Systems and Methods for Generating Large Arrays of Optical Traps in Neutral Atom Array Quantum Processors

by

Soroush Khoubyarian

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics (Quantum Information)

Waterloo, Ontario, Canada, 2025
© Soroush Khoubyarian 2025

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The content of this thesis reflects my original contributions within the experimental group led by Prof. Alexandre Cooper-Roy (ACR). The neutral atom array quantum processor used in this work was conceived and built prior to and during my studies, starting in 2019, through the efforts of several group members, including PhD students and postdoctoral researchers, under the supervision of ACR.

My specific contributions are as follows.

First, I designed and built the hardware and software for generating a stabilized, uniform two-dimensional trap array in our lab, which we use for daily experiments. This system enables long experimental runs without the need for periodic recalibration of trap depths, thereby increasing the duty cycle of the platform. In addition to constructing the optical path for the traps and the imaging system, I developed the software for a stabilization system that continuously measures trap powers and equalizes trap depths. I further integrated this system into the experimental setup, allowing the trap depths to be calibrated with the atomic signal and achieving in-situ uniformization of trap depths using atomic spectroscopy.

Second, I proposed a new method for circularizing laser beams using three standard cylindrical lenses and incorporated it into our trapping setup. This improved the beam circularity of our trapping beam, enabling the generation of more traps with the same optical power and equalizing the radial trapping frequency of our optical tweezers. The method is broadly applicable to other areas of optics, including laser machining, fiber coupling, and beam shaping. To our knowledge, this approach had not been proposed or implemented before and achieved performance previously attainable only with more advanced and expensive tools such as spatial light modulators, or with less power-efficient methods such as spatial filtering.

Third, I designed and implemented the software control and analysis infrastructure used in our lab for experiments. The control system integrates the different experimental apparatus in a seamless and time-efficient manner, enabling thermometry measurements and microwave horn spectroscopy. The analysis system captures and processes experimental data in real-time, allowing rapid generation of plots and automated analysis that guides the design of subsequent experiments.

Abstract

Neutral atom array quantum processors provide a scalable and controllable platform for quantum simulation, computation, sensing, and metrology. Strong Rydberg interactions enable entanglement and high-fidelity gates, while the ability to rearrange atoms allows for multi-qubit gates between spatially distant atoms and the implementation of quantum error-correcting codes with nonlocal stabilizers. These arrays also support electromagnetic field sensing and form the basis for next-generation atomic clocks with improved stability and precision. Recent progress has demonstrated platforms with thousands of qubits, the implementation of error correction codes, and experimental realization of quantum spin models and the study of quantum state phase transitions. Current efforts now focus on scaling to even larger arrays and reducing quantum gate errors, with the goal of achieving computations and simulations beyond the reach of classical devices.

However, increasing the scale and controllability of neutral-atom array platforms to address problems where quantum advantage can have practical impact remains challenging. Key difficulties include scaling to larger numbers of qubits, implementing fault-tolerant computation, and performing qubit gates between spatially distant atoms. A critical experimental challenge is ensuring that qubits form an indistinguishable spatiotemporal ensemble, so that atoms across the array and over multiple experimental shots can be treated as identical. Achieving this uniformity becomes more difficult as the array size increases, due to the added complexity of ensuring all atoms experience consistent control parameters, such as trap depth and magnetic fields. Without sufficient spatiotemporal uniformity, it is impossible to meaningfully average measurements across different qubits and experimental runs, undermining scalability and limiting the performance of quantum simulations, computations, and sensing applications.

In this thesis, I propose, design, and implement methods to scale up the number of qubits and to improve the spatiotemporal uniformity of qubit properties in acousto-optically generated trap arrays for neutral-atom quantum processors. These methods enable magnetic-field imaging across an array of 1,305 optical traps containing 690 qubits, as well as high-fidelity fluorescence imaging of single atoms in optical tweezers.

First, I demonstrate a novel method for circularizing and de-astigmatizing the trapping beam using three cylindrical lenses. This approach is cost-effective, power-efficient, and broadly applicable. By improving the beam shape, we reduce the per-trap laser power, enabling the use of larger qubit arrays in neutral atom quantum processors. Applied to a Ti:sapphire laser with an initial circularity of 0.69, this method achieves a circularity of 0.97 and a beam waist separation of 0.8% of the Rayleigh range, reducing the optical power required per trap by 5%.

After beam circularization, I develop a real-time closed-loop feedback system for an optical trap array generated by two orthogonal acousto-optical deflectors to enhance spatiotemporal uniformity of qubit properties. The system stabilizes trap depths using power measurements from a fast CMOS camera and in-situ depth estimates from an EMCCD camera based on the atomic signal. Noise is decomposed into uncorrelated temporal modes, each regulated by an independent PID loop. Unlike approaches that regulate only total laser power, our method stabilizes the depth of individual traps, eliminating the need for frequent recalibration of cooling beam parameters and enabling higher duty cycles. In a 45×29 array, the generated traps exhibit a standard deviation of 6% relative to the mean trap depth, limited by the number of independent control parameters actuating the acousto-optical deflectors.

