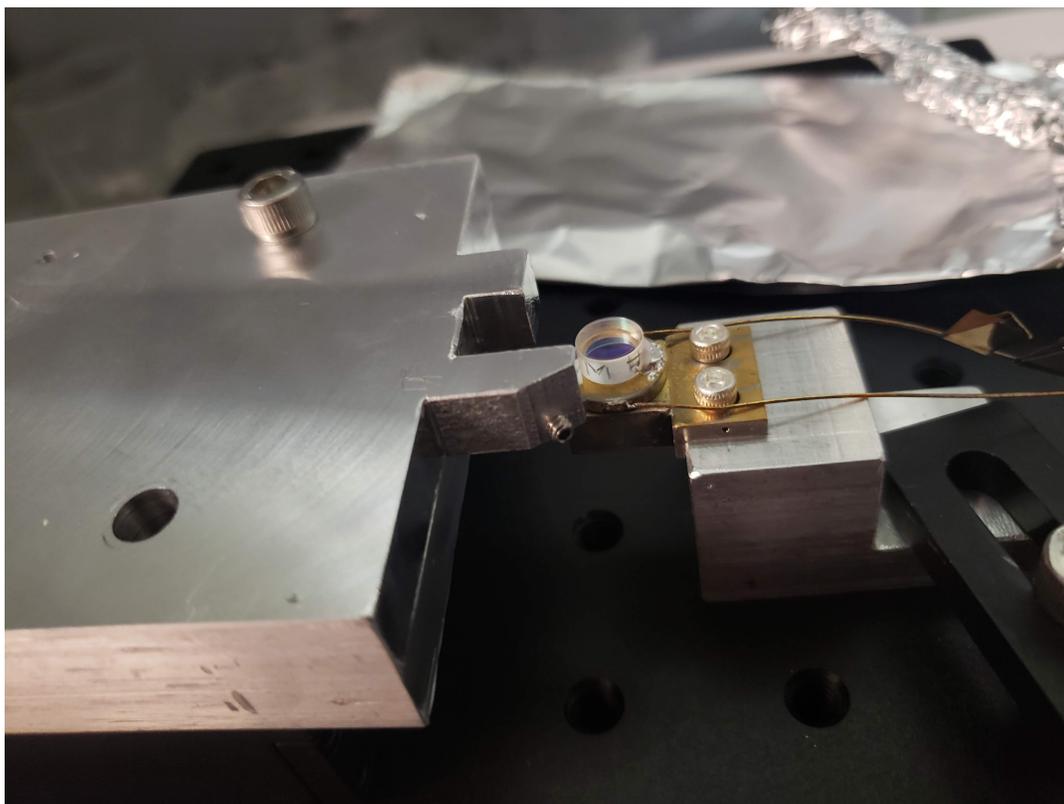


Figure 3-3: A piezo-mounted mirror immediately after being cured into place on its mount with silver epoxy. On the left of the image is a home-made clamp for our small mirrors. The clamp body is the same size as a standard Thorlabs translation stage, with through-holes for mounting directly on top of the stage. The center-left of the image shows the clamp jaw with a set screw at an angle. When tightened, this set screw holds a cavity mirror tightly into the far corner of the jaw clamp. The mirror mount is screwed into a metal holder that is clamped in place against a breadboard. The whole gluing jig holds all components in place while the mirror is centered on its mount and epoxy is applied to the edges of the mirror and surface of the mount.

that the mirror position was dependent on two factors: the accuracy of centering the mirrors onto their mounts during the epoxy process ($\sim 1\text{mm}$) and the size of the through-holes in the cavity mirror mounts relative to the screws (also $\sim 1\text{mm}$). Simply screwing the mirrors into place randomly was not sufficient to see cavity modes, and so we had to follow a more rigorous alignment procedure.

Mirror Placement

In this section, I will describe a specific alignment protocol that guaranteed an observable cavity transmission signal after one iteration. This section will also note unexpected difficulties that arose during the process along with their solutions. There are four main steps to aligning our cavities from scratch: 1) aligning the cavity output coupling mirror to an incident collimated beam, 2) aligning the mode matching lens to the cavity output coupling mirror, 3) aligning the cavity input coupling mirror to the cavity output coupling mirror, and 4) fine-tuning.

1. Cavity Output Coupling Mirror

First, we place the output coupling mirror of the cavity to set the position of the xy modes of the cavity. For optimal coupling to the TEM_{00} mode set by that mirror, we must ensure that the input light to the cavity is on-axis and converging at the correct rate to match the mode set by that curved mirror surface. To first align collimated light to the output coupling mirror, we mount the output coupling mirror roughly midway through its range and use a vacuum clean Allen wrench to tighten the two mounting screws with washers fully. We then set up a fiber launch and steering mirror that roughly directs light along where the cavity axis will be on the ring mount (see Fig. 3-5, 1.). For rough alignment, it is best to use the fiber launch to align the light to the position of the cavity input coupling mirror mount and alternately using the steering mirror to align the light to the position of the output coupling mirror mount.

Once the components are roughly aligned, we move on to fine alignment. This stage is easier if the light coming from the fiber is off-resonant with the cavity,

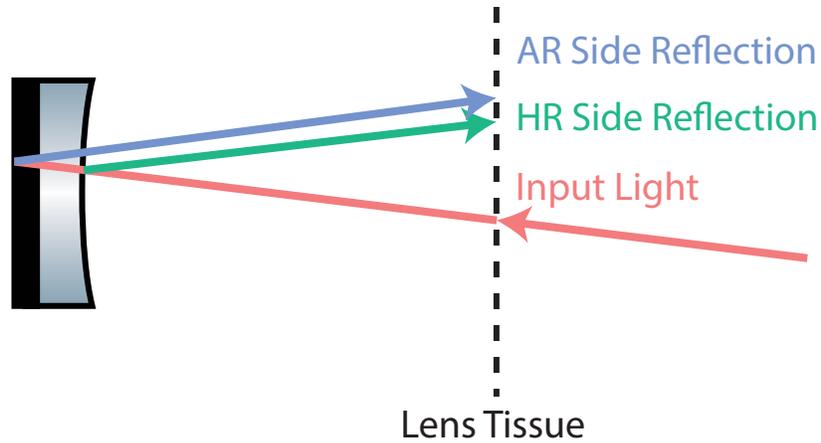


Figure 3-4: The three beams of light hitting the tissue paper while fine-aligning the cavity output mirror in step one. We only see all three spots of light on the tissue with off-resonant light, which allows for reflections off the AR coating.

because the AR coating is no longer effective. This makes it easier to see reflections off both surfaces of the mirror. Ideally, the alignment beam is also low-power to prevent dipole trapping of dust that can dirty the mirrors. We observed this process, and it significantly slowed our initial cavity alignment by blocking transmission through the aligned cavities. After installing a low-power, off-resonant beam, we place a single piece of lens tissue in the optical path between the steering mirror and the cavity ring. Three spots of light then become visible on the tissue: the input beam along with two beams of reflected light (one from the curved HR-coated surface of the output cavity mirror and one from the flat AR-coated surface of the output cavity mirror). If the reflected dots are not visible at this point, likely either the optical power transmitted through the lens tissue is too low or the rough alignment was not done well enough. After walking the three beams to overlap, the optical coupling path will be along the cavity axis. Finally, we place a beamsplitter, camera, and amplified photodetector (APD) at the output of the cavity mount (see Fig. 3-5, 2.). We center the alignment light on the camera and maximize the signal on the APD without touching the input light. We ensure that the chosen APD is faster than the ringdown time of the cavity. From this point on, we do not touch the cavity output mirror, steering mirror, fiber launch, camera, or APD.

2. Mode-matching Lens

The second step of the full alignment focuses the now aligned input beam to the waist of the cavity mode. We place a mode-matching lens on an XYZ translation mount at the proper distance from the center of the cavity (see Fig. 3-5, 3.). Just as we did while aligning the cavity output coupling mirror, we place a piece of lens tissue after the steering mirror. To ensure that the mode matching lens is perpendicular to the cavity axis and centered on the aligned input beam, we then move and tilt the lens until the reflections from the output coupling mirror are again overlapping with the incident beam. At this point we are especially careful, because multiple reflections are visible on the tissue from the two surfaces of the mode-matching lens itself, and can confuse alignment. To differentiate between the cavity mirror reflections and the mode-matching lens reflections, we can block the light between the mode matching lens and and cavity. The spots that disappear from the tissue are the reflections off the cavity output mirror, and are the correct spots to overlap for proper alignment. The reflections from the lens do not necessarily overlap with these reflections when properly aligned. From this point, we no longer touch the position of the mode-matching lens.

3. Cavity Input Coupling Mirror

Lastly, we attach and align the cavity input coupling mirror (see Fig. 3-5, 4.). The cavity input coupling mirror mounts in exactly the same way as the output coupling mirror, with two hex-head screws as shown in Fig. 2-2. The alignment is done by sliding the mirror around under finger-tight mounting screws. To determine proper alignment, one can align the cavity reflections with the input light using lens tissue as in the previous two steps. At this point in the alignment, however, we found that it is easier to align the cavity input mirror by simply monitoring the transmission camera of frequency-scanned resonant light. With our cavities, this means injecting our narrow-linewidth 780nm light into the cavity while scanning its frequency over about one free spectral range.

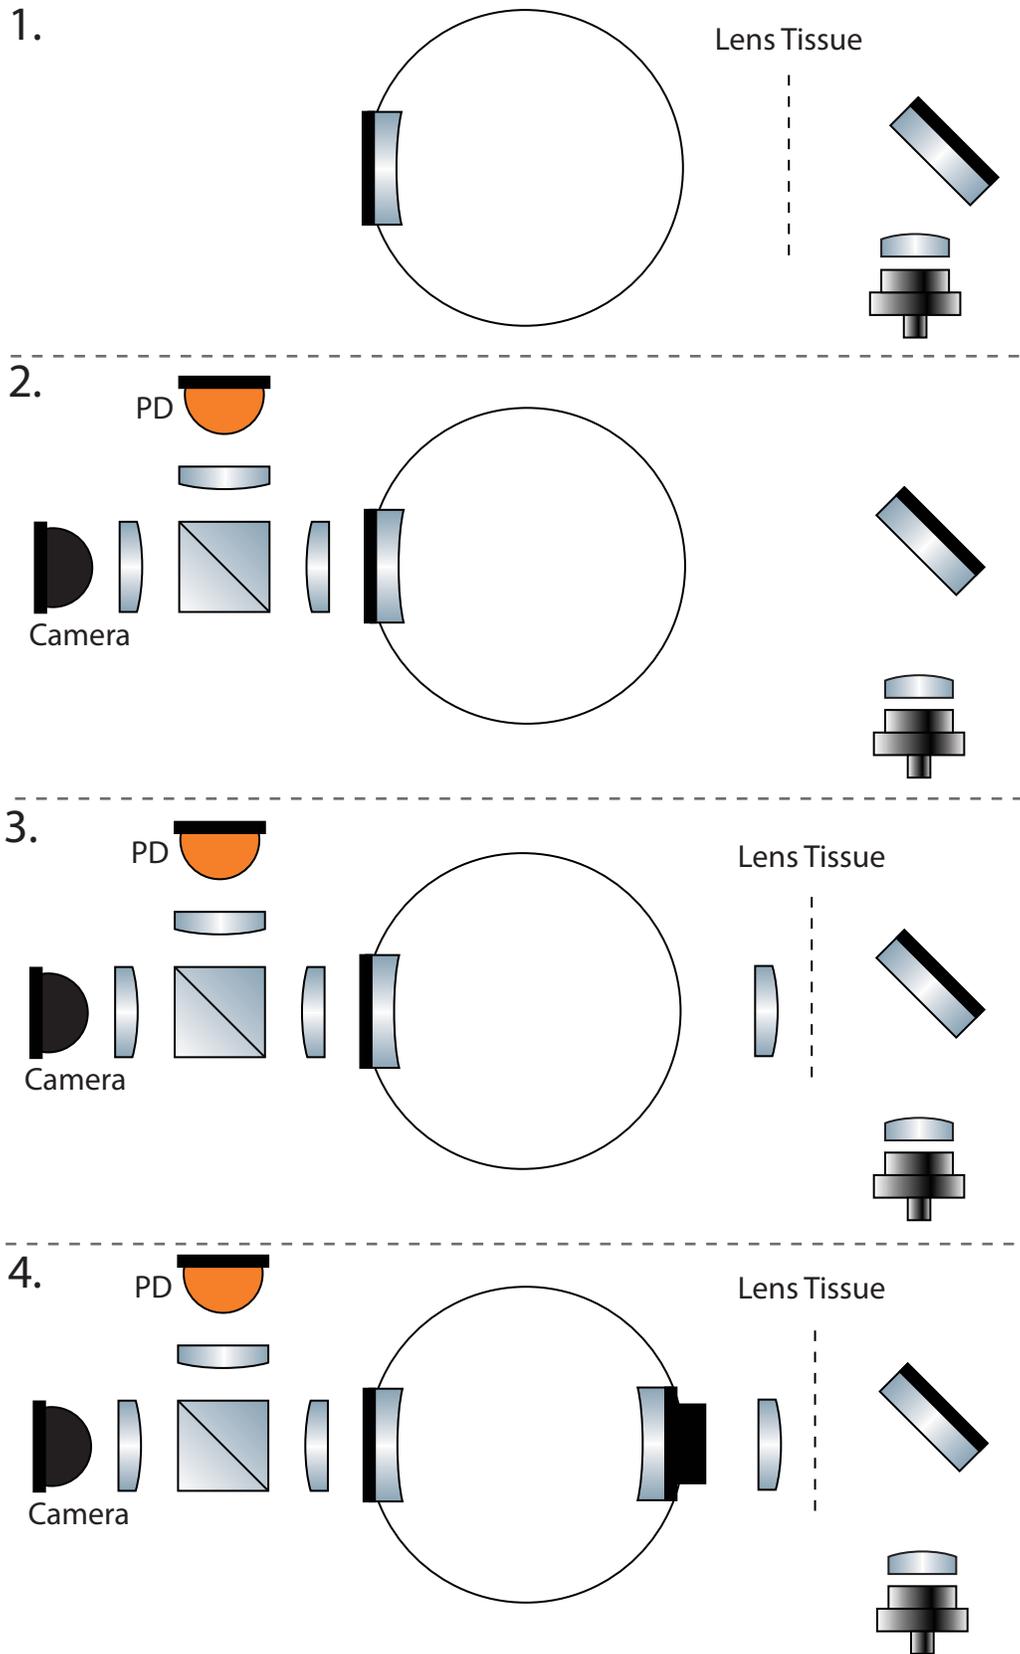


Figure 3-5: Schematic of full cavity mirror alignment, steps 1-4.

Assuming the previous two steps of the full alignment procedure were executed well and the mirrors are clean, at this point one can see various orders of cavity modes while sliding the input coupling mirror around. The mirror is in the properly aligned position when you can see the TEM_{00} mode at the center of the transmission camera. Because we aligned the camera to the center of the cavity output coupling mirror in step one, we now can be sure that the cavity input coupling mirror is properly aligned to the cavity axis. Figure 3-5, 4. shows the schematic of all optical components surrounding the cavity once these three steps are complete.

4. Fine-tuning Cavity Alignment

Once the above three steps are complete, the alignment should be well-matched to the TEM_{00} mode set by the mirrors of the cavity. If there is no transmission signal, however, there are a few common mistakes to look for. If the alignment was done with a high-power laser, it is possible that dust particles were trapped in the cavity mode and travelled to the mirrors, making them dirty. A good rule of thumb is to clean each mirror four times using a solvent wiping method with IPA to restore transmission (see App. A for details). If this does not work, try scanning the resonant laser frequency more slowly to allow more light to build up in the cavity before decaying (thus allowing a larger transmission signal). In practice, this means scanning the piezos of our 780nm Toptica laser or the cavity piezos themselves slowly (10Hz) over a fraction of a free spectral range ($\approx 100\text{m}V_{pp}$). If there is a transmission signal at the APD and the TEM_{00} mode is centered on the transmission camera, small adjustments of the fiber launch and steering mirror at this point can help increase the coupling to the TEM_{00} mode by up to $\approx 50\%$.

Several scenarios required us to do a partial alignment of the cavity instead rather than aligning from scratch as detailed above. For example, on some occasions we removed one mirror at a time to clean with a polymer solution. In this case, we cleaned the mirrors sequentially, re-aligning the cavity and optimizing to TEM_{00} after

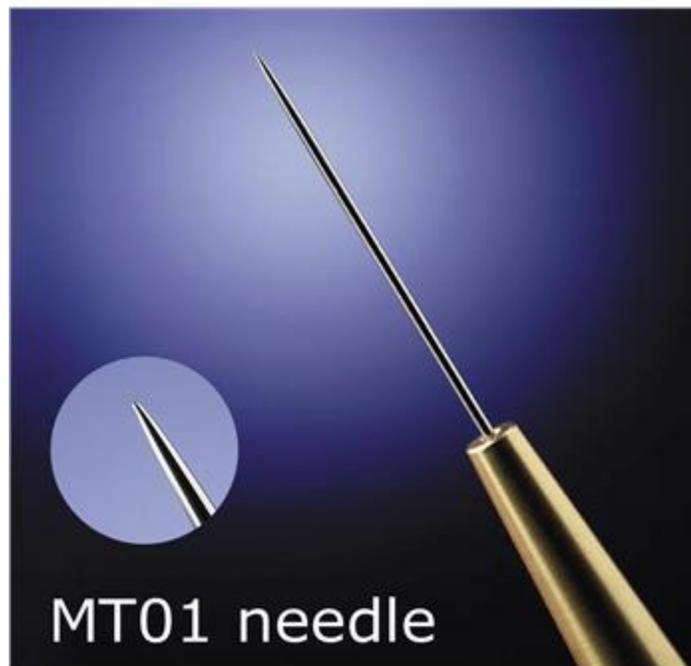


Figure 3-6: MT01 micro-needle tool from Thomas Scientific that was used to selectively block the cavity modes during the overlap procedure.

replacing each mirror. We followed the instructions in step 3 of full alignment above: mount the mirror mount with finger-tight screws, monitor the cavity transmission on the transmission camera, and tighten the mounting screws when we see TEM_{00} at the center of the camera. Whenever we remove both mirrors at once, we always need to repeat the full alignment procedure to find an initial signal and restore coupling to the TEM_{00} mode.

Overlapping Two Separate Cavity Modes

Although only one cavity was ultimately installed in the vacuum chamber for this experiment, we initially overlapped the centers of our two crossed cavities (described in Sec. 3.1) to within $8\mu\text{m}$ (see Fig. 3-8). To do so, we mapped out the location of both cavity modes with selective mode-blocking via a micro-needle (shown alone in Fig. 3-6) by mounting it into the cavity structure as shown in Fig. 3-7. By monitoring the cavity transmissions in the process of scanning the needle in Z , we were able to map out the two cavity modes and the separation between them. Quantitatively,

once we found the center of the cavity modes in XY, we scanned the needle in Z by the smallest distance that produced a noticeable change in at least one of the cavity transmissions, usually $\sim 10\mu\text{m}$. We then documented the transmitted power in terms of photodetector voltage at that needle position. After plotting and fitting this transmission data versus position in Z with a Gaussian function, we defined the center of the modes with the Gaussian mean, and the difference between the means as the distance between the cavities. After each full measurement of mode locations, we carefully adjusted mirror heights in the proper direction to decrease the distance between the cavities. This process was extremely stochastic and took a couple of weeks to complete. To make the smallest changes in mirror position at a time, we would imperceptibly loosen the mounting screws for the output mirror of the cavity that was higher in Z, and tap on the top of the mirror mount with a clean Allen wrench to lower it. We would then optimize that cavity's transmission by walking the input beam into the cavity to its new optimal position. The amount to which the cavity mirror mode would change during this procedure was extremely difficult to control, but the random walk eventually improved the alignment to the final $8\mu\text{m}$ separation. We stopped making changes at this point, because between each adjustment the relative position of the cavities often changed as much as $100\mu\text{m}$, so it was unlikely we could do better without spending significantly more time on the alignment.

The scanning translation stages and mount for the needle were similar to the apparatus used to hold a SNOM tip at the cavity modes to align the in-vacuum aspheric lenses, as discussed previously in Sec. 2.3.3. The main difference between the SNOM alignment method and the micro-needle method for mapping out the cavity modes was the angle of mounting. The micro-needle was mounted in the XY plane, bisecting the angle between the two different cavities. This allowed us to scan the needle in Z to map the heights. We roughly found the center of the two cavities in XY by translating the needle in that plane until a Z scan of the needle successfully extinguished the mode in both cavities at some point in the scan. For the SNOM alignment, the SNOM tip was directed in the Z direction, and scanned more carefully in XY to sit at the very center of the two overlapped cavity modes.

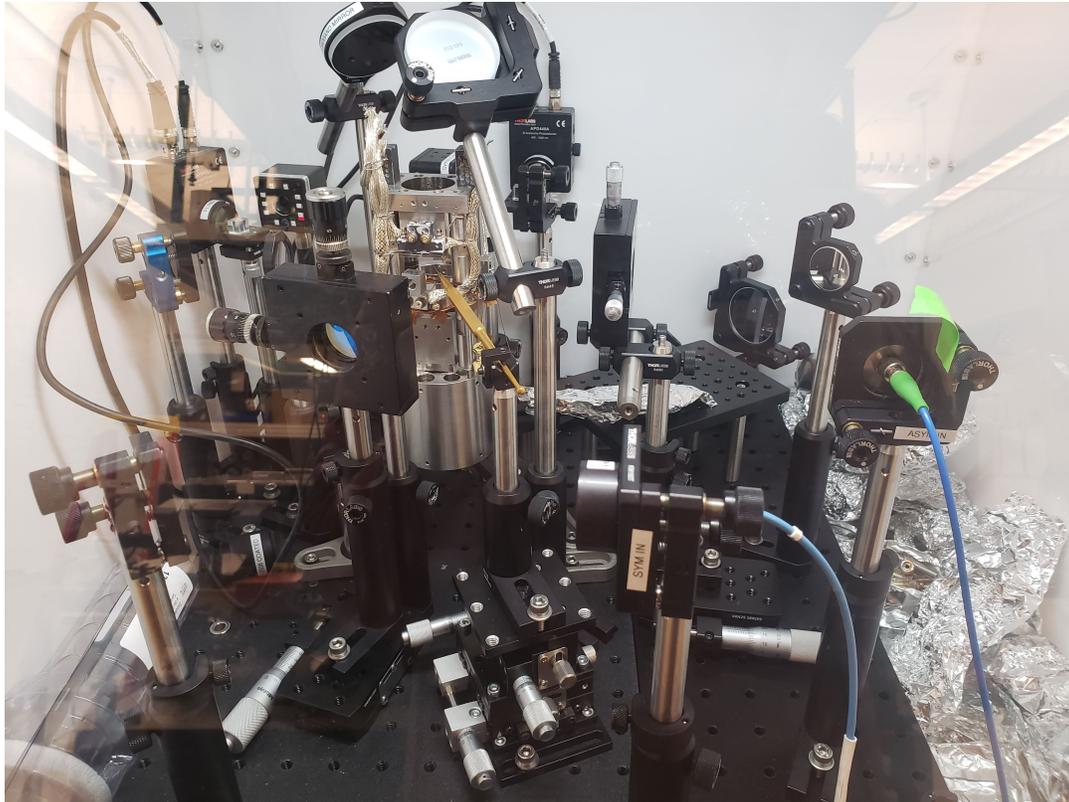


Figure 3-7: A view into the HEPA flow box during the cavity mode overlap process. The gold colored object clamped onto the three translation stages in the center of the image is the needle holder, extending into the region where the two cavity modes overlap.

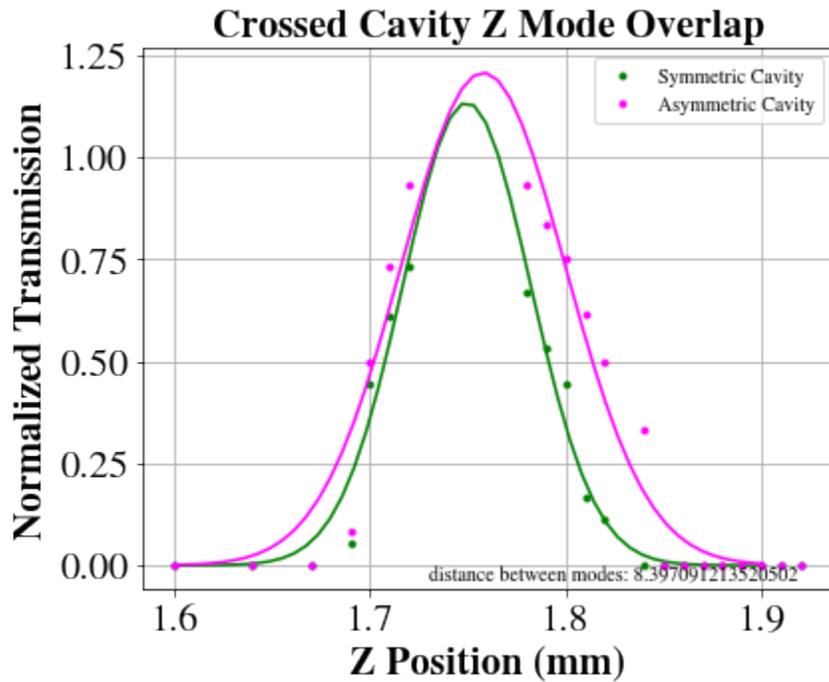


Figure 3-8: The final overlap of the two crossed cavity modes before the first vacuum installation of the cavity structure. By simultaneously scanning a needle through the mode of both cavities in the Z direction, we could reliably map out the mode positions relative to one another as we made small changes to the mirror positions to maximize overlap.

3.2.3 Characterization

Once the cavities were individually aligned and overlapped with one another, we moved on to fully understanding the optical properties. We noted several different parameters as a benchmark before installation into the vacuum chamber. We measured the ringdown time of each cavity, which is the inverse of the cavity linewidth and an indicator of the quality of the mirrors. To ensure we measured the best possible ringdown time, we cleaned the mirrors thoroughly. Ringdown time is also the final quantity we checked before installing the mirrors in vacuum to ensure cleanliness and cavity alignment. In addition to ringdown time, we measured the FSR of the cavities. As discussed in Sec. 1.4.1, ringdown time and FSR are the only parameters needed to fully define the behavior of a bare optical cavity. Before installation in vacuum, we only measured these parameters at 780nm. The properties at 1560nm were assumed to track with changes in the characteristics at 780nm, dependent only on mirror quality and cavity length. As such, the 1560nm properties were not measured until after vacuum installation. We discuss the pre-vacuum measurements of ringdown time and FSR below.