Having established these capabilities, we demonstrate controllability and practical utility of the platform by performing magnetic field sensing using microwave horn spectroscopy on a 45×29 trap array, mapping field gradients across the array. This establishes the foundation for advanced quantum sensing protocols, such as quantum lock-in amplification, and paves the way toward entanglement-based quantum sensors capable of offering a quantum advantage in sensing.

Lastly, to improve readout fidelity, I describe the optimization of the imaging beam parameters in neutral atom array quantum processors. This optimization maximizes atom classification accuracy while minimizing the probability of ejecting atoms from the traps. Applied to a 45×29 trap array, the method achieves a classification fidelity of 99.98% and a loss probability of 0.12% with a 75 ms imaging time. These improvements reduce measurement errors, enhance quantum state readout, and increase the sensitivity of neutral atom–based quantum sensors.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Alexandre Cooper-Roy, for his constant support, thoughtful mentorship, and encouragement throughout this project. His guidance has played a central role in my development as a researcher. I am especially grateful for the highly detailed and honest feedback he provided on all of my work in the lab. Alexandre spent considerable time reviewing even the smallest details of my experiments and analysis, consistently pushing me to reach higher standards. His mentorship significantly reshaped the way I think about scientific problems and helped me grow into a more rigorous and critical thinker. I also thank him for his unwavering support in connecting me with conference opportunities and encouraging me to engage with the broader research community, which is something I would not have done on my own.

I am also deeply thankful to my co-supervisor, David Cory, for his constructive feedback, which has helped me refine both the technical and conceptual aspects of this work. David consistently challenged me to think independently, to ask meaningful scientific questions, and to find creative ways to answer them using the tools at my disposal. His guidance pushed me to become a more thoughtful, resourceful, and rigorous scientist.

I owe special thanks to Anastasiia Mashko, Artem Zhutov, Kent Ueno, and Parth Padia for their guidance throughout my time in the lab. I learned much of what I know from their patient instruction and deep experience. I am especially grateful to Anastasiia and Artem for teaching me how to build production-quality optical systems, and for their *extraordinarily* patient endurance as I struggled (often repeatedly) to meet expectations while assembling and aligning the trapping system, not to mention their tolerance while helping me fix my mistakes along the way. I also thank Kent and Parth for helping me understand the software architecture of our platform and for their thoughtful and constructive feedback throughout the project.

In addition, I gratefully acknowledge the contributions of undergraduate students, whose efforts played a key role in developing the technical infrastructure underlying our experimental work. In particular, Amirhossein Dadpour made significant contributions to the implementation of our low-latency atom reconfiguration system, which enabled the controlled transport of atoms across different optical traps. His work substantially improved the quality, modularity, and maintainability of the underlying software, making the system more robust and easier to integrate into our experimental workflow.

I would also like to thank several postdoctoral researchers in our group, including Christopher Wyenberg, Anton Borissov and Meenu Kumari, for engaging discussions, thoughtful advice, and career suggestions. Their perspectives have significantly contributed to my growth as a researcher.

I would also like to sincerely thank my committee members Professor Dmitry Pushin, and Professor Kazi Rajibul-Islam, for their valuable feedback both during committee meetings and outside of them. Their insights and guidance have greatly improved the quality and clarity of this work, and their thoughtful suggestions helped me refine the scientific arguments throughout the project.

Finally, to my parents, thank you for your unwavering support and belief in me. Your emotional support—and, let's be honest, financial support—made this possible.

Chapter-by-Chapter Contributions

Chapter 2. This chapter is based on a manuscript co-authored with ACR and Anastasiia Mashko (AM). The paper involved ideation, theory, simulations, and experimental work. I proposed the circularization method, developed the theory, performed the simulations, and contributed to the experimental section. AM provided major assistance with data acquisition and contributed to data analysis. ACR led the project by giving valuable feedback on the development of the theory and simulations, and by preparing the manuscript.

Chapter 3. My contributions include formalizing the mathematical framework for the stabilization module, designing and implementing the real-time feedback system, integrating it into the experimental workflow, and characterizing its performance. ACR provided critical guidance for the system design, for both hardware and software.

Chapter 4. My contributions include developing the control system required for running spectroscopy experiments (Rabi, Ramsey, and spin-echo), creating the analysis infrastructure for these measurements, and designing an efficient control pipeline for actuating the arbitrary waveform generator (AWG) used to drive the microwave horn to perform arbitrary spectroscopy experiments. The hardware for the microwave horn, magnetic-field coils, and fine optical pumping was primarily constructed by AM and Artem Zhutov (AZ).

Chapter 5. I proposed the use of specific performance metrics, such as atom loss rate per scattered photon, to quantitatively evaluate imaging fidelity. Together with AM, we optimized the physical alignment of the polarization-gradient cooling beams to improve imaging and cooling efficiency. The optimization of the imaging beam parameters was carried out collaboratively with AZ and AM.

Table of Contents

Author's Declaration ii

Statement of Contributions iii

Abstract iv

Acknowledgements vi

List of Figures xii

List of Tables xxvii

1 Introduction 1

1.1 Neutral Atom Arrays as a Quantum Platform 3

1.1.1 Architecture and Physical Properties 3

1.1.2 Applications in Simulation, Computation, and Sensing 6

1.2 Challenges in Building Neutral Atom Array Quantum Processors 6

1.2.1 Scalability 7

1.2.2 Reliability 8

1.2.3 Controllability 8

1.2.4 Outline of the Thesis 9

2 Correcting astigmatism and ellipticity in Gaussian beams using a cylindrical lens pair with tunable focal lengths **14**

2.1 Background . . . 16

2.2 Theory – Gaussian Beam Circularization . . . 18

2.2.1 Three-Lens Solution to Gaussian Beam Circularization . . . 20

2.2.2 Robustness Against Finite Inter-Lens Distance . . . 26

2.3 Experiment – Astigmatism and Ellipticity Suppression . . . 29

2.3.1 Beam Circularization of a Commercial Laser Source . . . 30

2.4 Conclusion . . . 32

2.5 Appendices . . . 33

2.5.1 Matrix Formalism for Generalized Gaussian Beams . . . 33

2.5.2 Beam Circularization Using a Cylindrical Lens Pair . . . 39

2.5.3 Astigmatism Correction Using a Cylindrical Lens Pair . . . 44

2.5.4 Ellipticity and Astigmatism Correction with Three Cylindrical Lenses 47

2.5.5 Image Processing Method for Extracting Beam Properties . . . 50

3 A Feedback System for the Uniformization and Stabilization of Trap Depths in Optical Trap Arrays **52**

3.1 Trap Depth Uniformization and Stabilization Module . . . 55

3.1.1 Hardware Implementation . . . 55

3.1.2 Formulation of the Uniformization Problem . . . 57

3.1.3 Control Updates and Temporal Modes . . . 60

3.1.4 Stabilization Protocol . . . 67

3.1.5 Quantifying Trap Stability and Uniformity . . . 70

3.1.6 Time Dependence of the Jacobian Matrix . . . 72

3.2 Uniform Adiabatic Control of Trap Depths . . . 74

3.2.1 Adiabatic Trap Depth Control Using Waveform Modulation and Mixer-Based Power Stabilization . . . 75

3.2.2 Thermometry with Adiabatic Ramps . . . 79

4 Magnetic Field Sensing with Arrays of Neutral Atom Atoms **81**
4.1 Magnetic Field Sensing using Microwave Spectroscopy 82
4.1.1 State Preparation and Measurement 83
4.1.2 Microwave Spectroscopy in Shallow Traps 85
4.1.3 Future Steps – Quantum Sensing Applications 90
4.2 Colocated Optical Beams for Site-Selective Single-Qubit Control 91
4.2.1 Colocalization Procedure 94
4.2.2 Assessment of Performance 99

5 Optimizing Fluorescence Imaging for Quantum State Readout **101**
5.1 Derivation of the Figure of Merit 102
5.2 Optimization of Imaging Parameters 105
5.3 Characterization of the Imaging Classification Fidelity and Loss Probability 108

6 Conclusion and Outlook **111**

References **113**

Glossary **121**

List of Figures

1.1 **Quantum Simulation for Material Discovery.** Quantum processors can serve as quantum simulators by encoding and controlling the dynamics of quantum many-body systems. This approach enables the study of strongly correlated systems and accelerates the discovery of new materials with desired properties. 2

1.2 **Experimental Setup.** ^{87}Rb atoms are are dispensed from the atom source at the top of the vacuum cell. They are then trapped in a magneto-optical trap (MOT) generated by six MOT beams (orange arrows) and coils (red). After cooling, atoms are transferred into optical tweezers formed by 813 nm beams focused and collected by two microscope objectives (gray lenses). Chapters 2 and 3 describe trap engineering, including beam circularization (red) and 2D trap array multiplexing with power stabilization (purple boxes). Atoms are addressed with a microwave pulse from the microwave horn located beneath the vacuum cell (blue box). This setup is used for sensing electromagnetic fields as discussed in Chapter 4 across the trap array. Lastly, the six polarization-gradient cooling beams are used for imaging atoms with polarization-gradient cooling, optimized in Chapter 5 to maximize fidelity and minimize loss. 4