Photodetectors

To characterize our cavities during and after alignment, we used a combination of cameras and photodetectors at the cavities' transmissions. Because cameras have a much larger detection area than typical photodetectors, we used these as the initial indicators of alignment. We were able to see diffuse, high-order modes illuminate a camera sensor long before we had a detectable signal on a photodetector. After walking the cavity alignment and starting to see lower order modes on the camera, we switched to the photodetectors.

We considered two detector properties when choosing the proper photodetector for each cavity type: speed and sensitivity. Because our two cavities had significantly different properties, summarized in Tab. 3.2, we required a different photodetector for each. The high-finesse, symmetric cavity had an expected ring down time of

$12\mu\text{s}$ and the low-finesse, asymmetric cavity had an expected ringdown time of $4.5\mu\text{s}$. To determine if a photodetector was fast enough for these signals, we ensured the rise time of the detector was far below the expected ringdown time of the cavity. Thorlabs specifies the definition of rise time as $t_{rise} = \frac{0.3497}{BW}$ where BW is the detector bandwidth. Additionally, because of the lower losses and higher reflectivity in the symmetric cavity, the expected power at the transmission was lower than the asymmetric cavity's. Therefore the asymmetric cavity required a fast, lower sensitivity photodetector, and for this we used a Thorlabs APD410A (rise time: 35ns and maximum sensitivity: 500kV/A). The symmetric cavity required a slow, higher sensitivity photodetector, and for this we used a Thorlabs APD440A (rise time: $3.5\mu\text{s}$ and sensitivity: $5 \times 10^4\text{kV/A}$). Both are variable gain detectors, allowing us to max out the sensitivity during alignment and then turn it down during characterization. This helped to reduce the amplification noise for characterization and eventual locking. After switching to the single, OC-OC cavity, we continued using the Thorlabs APD410A photodetector, because the ringdown time was comparable to that of the original asymmetric cavity.

Only after fully installing the final cavity structure did we align and characterize the ringdown time at 1560nm. For this path, we used a Thorlabs PDA05CF2, which is an InGaAs detector with a rise time of 2.3ns and a maximum sensitivity of 5.25kV/W. This was sufficient for our expected ringdown time of 424ns, given previously measured values for the transmission (345ppm) and losses (0ppm) per output-coupling mirror. This Thorlabs PDA05CF2 photodetector was used both for initial alignment of the 1560nm light at transmission and for the PDH locking scheme via the reflection from the cavity. In addition to this photodetector, the 1560nm cavity path also has an InGaAs camera, Edmund Optics 56-567, which can be used to aid in alignment. However, because the sensitivity of InGaAs cameras is lower than that of Si detectors used at 780nm ($\sim 1\text{A/W}$ versus $\sim 50\text{A/W}$), this camera is difficult to align well enough to be useful.

Ringdown Measurements

Once we aligned our cavities and had a TEM_{00} signal both visible on our transmission cameras and aligned to our photodiodes, we characterized their quality via ringdown measurements. Through each step of mirror alignment and assembly, vacuum installation, and finally baking, this was the quantity that told us whether the cavity mirrors were stable or if they had degraded in some way. After the first installation of the two-cavity setup, we found the ringdown times decreased, which led us to the failure analysis detailed in Ch. 4. After the second installation of the single OC-OC cavity, we found the ringdown time to be stable and thus continued with bringing the experiment up and running: locking the cavity (Sec. 3.4) and aligning trapping and addressing optics (Ch. 5).

The ringdown measurement setup evolved over different periods of design, construction, and characterization. During the initial alignment and characterization stages, we measured the ring down times of the cavities in a fairly manual way. A diagram of the connections for the initial measurements at this stage are laid out in Fig. 3-9. We pass the resonant, scanning laser light through an EOM driven at about $0.5 \times \nu_{\text{FSR}}$ and at 18dBm before injecting it into the cavity. We connect the output of the cavity transmission APD to both an analog trigger box and an oscilloscope. We adjust the trigger level on the trigger box to fire when the EOM sideband is transmitted through the cavity in TEM_{00} , turning off the EOM RF source; in our case, a WindFreak RF Synthesizer. This trigger is also fed to an oscilloscope displaying the photodetector's output, allowing us to find and capture the fast exponential decay of light leaving the cavity. In this configuration, the correct point in frequency where the triggered TEM_{00} sideband is transmitted is not when the largest signal appears. Because the EOM can only finitely drive power into the sidebands, the carrier of multiple TEM_{mn} modes remains larger than the triggered TEM_{00} sideband. This fact sometimes causes confusion over which transmitted signal was the correct one to look at. However, it was easy to throw out improperly triggered traces after noticing their distinctly non-exponential decay. Instead of decaying exponentially near the

expected ringdown time of the cavity, the non-triggered bands slowly oscillate away from resonance instead of decaying exponentially. Once captured, we saved a .csv file of the properly triggered waveform and fit the exponential decay in Python to extract the time constant. The time constant of the exponential decay is the ringdown time of the cavity.

During the diagnosis of the initial cavity mirror degradation (see Ch. 4), we realized that performing ringdown measurements are simpler with some small changes to the scanning and triggering scheme described above. First, scanning the laser frequency is not the most dependable way to capture a good exponential ring down curve. Often, lasers mode hop or have other instabilities during the large frequency scan required to see multiple cavity modes, $\sim \nu_{FSR}$. This makes it difficult to accurately measure mode spacing. Ideally, we can instead hold the probe laser at a constant frequency and scan the cavity length using the mirror piezos instead. Second, using an EOM to trigger the turn-off of the cavity probe light is not ideal because of confusion between carrier and sidebands, as described above. A more convenient way to trigger the probe light on and off is to pass the probe light through a double-pass AOM (DPAOM) that is gated with an RF switch triggered by a high voltage at the transmission photodetector. Electronics for this more convenient ringdown measurement scheme are essentially the same as the other. The only added complexity comes from needing to align a DPAOM. Because we need the turn-off time of the DPAOM to be much faster than the cavity's ringdown time, the beam needs to be focused tightly and aligned as close to the RF transducer as possible. The only other addition was to use a commercial piezo driver to begin scanning our piezos across their entire 200V range. Finally, in the future, ringdown measurements of our cavity would be much more convenient with an automated measurement system that directly siphons captured exponential decays to our data analysis computer. Using our Labscript control system (discussed in Zachary Vendeiro's thesis), we could easily implement a routine that measures the cavity ringdown on a regular basis. This will help track any changes in the cavity's quality over time and even monitor the background pressure of rubidium.

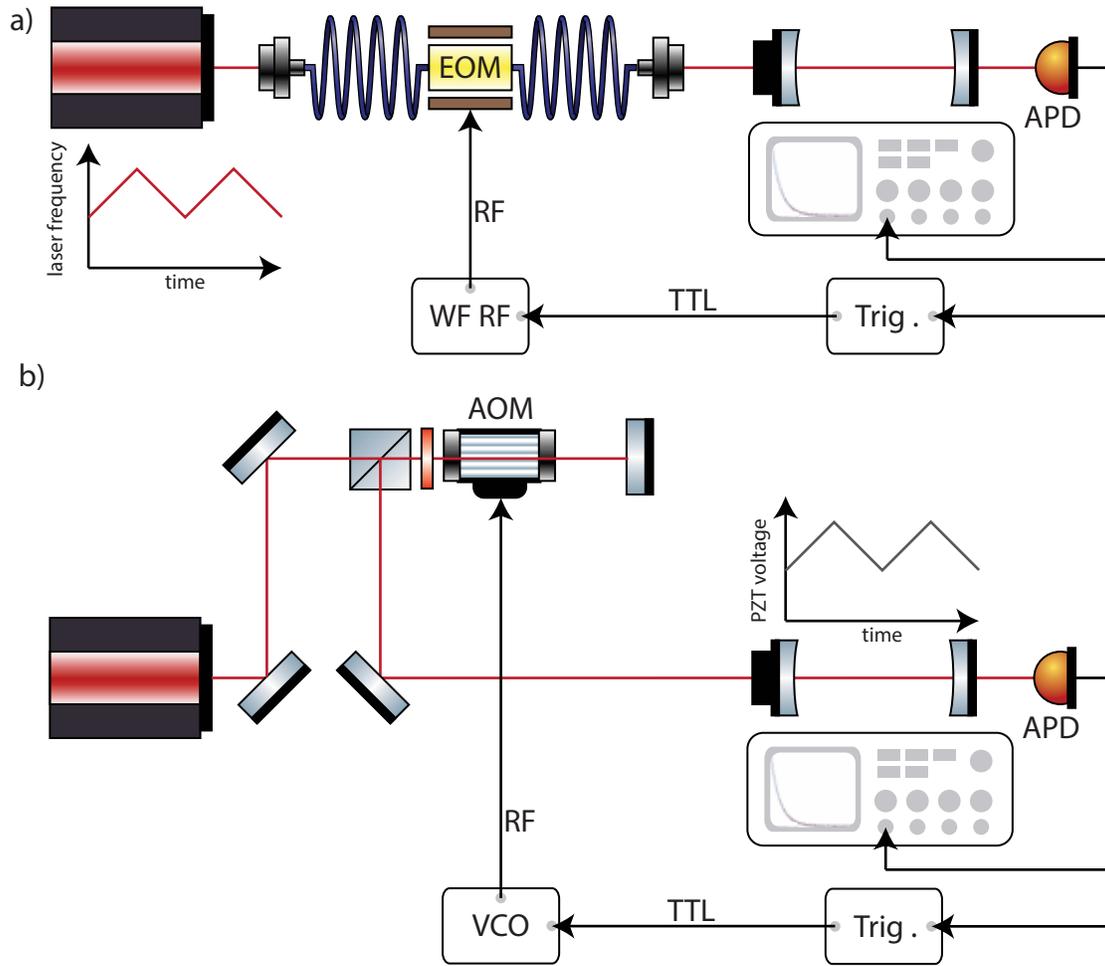


Figure 3-9: Controls and electro-optical components used for cavity ringdown measurements. The initial ringdown method using a triggered EOM while scanning laser frequency (a). Triggering an EOM was a less reliable method to measure the ringdown time than the DPAOM method with cavity mirror piezo scanning (b).

Free Spectral Range

The length of the cavities were originally designed to correspond to a specific FSR, one-half of the groundstate hyperfine splitting in rubidium-87, 3.42GHz. The goal of this frequency choice was to allow for two different rubidium-87 D_2 transitions to be on-resonance with the cavity at a given time. To confirm that the FSR of the assembled and aligned cavity matched the designed target of 3.42GHz (half the frequency of the ground state hyperfine splitting in rubidium), we first send the cavity input light through an electro-optic modulator (EOM). This creates two sidebands on the carrier frequency of the light with a frequency difference that matches the EOM RF input frequency. After adding these sidebands to the input laser light, scanning the overall laser frequency across cavity resonance reveals multiple TEM_{00} transmission peaks. By adjusting the EOM RF input frequency, we move the two sideband transmission peaks with respect to the carrier transmission peak. The points at which all three peaks are perfectly overlapped are where the EOM RF input frequency is an integer multiple of the cavity's free-spectral range, ν_{FSR} . By increasing the integer multiple, we can effectively "split the line" of this measurement making it more accurate by a factor of \sqrt{N} . These measurements confirmed our expected $\nu_{FSR} = \frac{c}{2L} = 3.42\text{GHz}$ from the lengths of our cavities.

3.3 Vacuum Installation

The full vacuum installation of the cavity apparatus involved 1) physically inserting it into the chamber and then 2) pumping and baking the full chamber to reach UHV. We did both of these steps twice, once with the originally designed crossed-cavity apparatus and once with the final OC-OC cavity. Many things went wrong during the first installation, so we will both describe the general installation procedure and the specifics of each of these installations to help others avoid our mistakes.

General Procedure

To keep close track of the cavity quality during the risky installation, we measured our cavities' ringdown times at all possible points: immediately before removing the apparatus from the HEPA flowbox, 2) after inserting the apparatus into the chamber, 3) after pumping the chamber before baking, and 4) after baking. Before deciding to remove the apparatus from the HEPA flowbox used for testing and characterization, we cleaned the mirrors until we achieved the longest possible ringdown time and used this number as our reference for the remainder of the process.

Dismounting the cavity apparatus from its testing and characterization home in the HEPA flowbox and inserting it into the UHV chamber was an intense process. We carefully choreographed each step and used four sets of hands during both installations. There were several concerns about what could go wrong in the process. As such, we worked to understand and minimize the risks while planning the process. For example, the HEPA flowbox was across the lab from the vacuum chamber, so walking the apparatus across the room opened up the possibility of 1) the mirrors getting dirty from dust in the air and 2) the apparatus being dropped through a collision or tripping. Both of these scenarios had the potential to drastically reduce the finesse of the cavity reversibly or otherwise. Before transporting the cavity apparatus, we rehearsed exactly where each person would stand and what motions they would take. In the end, the transport of the cavity structure across the room was uneventful in both installations.

After one person successfully carried the apparatus to the chamber, we adjusted grips and had three people lower the cavity apparatus through the upper chamber opening. Each person had a distinct job: 1) hold the full weight of the apparatus ($\sim 6\text{kg}$), 2) keep the apparatus parallel to the opening of the chamber, and 3) ensure that no wires got caught on the chamber opening. Because the cavity apparatus was designed to fill the full diameter of the upper flange tube, it was practically impossible to install the apparatus without some of the electrical components scraping the flange tube wall in the process. As such, during both installations, even though much care

was taken to safeguard the in-vacuum electrical connections with strain relief, some electrical connections were broken. These small tolerances are noted in the CAD image in Fig. 3-10, looking down on the apparatus in z through the upper flange of the main chamber.

In terms of the mechanical installation, the last step was to properly torque the top flange that holds the whole apparatus in place. Because the top flange of the cavity apparatus is entirely custom, if the knife edge were damaged in any way, it would have easily delayed the installation by months, especially because it was the middle of the Covid-19 pandemic. The flange tightening process was further complicated by the fact that the upper flange of our vacuum chamber body is rotatable, meaning that it does not stay in place on its own. Because this flange can fall down around the upper nipple of the chamber it was a constant struggle to keep it in the proper orientation. As such, during the lowering and flange torquing steps, one person was fully responsible for holding the upper flange into its highest position at the proper angle while ensuring that the copper gasket remained seated on the chamber's knife edge. This copper gasket was placed before we began transporting the cavity apparatus from the HEPA flowbox to the chamber. The angle at which the apparatus was finally installed in the XY plane was critical, because it determined at what angle the cavity axes would sit with respect to the chamber viewports. To ensure the correct angle, we continued to have one person hold up the rotatable flange while a second person held the weight of the almost-fully-lowered apparatus and a third person looked through the cavity viewports to center the cavity mirror mounts on the viewports. If this system were to be designed from scratch, a non-rotatable viewport would greatly simplify installation. Finally, we fully lowered the cavity apparatus onto the knife edge (still holding up the rotatable flange) and began to place and screw in the flange bolts in a star pattern, until all bolts reached a torque of 15ft-lbs, the recommended torque for a 4.5" CF flange. At this point, we checked the ringdown time of the cavities before moving on. During both installations, the ringdown times were acceptable at this point.

Once the apparatus was in place and secured inside of the vacuum chamber, we could begin the pump-down process. We followed a generally accepted procedure of

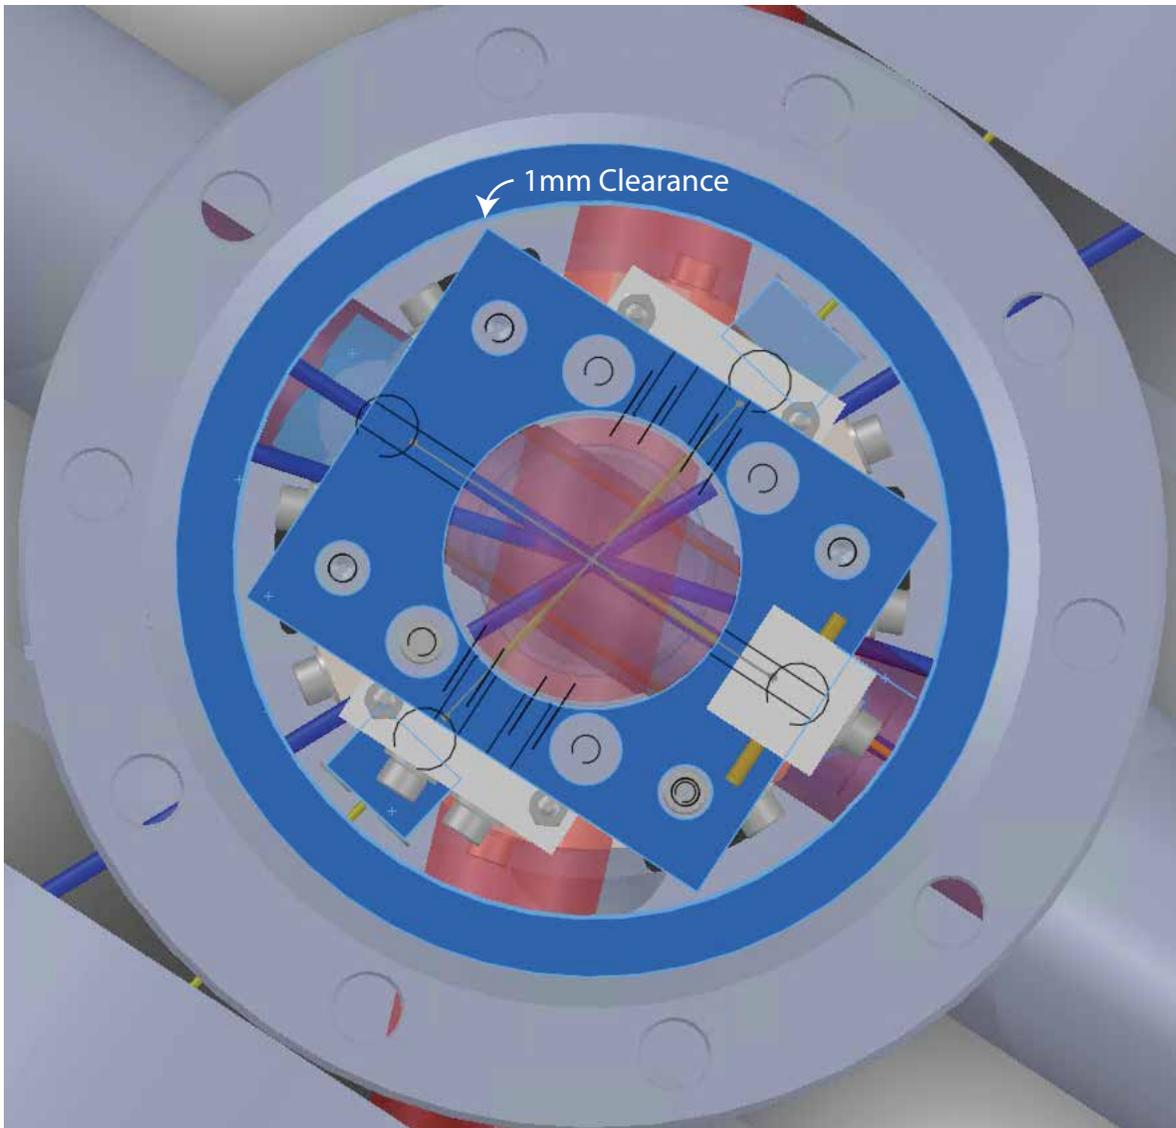


Figure 3-10: View of the cavity apparatus installed in the experimental chamber from below. The highlighted rectangle is the part of the apparatus with the largest diameter, only 2mm smaller than the diameter of the chamber tube inner diameter. If the cavity apparatus strayed from parallel to the chamber tube at any point during the installation, this portion of the apparatus and electronic connections hit the sides of the chamber.

1) turbo pumping, 2) rotary pumping, 3) baking with fiberglass heating tapes and foil covering, 4) ion pumping, 5) closing valve to rotary pump, 6) cooling, and 7) titanium sublimation (TiSub) pumping. The installation was only complete if the ringdown time of the cavities post-bake and pump was consistent with their ringdown times pre-installation. The first installation was not a success in terms of this metric.

First Installation

Because no one in Rubidium Lab had carried out a significant vacuum upgrade before the installation of our cavities, we spent considerable time preparing for the vacuum installation. Even so, we still had some gaps in our knowledge that led to mistakes summarized in this section. The first misstep in the first cavity installation was not entirely understanding how rotatable flanges operate. In the process of lowering the cavity apparatus into the chamber, we placed the full weight of its knife edge on the copper gasket and chamber flange knife edge before it was situated at the correct angle for cavity-viewport alignment. Because of this, we gouged out some of the copper on the gasket and had to pull the full apparatus out of the chamber, replace the gasket, and restart. On the second try, we successfully mounted the structure at the correct angle without damaging the copper gasket. We did not cause any damage to the upper flange knife edges during this process, so this error was not mission critical. Before we began the pumping process, however, we checked the impedance of all the in-vacuum electrical connections through the upper electrical feedthrough and found several issues. A few of the electrical connections were fully broken during the installation and almost all of the other connections demonstrated a significantly different impedance than they had prior to installation, including the mirror piezos. The piezo impedance was a major concern, because it could significantly reduce their scanning range if their electrodes are shorted. It was not clear why the electrical impedances changed until later, after we removed the apparatus from the vacuum. Eventually, however, we realized that all the in-vacuum component leads shorted through the Aquadag graphite layer on the in-vacuum DSub connectors. In effect, this put all of the in-vacuum components in parallel with one another to some degree. At

the time we pushed forward, however, having enough electrical components connected that we could have made the system work.

After confirming that the post-install cavity ringdowns were similar to the pre-install values, we moved on to pumping down the chamber. The full ion gauge pressure recording during this period is displayed in Fig. 3-11. At this point we unknowingly did not fully open the gate valve connecting the pumps to the chamber. We assumed the gate valve was similar to the Rb vapor cell valve, which can be opened with a small number of turns. However, the gate valve to our chamber actually must be loosened by many (~ 50) turns to be fully open, enough turns that the valve actually sticks out from the chamber by an inch or two. We did not realize this mistake until later in the bake when our ion pump stopped working due to shorts across its electrodes. The partially-opened valve likely contributed to the shorts in our ion pump by restricting flow out of the chamber to the rotary pump. This may have led to increased build up in the ion pump itself. Additionally, at this point we realized that we forgot to add the second layer of Faraday shielding over the cavity piezos, which was critical for future Rydberg atom experiments.

The failure of the ion pump was the most challenging and stressful hurdle in the first installation. Before we took any drastic actions when we saw that the ion pump was shorted, we made sure to check simple parts of the pumping setup that could have malfunctioned. We measured the resistance of the high voltage cable connecting the ion pump and its controller, and found no shorts. We also ruled out the ion pump controller as the issue by swapping it out for another controller, which gave the same error—indicating a short in the pump itself. After these sanity checks, we cooled the vacuum chamber back to room temperature. This allowed us to pursue two common ion pump short fixes. We first "hi-potted" the ion pump with a high-voltage, high-current source to burn shorts between the plates. We had difficulty finding a power supply with both high voltage and high current, but we ended up using a 12kV supply. When hi-potting alone did not immediately succeed, we then hit the ion pump with the rubber mallet to help breakdown internal debris while simultaneously hi-potting. The combination of hi-potting and hitting the ion pump was successful in burning off

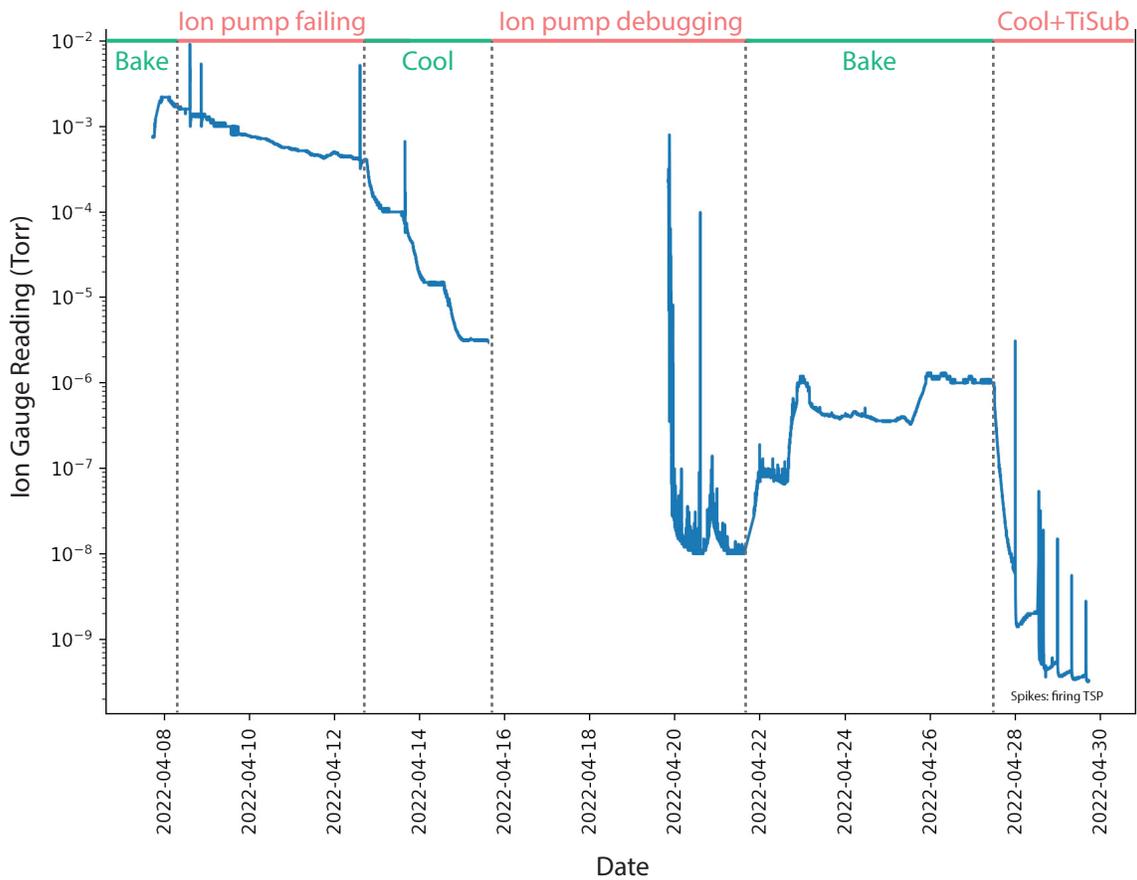


Figure 3-11: The ion gauge pressure reading in the experimental chamber as a function of time during the first installation of the cavity structure. We baked at a high temperature of 180°C . The spikes in current are the points at which we tried (unsuccessfully) to turn on and ramp up the voltage across the ion pump. After several rounds of hi-potting and hitting the ion pump with a rubber mallet, we were able to continue pumping down to UHV.

the shorts from the ion pump electrodes. After several rounds of these fixes and once the ion pump could turn on reliably without giving a short error on the controller, we ramped up the temperature of the entire system back to 180°C. From here, we continued the bake for another few days. Finally, we cooled the chamber down and ran the TiSub without incident, reaching UHV pressures.