1.3 **Architecture of a ^{87}Rb Rydberg Atom Array Quantum Processor.** Overview of a ^{87}Rb-based Rydberg atom array quantum processor. (a) Atomic transitions used to control ^{87}Rb atoms. The D_1 line enables hyperfine transitions via Raman pulses, as well as enabling Raman sideband cooling. The D_2 line is used for polarization-gradient cooling, pushout, and fluorescence imaging. Two-photon Rydberg excitation is performed using 420 nm and 1013 nm beams. (b) Hyperfine structure of the ground state and definition of the qubit states as $|0\rangle = |F = 1, m_F = -1\rangle$ and $|1\rangle = |F = 2, m_F = -2\rangle$. Transitions between the qubit states are driven with a 6.834 GHz microwave pulse. (c) Spatial configuration of the atoms and beam paths used for implementing optical or microwave-based single- and multi-qubit gates. The atoms are trapped in arrays of optical tweezers at a wavelength of 813 nm. The D_1 beams are incident on the atoms along the same direction as the traps, while the Rydberg beams (420 nm and 1013 nm) are incident from opposite directions along the orthogonal axis. The microwave horn addresses the atoms with pulses traveling perpendicular to the page at 6.8 GHz. (d) Configuration of the D_2 beams used for cooling, atom ejection, and imaging. Two counter-propagating beam pairs lie along the diagonal and anti-diagonal of the xz plane. A third counter-propagating pair lies along the diagonal of the yz plane. 11

1.4 **Qubits in a ^{87}Rb Rydberg Atom Array Quantum Processor.** Development and characterization of a ^{87}Rb-based Rydberg atom array platform. (a) Average fluorescence image from a 45×29 trap array (1305 traps) with a loading probability of 53%. Each dark spot corresponds to fluorescence from a single atom at 780 nm. (b) Single-shot image showing approximately 690 occupied traps and 615 empty traps. (c) Photon count histogram from a single trap during imaging, showing distinct peaks for background (vacant) and atomic signal (occupied). (d) Microwave Rabi spectroscopy performed on the array, showing population oscillations between the qubit states. A global π-pulse occurs near 25 μs, corresponding to a full population transfer to the $|1\rangle$ state (see Sec. 4.1.2). The achieved population is limited by the qubit gate fidelity. 12

1.5 **Outline of the Thesis.** Overview of the contributions of this thesis to a set of key challenges in building Rydberg atom array quantum processors. **Scalability:** Chapter 2 introduces a new beam circularization method using three cylindrical lenses to increase the number of qubits by 5%, and Chapter 5 optimizes imaging to reach a fidelity of 99.98% with a qubit survival probability of 99.9%. **Reliability:** Chapter 3 develops a uniformization and stabilization module for AOD-generated tweezers, producing trap arrays with homogeneous and stabilized trap depths, and enabling adiabatic ramps while maintaining homogeneity. **Controllability:** Chapter 4 demonstrates magnetic sensing in shallow optical traps using microwave spectroscopy, advancing toward dynamical decoupling and entanglement-enhanced sensitivity. 13

2.1 **Focusing elliptical Gaussian beams.** The figure shows the simulated intensity of an elliptical Gaussian beam focused with a Mitutoyo microscope objective of focal length $f = 4$ mm and aperture radius 4 mm, plotted as a function of the input beam circularity. For each circularity, the beam radii are chosen to maximize the intensity at the focal plane. The maximum focal intensity is achieved for a circular input beam (circularity = 1.0). As the input circularity decreases towards zero, the achievable focal intensity is reduced. 15

2.2 **Gaussian beams.** (a) A circular Gaussian beam has a circular intensity profile with radial symmetry. (b) An elliptical Gaussian beam has distinct divergence angles along its major (blue) and minor (red) axes. The beam radius along each axis is minimized in the same focal plane, where the intensity profile is elliptical. The intensity profile becomes circular in two planes symmetrically located on either side of the beam waist. (c) An astigmatic Gaussian beam is an elliptical beam whose beam waists do not overlap. There exists at least one plane where the intensity profile is circular. . . . 18

2.3 **Typical Two-Cylindrical-Lens Solution.** An astigmatic elliptical Gaussian beam can be converted into a circular Gaussian beam by placing two cylindrical lenses, L_1 and L_2, each oriented along one of the principal axes of the diverging beam. Each cylindrical lens is placed approximately one focal length away from its corresponding beam waist. 19

2.4 **Three-Lens Solution.** An astigmatic elliptical Gaussian beam is converted into a circular Gaussian beam using three cylindrical lenses. The first lens (L_1) is oriented along the major axis of the beam and placed in an arbitrary plane, not necessarily located a focal length away from the beam waist. The second and third lenses (L_{2-3}) are mounted one behind the other at an angle $\pm\theta$ in the plane where the intensity profile of the beam is circular. The relative angle is chosen to equalize the divergence angles of the beam along the major and minor axes. The solution is robust against small changes in the relative distance between the two lenses. The resulting circular beam can readily be collimated using a spherical lens. 21

2.5 **Equivalent Phase Profiles of a Cylindrical Lens Pair.** (a) The cumulative phase profile of two uniaxial cylindrical lenses of focal length f, counter-rotated by $\pm\theta$, is equivalent to the phase profile of a (b) biaxial lens, which is itself equivalent to the cumulative profile of (c) two uniaxial cylindrical lenses with tunable focal lengths $f_x(\theta) = f\sec^2(\theta)/2$ and $f_y(\theta) = f\csc^2(\theta)/2$. 22

2.6 **Astigmatism Correction.** (a) An elliptical beam is focused through a pair of cylindrical lenses with a tunable relative angle 2θ. The major and minor axes focus in different planes (blue and orange dashed lines). (b) The relative distance between the two focal planes (see Eq. (2.16)) changes with the lens angle and vanishes at three specific values. (c) Beam circularity (see Eq. (2.22)) is maximized close to one of the angles where the focal planes overlap; however, the biaxial lens pair alone is not sufficient to achieve perfect circularity. Note that the angle for which the circularity is maximized does not necessarily coincide with the angle for which astigmatism is suppressed, and the plots are not necessarily symmetric. 24

2.7 **Generalized Gaussian Beam Twisting.** The orientation of a generalized Gaussian beam rotates as it propagates through free space, while its transverse intensity profile in each plane is elliptical, similar to the intensity profile of an elliptical Gaussian beam (see insets). 25

2.8 **Dependence of Minimum Circularity on Nonzero Inter-Lens Spacing.** (a) The minimum circularity $\mathcal{C}_0$ (white-to-blue surface) changes with the rotation angle of the cylindrical lenses, θ_2 and θ_3. At $d_{23} = 0$ (yellow circle), the optimal minimum circularity is achieved for $\theta_2 = \theta_3 = \theta^*$ (gray dashed line), where $\theta^* = 41.4°$. For $d_{23} \neq 0$, the optimal angles deviates from the optimal predicted values (green line). The green disk indicates the optimal angles obtained for $d_{23} = 6$ mm, which approximately corresponds to the distance between the lenses when mounted in rotation mounts. (b) The minimum circularity of the beam depends on both the focal length of the lenses and their inter-lens distance. The circularity decreases with inter-lens spacing. For a fixed inter-lens spacing, increasing the focal length improves the minimum achievable circularity. These results were obtained for a Gaussian beam with a wavelength of $\lambda = 415$ nm whose beam waist radii are $w_{0x} = 260$ μm and $w_{0y} = 1.015$ mm at the beam waists located $z_{0x} = -6$ m and $z_{0y} = -22$ m, respectively, where $z = 0$ denotes the plane where the first lens of the lens pair is located. 27

2.9 **Model Validation.** (a) Divergence angles and (b) beam waists of an astigmatic elliptical Gaussian beam measured along its principal axes, which are aligned with the x (red) and y axes (blue). The beam is generated by transmitting a nearly circular Gaussian beam at 780 nm through a pair of cylindrical lenses with a tunable relative angle 2θ, oriented asymmetrically about the x axis. Each data point represents a fit of the longitudinal intensity profile to a Gaussian beam model. Solid lines indicate theoretical predictions based on the input beam and the geometry of the lens system. The cylindrical lenses have a focal length of $f = 1000$ mm. The dashed lines show the expected behavior obtained when choosing the lenses to have an effective focal length along the transverse direction $f_\perp = 51.4$ m. 29

2.10 **Experimental Demonstration of Beam Circularization.** (a-left) A commercial laser source emits a diverging beam with an elliptical intensity profiles (insets). The beam radii (squares) are measured along the x (red) and y (blue) axes and fitted to Gaussian beam profiles (solid lines). (a-middle) Focusing the beam along its major axis using a cylindrical lens ($f_1 = 500$ mm) placed at P_1 (thick green line) creates a plane P_2 at $z = 176$ mm where the intensity profile becomes circular. (a-right) Focusing the beam along both its major and minor axes using a cylindrical lens pair ($f_2 = f_3 = 500$ mm) with an optimized relative angle suppresses astigmatism and restores circularity. (b) Focusing the beam after P_2 with an achromatic doublet ($f = 150$ mm) further demonstrates the suppression of astigmatism, with the beam waists overlapping at $z = 298$ mm. (c) The minimum beam circularity near the beam waist averages to 0.95(1), bounded between 0.92(1) and 0.99(1). 31

2.11 **Imperfect Lens Orientation.** A cylindrical lens of focal length $f = 500$ mm is placed at $z = 0$ mm at an angle θ with respect to the major axis of an incoming elliptical beam. The maximum circularity of the beam decreases away from 1.0 when the cylindrical lens is rotated away from $\theta = 0°$. At $\theta = 0°$ the beam radii intersect at $z = 200$ mm and the beam orientation remains constant. At $\theta = 1.5°$ the beam radii of the beam do not intersect, and the beam rotates by almost 90°. The beam reaches a maximum circularity of 0.98 at $z = 200$ mm. 49

2.12 **Image Processing.** (left) The raw image captured by the camera exhibits interference effects caused by its protective glass cover. The beam radii along the principal axes are extracted from a rescaled best-fit ellipse (solid red) applied to the filtered image (dashed red). (right) The filtered image is obtained by applying a Gaussian filter with $\sigma = 0.55$ mm to the raw image. The principal axes are extracted by fitting the filtered image to the best-fit ellipse (dashed red). 50