After unwrapping the chamber and re-aligning the cavity optics, however, it was immediately obvious that our cavities did not survive the bake in the way we had hoped. Both cavities were misaligned with significant astigmatism, which we could see qualitatively in the asymmetry of the shapes of transmitted modes. For example, the TEM₀₀ mode was an elongated elliptical Gaussian instead of a circular Gaussian. Further, the ringdown times were significantly faster than the response times of the photodetectors, indicating that the quality of the cavities had degraded severely. After confirming this degradation for both cavities, we concluded that the only way to move forward was to remove the cavity structure from the vacuum chamber, recover the cavities, fix the first installation mistakes, and attempt a second installation. Initially, we thought the reduction in cavity ringdown times could be fully explained by misalignment during the ion pump hammering. However, after removing the cavities from vacuum and realigning the mirrors, the ringdown times were still several times faster than before the bake. Cleaning the mirrors with progressively more aggressive solvents also did not recover the original ringdown times, and in the end we were not able to recover the HR mirrors. A detailed look into the failure analysis of these mirrors can be found in Ch. 4. During this failure analysis, we realized that the one output coupling (OC) mirror that went through the bake had minimal, if any, degradation. As such, we made the difficult decision to move forward with installing a cavity made from two of the spare OC mirrors rather than delaying the installation further by ordering new mirrors. At that point in time, the lead times to get new HR mirrors was months to a year.

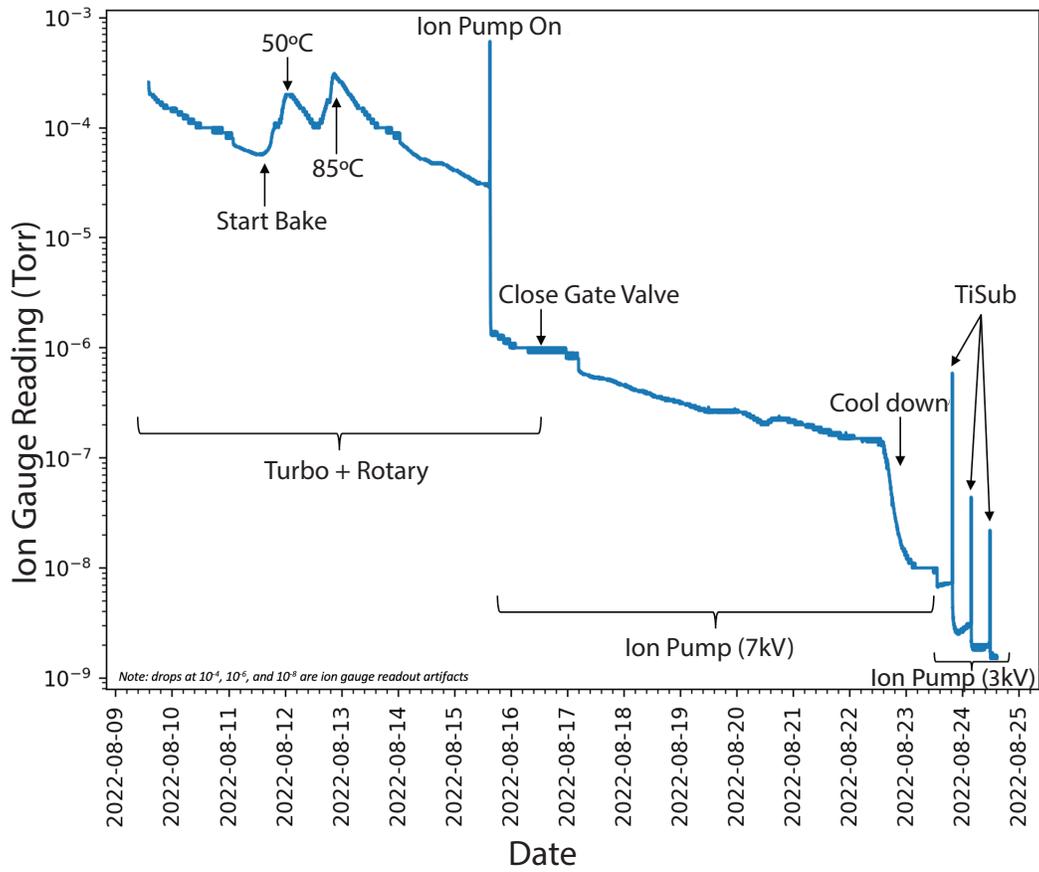


Figure 3-12: The ion gauge pressure reading in the experimental chamber as a function of time during the second and final cavity installation.

Step	Date	ringdown Time [μ s]
HEPA Box	08/08/2022	2.1
Pre-Bake	08/09/2022	2.4
Post-Bake	08/23/2022	2.1
Data Taking	01/31/2023	1.7

Table 3.3: Ringdown times of the installed OC-OC cavity at various stages in installation. The ringdown time has fluctuated over time, potentially due to dust falling on and off the mirrors. It is also possible that in the time since installation these mirrors have experienced vacuum degradation from oxygen vacancy formation.

Second Installation

Before we moved on to installing a new cavity in the vacuum chamber, we had to spend considerable time carrying out failure analysis. We tabulated all mistakes made, repaired electrical connections, and re-aligned and characterized the new cavity. In terms of electronics, this meant soldering together wires that had broken apart, scraping off all Aquadag graphite from the DSub connectors, removing shorts in the mirror piezos by scraping away more of the gold coating, and replacing all the heaters and thermistors. We determined that we could not recover the baked HR mirrors as discussed in Ch. 4 and so focused on gluing the spare OC mirrors onto mounts for alignment. Once the OC-OC cavity was aligned and characterized, we also remembered to install the second layer of Faraday shielding over the mirror piezo mount, putting us in good shape for our future Rydberg experiments.

By the time we made it to the second vacuum installation, we were able to avoid nearly all of the mistakes from the first installation, with the one exception of breaking electrical connections. While sliding the cavity apparatus back into the chamber, one of the new thermistor’s electrical connections fully broke. Luckily, we added two new thermistors for redundancy during the repair, leaving us with a functioning heater and thermistor. Otherwise, all of the electrical connections remained intact. During the pump and bake, we fully opened the gate valve and kept the chamber below 100°C to reduce the risk of mirror degradation. The ion gauge pressure readings during the second vacuum installation are plotted in Fig. 3-12. We did not experience issues with the ion pump shorting during this installation.

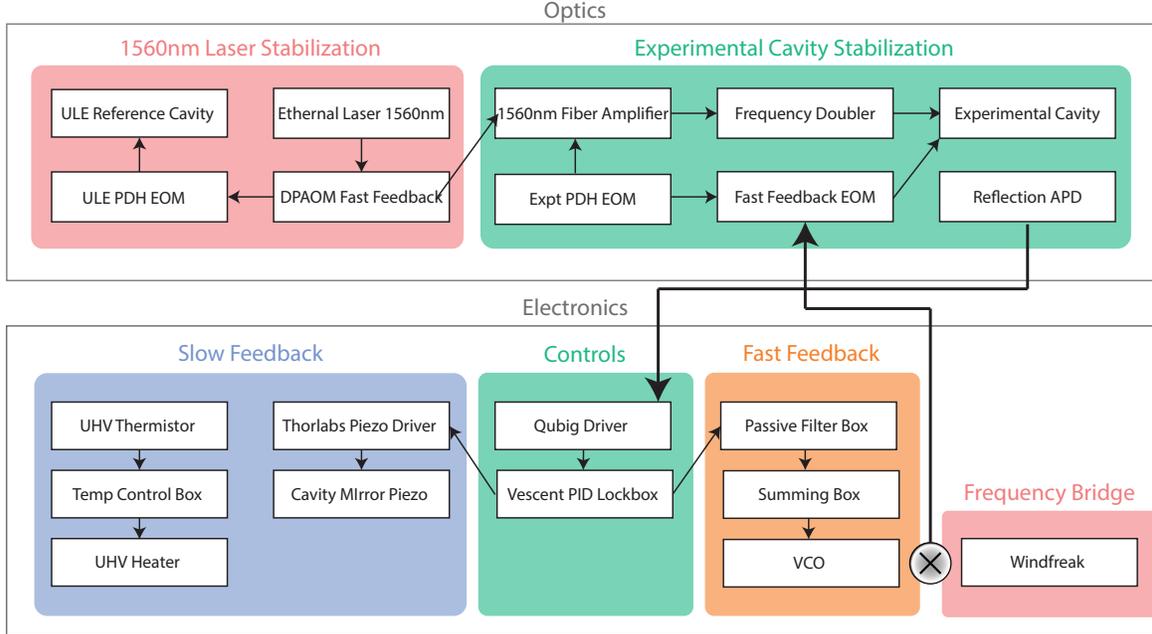


Figure 3-13: Block diagram of all major components of the experimental cavity stabilization scheme including passive, active, fast, and slow feedback.

After successfully installing the cavity apparatus in vacuum and achieving UHV pressures $\sim 10^{-9}$ Torr, we aligned the mode matching optics through the cavity chamber view ports. We measured the ringdown time of the cavity, and it was roughly the same, if not slightly longer, than the measurement we took immediately before installing the cavity in the vacuum chamber. Table 3.3 summarizes OC-OC experimental cavity ring down measurements as a function of date. There have been fluctuations in the ringdown times of this cavity over time. Specifically, post-installation the finesse of the cavity seems to be slowly decreasing. We did not realize until after installation, but these mirrors are capped in a Ta_2O_5 layer, just as the HR mirrors were. This means that the drop in finesse we are witnessing could be a result of oxygen vacancy formation over time in the cap layer (see Ch. 4 for more details). However, at a ring down time of $1.7\mu\text{s}$, our single atom cooperativity is still $\eta = 2.1$ and squarely within the strong-coupling regime. Thus, this installation was successful, and we moved on to stabilizing the length of the cavity and realizing atom-cavity coupling.

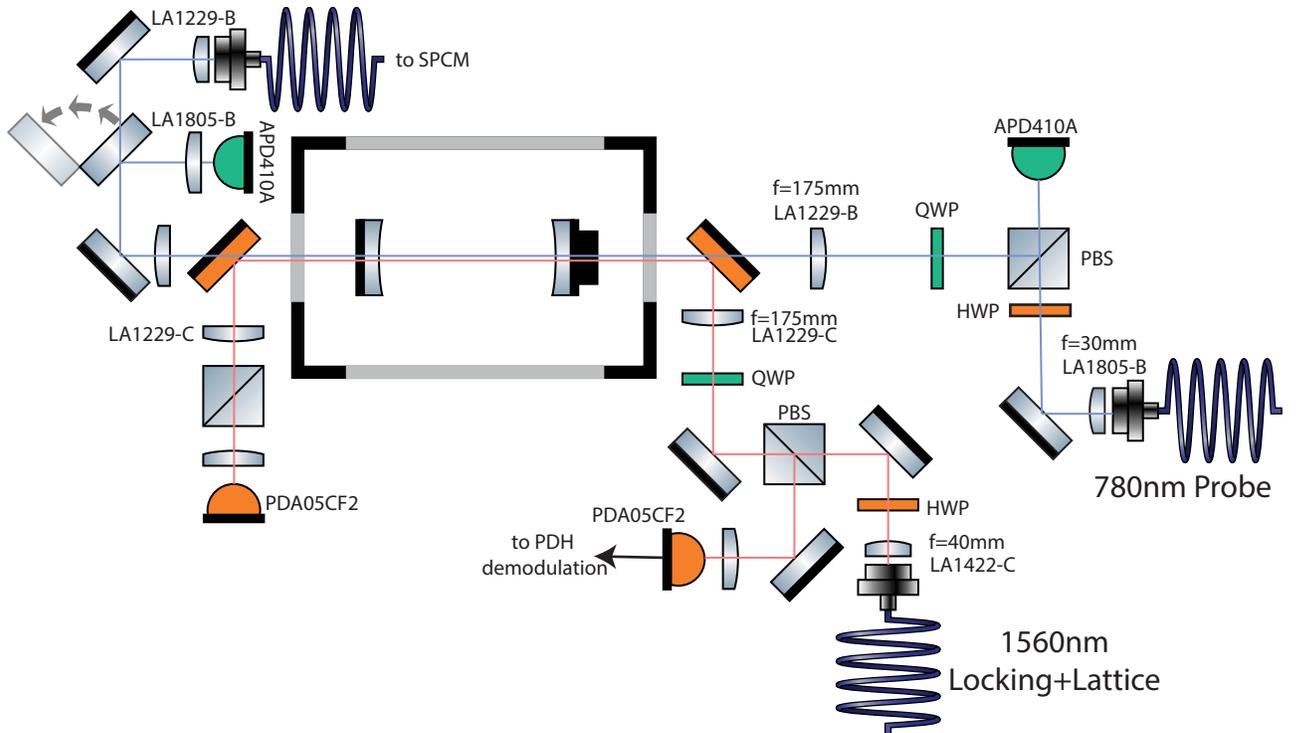


Figure 3-14: All cavity optics paths: 1560nm locking and lattice (pink beam) and the 780nm probe (purple beam). The beams are offset for visual clarity—in reality they overlap.

3.4 Stabilization

To stabilize our experimental cavity, we use the standard Pound-Drever-Hall (PDH) locking scheme to extract an intensity-noise independent error signal. The mechanism and math behind the PDH locking scheme is described in detail in many sources [15, 31]. Conceptually, however, this locking scheme operates by creating an optical error signal that is asymmetric about cavity resonance at the reflection of an optical cavity. Using the reflected signal decouples frequency fluctuations and intensity fluctuations in the probe light. Creating an asymmetric error signal about resonance provides information about which direction the probe light frequency needs to change with respect to the cavity frequency to reach resonance. PDH achieves these two important capabilities by taking advantage of constructively and destructively interfering frequency sidebands in the cavity-reflected light. This signal has a steep slope set by the cavity linewidth centered on the cavity resonance.

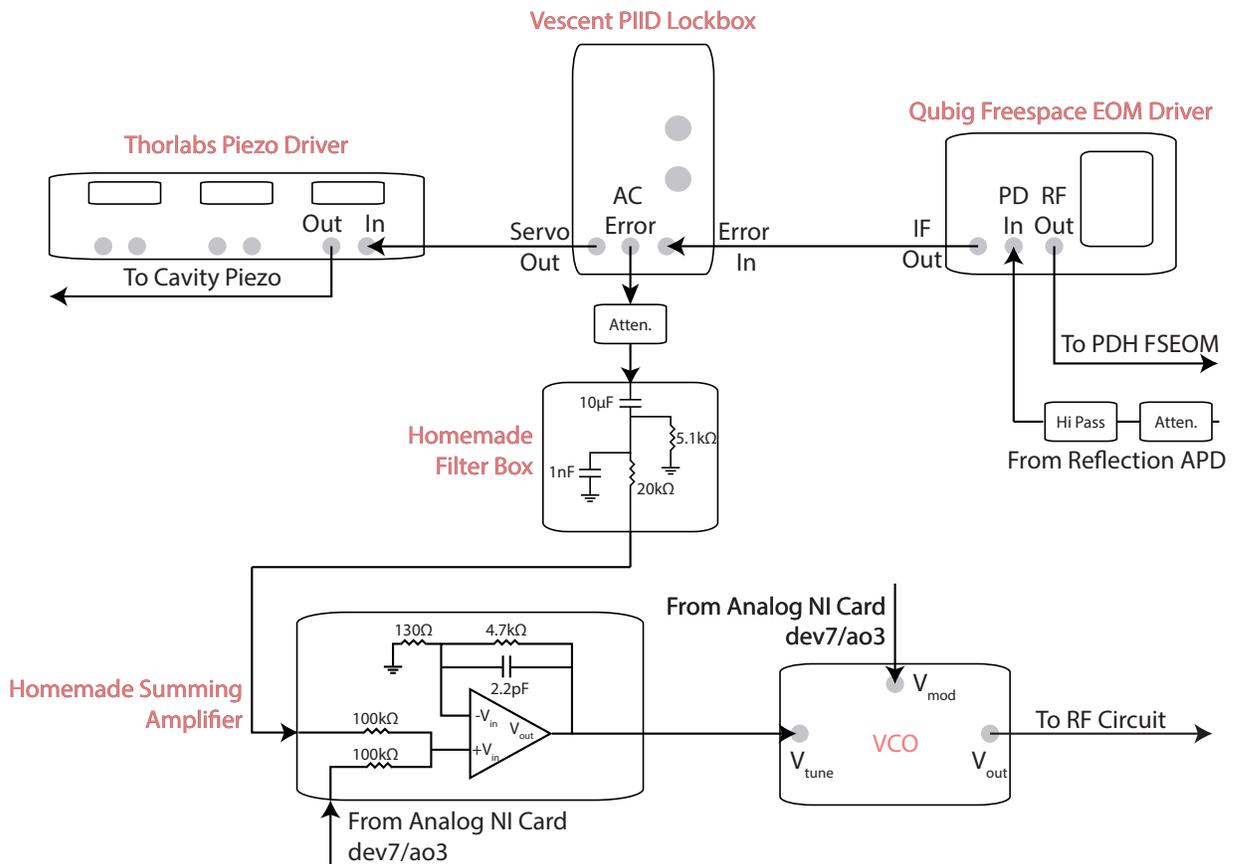


Figure 3-15: A detailed schematic of the first stage of PDH locking electronics, where both the fast and slow servo outputs are created. The slow servo output is taken directly from the lockbox output and sent to the cavity mirror piezo. The fast servo is created from the AC Error signal from the lockbox after passing through an attenuator, homemade filters, and a homemade summing amplifier.

An overview of our implemented PDH scheme for stabilization, including all electronics, RF circuits, and optics, is shown in Fig. 3-13. We have three knobs with which to set the relative frequency between our cavity and 1560nm locking laser: cavity length via temperature, cavity length via piezo, and locking laser frequency. Because temperature control of the cavity length is extremely coarse and slow, we chose to use it as a passive offset that is not controlled by the PDH scheme. Instead, we send the slow feedback from the PDH error signal lockbox (Vescent D2-125 Laser Servo in side-lock mode) in to the cavity piezo, and we send the fast feedback from the AC error of the lockbox to the cavity locking laser through an EOM in what is called a "feed-forward" loop. In this way, the frequency of the cavity follows the locking light at low frequencies, and the frequency of the locking light follows the cavity at high frequencies. This is an appropriate choice of feedback because mechanical transducers like piezos have a larger response at low frequencies and EOMs have a larger response at high frequencies. To manage the relative gains of our fast and slow feedback loops, we use several stages of filtering, amplification, and attenuation, as shown in the block diagram of the stabilization system in Fig. 3-13. A more detailed look at the circuit diagram and components is laid out in Fig. 3-15.

In addition to stabilizing the cavity and cavity lock laser to one another, we need to stabilize 780nm probe light with the same fast feedback to ensure the 1560nm locking light and 780nm probe light could both transmit through the cavity. The technically simplest way to do this is to double the 1560nm PDH light with a nonlinear crystal. This ensures that the resulting 780nm light is not only the frequency-locked but is also the phase-locked to the 1560nm stabilization light, without the added complexity of a separate phase-locked loop. However, because experiments may require different atom-cavity, atom-probe, and probe-cavity detunings, we ensured that our RF frequency system in Fig. 3-18 allows for all three types of detunings. We feedforward the same (frequency doubled) fast feedback from the 1560nm cavity lock laser to the 780nm probe laser stabilization loop in this RF circuit.

Finding the correct gains for both the fast and slow feedback loops took trial and error. After attempting many combinations of gains on our Vescent PIID box and



Figure 3-16: The PDH error signal for our installed experimental cavity without (a) and with (b) the fast feedback loop engaged. The sensitivity of the lock can be estimated via the slope of the central feature of the error signal without fast feedback engaged.

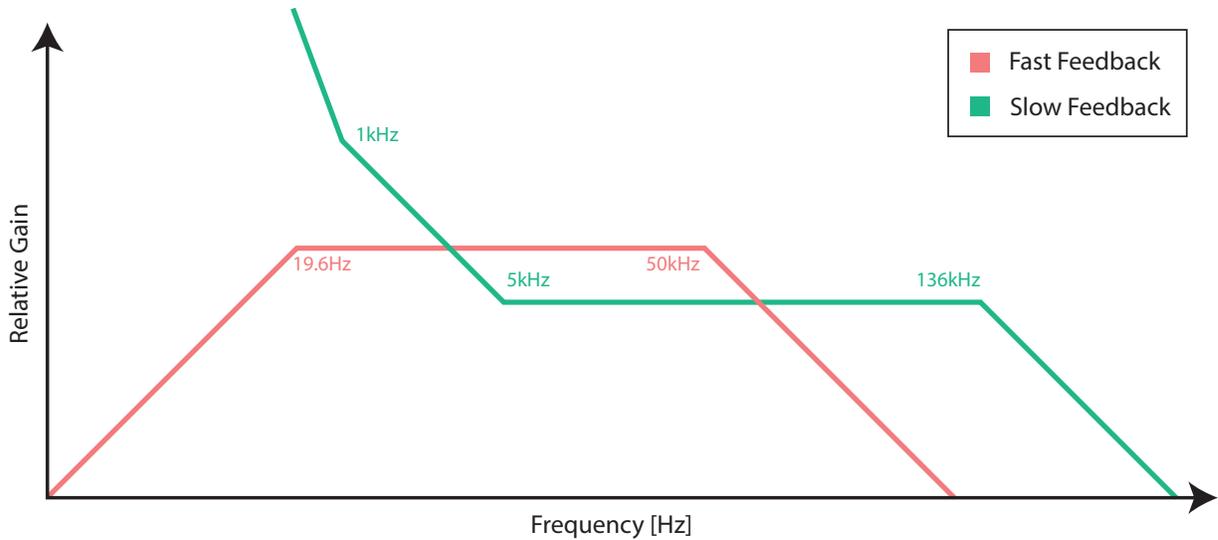


Figure 3-17: Estimated theoretical Bode plots for both the fast and slow cavity feedback loops. The slow feedback corner frequencies correspond to the first and second lockbox integrator frequencies and the piezo roll-off frequency, respectively. The fast feedback corner frequencies correspond to the cutoff frequencies of the homemade filter box shown in Fig. 3-15. This diagram is schematic and the relative gains and frequencies of the slow versus fast feedback are not plotted to scale.

switching out attenuators, amplifiers, and homemade filter boxes on each the fast and slow feedback loops, we found our system was most stable (i.e. responsive to high and low frequency perturbations without oscillating) with a feedback loop described by the approximate Bode plot in Fig. 3-17. We depended on noise measurements to guide us in optimization as we made changes to the various gains across the feedback loops. Initially, we worked to minimize the rms deviations in the error signal voltage as measured on an oscilloscope. Once the lock was fairly stable and could remain locked for tens of minutes, we moved on to analyze the AC error signal on a spectrum analyzer. Here, we adjusted gains to minimize resonances on the spectrum. After our initial gain adjustments monitoring the noise on the AC Error signal from the Vescent lockbox, we settled on a first integrator frequency of 500Hz, a second integrator frequency of 1kHz, and an overall coarse gain of -40dB. These are not the gains currently used in the lab, because we were able to further optimize the cavity feedback once we began taking data on the system.

After reaching a successful first cavity lock as described above, the next steps were

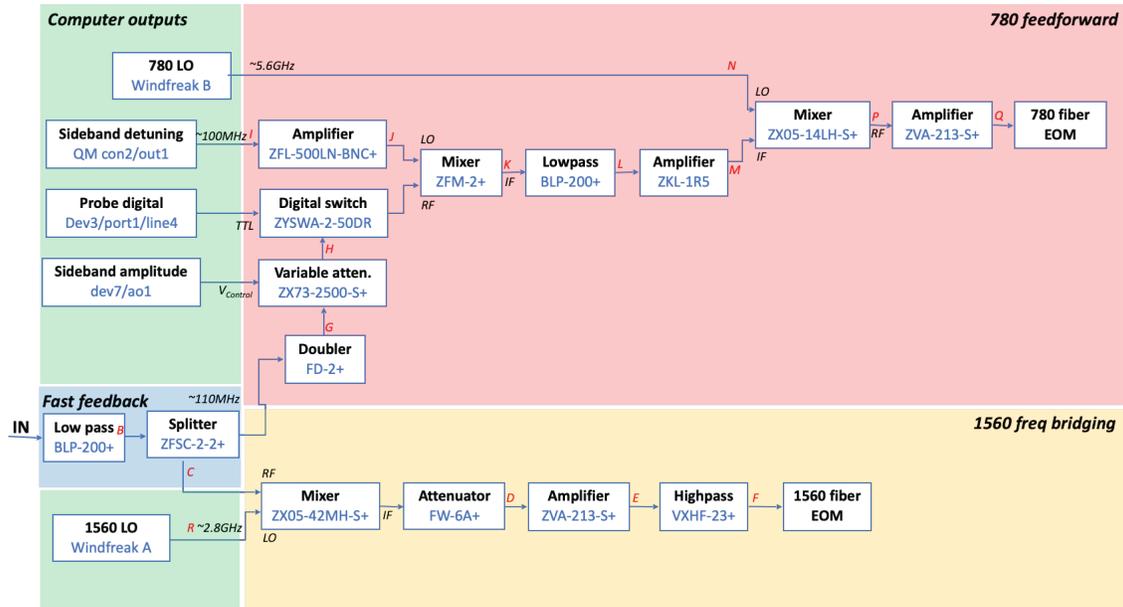


Figure 3-18: The RF circuit diagram for the fast feedback portion of the cavity locking scheme. The error signal resulting from the initial locking electronics in Fig. 3-15 is fed into a VCO. The resulting RF signal is mixed with an offset RF tone from an additional RF synthesizer, in this case a WindFreak. After mixing, the summed signal is filtered, attenuated, and amplified in order to create a useful gain. This final signal is sent to the cavity locking laser’s fiber EOM, which creates an optical sideband at the desired frequency (i.e. at the addition of the fast feedback signal and the offset RF). In parallel, the same fast feedback signal that is sent to the mixer for the 1560nm cavity locking laser is doubled and mixed with a second offset RF tone from a second WindFreak RF synthesizer. The resulting signal is sent to the 780nm probing laser EOM to ensure that both lasers see the same fast feedback and can be detuned separately.