3.1 **Drift in the Relative Power of Optical Traps in an 11×11 Tweezer Array.** (a) Fraction of the total power contained in four selected columns of an initially uniformized 11×11 trap array over a duration of 50 seconds. Each column exhibits independent drift and fluctuations. For example, column 2 shows a steady increase in relative power while column 10 decreases. (b) Temporal evolution of the non-uniformity metric σ/μ, where σ is the standard deviation and μ is the mean of the trap depths. The metric increases over time, indicating a loss of uniformity in trap depths due to drift in the relative optical power distribution. 54

3.2 **Hardware Architecture for Closed-Loop Stabilization and Uniformization.** Schematic of the closed-loop feedback system used to stabilize and uniformize trap depths in a 45×29 ^{87}Rb atom array. Polychromatic RF waveforms are generated by a pair of AWG channels, amplified, and sent to orthogonal AODs to create a multiplexed tweezer array. The laser beam is focused by a microscope objective to trap single atoms and is collected by a second objective. The outgoing light is split: trapping light at 813 nm is detected by a CMOS camera to estimate relative tweezer powers, while fluorescence at 780 nm is imaged by an EMCCD camera to extract atom-dependent light shifts. These measurements feed into an estimator module (see Fig. 3.3) that reconstructs the trap depth vector $\underline{D}^{\text{estimate}}$. A controller module (see Fig. 3.6) computes a correction vector $\Delta\underline{\alpha}$ to update the AWG waveforms and minimize the deviation $\|\underline{D}^{\text{estimate}} - \underline{D}^{\text{target}}\|$ (see Sec. 3.1.2). 56

3.3 **Trap Depth Estimator for Closed-Loop Stabilization and Uniformization.** Implementation of the estimator module within the stabilization and uniformization feedback system (see Fig. 3.2). (a) The differential light shift $\delta\nu$ of the D_2 transition is measured using pushout spectroscopy: a near-resonant beam is swept across frequencies, and the frequency that maximizes atom loss is taken as the light-shifted resonance. (b) The measured $\delta\nu$ is related to the trap intensity via atomic polarizabilities of the $5S_{1/2}$ and $5P_{3/2}$ states, allowing extraction of trap depth and peak intensity. Assuming Gaussian beam profiles, the trap power can be inferred. A diagonal weight matrix $\underline{\underline{W}}$ is introduced to map CMOS-measured power $\underline{P}^{\text{camera}}$ to trap depth $\underline{D}$. (c) The estimator combines fast CMOS measurements of $\underline{P}^{\text{camera}}$ with slowly refreshed $\delta\underline{\nu}$ estimates from the EMCCD to compute $\underline{D}^{\text{estimate}} = \underline{\underline{W}} \cdot \underline{P}^{\text{camera}}$. A slow PID loop updates $\underline{\underline{W}}$ to reduce error between $\underline{D}^{\text{estimate}}$ and the EMCCD-inferred $\underline{D}^{\text{traps}}$ (see Sec. 3.1.2). 58

3.4 **Update Rule for Control Parameters to Correct for Trap Depth Noise.** (a) Effect of perturbing the control parameters $\underline{\alpha}^x$ along individual basis directions $\underline{e}^i$ in a 5×1 trap array. Each change primarily affects the corresponding column, but also induces residual changes in other traps due to nonlinear intermodulation from the RF amplifier. The varying color intensity on the ith trap indicates non-uniform sensitivity to each control parameter. (b) Applying the Jacobian-based update rule (see Eq. (3.11)) using the system's Jacobian matrix $\underline{\underline{J}}$ (see Eq. (3.8)) compensates for intermodulation and non-uniform response. This enables targeted, uniform updates of individual columns. In this 1D example, $\underline{\underline{J}}$ is invertible. (c) Conceptual illustration for larger arrays where $\underline{\underline{J}}$ is not invertible. The control space (left) and the higher-dimensional trap depth space (right) are connected via $\underline{\underline{J}}$. The accessible set of trap depths (black ellipse on the right) is a submanifold of the full target space. The current trap depths (black asterisk) and desired trap depths (blue circle on the right) are generally not directly connected via the control parameters. The best approximation of the target (orange triangle on the right) lies on the accessible manifold. Its pre-image under $\underline{\underline{J}}$ is a submanifold in control space (orange curve on the left). The update rule computes the smallest-norm change $\Delta\underline{\alpha}$ that brings the current control parameters to the closest point on this pre-image, moving along the shortest path (solid blue arrow), while naive rules may follow suboptimal directions (dashed green arrow). 63

3.5 **Decomposition of Trap Depth Noise into Uncorrelated Temporal Modes.** (a) Noise in a 2×1 trap array is decomposed into uncorrelated temporal modes by diagonalizing the correlation matrix of the estimated trap depths over a 4 s period (see Eq. (3.14)). The first mode corresponds approximately to a global fluctuation $(1\ 1)^T/\sqrt{2}$, and the second to a differential fluctuation $(1\ -1)^T/\sqrt{2}$, indicating two independent sources of noise: fluctuations in total laser power and in AOD diffraction efficiency, which affects the power balance between traps. The two eigenvalues of the correlation matrix $\underline{\underline{\mathcal{C}}}$ are shown to the right. (b) Decomposition of trap depth fluctuations in a 2×2 array into four orthogonal modes. The first mode is nearly global, while the second and third correspond to row-wise and column-wise variations. The column mode has a larger eigenvalue, indicating more significant column fluctuations. The four eigenvalues of $\underline{\underline{\mathcal{C}}}$ are shown in descending order. (c) Temporal modes for an 11×11 array, illustrating spatial structure in the trap depth noise and the relative importance of collective fluctuation modes. 66

3.6 **Controller Module Used in the PID Protocol to Update Control Parameters.** The controller module used in the uniformization and stabilization feedback system for an $N \times M$ trap array (see Fig. 3.2). The measured trap depth $\underline{D}^{\text{estimate}}$ is subtracted from the target depth $\underline{D}^{\text{target}}$ to compute the trap depth error $\Delta\underline{D}$. This error is decomposed into a truncated basis of $L+1$ uncorrelated temporal modes $\{\tilde{\underline{\mathcal{C}}}^i\}_{i=0}^{L}$, where $L+1 \leq NM$ and $L+1$ is the number of independent modes the controller is designed to correct. Each mode amplitude $\tilde{E}^i(t)$ is extracted via projection and sent to a dedicated PID loop, which computes the magnitude of the corrective response. The update direction is given by the update rule (see Eq. (3.11)) and the temporal mode vector $\tilde{\underline{\mathcal{C}}}^i$, resulting in a control update $\Delta\underline{\alpha}^i$. The total update to the control parameters is $\Delta\underline{\alpha} = \sum_{i=0}^{L} \Delta\underline{\alpha}^i$. 69

3.7 **Performance of the Stabilization and Uniformization System.** (a) Histogram of the relative deviation of time-averaged trap depths from the target trap depth for a 21×21 array. The red distribution (with feedback) is significantly narrower and centered around zero compared to the blue distribution (without feedback), indicating substantial improvement in time-averaged uniformity. (b) A high-resolution view of the red distribution in (a), showing its sharp peak around zero. (c) Evolution of the non-uniformity metric (σ/μ) over a 60 s period, showing long-term stability of the uniformization system (c.f. Fig. 3.1b). (d) Histogram of the relative temporal fluctuation σ/μ, where σ is the standard deviation of the trap depth over time and μ is the mean. Data is measured over 4 s for each trap. The red distribution (with feedback) is more compressed and closer to zero compared to the blue distribution (without feedback), indicating significant suppression of temporal fluctuations by the feedback system. 71

3.8 **Time Dependence of the Jacobian Matrix.** (a) Pictorial representation of the measured relative Jacobian matrix, $\underline{\underline{J}}_{\text{relative}} = \underline{\underline{J}}/\|\underline{\underline{J}}\|_\infty$, over ≈ 12000 s for an 11×11 trap array (121 traps). Traps are indexed in row-major order. The first 11 control parameters correspond to $\underline{\alpha}^y$ and predominantly affect their corresponding rows (dark blue blocks in the first 11 columns), while the second 11 control parameters correspond to $\underline{\alpha}^x$ and predominantly affect their corresponding columns. Notably, for the second set of parameters, there are non-zero correlations with trap powers outside the corresponding columns due to intermodulations between frequency tones applied to the AODs. (b) Relative variations of the Jacobian matrix as a function of time over a 12000 s period, with data points spaced by 23 s. The relative variation stabilizes after approximately 1500 s and remains roughly constant around 30×10^{-3}, with periodic peaks every 1–2 minutes. (c) Power spectral density (PSD) of the variations. Major frequency components are below 10 mHz, corresponding to periods of $\sim$ 100 s, likely related to the periodic peaks observed in (b). (d) Norm of the Allan variance vector of the Jacobian as a function of the time bin size τ. The norm reaches a minimum at ≈ 3132 s, indicating drift in the Jacobian and suggesting a recalibration period of approximately 3132 s, i.e., 52 min. 73

3.9 **Non-Uniformity in Output RF Power from a Mixer at Different Frequencies.** Output RF power of a double-balanced mixer (Mini-Circuits, ZFM-2+) as a function of DC voltage applied to the I input, with a fixed RF input power of approximately −13 dBm supplied to the L input. The output is shown for two RF frequencies: 80 MHz (red) and 120 MHz (blue). As the DC voltage approaches 0, a significant frequency-dependent discrepancy in output power emerges, with a maximum difference of up to 5.2 dBm. This non-uniformity implies that using a single mixer for global control of multiple traps can introduce large variations in trap depth during ramp-down. . . . 74