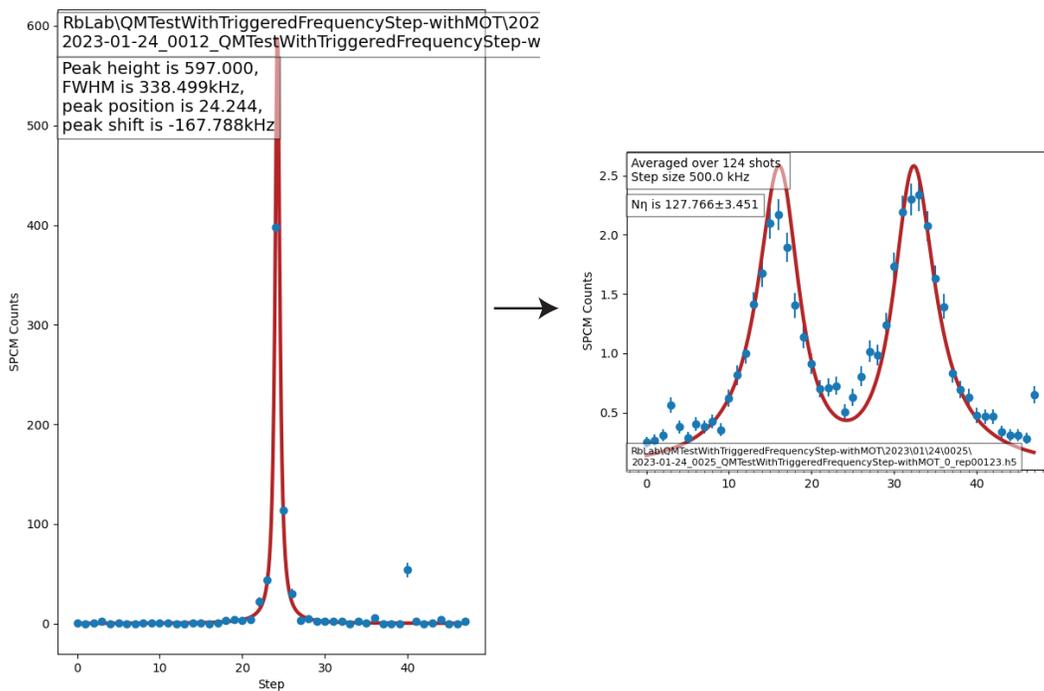


Figure 3-19: The cavity transmission spectrum without a MOT (left) and the initial Rabi splitting of the atom-cavity coupled system with a large atomic ensemble in the cavity after releasing a MOT (right).

to 1) resuscitate the MOT, 2) trap atom ensembles within the cavity mode, and 3) trap individual atoms. By adjusting the bias magnetic coils on the outside of our vacuum chamber and monitoring the cavity transmission on a camera, we were able to center the MOT in the cavity mode. When the alignment is close, we see the MOT atoms fluorescing into high-order cavity modes as we ramp the cavity length via the mirror piezo. As the alignment improves and on resonance, we see the MOT atoms fluorescing more and more into the TEM00 mode of the cavity. With the MOT in the cavity, we see Rabi splitting in the atom-cavity transmission spectra—a good indication of coupling shown in Fig. 3-19. However, as we monitored the photon counts at the transmission of the cavity after releasing the MOT, we noticed a large oscillation in our signal, shown in Fig. 3-20. The signal displayed fast oscillations at 430Hz and slower oscillations at 26Hz. We were able to find matching oscillations in the error signal of the cavity lock by triggering the oscilloscope on the MOT coil turn-off time. This was a strong indication that the oscillations were due to gain imbalance in the cavity lock. That is, our cavity lock could not compensate for mechanical oscillations cause by shutting off the MOT coils. To further damp the oscillations, we adjusted the Bode plot of our lock, shifting the first integrator frequency up to 1kHz, the second integrator frequency up to 5kHz, and attenuating the fast feedback by 2dBm. This reduced the oscillations' amplitude to below our noise floor, and set the final cavity lock loop gains and frequency cut-offs to those documented in Fig. 3-17.

The resulting error signal from our PDH circuit is shown in Fig. 3-16 both without (a) and with (b) the fast feedback loop engaged. Because the fast feedback loop is entirely homemade and takes the AC Error output signal from the Vescent PIID lockbox, the fast feedback is always engaged, even when the lockbox is in the "Ramp" or "Unlock" modes. The fast feedback has the effect of stretching out the appearance of the PDH error signal when ramping across resonance as it tries to keep the laser and cavity on resonance as long as possible. Our cavity has a much narrower linewidth at the probe wavelength of 780nm than at the locking wavelength of 1560nm. Thus, to characterize the performance of our lock and judge if it is "good enough," we can simply look at the noise on the transmitted 780nm probe light when locked. With the

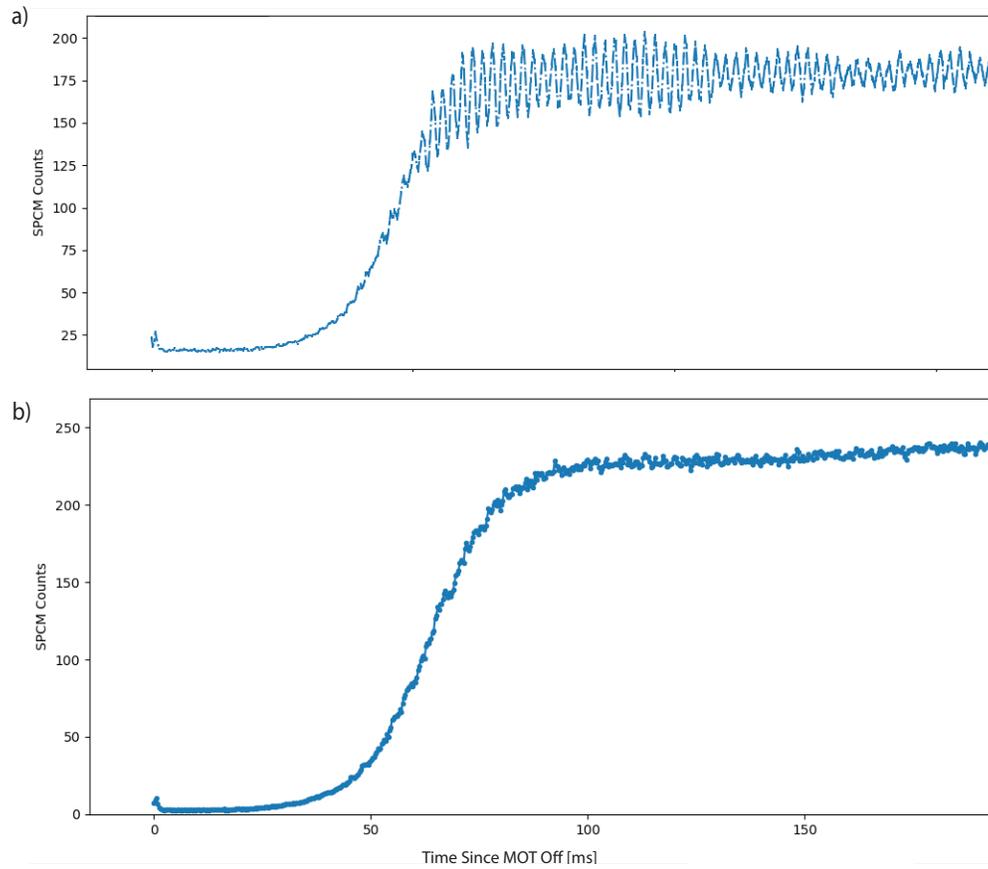


Figure 3-20: Counts on the SPCM at the cavity transmission as a function of time after a loaded MOT was released by turning off the MOT light and coils both a) before and b) after adjusting the cavity lock parameters to reduce oscillations induced by MOT turn-off.

Property	Value
κ	127kHz
τ	1.7 μ s
η	2.10
κ_{lock}	20.4kHz
ν_{FSR}	3.42GHz

Table 3.4: Properties of the installed and stabilized experimental cavity at the probe wavelength, 780nm.

adjustments to minimize the oscillations in the SPCM signal after MOT turn-off, we measure a $55.5V_{rms}$ error in the transmitted optical power on a signal with a peak-to-peak amplitude of 2.2V. Given the known cavity linewidth of 127kHz at 780nm, this implies our lock linewidth is 20.4kHz, or 16% of the cavity probe linewidth. In future iterations of this experiment, we can improve the locking substantially by designing robust cavity mirrors that have a narrower linewidth at the locking frequency than at the probing frequency. Having a narrower linewidth at the locking frequency will mean that the electronics of the lock will not have to be tuned as finely, hugely simplifying the stabilization process. The original cavity design did have this property, but unfortunately the required HR mirrors were lost in the first bake.

Chapter 4

Cavity Mirror Materials

High-reflectivity, low-loss optical mirrors are the foundation for high-finesse optical cavities, indispensable tools for studying fundamental light-atom interactions and quantum information science [77, 103, 53, 82]. Cavity QED systems can mediate entanglement generation in many-atom systems on demand for both metrological and quantum simulation applications [68, 49]. Indeed, cavity-mediated entanglement allows for large-atom-number ensembles to be entangled, enabling measurement resolutions beyond the standard quantum limit (SQL) [46]. Further, optical cavities can act as nodes for coherent photon transfer in quantum networks of various hardware platforms, from NV-centers, to trapped ions, to neutral atoms [83, 20, 28]. Cavities even can help on the path to a fault-tolerant quantum computing. At some point, quantum computing architectures like neutral atom arrays will reach a critical point in system size and atomic qubit state readout speed. By connecting smaller modules and by mediating nondestructive microsecond-scale state readout through cavity measurements, cavities can address both of these challenges [82, 99]. To achieve these proposed benefits across metrology, quantum simulation, and quantum information science with cavities, the cavity-atom system must be well within the strong-coupling regime. To be in this regime, the system must have a large single-atom cooperativity, $\eta = 4g^2/\kappa\Gamma \geq 1$, which depends both on atomic (Γ) and cavity (κ) losses [98]. Thus, for a fixed atomic species, atom number, and cavity mode volume, reducing cavity mirror loss and transmission is critical to reaching the strong-coupling regime and

taking full advantage of optical cavities' capabilities [61].

High-reflectivity optical cavity mirrors near infrared wavelengths are typically implemented as distributed Bragg reflectors (DBRs) with alternating layers of Ta₂O₅ (n=2.04) and SiO₂ (n=1.45). To reach the required reflectivities (1-50ppm transmission), these stacks consist of twenty to fifty layers of these alternating materials, with each layer's thickness determined by the design optical wavelength(s) [44]. These stacks tend to be capped in the higher index of refraction material Ta₂O₅ by default to reach the largest index differential between the dielectric stack and the environment, and therefore the highest reflectivities. To build an atomic system within an optical cavity, the mirrors need to be UHV-compatible, often down to pressures as low as 10⁻¹⁰ Torr. In order to reach such a low pressure, the entire system must be annealed at temperatures approaching 200°C. Unfortunately, the optical properties of the standard mirror top layer Ta₂O₅ degrade in ultra-high vacuum (UHV) conditions, supposedly because the UHV pressure drops below the vapor pressure of oxygen in the material; thus resulting in oxygen reduction [43, 22, 95, 17]. Previous studies show that this vacuum-induced oxygen loss can be reversed by immersion in a high-O₂ environment and prevented by depositing a thin (1-2 nm) cap-layer of SiO₂ in most scenarios. However, if a mirror is heated above 150° while in vacuum, the degradation is irreversible. The material mechanism that translates oxygen reduction in Ta₂O₅ to optical loss and the reason for irreversible loss after a higher temperature anneal is not well-understood.

In this chapter, we our observation of an increase in HR cavity mirror losses after a twelve-day UHV anneal at 180°C. To identify the cause of the degradation, we complete several diagnostic and attempted recovery steps: 1) mirror re-alignment, 2) mirror cleaning, 3) optical metrology, 4) atomic-force microscopy (AFM), 5) X-ray fluorescence (XRF) spectroscopy, and 6) selective wet-etching. Mirror re-alignment and cleaning does not recover the initial finesse of our cavity, and optical metrology does not yield an explanation for the increase in loss. AFM images of the mirror surfaces reveal a large increase in surface roughness of high-reflectivity cavity mirrors after UHV anneal. To identify the chemical nature of the roughness increase, we use

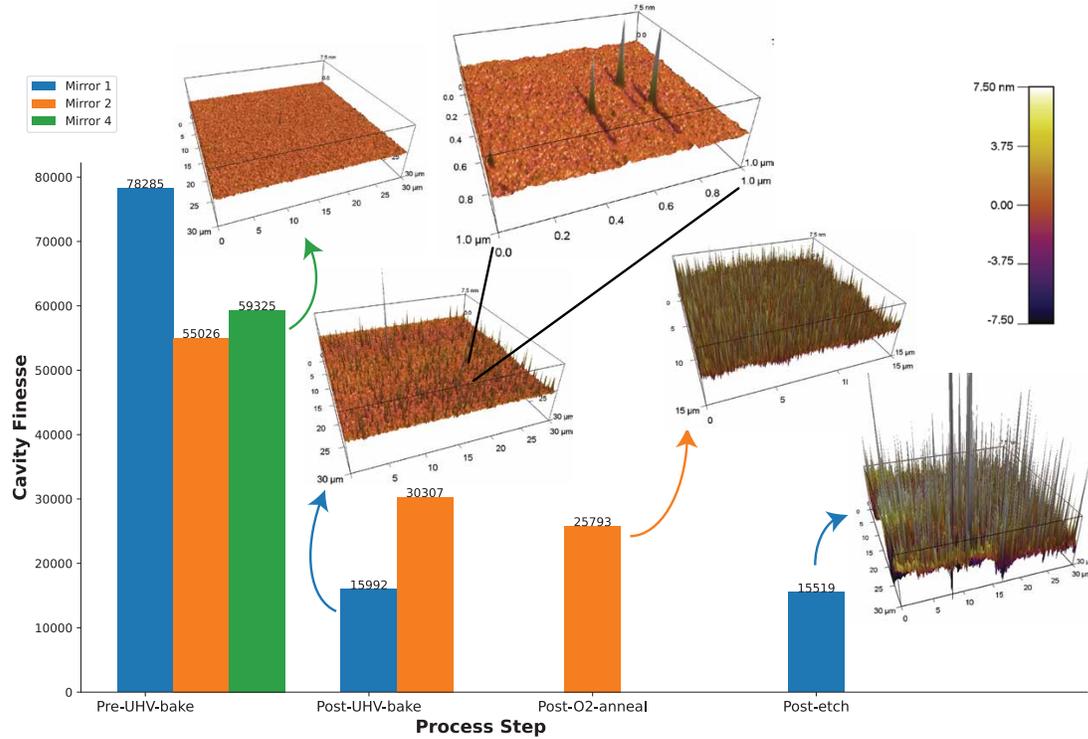


Figure 4-1: Cavity finesse as a function of mirror processing step. The cavity finessses decreased after the initial vacuum bake and after each subsequent recovery attempt that included a bake step. Corresponding 3D AFM images of the top surface of the mirror coatings at each step are also pictured.

XRF and selective wet-etching, which identify the elemental composition and presence of SiO_2 , respectively. Finally, we propose an explanation for the degradation of high-finesse DBR optical cavity mirrors post-high-temperature UHV anneal and make recommendations for the prevention of degradation. We conclude the degradation is caused by oxygen reduction in the top Ta_2O_5 layer followed by nucleation and growth of Ta-oxide, pure Ta, or pure Si defects at the resulting oxygen vacancies. We discuss the dependence of mirror losses on surface roughness and show the optical loss scales as expected. Finally, we make recommendations for avoiding these vacuum-induced losses in the future.

4.1 Materials and Methods

4.1.1 Mirror Coating Runs

The mirrors in this study come from two distinct coating runs (from Advanced Thin Films (ATF) in Boulder, CO in 2015). Both dielectric stack coatings were designed at 780nm and 1560nm wavelength light. Mirror Coating One was designed for transmission of 2ppm at 780nm and has a total thickness of $11.6\mu\text{m}$. Mirror Coating Two was designed for transmission of 50ppm at 780nm with a total thickness of $7.8\mu\text{m}$. To meet these goals, the two mirrors have a different number of layers in the DBR dielectric stack, but the substrates (UV Fused Silica), layer types (Ta_2O_5 / SiO_2), and layer orders (capped in Ta_2O_5) are otherwise identical. Both coating runs were completed using ion beam sputtering (IBS) followed by a high temperature anneal.

4.1.2 Annealing and Optical Methods

The initial UHV anneal of the DBR mirrors in this study was conducted in a stainless steel and fused silica vacuum chamber used for experiments with rubidium atoms. After the cavity mirrors (four separate mirrors, from two separate optical cavities, and two mirror coating runs) were installed in the chamber through the upper view port of the system, all conflat flanges were sealed with the recommended torque and the system was pumped out following a standard pumping procedure with stages including turbo, ion, and titanium sublimation pumping. During this process, the chamber was heated to a temperature of 180°C using fiberglass heating tapes wrapped around the metal surfaces. The chamber was held at this temperature for twelve days. The pressure at the beginning of this period was 2×10^{-3} Torr and at the end of the period was 3×10^{-10} Torr as indicated by the ion pump current.

After annealing the mirrors during this vacuum bake-out, initial ringdown time measurements in the chamber were less than the cavity transmission photodetectors' response times. Additionally, the modes of both cavities were visibly asymmetric on transmission cameras, implying mirror misalignment of both cavities. After removing

the cavities from the chamber, re-aligning, and cleaning the mirrors, the ringdown times and thus finesse were still significantly less than before the anneal, as depicted in the "post-UHV-bake" bar in Fig. 4-1. Following the procedure in [43], we took a subsequent anneal processing step aimed at recovering mirror quality by removing oxygen vacancies. This anneal was carried out in a 100% O₂ tube furnace at 400°C.

Before and after each anneal and processing step, we optically characterized the cavities via ringdown measurements with resonant light. Resonant light at 780nm was focused onto the cavity mode using mode matching lenses and steering mirrors. An acousto-optic modulator (AOM) triggered by a high optical signal at the output of the cavity switches off input light to the cavity as a cavity piezoelectric transducer (PZT) scans the length of the cavity across laser resonance. The resulting decay of light power at the cavity transmission photodiode is fit by an exponential and the characteristic 1/e ringdown time (τ) is extracted from that fit. With the cavity free-spectral range defined as $\nu_{FSR} = \frac{c}{2L}$ with L the distance between cavity mirrors, this ringdown time allows us to calculate cavity finesse using $\mathcal{F} = \nu_{FSR}/\kappa$ where $\kappa = 1/(2\pi\tau)$ is the cavity linewidth in Hz [98].

4.1.3 Non-Optical Characterization Methods

We utilized three main non-optical characterization methods to diagnose and understand the loss mechanism in our mirrors: atomic force microscopy (AFM), X-ray fluorescence (XRF), and selective etching. After several rounds of cleaning and re-aligning, the cavity mirrors did not recover their initial, pre-anneal performance. Thus, we postulated that physical changes to the surface of the mirror could be the source of loss. To verify this hypothesis, we measured mirror roughness using AFM, a method that maps the topology of a surface using deflections of and active feedback to a microscopic cantilever. In our case, we used an Asylum Research Cypher AFM machine in repulsive tapping-mode/AC with a 160AC-NA aluminum-coated tip. After mapping all of our mirrors' HR surfaces and noting a roughness increase after the anneal, we then moved to identify the defects. To do so, we carried out two material analyses: XRF and selective wet etching. We used a Bruker Tracer-III

handheld XRF device to qualitatively determine which elements were present in the reflective coating of our mirrors, both before and after UHV anneal. The XRF results showed that the increase in roughness was not due to contamination of the mirror dielectric stack with an external material. We then further narrowed down the identity of the growths by using a selective wet etch of the surface. In the semiconductor industry, a common chemical etch, or method to remove material from a surface in a controlled manner, is a buffered oxide etch (BOE). BOE etches SiO_2 at a rate orders of magnitude faster than our mirrors' other semiconductor materials. For this reason, we used selective wet-etching with a 7:1 buffered oxide etch (BOE) in the MIT.nano cleanroom to determine whether the defects were made of SiO_2 or another material.

4.2 Results and Discussion

4.2.1 Process Effects on Cavity Finesse

The initial bake in UHV caused a severe decline in cavity finesse, as shown in the first two data points of Fig. 4-1. The drop in mirror finesse was only observed on mirrors from Mirror Coating Run I. The cavities were designed to be in the strong coupling regime with a single atom cavity cooperativity $\eta \geq 1$ for the rubidium D_2 line. Before the bake, $\eta = 3.4$ and with the drop in finesse of mirrors from Mirror Coating Run I, $\eta = 0.9$ and is no longer in the strong-coupling regime. In an attempt to recover the initial finesse of the cavities, we followed the recommendation of a previous study that claimed oxygen vacancy formation in the Ta_2O_5 cap [43]. We anneal one degraded mirror from Mirror Coating Run I ("Mirror 2") in a pure oxygen environment at 400°C for four hours. Although the measured finesse was slightly higher than the other degraded mirror after this O_2 anneal, it did not recover even 50% of the initial finesse, implying that an additional mechanism is responsible for the degradation.

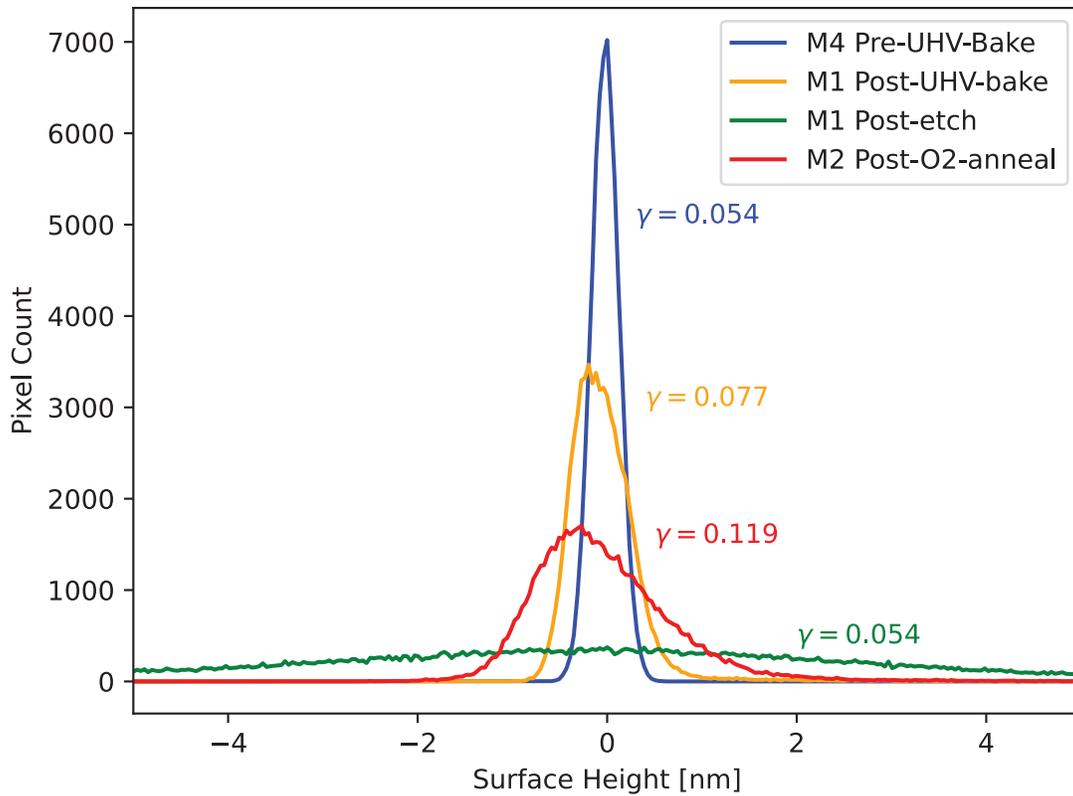


Figure 4-2: Histogram of mirror surface heights after the various processing steps. The distribution in surface height of a single mirror becomes less Gaussian with each step, indicated by increasing skew, γ . Additionally, each processing step intended to return the mirror surface roughness to the pre-UHV-bake value actually increased the deviation in mirror roughness.

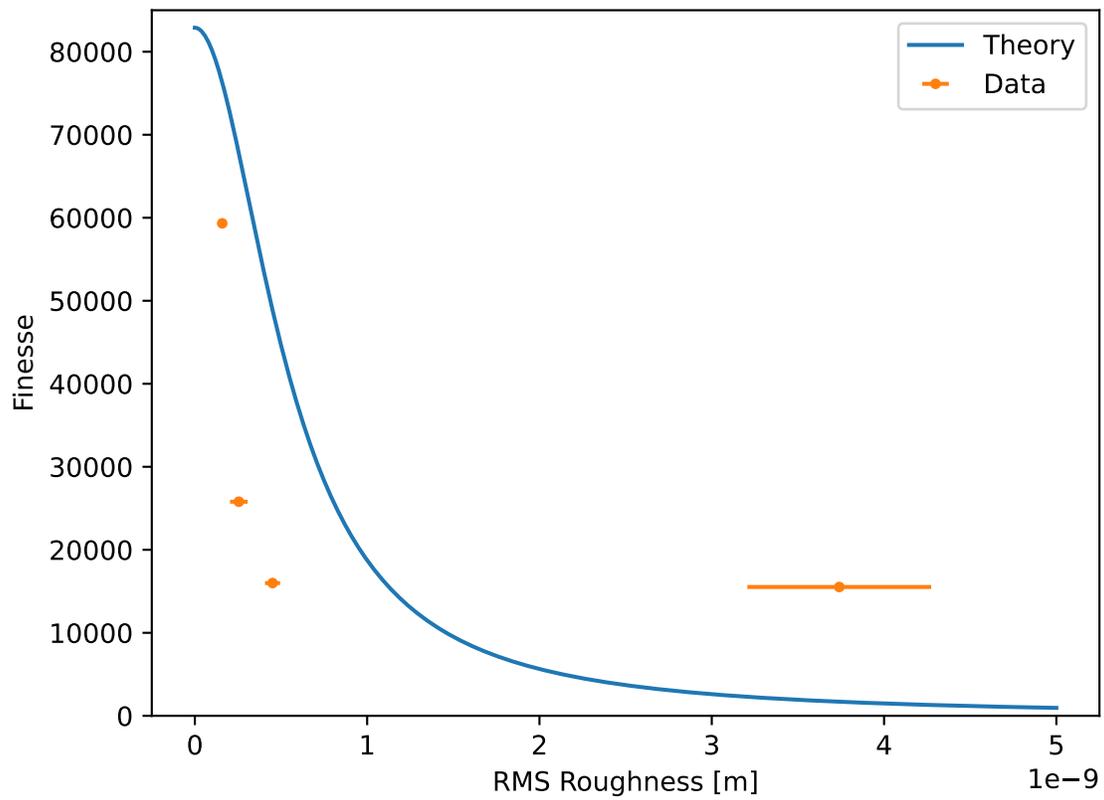


Figure 4-3: Expected theoretical reduction in cavity finesse with mirror surface roughness compared with measured reduction in cavity finesse with mirror surface roughness. The theory curve is calculated with no free parameters.

4.2.2 Process Effects on Mirror Roughness

AFM images of the DBR surface of UHV annealed cavity mirrors from Coating Run I show a drastic change in surface roughness as compared with images of non-annealed mirrors, shown in Fig. 4-1. The roughness is attributable to an abundance of high-aspect-ratio point defects after the anneal. Unannealed mirrors from Coating Run I do exhibit these point defects, but at a much lower density than their annealed counterparts. Mirrors from Coating Run II never display these point defects, before or after the anneal. This discrepancy between the behavior of the two coating runs implies that manufacturing errors in Coating Run I could be partially responsible for the severe degradation under UHV anneal. In addition to the deviation in mirror defect heights increasing with each processing step, as shown in Fig. 4-2, the distribution of defect heights also becomes more skewed and less Gaussian after each processing step. The post-BOE curve does not follow this increase in skew, however, the distribution of defect heights is so uniform that the error in calculating skew is much higher in this case than the others. This redistribution of defect sizes between processing steps could indicate a higher level of periodicity in the surface topology. Such periodicity could introduce interference phenomena, complicating our simple model of the cavity finesses' dependence on surface roughness plotted in Fig. 4-3. The root-mean-square (rms) surface roughness extracted from the AFM measurements of mirrors from Coating Run I plotted in Fig. 4-3 shows that the increase in roughness from pre- to post-anneal qualitatively tracks the decrease in cavity finesse during the same period. Indeed, one would expect the loss of the mirrors to increase with surface roughness as [14]

$$L = L_0 + \frac{(4\pi\sigma_{rms})^2}{\lambda^2}. \quad (4.1)$$

We can relate the losses of individual mirrors to the finesse of an optical cavity with $\mathcal{F} = 2\pi/(L_1 + L_2 + T_1 + T_2)$ [66]. Combining these equations describes the relationship between our cavity finesse and mirror roughness as

$$\mathcal{F} = \frac{2\pi}{L_1 + L_2 + \exp[-(4\pi\sigma_{rms})^2/\lambda^2] + T_1 + T_2}, \quad (4.2)$$

where $L_1 = 10\text{ppm}$ and $L_2 = 10\text{ppm}$ are the initial losses of the two cavity mirrors before annealing, $T_1 = 50\text{ppm}$ and $T_2 = 5.8\text{ppm}$ are the transmission coefficients of the two cavity mirrors, and we have the added loss term that depends on σ_{rms} [96]. Our finesse and roughness data from the optical and AFM measurements are plotted in Fig. 4-3 along with this theoretical curve. The shaded region around the theory curve indicates one standard deviation of error in our measurements of both roughness and finesse. We find good qualitative agreement between data and theory and conclude our data are consistent with mirror roughness as the dominant loss mechanism in the degraded mirrors. The final data point at highest roughness in this plot strays from the theory the most, lying far outside one standard deviation of error in roughness and finesse. It is possible this can be explained by our inability to perform AFM and ring-down measurements on the exact same point on the mirror surface due to a lack of location-identifying marks on the mirror surface.

4.2.3 Elemental Analysis

To further understand the degradation mechanism, we set out to identify the chemical composition of the defect growths.

X-Ray Fluorescence Spectroscopy

To determine the identity of individual elements in the surface point defects in the AFM images in Fig. 4-1, we analyzed each mirror surface with XRF, yielding an identical set of characteristic peaks as shown in Fig. 4-4. The most prominent fluorescence peaks were those characteristic of tantalum, followed by silicon, palladium, rhodium, and iron, the last three of which are elements found in the XRF detector itself. This measurement confirms that the defects must originate from the materials within the DBR dielectric stack instead of from a contaminant. Thus, the XRF implies the defects are pure silicon, pure tantalum, or oxides of those elements.

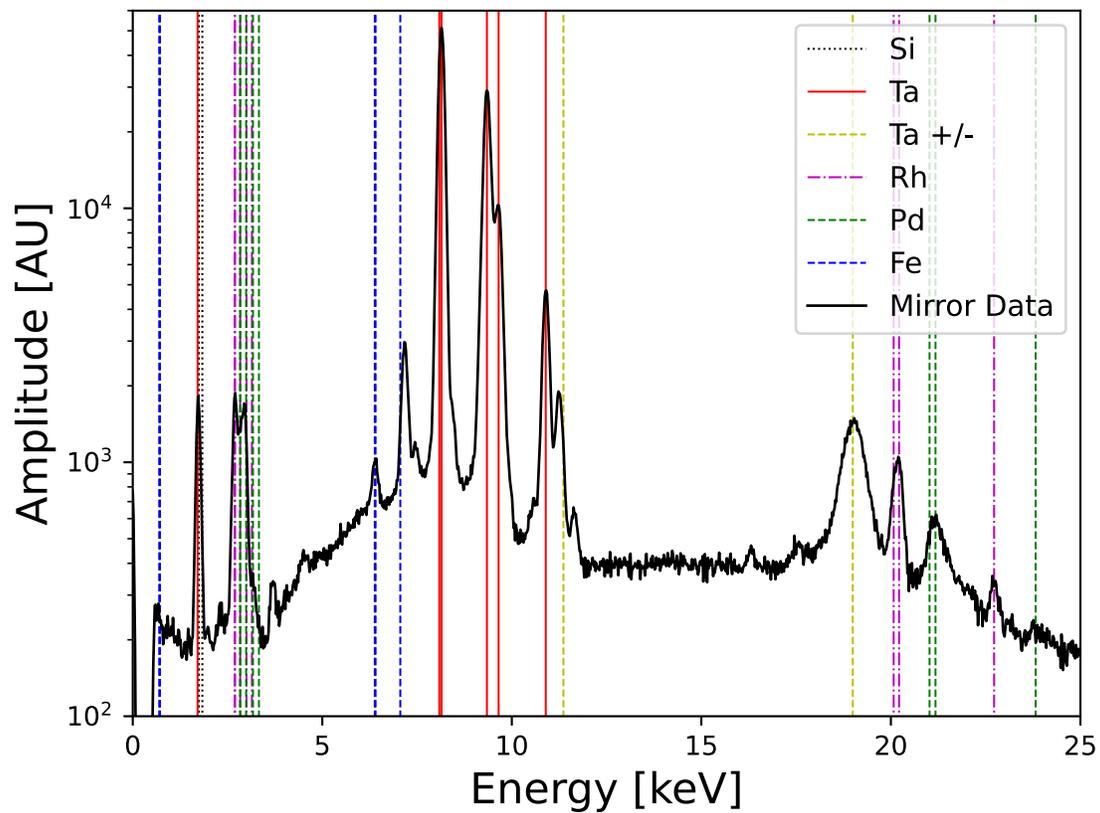


Figure 4-4: X-ray fluorescence rules out the possibility that a contaminant from the vacuum chamber caused the mirror surface roughness increase. Only Si, Ta, and artifacts of the XRF device (Rh, Pd, and Fe) were resolvable in the measurement.

Selective Wet-Etching

A selective BOE further narrows down the identity of the defect: BOE is highly selective for SiO_2 over the other material possibilities. Specifically, BOE etches SiO_2 at a rate of 2nm/s, Si at a rate of 0nm/s, and Ta_2O_5 at a rate of 3×10^{-4} nm/s at room temperature and pressure [18, 27]. Thus, etching our HR mirror surface with BOE revealed whether the roughness increase was due to SiO_2 or one of the other dielectric materials in our mirrors. At room temperature, a BOE of fourteen seconds should etch away at least 16nm of SiO_2 , which is not what we observe [80]. Because we would expect no change to the mirror surface if the defects were not SiO_2 , and this is what we observed, the defects must be pure Si, pure Ta, or Ta_2O_5 . However, as is shown in Fig. 4-1, the roughness of Mirror One increased significantly after the BOE step. This un-intuitive result is explained by additional heating during the placement of a photomask. Before etching, an anneal process is required to set a protective photoresist layer on the anti-reflective side of the mirrors. This additional annealing could have caused further defect growth. The photoresist is applied manually because the concave geometry of the high-reflective (HR) mirror surface prevents spin coating. The mirror with photoresist bakes at 95°C for twenty to forty five minutes. Given that the roughness of each mirror increased during any processing step that elevated temperatures, the increase of roughness after etching is consistent with our other observations.

4.3 Conclusion

In conclusion, we find that the UHV degradation of Ta_2O_5 -capped DBR cavity mirrors is not only attributable to oxygen vacancy formation in the tantalum oxide layer. In fact, we find that when annealed above 150°C in UHV, these mirrors can undergo an irreversible nucleation and growth process of the DBR stack materials. This material change is directly observable in optical ring down time characterizations of the cavity and under AFM inspection. This observation confirms and extends the result from a previous study where optical quality of cavity mirrors baked at 150°C

could not be recovered with oxygen flow [43]. It is likely that oxygen defects form at low temperatures and then act as nucleation points for defect growth at higher temperatures.

We also find that the presence of low-density defects on unannealed mirrors is correlated with degradation post-anneal. This implies that small changes in the commercial IBS and annealing processes used to create these DBR dielectric stacks has immense impact on the longevity and performance of the mirrors. The process may become more unstable as dielectric stack thickness increases, as we only observed unannealed defects in our thickest mirror coating, Coating Run I. The formation of "whisker" defects with high aspect ratio are well-documented in other species of metal-oxide materials during the oxygen reduction process [23]. Although the whisker formation process has not previously been observed in Ta_2O_5 specifically, it is generally accepted that the presence of surface defects (oxygen vacancies or external defects) in transition metal oxides often are responsible for subsequent catalytic activity, which supports our hypothesis of oxygen vacancies as defect nucleation sites [55].

Chapter 5

Atom Trapping and Addressing

A useful cavity-coupled neutral atom array requires optics both for array atom trapping and the atomic state addressing to manipulate states and create interactions. The cavity and vacuum apparatus is at the heart of the experimental upgrade completed in this thesis, but our trapping and addressing systems will be just as critical in moving toward new quantum information directions. The atom trapping and addressing systems designed for and implemented into our experiment take and combine many recent technical improvements from other atom array systems. Thus, our system has the ability for atom motion and Raman beam generation on top of standard trapping and single-photon addressing [16, 64]. The upgraded experiment has the ability to trap atoms in individual optical tweezers, in a cavity lattice, or both. The internal states of the atoms will be manipulated with single-site Raman beams, two-photon Rydberg excitations, and global manipulations via our MOT optics. In this section, we discuss the theoretical goals and resulting technical requirements of each of these trapping and addressing systems. We then document the initial in-lab realizations and finally summarize what is yet to be designed and built. At the time of this writing, each of the trapping and addressing systems are in various stages of installation, alignment, and testing. The optical tweezers and cavity lattice were first to be incorporated into the experiment. The Raman hyperfine shelving beams, Rydberg transition optics, and stationary SLM tweezer array were all in progress at the time of thesis submission.

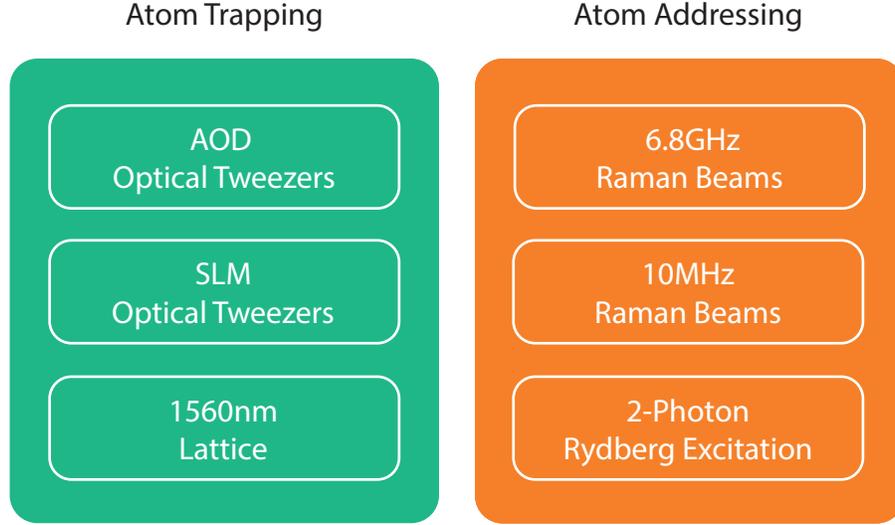


Figure 5-1: Overview of the atom trapping and addressing systems in the new experiment platform. Not all of these systems will be integrated in the first iteration of this experiment. Currently, only the AOD tweezers are fully built and integrated with the vacuum chamber. The 6.8GHz Raman beams and 2-photon Rydberg Excitation optics have been partially built, but not yet integrated with the chamber. The SLM optical tweezers are in the design and testing process and the 10MHz Raman beams have not been started.

For a visual summary of the atom trapping and addressing systems, see the block diagram in Fig. 5-1. Ultimately, we plan to have three distinct trapping systems: AOD optical tweezers with motion, a cavity lattice, and a stationary two-dimensional spatial light modulator (SLM) tweezer array. Similarly, beyond the global MOT addressing, we will have three new individual atom addressing systems: 6.8GHz Raman beams, 10MHz Raman beams, and two-photon Rydberg excitation. Because basic MOT and cooling interactions are extensively discussed elsewhere, both generally [108] and for this experiment specifically [101] and because atom array trapping is similarly well understood, we will focus this chapter on the specifics of the new components in our experiment.

5.1 Optical Tweezers

To trap a single rubidium atom in an optical tweezer, the potential the atom experiences via the dipole interaction with the tweezer light must be larger than the energy

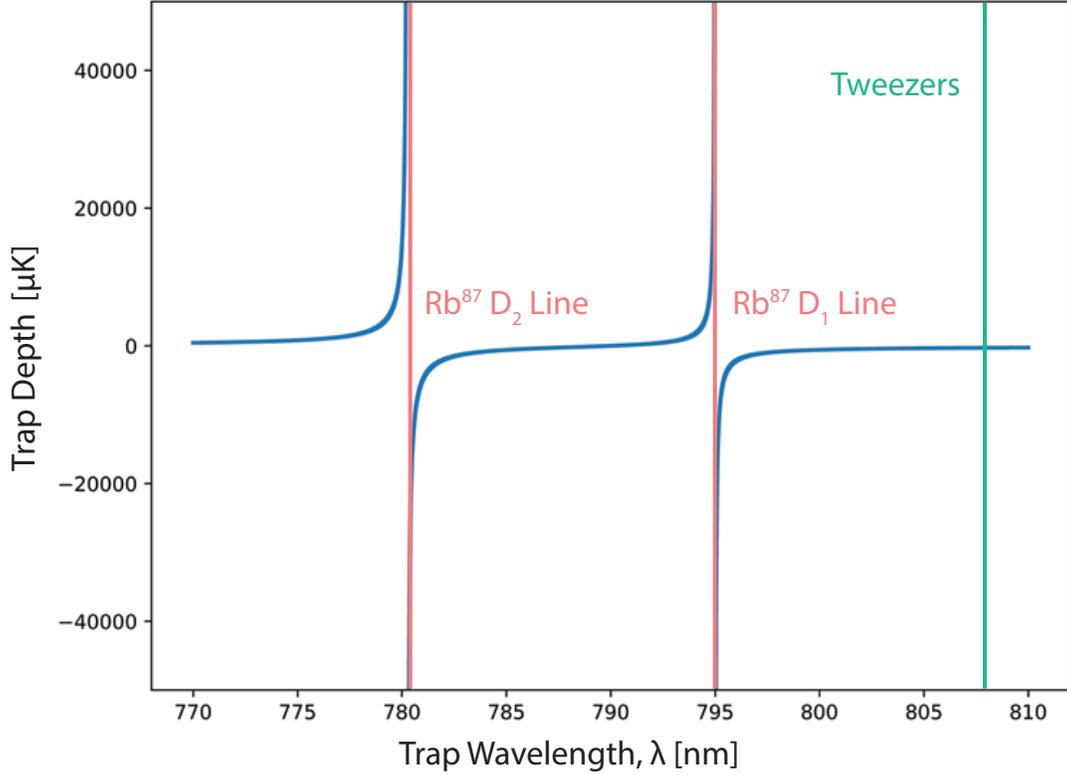


Figure 5-2: Dipole trap depth in μK as a function of trap wavelength for a rubidium-87 atom.

of the atom. In our lab the first stage of cooling and trapping is in our magneto-optical trap (MOT), which can cool atoms to the Doppler limit, $T_D = \frac{\hbar\Gamma}{2k_B}$ [94]. After creating our MOT on the ^{87}Rb D_2 line from $F=2 \rightarrow F'=3$, we expect our atoms to be cooled to $145\mu\text{K}$. To reach lower temperatures, we can perform polarization gradient cooling (PGC) to reach sub-Doppler temperatures by detuning the MOT beams and turning off the magnetic coils [26]. The theoretical limit of this cooling is given by the recoil limit, $T_{recoil} = \frac{\hbar^2 k^2}{Mk_B}$ [29]. With rubidium 87 atoms, this corresponds to a temperature of 362nK [94].

For a tightly-focused optical dipole trap, also called an optical tweezer, the trap depth must be several times larger than the atom's temperature to trap it. Assuming a wavelength of 808nm and linear polarization in the limit of large detunings, negligible saturation, and the rotating wave approximation (RWA), we can calculate our tweezer trap depth via Equation 5.3 [47]. In our case, if we cool with only a MOT, we require

a $1\mu\text{m}$ waist tweezer with a power of $\sim 330\mu\text{W}$ to give us a trap depth of twice the Doppler limit, $290\mu\text{K}$. With an added PGC stage, the requirements become more lax and each $1\mu\text{m}$ tweezer would only need a power of 800nW to give us a trap depth of twice the recoil limit, 660nK . We assume the beam waist to be $1\mu\text{m}$ for these calculations for two reasons. First, at this trap size, collisional losses dominate inside the trap and ensure only one or zero atoms will be loaded [100, 48]. Second, this size is a few times larger than the diffraction limit of our in-vacuum aspheric lens that is used to focus the tweezers onto the atoms, as discussed in Ch. 2, Sec. 2.3. As such, $1\mu\text{m}$ is approximately the smallest trap waist size we expect to achieve after our beam passes through several optics along the way to the atoms.

$$U_{dipole}(r) = \frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) I(r) \approx \frac{1}{k_B} \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\Delta} I(r) \quad (5.1)$$

$$\Gamma_{sc}(r) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right)^2 I(r) \approx \frac{3\pi c^2}{2\hbar\omega_0^2} \left(\frac{\Gamma}{\Delta} \right)^2 I(r) \quad (5.2)$$

$$T_{trap}(r) = \frac{U_{dipole}(r)}{k_B} \quad (5.3)$$

Trap depth is not the only figure of merit for a tweezer systems, however, as the scattering rate of the atom in the trap also determines heating over time and limits the lifetime of trapping. Minimizing the heating rate of our trapped atoms is what led us to choose 808nm as our main trapping wavelength. Eq. 5.2 gives the expression to calculate the scattering rate of an atom in a tweezer for the same approximations as above: large detunings, low-saturation, and RWA [47]. This expression can be related to the trap depth in Eq. 5.3 simply by $\hbar\Gamma = \frac{\Gamma}{\Delta} T_{trap} k_B$, immediately revealing that for a fixed trap depth, one can minimize the scattering by detuning the trap light from the atomic resonance. This process is not as simple as considering just one atomic transition, however. If we consider the D_2 line of rubidium-87 at 780nm , as we further red-detune a trap laser, we get closer to resonance with the D_1 line at 795nm . This means we must balance the energy of our trapping laser to be in between the two lines or red-detuned from them both. A plot of this dependence is shown in Fig. 5-3.

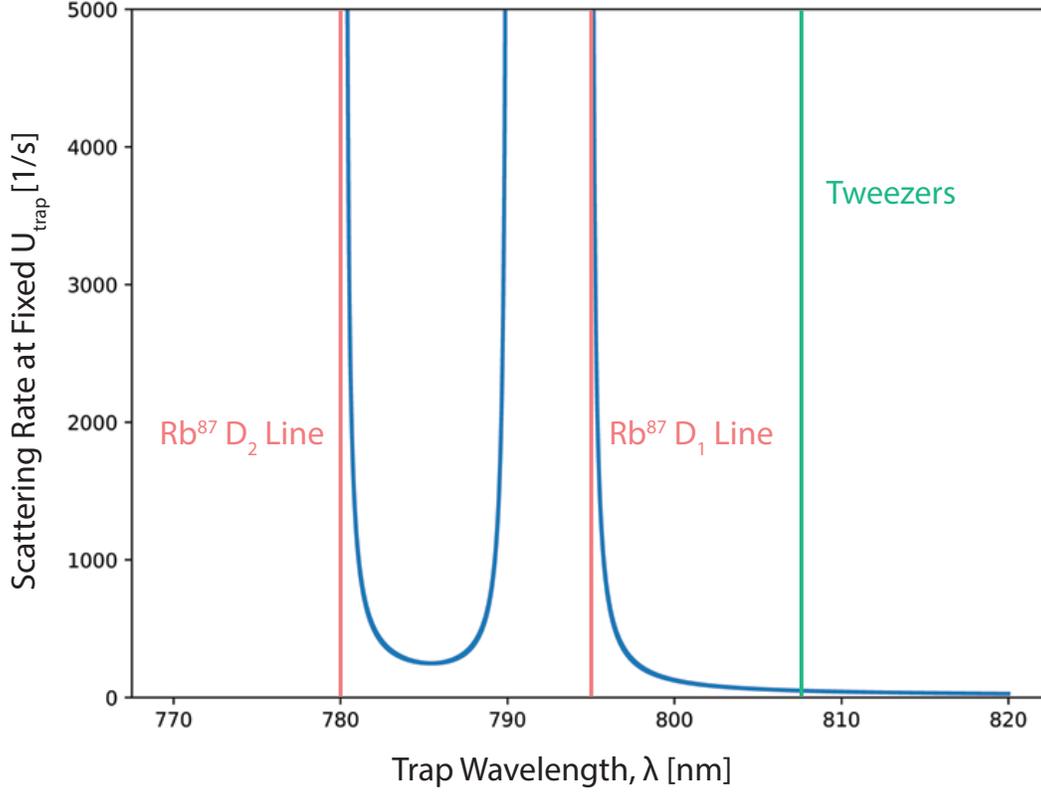


Figure 5-3: The dependence of photon scattering rate on the tweezer trap wavelength in rubidium-87. The heating rate is directly proportional to the scattering rate of the atoms in the trap, so we chose a trap wavelength of 808nm to minimize scattering.

At a trap wavelength of 808nm and a trap depth of $500\mu\text{K}$, we have a scattering rate of 48 photons/second. Using the relation between heating rate and scattering rate, $P_{heat} = k_B T_{recoil} \Gamma_{sc}$, this is equivalent to a heating rate of $17.5\mu\text{K/s}$.

To create the requisite small-waisted, high power tweezer beams in our system, we rely on high-numerical aperture lenses to magnify a collimated beam down to a much smaller image. Specifically, we deploy our optics in a setup show in Fig. 5-4. We first fiber launch the light with a Schäfter+Kirchhoff collimator into a 5x magnification telescope before relaying the resulting beam through a crossed AOD and into the vacuum chamber. Approximately two centimeters from the cavity axis in vacuum, the beam is finally focused down using the high-NA (0.55) in-vacuum aspheric lens discussed in Ch. 2, Sec. 2.3. Taking into account all of the lenses, this system takes a collimated beam with radius 3.3mm, splits it into multiple beams in XY, magnifies

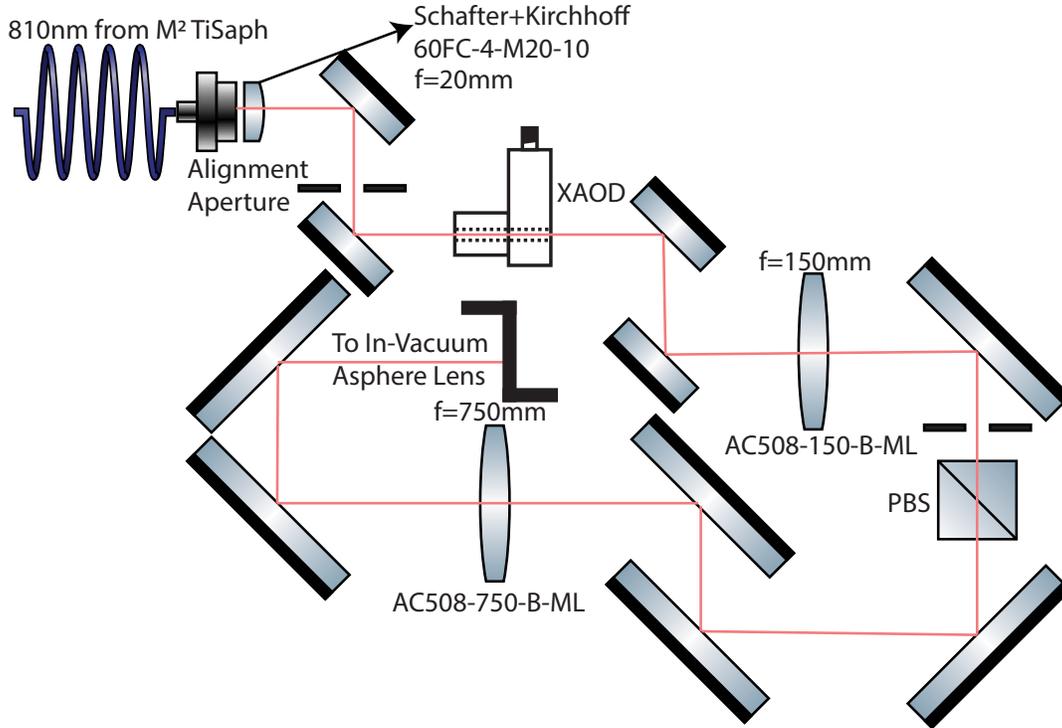


Figure 5-4: Optical schematic of the optics that create the optical tweezers immediately before the vacuum chamber lower viewport. The final lens is the in-vacuum aspheric lens, AL3026, as discussed in Chapter 2.

it by five to a radius of 1.65cm, and finally focuses it down to a waist of $1\mu\text{m}$ with the in-vacuum aspheric lens.

In addition to the correct magnification for our required trap depth, we need a system that creates not just one, but many, optical tweezers. To do this, we chose to create the tweezer traps with a two-dimensional AOD. This allows us to trap multiple atoms along the cavity mode in one-dimension while maintaining the ability to move rows of atoms in and out of the cavity mode with the other dimension of control, as depicted in Fig. 1-3 of the introduction. In addition to the AOD tweezer array, we also intend to implement a stationary two-dimensional tweezer array with a spatial light modulation (SLM) in the future. Essentially, having both traps will allow for AOD tweezer motion to pick-and-place atoms where they are needed in a static SLM tweezer array. The benefits of having both types of traps in the same experiment will be explained in more detail in Ch. 6. The optics for the SLM path will be similar to those in Fig. 5-4 for the AOD tweezers. The only foreseen changes will be to

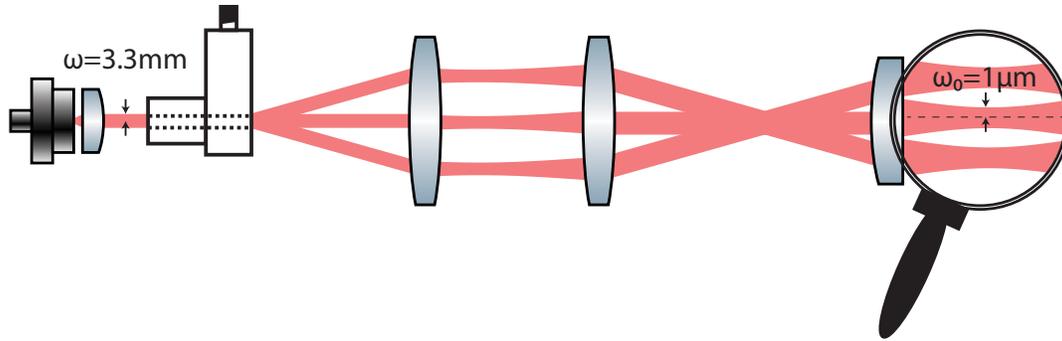


Figure 5-5: Detailed picture of how the many orders of light created in a crossed AOD converge/diverge and focus/defocus through the tweezer array optical path.

swap out one telescope lens to adjust the magnification of these traps to the size of the image created by the SLM in the Fourier plane. We invested in a new type of control system (the FPGA-based OPX) from Quantum Machines to create the RF tones necessary to control the motion of the tweezers in both dimensions with and without this added SLM static trap. Beyond RF synthesis, the OPX is able to, in real time, do calculations using analog inputs and change the analog outputs based on the results. This ability is a critical addition to our shot-based Labscript control software, because Labscript is unable to perform feedback within an experimental shot. Being able to make a measurement and then move or turn off traps or address an atom with resonant light in response is critical to many of our proposed future experiments.

5.2 Cavity Lattice

Because our cavity stabilization laser is red-detuned from the D_1 and D_2 lines in rubidium-87, the light can also be used to create a standing lattice trap along the cavity mode. This trap does not require any components in addition to the cavity apparatus, so we briefly considered doing our first experiments in the lattice. This gave us some redundancy in the event that aligning the optical tweezers to the cavity mode was more time consuming than anticipated. However, tweezer alignment to the cavity mode was relatively quick, so we moved directly to trapping atoms in tweezers.

Even though the lattice trap is not required in our experiment to trap atoms, it can still provide benefits as a secondary trap potential. For instance, the lattice can naturally increase atom localization to the anti-nodes of the cavity probe light, as shown in Fig. 5-6, nudging the atoms to spend more time maximally coupled to the cavity mode. This periodic potential along the cavity axis, Z , is expressed in Eq. 5.4 [47] where U_0 is $4U_{dipole}$ (factor of four coming from the intensity of the standing wave), z is the distance along the cavity axis, and r is the distance from the cavity axis radially. On axis and in an antinode of the cavity ($r = 0$ and $z = 0$), we maximize the trap depth at $-U_0$. However, because our lattice trap is not near-detuned to the atom transitions, we can no longer use the RWA, meaning that now we must take the full form of Eq. 5.3 [47]. Beyond simple trap depth, however, it is also important to quantitatively confirm whether the benefits of added localization outweighs the effects of differential AC Stark shifts between tweezer sites caused by differences in the lattice power. An AC Stark shift differential across tweezer sites complicates single-site addressing and would reduce the fidelity of single and multi-qubit gates. Additionally, a time-varying AC stark shift could cause significant heating in the trapped atoms and decrease the vacuum lifetime. In this section, we will discuss the trapping capabilities of the cavity lattice on its own and will finish with estimations of the lattice effect on tweezer-trapped atoms.

$$U_{lattice}(z) \approx -U_0 \cos kz^2 \left[1 - 2 \left(\frac{r}{\omega_0} \right)^2 - \left(\frac{z}{z_R} \right)^2 \right] \quad (5.4)$$

We implement the lattice using the same optics and frequency stabilization electronics from the experimental cavity PDH lock, as described in detail in Sec. 3.4. The only addition to this path is more optical power, provided by a C-band fiber amplifier, specifically a MARS Series High Power Optical Fiber Amplifier made by Connet Laser Technology. This amplifier takes ~ 5 mW of light from our 1560nm Eternal Fiber laser and amplifies it to up to 1W without broadening the linewidth. Calculating the strength of the lattice potential via the dipole force from Eq. 5.4, we find we have the ability to tune the depth of our lattice trap along the cavity up to

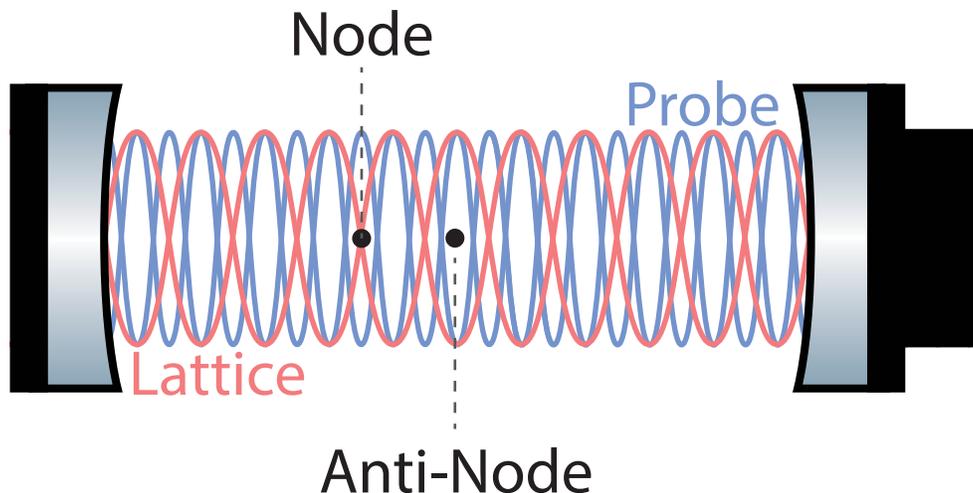


Figure 5-6: Schematic of the inter-cavity trapping lattice (pink) overlapped with the probe light standing wave (purple). Atoms feel the strongest trapping potential at the anti-nodes of the red-detuned 1560nm light. The antinodes of the trapping light line up with the antinodes of the 780nm cavity probe standing wave, allowing for the trapped atoms to be held at the point of strongest coupling along the cavity axis. Atomic coupling to the cavity mode is variable based on the position along the cavity mode.

$T_{lattice} = 28\mu K$ —assuming perfect coupling to the cavity and sending in all 1W of 1560nm power from the fiber amplifier. However, this trap mostly only confines the atoms strongly along the cavity axis, and only weakly in the tangential directions. The trap strength tangential to the cavity axis will be $\sim U_0/\omega_0$ giving us a negligible trap depth in that direction [47]. Thus, we will not be able to trap our atoms in the 1560nm lattice alone. If we still choose to use the 1560nm lattice for additional localization along the cavity axis, time-varying light shifts could become an issue. The power instability of the 1560nm light after the fiber amplifier is $\sim 1\%$. Using Eq. 5.5, which assumes a large detuning between atom and light, we find that this would correspond to a max difference in AC Stark shifts over time of about 56kHz, not enough to significantly impact the fidelity our 10MHz Raman addressing beams discussed below in Sec. 5.3 [50, 47, 94, 38].

$$\Delta E_{Stark} \approx \frac{\Omega_{eg}^2}{4(\omega - \omega_0)} \quad (5.5)$$

5.3 Raman Addressing and Rydberg Transitions

Neither the Raman addressing nor the Rydberg transition systems have yet been installed or tested with atoms in the vacuum chamber. However, the basic optics creating the 6.8GHz Raman beams and the two-photon 480nm Rydberg beam have been set up. Both systems have been designed to allow for internal state manipulation of individual atoms within our array, but they can also be used for global excitation depending on the final optics delivering the beams to the vacuum chamber. This section will give a brief overview of the Raman and Rydberg setups and their intended future uses.

To achieve individual atomic quantum gates, we need hyperfine state control for each atom. A single-site addressing system analogous to our optical tweezers will enable this capability, with the main difference between the two systems being the number of frequency components in the beam. For a Raman transition within the ground state manifold, we need the ability to direct two different optical frequencies with a well-defined frequency difference to each atom. The goal here is to drive transitions like those in Fig. 1-1c). To do this, we replicate the optical setup in Fig. 5-4, but instead using a one-dimensional (1D) AOD and using the lower-NA 480nm in-vacuum aspheric lens. A 1D AOD will allow us to address one atom along the cavity mode at a time, so all operations would need to be performed serially. To create the dual-frequency light for this single-site addressing system, we use a free-space EOM and a chirped Volumetric Bragg Grating (VBG). A diagram of this optical setup is in Fig. 5-7. The phase modulation from the EOM, which creates sidebands at the modulation frequency, is converted to an amplitude modulation in the highly dispersive VBG, which changes the relative phases between the sidebands via a periodically varying index of refraction. Ultimately, this gives us pairs of frequency components in the light with each difference in frequency set by the initial phase modulation frequency—in our case, 6.8GHz to enable Raman transitions between the ground state hyperfine manifolds in rubidium-87. We chose this method of frequency generation because it is passively stable and allows for higher powers than more tra-

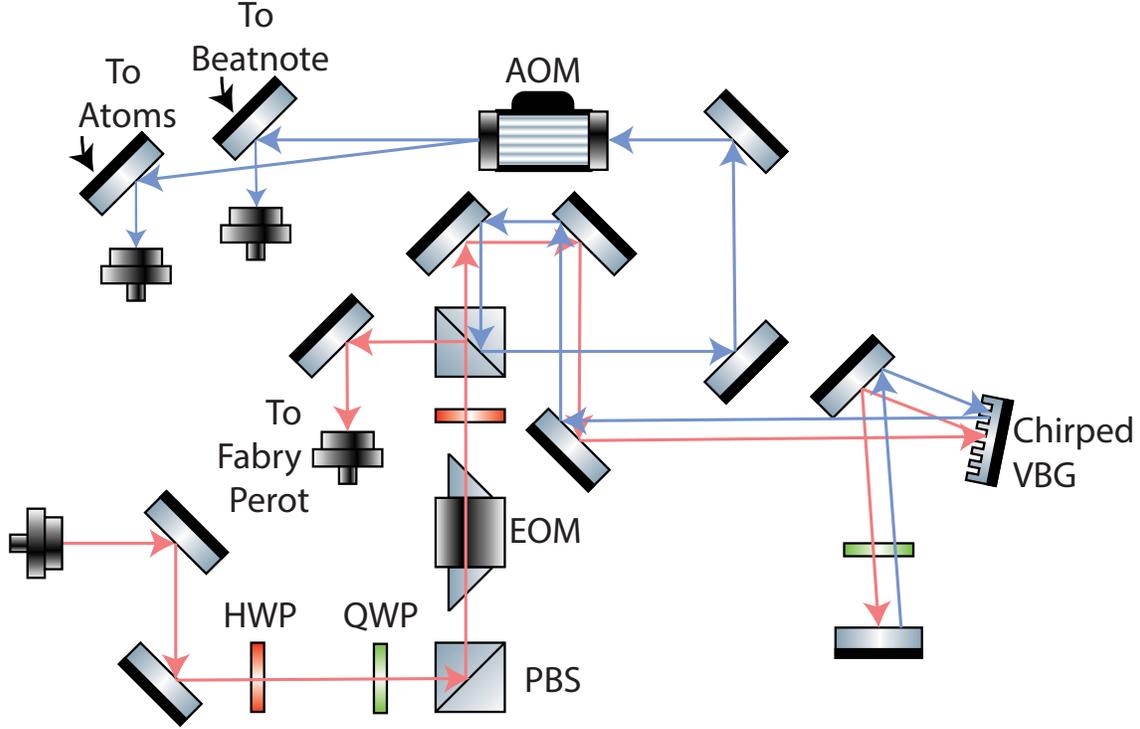


Figure 5-7: Optics for creating the Raman beam pairs at 6.8GHz around 795nm with two passes through a chirped Volumetric Bragg Grating (VBG).

ditional methods of direct amplitude modulation [64]. In addition to an addressing system that allows transitions between the ground state $F=1$ and $F=2$, we would also eventually like to have a second system that allows for transitions between m_F levels within the same hyperfine manifold when the degeneracy is Zeeman split by an external magnetic field. Such a system could be constructed in two main ways: 1) passively using a VBG with a set frequency, likely $\sim 10\text{MHz}$ or 2) actively by using EOMs and AOMs. Because the frequency difference between the two Raman components is smaller in this case, the technical requirements are less stringent.

To reach high-lying Rydberg states in our atoms, we chose a two-photon transition using 780nm and 480nm light, shown schematically in Fig. 5-8. The 780nm light will either be a global beam from the MOT optics or will be sent through one of our Raman addressing systems. This would require setting up a switch to selectively send different either 780nm light or the Raman 795nm light through the 1D AOD. To direct the 480nm light to individual atoms and make the Rydberg excitation

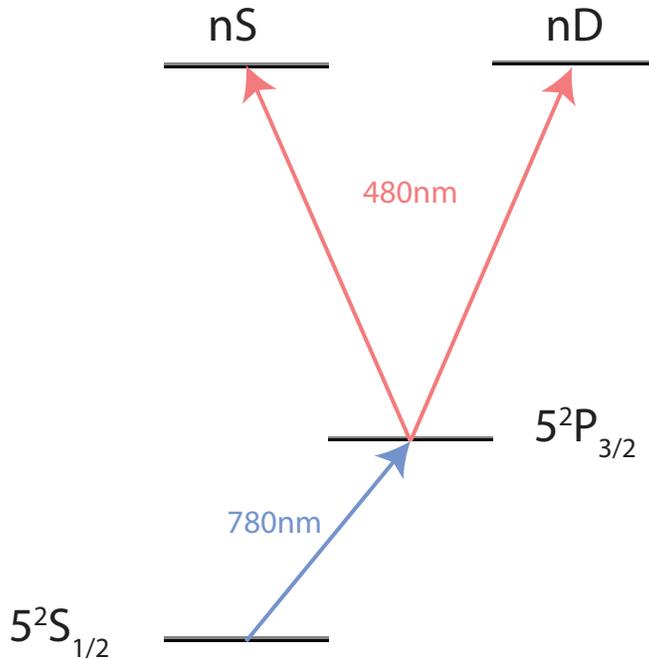


Figure 5-8: Energy diagram showing the two-photon Rydberg transition with 780nm and 480nm light. [84].

local, we plan to build a third AOD system, analogous to the tweezer and Raman addressing systems, but at 480nm and sent through the smaller-NA in-vacuum lens. The 480nm light is provided by a Toptica SHG system, specifically a TA-SHG Pro. This system consists of a 960nm external-cavity diode laser (ECDL) that is sent through a tapered amplifier and doubled in a bowtie cavity to 480nm. We separately stabilize the 960nm light using a PDH lock to our ULE cavity from Stable Laser Systems. Currently, however, we have been unable to lock the 960nm light to the ULE cavity and stabilize the doubling cavity at the same time. The doubling cavity feedback loop feeds back to the 960nm laser current, coupling the two systems in an unpredictable way. Before this system is used on the atoms, an electronics solution to this feedback coupling will need to be implemented and the stability of the 480nm output characterized. When the doubling cavity is stabilized, however, we can achieve a maximum output power of about 225mW at a current setting of 77mA on the 960nm laser and 2300mA on the TA. We expect to be able to create 480nm spots with waists $\sim 5\mu\text{m}$ using the upper, smaller-NA aspheric lens (see Sec. 2.3).

Beyond the optics for the Rydberg excitations, we also carefully considered the large polarizability of Rydberg atoms while building many of the vacuum components in this new experiment. The largest external electric field in our chamber is created by the cavity mirror piezo, which can be driven at up to a 200V difference across the electrodes. The piezo is only $\sim 2.2\text{cm}$ away from the center of the cavity mode, meaning that the atoms could see an electric field of hundreds of volts per meter without any intervention on our part [101]. We took several steps to reduce the amplitude of the field at the atoms: plating the piezo in chrome and gold via vacuum sputtering in MIT ONELab, enclosing the piezo in an additional stainless steel shield, adding temperature control to the cavity mounting ring to help with zeroing the voltage at the piezo, and incorporating four electrodes to allow for active cancellation of remaining fields. Based on past tests [101], we believe that all of these interventions may be necessary to reduce electric fields at our Rydberg atoms to below $\sim 10\text{V/m}$, which is the point at which the Stark Shifts on the $100\text{S}_1/2$ Rydberg state approach $\sim 10\text{MHz}$ and it starts mixing with nearby states [101].

Chapter 6

Outlook

Much of the work documented in this thesis has been technical in nature: designing, simulating, building, and characterizing complex hardware systems. However, many hours were also spent in collaboration walking through the future scientific capabilities of the system and how to best achieve them. In fact, during the height of the Covid-19 pandemic, this was the most productive work we could do, separated from the lab but able to discuss over Zoom. This entire upgrade process has been a conceptual bridge taking Rubidium Lab from fundamental atomic physics research in laser cooling and quantum degeneracy to more application-focused with new directions in quantum information science and quantum computing. As such, experiments allowing proof-of-principle demonstrations of key requirements for quantum computing, such as error correction, featured heavily in our discussions. In this final chapter, we will discuss the high-level outlook for the experiment and describe a couple of exciting demonstrations we can achieve with our cavity-coupled neutral atom array system: cavity-mediated surface code error correction and fast non-destructive atom array loading.

6.1 Error Correction

Once all supporting trapping and addressing optics are up and running as described in Ch. 5, the lab can move on to proof-of-principle error correction experiments. As discussed in the introduction, Ch. 1, the standard atomic state readout methods of

fluorescence and blowing away atoms in neutral atom arrays are either too slow to allow for sustained computation within atomic coherence times or are destructive, causing the loss of trapped atoms in the process. In order for Rydberg atom arrays to compute fault-tolerantly one day, they must be error corrected and to successfully error correct a system, we need to be able to quickly readout qubit states and then feedback to correct those errors. With the integration of our cavity for readout and our real-time feedback to individual addressing of atoms, we will be able to do just that.

A prototypical and near-term example of an error correcting scheme that could be implemented with such readout is a surface error correction code. In such a code, one logical qubit is encoded in a two-dimensional surface of physical qubits of two types: data and syndrome. Each cycle of the code involves checking the parity of adjacent data qubits via two-qubit measurements ("stabilizers") in each the X and Z bases of the logical Hilbert space. In a Rydberg atom array, these stabilizer measurements are implemented with Rydberg-blockade-mediated gates. The results of these measurements are then stored in the syndrome qubits' states until the errors can be read, corrected, and the system can continue on to the next step in the calculation. So far, only one cycle of such a code has been physically implemented in a Rydberg atom array system, and the limiting factor in continuing the correction onto larger numbers of cycle was readout [16].

Thus, we must devise a way to determine atomic states without perturbing them. As we explained in the introduction, we can achieve this via off-atomic-resonance probes through a cavity mode, either in a fluorescence or transmission configuration. Simply put in Fig. 6-1, if the atom is on-resonance with the cavity resonance, the coupled system's resonance will shift and if the atom is not on-resonance with the cavity, the coupled system's resonance will be the bare cavity resonance. For a real level-structure scheme that works with Rb87, see Fig. 6-2.

However, the problem then becomes a lack of parallelism in the measurement. Any light sent through a cavity mode will interact with all atoms sitting within the Gaussian mode. This means that if *any* atom is in the $|1\rangle$ state, the atom-cavity res-

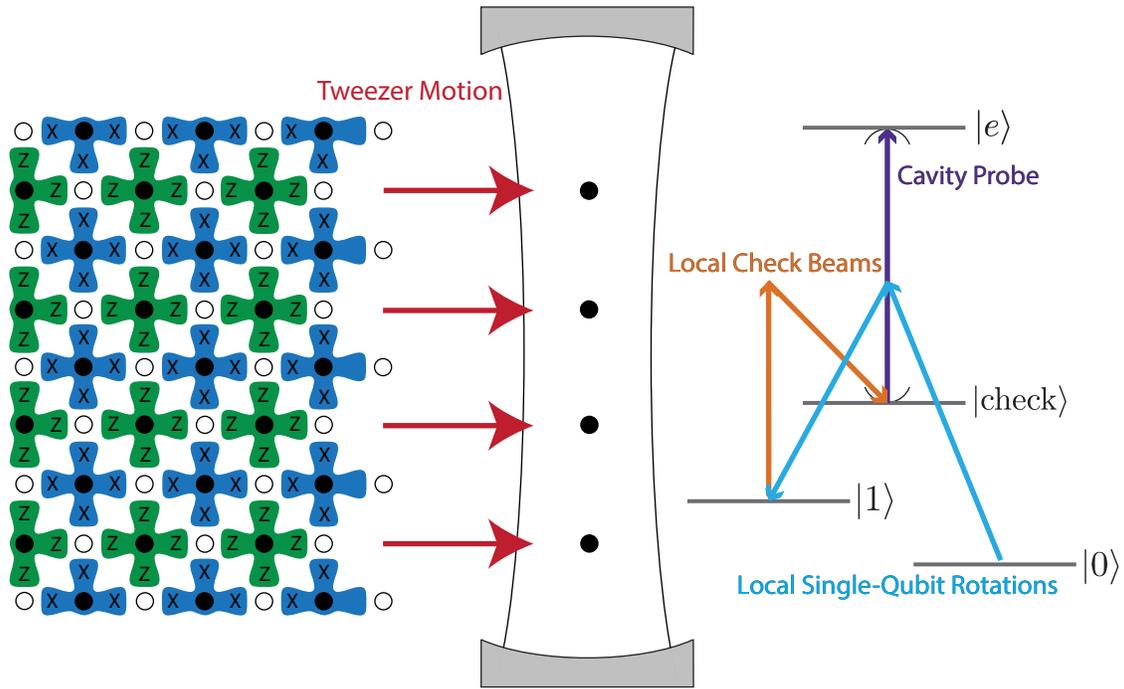


Figure 6-1: A schematic representation of the future of Rydberg Atom Arrays, coupled to an optical cavity. A static optical tweezer array (SLM) holds the surface code patch physical qubits outside of the cavity mode while moving tweezer traps (crossed AODs) constantly move syndrome qubits in and out of the cavity mode for error correction and the reloading of traps that lost an atom due to vacuum lifetime. On the left is a graph representation of a two-dimensional surface error correcting code, with data qubits (white) and syndrome qubits (black) connected via stabilizer measurements in the X (blue) and Z (green) bases. On the right is a simplified four-level atom in a cavity. The atom will shift the resonance of the coupled atom-cavity system if it's in state $|check\rangle$, but not if it's in state $|0\rangle$ or state $|1\rangle$, allowing for a determination on atomic state based on the probe transmitted power.

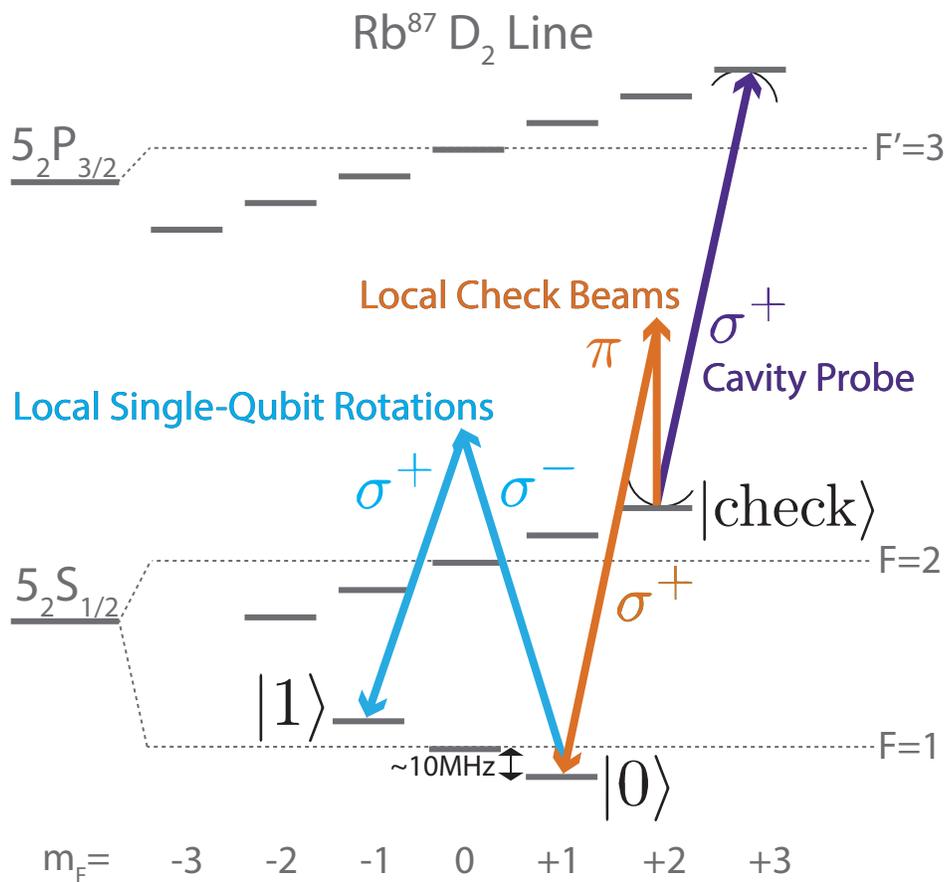


Figure 6-2: Physical internal state scheme for logical states in rubidium-87 that can be used for error diagnosis with a cavity. This scheme requires a small external magnetic field to raise the hyperfine degeneracy, allowing access to different m_F levels with Raman transitions.

onance will be shifted, but we will have no information on *which* atom(s) caused the shift. Naively, one could measure each atom individually, working through all syndrome qubits in a purely sequential manner. However, this means the measurement scales with N , making it impractical as atom number continues to increase in neutral atom array systems. A more clever way to speed up the measurements and make them more scalable is to implement a binary search algorithm to find the location of individual errors within the syndrome qubits. To do this, we must have at least one method with which to "turn-off" our atoms' couplings to the cavity mode. In our system, we have two: hyperfine shelving and tweezer motion. Depending on which is technically simpler to implement, we can choose to individually change atoms' internal states to those outside of the syndrome qubit logical manifold and thus also out of resonance with the readout optical cavity. Similarly, we can simply move the tweezers of atoms that we want to turn off to a position several standard deviations away from the center of the optical cavity mode. If in the future we find any cross talk or noise induced in the measurement when only one of these procedures is followed, we can also use both shelving and motion to bring the atom-coupling even closer to zero.

With the ability to dynamically change cavity coupling for each individual syndrome qubit in our neutral atom array, we can fully implement a binary search by probing the cavity system with increasingly smaller groups of atoms coupled to the cavity at a given time, as depicted in Fig. 6-3. The binary search immediately gives us a logarithmic speedup over sequential atomic state measurements, but a separate factor entirely makes this readout method faster than other options. In order for a surface error correction code to work, for it to remove more errors from a logical qubit than it creates, the total errors occurring across the system must happen $p \lesssim 1\%$ of the time. This means that any functioning surface code will have what are called biased errors. Most of the time during the first measurement of the syndrome qubits, there will be no errors. This gives the overall measurement scheme a significant reduction in the average time to diagnose an error to

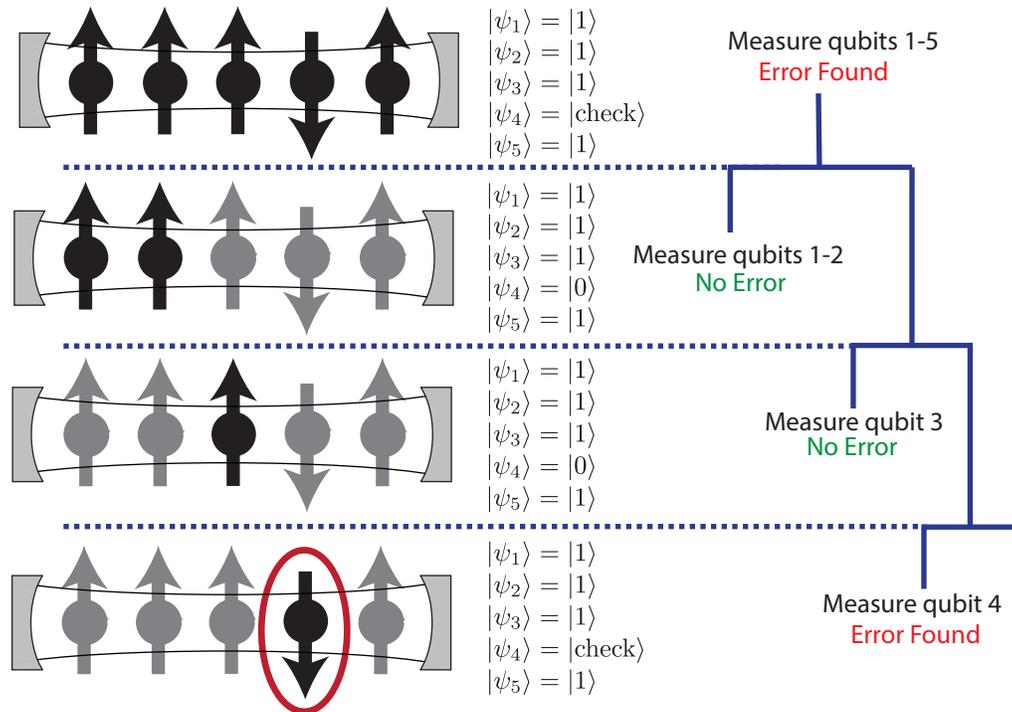


Figure 6-3: The procedure for finding the location of an error (depicted here as the atom in the "down" state) in faster than serial times using the nondestructive, off-resonant cavity readout in a binary search scheme. This requires moving atoms out of the logical basis with Raman addressing beams to "shelve" them when the search is not in that branch of the tree, and then moving them back into the logical basis to continue the search until the error has been found.

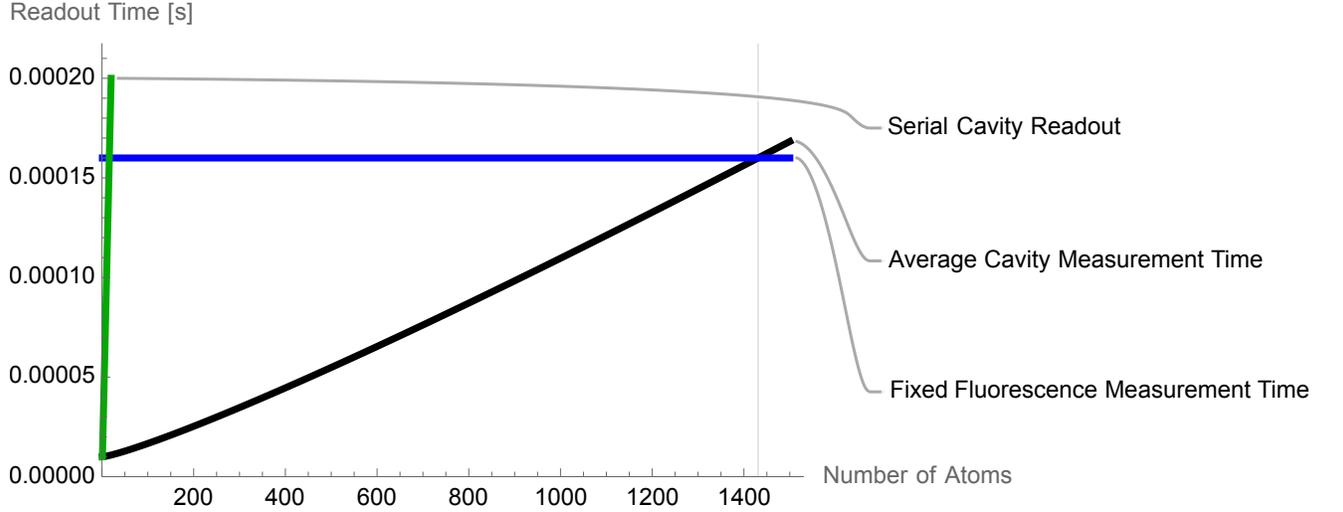


Figure 6-4: The scaling of cavity error diagnosis speed with atom number for a cavity binary search method, sequential cavity search, and a parallel fluorescence readout.

$$T_{avg} = 3\tau[1 + Np \log_2(N)]. \quad (6.1)$$

That is, most of the time ($1 - p$ of the time), we will only need to complete one probe through the cavity to confirm that the code patch has no errors. With one measurement being a function of the ringdown time of the cavity, this means that most code cycles of the surface error correction code will only require a diagnosis step with $T_{diag} \sim 10\mu s$. Taking into account both the binary logarithmic speedup and the biased error speedup, we find that binary cavity error diagnosis scales better with atom number than any other current readout method until we reach ~ 1500 atoms, a value that Rydberg atom array systems are currently far from reaching. Fig. 6-4 demonstrates this crossover point and also compares to how the binary search with biased errors significantly out-performs the naive sequential state measurement with a cavity.

The ability to more quickly find errors in a group of syndrome qubits removes one of the largest barriers to surface code error correction in Rydberg atom array systems. However, the capabilities discussed in this section assumed that all errors associated with creating the surface code and performing calculations on it are already below the threshold error for a surface code, $\sim 1\%$. In reality, most Rydberg atom array

systems are close to—but not quite at—this level of fidelity. Although single-qubit gates in the form of π and $\pi/2$ pulses have extremely high fidelities, approaching 99.9%, state-of-the-art Rydberg CNOT gates have errors $\sim 3\%$. This is a reminder of the complexity of these systems and that consistent progress needs to be made on all fronts to reach the field's goal of fault-tolerant computing.

6.2 Atom Array Loading

Seven years ago, deterministically loading a filled two-dimensional array of neutral atoms kicked off an explosion of developments in the Rydberg atom array community [10, 35]. The ability to create a fully-filled array of trapped atoms enabled subsequent physics experiments to have tunable geometries rather than relying on the statistical loading of traps—a dream system for both condensed matter and quantum computing applications. However, although some progress has been made in loading atoms into set traps from a reservoir [78], current tweezer loading methods all rely on fluorescence imaging to determine the location of trapped atoms before re-arranging or moving into their final trapped location. For the same reasons given in the introduction for atomic state readout speeds, a scaled up system will require faster loading, and a cavity can provide it.

Loading cannot be sped up as much as error correction can be as described in the previous section, because a binary search does not help. This is due to the "error" state of an unfilled trap being no atom, rather than an atomic state that can shift the energy of the cavity-atom system. Thus, instead of applying "check" pulses onto trap sites to ensure the error state is on resonance with the cavity and the non-error state is not, loading will require the opposite order. Before moving a tweezer into the cavity to check for an atom, we will apply pulses to ensure the present atoms will shift the cavity resonance. Thus, the full procedure is schematically represented in Fig. 6-5. We will require a constant reservoir of atoms outside of the cavity mode, likely in a large-waist dipole trap. We will use tightly focused beams through a two-dimensional AOD to pick up individual atoms from this reservoir and move through

the cavity, one at a time. Before the tweezer reaches the cavity, we will need to ensure that the atom is on-resonance with the cavity, meaning that all the atoms are in $|5^2S_{1/2}, F = 2\rangle$, which should be the case in our system after our MOT. While moving tweezers through the cavity to check occupation, we must ensure that only one tweezer is coupled to the mode at a time, so the separation between the tweezers must be $\gtrsim 3 \times \omega_0 = 135\mu m$. For each time period that one tweezer is coupled to the cavity mode, $t = \frac{135\mu m}{v}$, we will send a probe light pulse through the cavity at the bare cavity resonance frequency while gating open our SPCM and counting the transmission through the cavity. A low count at the cavity transmission after a few ring down times and thus a shifted resonance indicates a filled trap while a high count at the cavity transmission and thus no shift indicates an empty trap. Based on the results of this measurement for each tweezer, we can turn off the empty traps, removing them from the train of atoms. The remaining traps in the train will be moved to a static two-dimensional SLM tweezer array, also sitting well outside the cavity mode.

By assuming tweezer speeds demonstrated in other atom array experiments, we can produce a good upper bound approximation of loading times using our cavity checking method. Recent atom array experiments have demonstrated atom motion across a distance of $\sim 100\mu m$ at speeds $\lesssim 0.55\mu m/\mu s$ [16]. At these speeds, we could pull a single tweezer through the full cavity mode (defined conservatively as $6 \times \omega_0$, or three standard deviations on either side of the center of the cavity mode) in about $500\mu s$. To fill a 10×10 array with this method, then, it would take $\sim 100ms$, on the order of parallel fluorescence readout methods. However, because the cavity readout method for filled traps is purely sequential, there is an atom number crossover point at which parallel fluorescence readout always wins when filling an array from scratch. Even so, filling an array from scratch is not the primary benefit to using fast cavity readout for array loading. During error correction, if we determine an atom has been lost, we can actively replace the atom and continue operating the array without the need to re-load a full array. Assuming a vacuum lifetime $\sim 1s$, this means that on average, we would need to replace 100 atoms/s, which can easily be done within

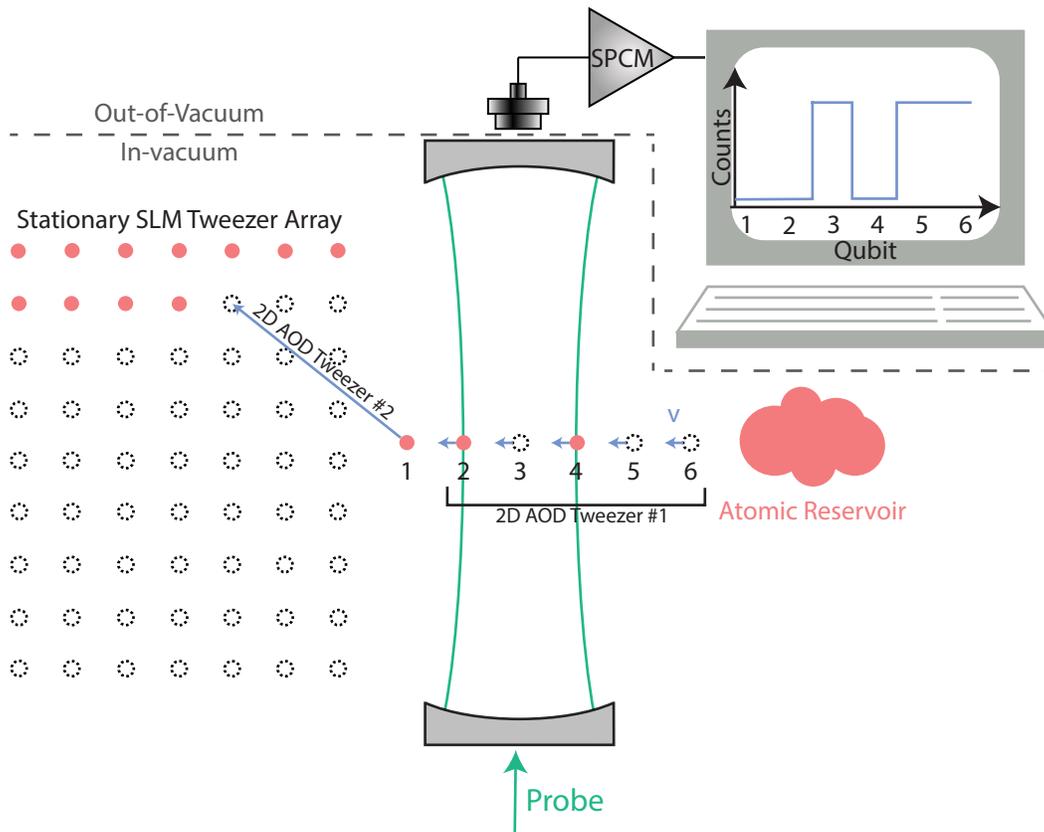


Figure 6-5: Schematic of the cavity-mediated array loading procedure from scratch. An atomic reservoir sits outside the cavity mode, either in a large ensemble dipole trap or a MOT. A combination of 1D and 2D AOD tweezers pick up atoms from the reservoir, drag them through the cavity mode while a probe beam turns on. Based on the transmission counts at the output of the cavity on an SPCM, the tweezer is moved to the next empty spot in a stationary SLM tweezer array on the other side of the cavity mode or the tweezer is turned off, and the process continues until the 2D array is 100% filled.

hundreds of milliseconds with our cavity readout and motion time of $\sim 500\mu\text{s}/\text{atom}$.

Appendix A

Cavity Mirror Cleaning

When building a cavity experiment, the cleanliness of high-finesse cavity mirrors is critical. One piece of dust on a mirror that overlaps with the cavity mode can decrease ring down times by orders of magnitude. Additionally, during vacuum chamber baking it is possible, though rare, for out-gassed substances to coat the mirror surface and disrupt the high-reflectivity coating. Depending on the contaminant, different procedures are used to clean the mirrors. This section will detail the chemicals and methods we used to maintain the high finesse of our crossed cavities at various points during installation.

Before cleaning, it is best to properly align the cavity at hand and optimize the TEM₀₀ coupling. This allows us to monitor the transmission and finesse (by ring down time measurements) of the cavity after each stage of cleaning. The properties of the transmission signal will indicate whether the mirrors are getting cleaner or less clean.

The most common way to clean precision optics is with solvents. In optics labs this usually means isopropyl alcohol, acetone, methanol, or even distilled water. These different solvents each have different benefits and drawbacks, from strength and drying speed, to safety. For our cavities, when cleaning the mirrors for dust removal, we use isopropyl alcohol, which dries quickly with little residue and is not as harmful to our health as acetone and methanol are. We use acetone, methanol, and water if we suspect that our mirrors are coated in a substance (immediately following our

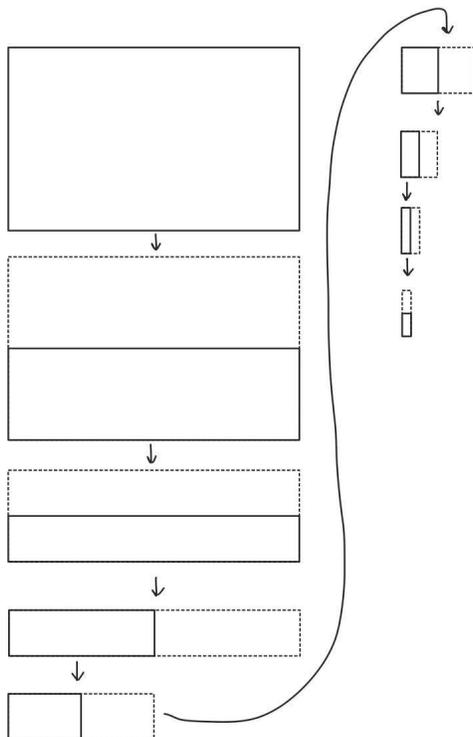


Figure A-1: How to fold a piece of lens tissue for cleaning our small cavity mirrors.

first installation attempt we assumed this was the case). At one point of the mirror recovery effort after the first vacuum installation attempt, we also cleaned the mirrors by soaking and/or carefully sonicating them in dichloromethane (DCM) and toluene. After these cleaning methods did not work, we knew that cleanliness of the mirror surface was not the reason for the mirror degradation.

For future cavity installations or adjustments, we have tabulated more detailed instructions on how to clean our cavity mirrors with solvents below. There may be other technical procedures that work for cleaning, but these procedures worked best after much trial and error to install the current system.

1. Wearing clean, powder-free gloves, ensure the mirror to be cleaned is secured on its mount and the mount is screwed in tightly to the cavity ring with a vacuum-clean Allen wrench.
2. Prepare a piece of clean, dry lens tissue paper by folding it as illustrated in figure A-1 and securing it in the jaws of a vacuum-clean hemostat, holding the

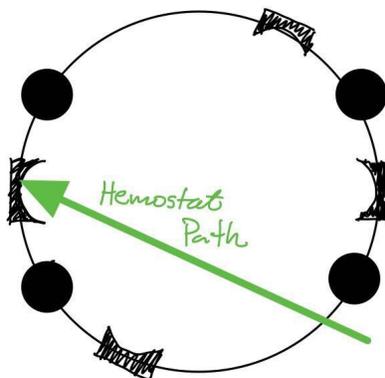


Figure A-2: A potential path through the crossed cavity apparatus for mirror cleaning.

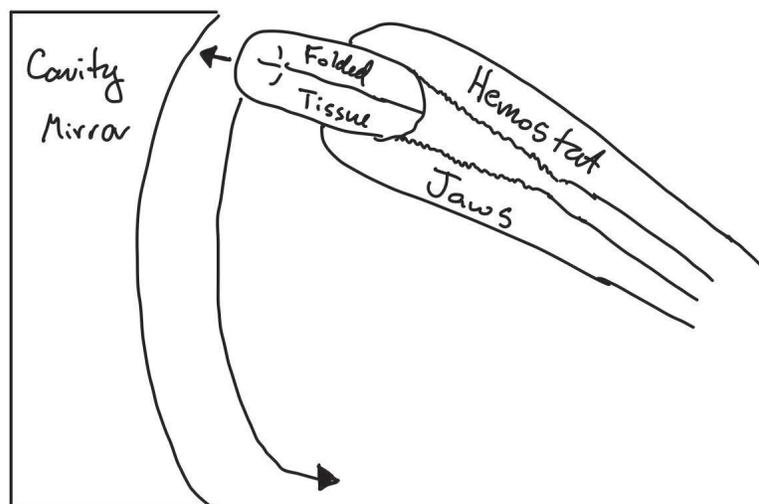


Figure A-3: How to wipe a cavity mirror with a properly folded tissue held by a hemostat.

final fold together. This folding technique ensures that the lens tissue is stiff for cleaning with moderate force and is small enough to stay in contact with the curved surface of the mirror during the entirety of the wipe.

3. Using a chemical squeeze bottle (wrapped in foil to make the process cleaner and easier), drip \approx five drops of the solvent onto the end of the piece of tissue. If the tissue is dripping, saturated, or shiny, allow some solvent to evaporate before cleaning the mirror. Waving the hemostat in the air several times can help speed up this process. If too much solvent is on the tissue, it will leave behind large streaks during the wipe that will likely make the mirror dirtier. If not enough solvent is on the tissue, there is a high risk of scratching the mirror.
4. Find a stable position where the hemostat can reach the mirror to be cleaned with room for the full wiping motion without hitting or disturbing other components of the crossed cavity apparatus. Using two hands to stabilize the hemostat is a good idea. See Fig. A-2 for an aerial view of such a position.
5. Starting about 1mm below the top of the mirror (where the curve of the glass has already begun), press the wet, folded piece of tissue firmly to the mirror surface and drag downward as slowly as possible in a continuous arc shape matching the curvature of the mirror. Be extremely careful not to touch any part of the mirror surface with the metal hemostat. This will severely damage the high-reflectivity coating. See Fig. A-3 for a pictorial description of this wiping motion.
6. Slowly retract the hemostat from the crossed cavity apparatus and dispose of the used tissue in a chemical hazardous waste container, as appropriate. Check the transmission of the cavity (signal amplitude and ring down time) and repeat the full cleaning procedure as necessary.

Another common way to clean cavity mirrors is with a polymer or "nail polish" cleaner. These polymer cleaners come in a small bottle with a brush lid for application. Generally, the polymer is dripped and dragged onto the optical surface with the

brush (without ever touching the optical surface with the brush bristles themselves) and then allowed to dry before being carefully peeled off the optical surface. This polymer solution contains strong solvents for dissolving contaminants and reduces the risk of scratching the optical surface as one might while wiping with a solvent. More detailed instructions on how to clean with a polymer optics cleaner are enumerated below.

1. Place the mirror on a flat, vacuum clean work surface with the optical surface to be cleaned facing upward. If the cavity mirror to be cleaned is connected to the cavity ring by piezo wires, build a platform out of Thorlabs components that is the proper height to keep the mount as close to the cavity ring as possible during cleaning. Make sure to wrap this platform in clean foil before placing the mirror on it. Be gentle with the piezo wires, as they can easily rip off the piezo electrodes (they are plated on the surface of the piezo ceramic).
2. Hold the mirror mount in place with one gloved hand while using the polymer cleaner brush to place a drop of polymer solution on the middle of the cavity mirror.
3. Leave the drop to pre-dry for about twenty minutes.
4. Place the end of a string of unwaxed dental floss on the drop of polymer solution. Press on top of the floss gently with a gloved finger to just submerge it in the polymer solution.
5. Using the polymer applicator brush, add additional polymer solution to secure the floss and cover the majority of the optical surface of the mirror.
6. Leave the solution to dry on the mirror, ~ 2.5 hours.
7. Periodically test to see if the polymer is dry by touching the surface extremely lightly with a gloved finger. If the polymer easily depresses, it is not yet dry.

8. When the polymer solution is fully dry, hold the mirror edges with one gloved hand and use the other gloved hand to peel off the dried polymer solution. If properly applied and dried, it should all come off easily in one motion.
9. Re-mount the cavity mirror (being careful not to trap any piezo wires between the mirror mount and cavity ring), and re-align to the TEM00 mode as per the instructions for full cavity alignment described in Sec. 3.2.2.
10. Measure the ring down time of the cavity, determine whether the mirror requires more cleaning, and repeat these steps as necessary.

Bibliography

- [1] Ibm builds the world’s largest dilution refrigerator for quantum computers – engineering exploration.
- [2] Ibm unveils world’s largest quantum computer at 433 qubits - new scientist.
- [3] It’s colossal: Creating the world’s largest dilution refrigerator.
- [4] The first hundred years. *Nature Publishing Group*, 2013.
- [5] Rajeev Acharya, Igor Aleiner, Richard Allen, Trond I. Andersen, Markus Ansmann, Frank Arute, Kunal Arya, Abraham Asfaw, Juan Atalaya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Joao Basso, Andreas Bengtsson, Sergio Boixo, Gina Bortoli, Alexandre Bourassa, Jenna Bovaird, Leon Brill, Michael Broughton, Bob B. Buckley, David A. Buell, Tim Burger, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Josh Cogan, Roberto Collins, Paul Conner, William Courtney, Alexander L. Crook, Ben Curtin, Dripto M. Debroy, Alexander Del Toro Barba, Sean Demura, Andrew Dunsworth, Daniel Eppens, Catherine Erickson, Lara Faoro, Edward Farhi, Reza Fatemi, Leslie Flores Burgos, Ebrahim Forati, Austin G. Fowler, Brooks Foxen, William Giang, Craig Gidney, Dar Gilboa, Marissa Giustina, Alejandro Grajales Dau, Jonathan A. Gross, Steve Habegger, Michael C. Hamilton, Matthew P. Harrigan, Sean D. Harrington, Oscar Higgott, Jeremy Hilton, Markus Hoffmann, Sabrina Hong, Trent Huang, Ashley Huff, William J. Huggins, Lev B. Ioffe, Sergei V. Isakov, Justin Iveland, Evan Jeffrey, Zhang Jiang, Cody Jones, Pavol Juhás, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Tanuj Khattar, Mostafa Khezri, Mária Kieferová, Seon Kim, Alexei Kitaev, Paul V. Klimov, Andrey R. Klots, Alexander N. Korotkov, Fedor Kostritsa, John Mark Kreikebaum, David Landhuis, Pavel Laptev, Kim Ming Lau, Lily Laws, Joonho Lee, Kenny Lee, Brian J. Lester, Alexander Lill, Wayne Liu, Aditya Locharla, Erik Lucero, Fionn D. Malone, Jeffrey Marshall, Orion Martin, Jarrod R. McClean, Trevor McCourt, Matt McEwen, Anthony Megrant, Bernardo Meurer Costa, Xiao Mi, Kevin C. Miao, Masoud Mohseni, Shirin Montazeri, Alexis Morvan, Emily Mount, Wojciech Mruczkiewicz, Ofer Naa-man, Matthew Neeley, Charles Neill, Ani Nersisyan, Hartmut Neven, Michael Newman, Jiun How Ng, Anthony Nguyen, Murray Nguyen, Murphy Yuezhen Niu, Thomas E. O’Brien, Alex Opremcak, John Platt, Andre Petukhov, Rebecca Potter, Leonid P. Pryadko, Chris Quintana, Pedram Roushan, Nicholas C.

- Rubin, Negar Saei, Daniel Sank, Kannan Sankaragomathi, Kevin J. Satzinger, Henry F. Schurkus, Christopher Schuster, Michael J. Shearn, Aaron Shorter, Vladimir Shvarts, Jindra Skruzny, Vadim Smelyanskiy, W. Clarke Smith, George Sterling, Doug Strain, Marco Szalay, Alfredo Torres, Guifre Vidal, Benjamin Villalonga, Catherine Vollgraft Heidweiller, Theodore White, Cheng Xing, Z. Jamie Yao, Ping Yeh, Juhwan Yoo, Grayson Young, Adam Zalcman, Yaxing Zhang, and Ningfeng Zhu. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* 2023 614:7949, 614:676–681, 2 2023.
- [6] Fangzhao Alex An, Anthony Ransford, Andrew Schaffer, Lucas R Sletten, John Gaebler, James Hostetter, and Grahame Vittorini. High fidelity state preparation and measurement of ion hyperfine qubits with $i > 1$. 2.
- [7] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G.S.L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature* 2019 574:7779, 574:505–510, 10 2019.
- [8] C H Baldwin, B J Bjork, J P Gaebler, D Hayes, and D Stack. Subspace benchmarking high-fidelity entangling operations with trapped ions. 2020.
- [9] Katrina Barnes, Peter Battaglino, Benjamin J. Bloom, Kayleigh Cassella, Robin Coxe, Nicole Crisosto, Jonathan P. King, Stanimir S. Kondov, Krish Kotru, Stuart C. Larsen, Joseph Lauigan, Brian J. Lester, Mickey McDonald, Eli Megidish, Sandeep Narayanaswami, Ciro Nishiguchi, Remy Notermans, Lucas S. Peng, Albert Ryou, Tsung Yao Wu, and Michael Yarwood. Assembly and coherent control of a register of nuclear spin qubits. *Nature Communications* 2022 13:1, 13:1–10, 5 2022.
- [10] Daniel Barredo, Sylvain De Léséleuc, Vincent Lienhard, Thierry Lahaye, and Antoine Browaeys. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354:1021–1023, 11 2016.

- [11] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Deterministic quantum teleportation of atomic qubits. *Nature* 2004 429:6993, 429:737–739, 6 2004.
- [12] David Beckman, Amalavoyal N. Chari, Srikrishna Devabhaktuni, and John Preskill. Efficient networks for quantum factoring. *Physical Review A*, 54:1034, 8 1996.
- [13] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22:563–591, 5 1980.
- [14] H. E. Bennett and J. O. Porteus. Relation between surface roughness and specular reflectance at normal incidence. *JOSA, Vol. 51, Issue 2, pp. 123-129*, 51:123–129, 2 1961.
- [15] Eric D Black. An introduction to pound-drever-hall laser frequency stabilization. 2001.
- [16] Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T. Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. *Nature* 2022 604:7906, 604:451–456, 4 2022.
- [17] B Brandstätter, A Mcclung, and K Schüppert. Integrated fiber-mirror ion trap for strong ion-cavity coupling articles you may be interested in. *Cite as: Rev. Sci. Instrum*, 84:123104, 2013.
- [18] Boris Braverman. Cavity quantum electrodynamics with ensembles of ytterbium-171. 2018.
- [19] Sergey Bravyi, Oliver Dial, Jay M. Gambetta, Darió Gil, and Zaira Nazario. The future of quantum computing with superconducting qubits. *Journal of Applied Physics*, 132:160902, 10 2022.
- [20] Manuel Brekenfeld, Dominik Niemietz, Joseph Dale Christesen, and Gerhard Rempe. A quantum network node with crossed optical fibre cavities. *Nature Physics* 2020 16:6, 16:647–651, 4 2020.
- [21] Colin D. Bruzewicz, Robert McConnell, John Chiaverini, and Jeremy M. Sage. Scalable loading of a two-dimensional trapped-ion array. *Nature Communications* 2016 7:1, 7:1–6, 9 2016.
- [22] Marko Cetina, Alexei Bylinskii, Leon Karpa, Dorian Gangloff, Kristin M Beck, Yufei Ge, Matthias Scholz, Andrew T Grier, Isaac Chuang, and Vladan Vuletić. One-dimensional array of ion chains coupled to an optical cavity. *New Journal of Physics*, 15:53001, 2013.

- [23] Mei Chang and Lutgard C. De Jonghe. Whisker growth in reduction of oxides. *Metallurgical Transactions B*, 15:685–694, 12 1984.
- [24] Disheng Chen, Nikolay Zheludev, and Wei bo Gao. Building blocks for quantum network based on group-iv split-vacancy centers in diamond. *Advanced Quantum Technologies*, 3, 2 2020.
- [25] Lilian Childress and Ronald Hanson. Diamond nv centers for quantum computing and quantum networks. *Mrs Bulletin*, 38:134–138, 2 2013.
- [26] Yue-Sum Chin, Matthias Steiner, and Christian Kurtsiefer. Polarization gradient cooling of single atoms in optical dipole traps. *PHYSICAL REVIEW A*, 96:33406, 2017.
- [27] Carsten Christensen, Roger De Reus, and Siebe Bouwstra. Tantalum oxide thin films as protective coatings for sensors. *Journal of Micromechanics and Microengineering*, 9:113, 6 1999.
- [28] Costas Christoforou, Corentin Pignot, Ezra Kassa, Hiroki Takahashi, and Matthias Keller. Enhanced ion–cavity coupling through cavity cooling in the strong coupling regime. *Scientific Reports 2020 10:1*, 10:1–7, 9 2020.
- [29] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the doppler limit by polarization gradients: simple theoretical models. *JOSA B, Vol. 6, Issue 11, pp. 2023-2045*, 6:2023–2045, 11 1989.
- [30] David P Divincenzo. The physical implementation of quantum computation.
- [31] R. W.P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B Photophysics and Laser Chemistry*, 31:97–105, 1983.
- [32] Rainer Dumke, Zehuang Lu, and John Close. Paper • open access you may also like roadmap on quantum optical systems. 2014.
- [33] John Dunlavy. Wide range tunable transmitting loop antenna. 10 1967.
- [34] M. Ebert, M. Kwon, T. G. Walker, and M. Saffman. Coherence and rydberg blockade of atomic ensemble qubits. *Physical Review Letters*, 115:093601, 8 2015.
- [35] Manuel Endres, Hannes Bernien, Alexander Keesling, Harry Levine, Eric R. Anschuetz, Alexandre Krajenbrink, Crystal Senko, Vladan Vuletic, Markus Greiner, and Mikhail D. Lukin. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354:1024–1027, 11 2016.

- [36] Alexander Erhard, Hendrik Poulsen Nautrup, Michael Meth, Lukas Postler, Roman Stricker, Martin Stadler, Vlad Negnevitsky, Martin Ringbauer, Philipp Schindler, Hans J. Briegel, Rainer Blatt, Nicolai Friis, and Thomas Monz. Entangling logical qubits with lattice surgery. *Nature* 2021 589:7841, 589:220–224, 1 2021.
- [37] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21:467–488, 6 1982.
- [38] C.J. Foot. Atomic physics - c.j. foot - oxford university press. *Oxford University Press*, 2004.
- [39] Luigi Frunzio, Andreas Wallraff, David Schuster, Johannes Majer, and Robert Schoelkopf. Fabrication and characterization of superconducting circuit qed devices for quantum computation. *IEEE Transactions on Applied Superconductivity*, 15:860–863, 6 2005.
- [40] Zhuo Fu, Peng Xu, Yuan Sun, Yang Yang Liu, Xiao Dong He, Xiao Li, Min Liu, Run Bing Li, Jin Wang, Liang Liu, and Ming Sheng Zhan. High-fidelity entanglement of neutral atoms via a rydberg-mediated single-modulated-pulse controlled- phase gate. *Physical Review A*, 105:042430, 4 2022.
- [41] J. Gallego, W. Alt, T. Macha, M. Martinez-Dorantes, D. Pandey, and D. Meschede. Strong purcell effect on a neutral atom trapped in an open fiber cavity. *Physical Review Letters*, 121:173603, 10 2018.
- [42] Jay M. Gambetta, Jerry M. Chow, and Matthias Steffen. Building logical qubits in a superconducting quantum computing system. *npj Quantum Information* 2017 3:1, 3:1–7, 1 2017.
- [43] Dorian Gangloff, Molu Shi, Tailin Wu, Alexei Bylinskii, Boris Braverman, Michael Gutierrez, Rosanna Nichols, Junru Li, Kai Aichholz, Marko Cetina, Leon Karpa, Branislav Jelenkovic Jelenkovic, Isaac Chuang, Vladan Vuleticvuletic, 11 H Loh, Y.-J Lin, I Teper, M Cetina, J Simon, and J K Thompson. Preventing and reversing vacuum-induced optical losses in high-finesse tantalum (v) oxide mirror coatings. 23, 2015.
- [44] Boxuan Gao, Jeroen Beeckman, Kristiaan Neyts, and John Puthenparampil George. Design, fabrication and characterization of a distributed bragg reflector for reducing the x00e9;tendue of a wavelength converting system. *Optics Express, Vol. 28, Issue 9, pp. 12837-12846*, 28:12837–12846, 4 2020.
- [45] T M Graham, Y Song, J Scott, C Poole, L Phuttitarn, K Jooya, P Eichler, X Jiang, A Marra, B Grinkemeyer, M Kwon, ‡ M Ebert, J Cherek, M T Lichtman, M Gillette, J Gilbert, D Bowman, T Ballance, C Campbell, E D Dahl, O Crawford, N S Blunt, B Rogers, T Noel, and M Saffman. Demonstration of multi-qubit entanglement and algorithms on a programmable neutral atom quantum computer. 2022.

- [46] Graham P. Greve, Chengyi Luo, Baochen Wu, and James K. Thompson. Entanglement-enhanced matter-wave interferometry in a high-finesse cavity. *Nature* 2022 610:7932, 610:472–477, 10 2022.
- [47] Rudolf Grimm, Matthias Weidemüller, and Yurii B Ovchinnikov. Optical dipole traps for neutral atoms. 1999.
- [48] T. Grunzweig, A. Hilliard, M. McGovern, and M. F. Andersen. Near-deterministic preparation of a single atom in an optical microtrap. *Nature Physics* 2010 6:12, 6:951–954, 9 2010.
- [49] Fernando Guadalupe and Santos Lins Brandão. Entanglement theory and the quantum simulation of many-body physics. 10 2008.
- [50] David Guery-Odelin and Claude Cohen-Tannoudji. Advances in atomic physics: An overview.
- [51] Aljaafari Hamza and Basant Kumar. A review paper on des, aes, rsa encryption standards. *Proceedings of the 2020 9th International Conference on System Modeling and Advancement in Research Trends, SMART 2020*, pages 333–338, 12 2020.
- [52] T P Harty, M A Sepiol, D T C Allcock, C J Ballance, J E Tarlton, and D M Lucas. High-fidelity trapped-ion quantum logic using near-field microwaves. 2016.
- [53] Markus Hijkema, Bernhard Weber, Holger P. Specht, Simon C. Webster, Axel Kuhn, and Gerhard Rempe. A single-photon server with just one atom. *Nature Physics* 2007 3:4, 3:253–255, 3 2007.
- [54] C W Hoyt, Z W Barber, C W Oates, T M Fortier, S A Diddams, and L Hollberg. Observation and absolute frequency measurements of the $1\text{ s } 0\text{-}3\text{ p } 0$ optical clock transition in neutral ytterbium. 2005.
- [55] Chris Huggins. Growth surface characterisation of metal oxide thin films under uhv conditions.
- [56] M. Karski, L. Förster, J. M. Choi, W. Alt, A. Widera, and D. Meschede. Nearest-neighbor detection of atoms in a 1d optical lattice by fluorescence imaging. *Physical Review Letters*, 102:053001, 2 2009.
- [57] Hyosub Kim, Woojun Lee, Han Gyeol Lee, Hanlae Jo, Yunheung Song, and Jaewook Ahn. In situ single-atom array synthesis using dynamic holographic optical tweezers. *Nature Communications* 2016 7:1, 7:1–8, 10 2016.
- [58] H. J. Kimble. Strong interactions of single atoms and photons in cavity qed. *Physica Scripta*, 76:127, 1998.

- [59] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. 2020.
- [60] Daniel Kleppner. A short history of atomic physics in the twentieth century. *Reviews of Modern Physics*, 71:S78, 3 1999.
- [61] H R Kong and K S Choi. Physical limits of ultra-high-finesse optical cavities: Taming two-level systems of glassy metal oxides.
- [62] Minho Kwon, Matthew F Ebert, Thad G Walker, and M Saffman. Parallel low-loss measurement of multiple atomic qubits. 2017.
- [63] Zhangli Lai, Shichao Zhang, Qingdong Gou, and Yong Li. Polarizabilities of rydberg states of rb atoms with n up to 140. *Physical Review A*, 98:52503, 2018.
- [64] Harry Levine, Dolev Bluvstein, Alexander Keesling, Tout T Wang, Sepehr Ebadi, Giulia Semeghini, Ahmed Omran, Markus Greiner, Vladan Vuleticvuletic, and Mikhail D Lukin. Dispersive optical systems for scalable raman driving of hyperfine qubits. *Physical Review A*, 105:32618, 2022.
- [65] Bo Li, Meng Li, Xiaojun Jiang, Jun Qian, Xiaolin Li, Liang Liu, and Yuzhu Wang. Optical spectroscopy of np rydberg states of $\{87\}\{\{\text{rb}\}\}$ atoms with a 297-nm ultraviolet laser. 2019.
- [66] Zeyang Li, Boris Braverman, Simone Colombo, Chi Shu, Akio Kawasaki, Albert F Adiyatullin, § Edwin Pedrozo-Peñafiel, Enrique Mendez, and Vladan Vuleticvuletic. Collective spin-light and light-mediated spin-spin interactions in an optical cavity. *Physical Review Applied*, 10:20308, 2030.
- [67] Markus Mack, Florian Karlewski, Helge Hattermann, Simone Höckh, Florian Jessen, Daniel Cano, and József Fortágh. Measurement of absolute transition frequencies of 87 rb to ns and n d rydberg states by means of electromagnetically induced transparency. *PHYSICAL REVIEW A*, 83:52515, 2011.
- [68] Stuart J Masson and Scott Parkins. Rapid production of many-body entanglement in spin-1 atoms via cavity output photon counting. 2019.
- [69] Adam R. Mills, Charles R. Guinn, Michael J. Gullans, Anthony J. Sigillito, Mayer M. Feldman, Erik Nielsen, and Jason R. Petta. Two-qubit silicon quantum processor with operation fidelity exceeding 99 *Science Advances*, 8:5130, 4 2022.
- [70] C Monroe and J Kim. Scaling the ion trap quantum processor. 8, 2016.
- [71] A H Myerson, D J Szwer, S C Webster, D T C Allcock, M J Curtis, G Imreh, J A Sherman, D N Stacey, A M Steane, and D M Lucas. High-fidelity readout of trapped-ion qubits. 2008.

- [72] Warren Nagourney. Quantum electronics for atomic physics.
- [73] Minh-Thi Nguyen, Jin-Guo Liu, Jonathan Wurtz, Mikhail D Lukin, Sheng-Tao Wang, and Hannes Pichler. Quantum optimization with arbitrary connectivity using rydberg atom arrays.
- [74] Michael A. Nielsen and Isaac L. Chuang. Quantum computation and quantum information: 10th anniversary edition. *Quantum Computation and Quantum Information*, 12 2010.
- [75] S. Olmschenk, R. Chicireanu, K. D. Nelson, and J. V. Porto. Randomized benchmarking of atomic qubits in an optical lattice. *New Journal of Physics*, 12:113007, 11 2010.
- [76] Karmela Padavic-Callaghan. The world’s largest quantum computer fridge. *New Scientist*, 255:18, 9 2022.
- [77] A J Park, J Trautmann, N Šanticšantic, V Klüsener, A Heinz, I Bloch, and S Blatt. Cavity-enhanced optical lattices for scaling neutral atom quantum technologies to higher qubit numbers. *Physical Review Applied*, 10:30314, 2022.
- [78] Lars Pause, Tilman Preuschoff, Dominik Schäffner, Malte Schlosser, and Gerhard Birkel. Reservoir-based deterministic loading of single-atom tweezer arrays.
- [79] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum ccd computer architecture. *Nature 2021 592:7853*, 592:209–213, 4 2021.
- [80] Harald Proksche, Giinter Nagorsen, and Detlef Ross. The influence of nh4f on the etch rates of undoped si02 in buffered oxide etch.
- [81] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance absorption by nuclear magnetic moments in a solid [7]. *Physical Review*, 69:37–38, 1 1946.
- [82] Joshua Ramette, Josiah Sinclair, Zachary Vendeiro, Alyssa Rudelis, Marko Cetina, and Vladan Vuletić. Any-to-any connected cavity-mediated architecture for quantum computing with trapped ions or rydberg arrays. *PRX Quantum*, 3:010344, 3 2022.
- [83] Stephan Ritter, Christian Nölleke, Carolin Hahn, Andreas Reiserer, Andreas Neuzner, Manuel Uphoff, Martin Mücke, Eden Figueroa, Joerg Bochmann, and Gerhard Rempe. An elementary quantum network of single atoms in optical cavities. *Nature 2012 484:7393*, 484:195–200, 4 2012.
- [84] I I Ryabtsev, I I Beterov, D B Tret’yakov, V M Èntin, and E A Yakshina. Spectroscopy of cold rubidium rydberg atoms for applications in quantum information. *Physics-Uspexhi*, 59:196–208, 2 2016.

- [85] M Saffman. Review to cite this article: M saffman. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49:202001, 2016.
- [86] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with rydberg atoms. *Reviews of Modern Physics*, 82:2313–2363, 8 2010.
- [87] Malte Schlosser, Daniel Ohl De Mello, Dominik Schäffner, Tilman Preuschoff, Lars Kohfahl, and Gerhard Birkl. Assembled arrays of rydberg-interacting atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 53:144001, 6 2020.
- [88] Kai Niklas Schymik, Bruno Ximenez, Etienne Bloch, Davide Dreon, Adrien Signoles, Florence Nogrette, Daniel Barredo, Antoine Browaeys, and Thierry Lahaye. In situ equalization of single-atom loading in large-scale optical tweezer arrays. *Physical Review A*, 106:022611, 8 2022.
- [89] Margaret E Shea, Paul M Baker, James A Joseph, Jungsang Kim, and Daniel J Gauthier. Submillisecond, nondestructive, time-resolved quantum-state readout of a single, trapped neutral atom. *PHYSICAL REVIEW A*, 102:53101, 2020.
- [90] Kevin Singh, Shraddha Anand, Andrew Pocklington, Jordan T Kemp, and Hannes Bernien. Dual-element, two-dimensional atom array with continuous-mode operation. *Physical Review X*, 12, 2022.
- [91] Sergei Slussarenko and Geoff J. Pryde. Photonic quantum information processing: A concise review. *Applied Physics Reviews*, 6:41303, 12 2019.
- [92] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T.C. Allcock, and D. H. Slichter. High-fidelity laser-free universal control of trapped ion qubits. *Nature 2021 597:7875*, 597:209–213, 9 2021.
- [93] Daniel A Steck. Classical and modern optics.
- [94] Daniel A Steck. Rubidium 87 d line data. 2001.
- [95] J D Sterk, L Luo, T A Manning, P Maunz, and C Monroe. Photon collection from a trapped ion-cavity system. *PHYSICAL REVIEW A*, 85:62308, 2012.
- [96] Nicolas Straniero, Jérôme Degallaix, Raffaele Flaminio, Laurent Pinard, Gianpietro Cagnoli, F Acernese, M Agathos, K Agatsuma, D Aisa, N Allemandou, A Allocca, J Amarni, P Astone, G Balestri, G Ballardini, F Barone, J p Baronick, M Barsuglia, A Basti, F Basti, T S Bauer, V Bavigadda, M Bejger, M G Beker, C Belczynski, D Bersanetti, A Bertolini, M Bitossi, M A Bizouard, S Bloemen, M Blom, M Boer, G Bogaert, D Bondi, F Bondu, L Bonelli, R Bonnard, V Boschi, L Bosi, T Bouedo, C Bradaschia, M Branchesi, T Briant, A Brillet, V Brisson, T Bulik, H J Bulten, D Buskulic, C Buy, G Cagnoli, E Calloni, C Campeggi, B Canuel, F Carbognani, F Cavalier, R Cavalieri, G Cella,

- E Cesarini, E C Mottin, A Chincarini, A Chiummo, S Chua, F Cleva, E Coccia, P f Cohadon, A Colla, M Colombini, A Conte, J p Coulon, E Cuoco, A Dalmaz, V Dattilo, M Davier, R Day, G Debreczeni, J Degallaix, S Deléglise, W D Pozzo, H Dereli, R D Rosa, L D Fiore, A D Lieto, A D Virgilio, M Doets, V Dolique, M Drago, M Ducrot, G Endrőczy, V Fafone, S Farinon, I Ferrante, F Ferrini, F Fidecaro, I Fiori, R Flaminio, J d Fournier, S Franco, S Frasca, F Frasconi, L Gammaitoni, F Garufi, M Gaspard, A Gatto, G Gemme, B Gendre, E Genin, A Gennai, S Ghosh, L Giacobone, A Giazotto, R Gouaty, M Granata, G Greco, P Groot, G M Guidi, J Harms, A Heidmann, H Heitmman, P Hello, G Hemming, E Hennes, D Hofman, P Jaranowski, R J G Jonker, M Kasprzack, F Kéfélian, I Kowalska, M Kraan, A Królak, A Kutynia, C Lazzaro, M Leonardi, N Leroy, N Letendre, T G F Li, B Lieunard, M Lorenzini, V Lorientte, G Losurdo, C Magazzù, E Majorana, I Maksimovic, V Malvezzi, N Man, V Mangano, M Mantovani, F Marchesoni, F Marion, J Marque, F Martelli, L Martellini, A Masserot, D Meacher, J Meidam, F Mezzani, C Michel, L Milano, Y Minenkov, A Moggi, M Mohan, M Montani, N Morgado, B Mours, F Mul, M F Nagy, I Nardecchia, L Naticchioni, G Nelemans, I Neri, M Neri, F Nocera, E Pacaud, C Palomba, F Paoletti, A Paoli, A Pasqualetti, R Passaquieti, D Passuello, M Perciballi, S Petit, M Pichot, F Piergiovanni, G Pillant, A Piluso, L Pinard, R Poggiani, M Prijatelj, G A Prodi, M Punturo, P Puppo, D S Rabeling, I Rácz, P Rapagnani, M Razzano, V Re, T Regimbau, F Ricci, F Robinet, A Rocchi, L Rolland, R Romano, D Rosińska, P Ruggi, E Saracco, B Sassolas, F Schimmel, D Sentenac, V Sequino, S Shah, K Siellez, N Straniero, B Swinkels, M Tacca, M Tonelli, F Travasso, M Turconi, G Vajente, N van Bakel, M van Beuzekom, J F J van den Brand, C Van Den Broeck, M V van der Sluys, J van Heijningen, M Vasúth, G Vedovato, J Veitch, D Verkindt, F Vetrano, A Viceré, J y Vinet, G Visser, H Vocca, R Ward, M Was, L w Wei, M Yvert, A Z ́ny, and J p Zendri. 4040) mirrors; (220.5450) polishing; (050.2230) fabry-perot; (220.2740) geometric optical design. 23, 2015.
- [97] R T Sutherland, R Srinivas, S C Burd, H M Knaack, A C Wilson, D J Wineland, D Leibfried, D T C Allcock, D H Slichter, and S B Libby. Laser-free trapped-ion entangling gates with simultaneous insensitivity to qubit and motional decoherence. *PHYSICAL REVIEW A*, 101:42334, 2020.
- [98] Haruka Tanji-Suzuki, Ian D Leroux, Monika H Schleier-Smith, Marko Cetina, Andrew T Grier, Jonathan Simon, and Vladan Vuleticvuletic. Interaction between atomic ensembles and optical resonators: Classical description.
- [99] Igor Teper, Yu-Ju Lin, and Vladan Vuleticmit Vuletic. Resonator-aided single-atom detection on a microfabricated chip. 2006.
- [100] J D Thompson, T G Tiecke, A S Zibrov, Vuletic V Vuletic, and M D Lukin. Coherence and raman sideband cooling of a single atom in an optical tweezer. 2013.
- [101] Zachary Vendeiro. Raman cooling to high phase space density. 2022.

- [102] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger. Trapped-ion quantum logic with global radiation fields. *Physical Review Letters*, 117:220501, 11 2016.
- [103] Stephan Welte, Bastian Hacker, Severin Daiss, Stephan Ritter, and Gerhard Rempe. Photon-mediated quantum gate between two neutral atoms in an optical cavity. *PHYSICAL REVIEW X*, 8:11018, 2018.
- [104] Steve. Winder, Joseph J. Carr, and John Davies. Newnes radio and rf engineering pocket book. page 344, 2002.
- [105] Nicolas Wittler, Federico Roy, Kevin Pack, Max Werninghaus, Anurag Saha Roy, Daniel J. Egger, Stefan Filipp, Frank K. Wilhelm, and Shai Machnes. Integrated tool set for control, calibration, and characterization of quantum devices applied to superconducting qubits. *Physical Review Applied*, 15:034080, 3 2021.
- [106] Hao Wu, Michael Mills, Elizabeth West, Michael C Heaven, and Eric R Hudson. Increase of the barium ion-trap lifetime via photodissociation. *Physical Review A*, 104:63103, 2021.
- [107] Amnon Yariv. Amnon yariv-quantum electronics-wiley (1989). 1988.
- [108] Jerzy Zachorowski, Tadeusz Pałasz, and Wojciech G Awlik. Magneto-optical trap for rubidium atoms. *Optica Applicata*, 1998.
- [109] Yifu Zhu, Daniel J. Gauthier, S. E. Morin, Qilin Wu, H. J. Carmichael, and T. W. Mossberg. Vacuum rabi splitting as a feature of linear-dispersion theory: Analysis and experimental observations.