3.10 **Hardware Setup for Adiabatic Ramps with Uniformized Trap Depths.** An updated hardware setup extending the uniformization and stabilization system previously introduced in Fig. 3.2 to enable adiabatic ramping while maintaining uniform trap depths. In addition to the CMOS and EMCCD cameras used to measure trapping power and trap depth, respectively, a photodiode is added to monitor the total trapping power. A double-balanced mixer is inserted between the second AWG channel and AOD X to allow analog control over the global RF power via its I input. The uniformization and stabilization module continues to update the AWG waveforms using camera measurements, but it leaves the global power mode unchanged, i.e. $\tilde{\mathcal{C}}^0$. Instead, the photodiode signal is sent to an analog PID controller (Moku:Lab), which compares it to an external analog input that defines the desired adiabatic power ramp. The controller outputs an analog signal that adjusts the mixer voltage through the analog PID controller, scaling the RF drive to achieve the desired global power profile. Simultaneously, a digital trigger initiates a coordinated sweep of the control parameters $\underline{\alpha}$ within the uniformization module, ensuring trap depth uniformity is preserved throughout the ramp. 76

3.11 **Characterizing Post-Ramp Survival Probability with and without Uniformization of Trap Depths.** Comparison of atom survival probability distributions after a trap depth ramp in a 45×29 array. (a) When using only the mixer to perform the ramp, a strong column dependence is observed in the survival probability, with a high non-uniformity of $\sigma/\mu = 1.0$. (b) When using the mixer together with the uniformization and stabilization module to update the control parameters $\underline{\alpha}$ during the ramp, the survival probability becomes more homogeneous, with non-uniformity reduced to 0.5 and no significant column dependence. Residual spatial variation is attributed to initial non-uniformity in trap depths prior to the ramp, a consequence of the reduced control space dimensionality in a two-AOD system (see Fig. 3.4c). 78

3.12 **Thermometry of Single Atoms in Optical Tweezer Arrays.** Atoms are loaded into traps with initial depth U_0 and adiabatically ramped down to a final depth U_i. The survival probability of atoms is measured as a function of U_i and fitted using the model provided in Eq. (3.25). The fit yields an estimated atomic temperature of 24 μK. 79

4.1 **Microwave addressing of hyperfine transitions.** (a) Microwave radiation from a horn antenna drives hyperfine transitions between the $F = 1$ and $F = 2$ manifolds of the $5s_{1/2}$ ground state. In an external magnetic field, $\sigma_\pm$ components couple to seven allowed transitions, grouped into four distinct frequencies (colored). (b) Microwave spectroscopy is performed by sweeping the IQ mixer frequency and applying a π-pulse. With coarse optical pumping, four peaks corresponding to the transitions are observed, while fine pumping prepares atoms in $|F = 2, m_F = -2\rangle$, leaving only a single peak. (c) Experimental setup: a local oscillator at ~ 6.734 GHz is IQ-modulated by an AWG, amplified, and delivered to the atoms with a horn antenna. 88

4.2 **Magnetometry using Ramsey spectroscopy.** (a) Ramsey signal measured across a 45×29 trap array. The observed oscillations correspond to a Rabi frequency of 21 kHz, and the fitted decay envelope yields a coherence time of $T_2^* = 896$ μs. (b) Pulse sequence: a $\pi/2$ pulse is applied along the y axis, i.e. $\phi = \pi/2$, followed by a variable delay τ, and a second $\pi/2$ pulse with phase $\phi = \pi/2 + \nu\tau$, where $\nu = 15$ kHz. (c) Spatially resolved magnetometry. Atoms are grouped into 5×5 bins (with the final row containing 4 traps). For each bin, the Ramsey signal is averaged and the oscillation frequency is extracted. In the absence of a magnetic field, the frequency should match the reference $\nu = 15$ kHz. Deviations from this value reveal spatial variations in the magnetic field introduced by external coils. The observed detuning pattern indicates the presence of a magnetic field gradient along both the x and y axes, which we have intentionally introduced using magnetic coils. 89

4.3 **Quantum lock-in amplification sequence for electromagnetic field imaging.** A Ramsey spectroscopy sequence incorporating a train of equidistant π pulses between two $\pi/2$ pulses. The signal to be measured, in our case a near-infrared (NIR) electromagnetic field, is modulated at a frequency equal to half the repetition rate of the π pulses. This modulation ensures that the signal constructively interferes with the Ramsey sequence, while broadband noise is dynamically decoupled and suppressed. The technique acts as a quantum lock-in amplifier, enabling sensitive detection of weak, narrowband signals in the presence of dominant noise. Adapted from Ref. [1], Fig. 2b. 92

4.4 **Hardware setup for implementing local single-qubit gates using the Raman D_1 transition.** (a) Atomic transitions used for coherent control between $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 1\rangle$ via a Raman D_1 beam composed of two σ_+ polarized frequency components addressing the $F = 1 \to F' = 2$ and $F = 2 \to F' = 2$ transitions, both red-detuned from resonance by Δ. (b) Experimental setup for generating the optical tweezers and multiplexed D_1 Raman beams. The traps and Raman beams are each created using pairs of orthogonal AODs. The D_1 beam passes through a $4f$ imaging system before the AODs; the lens before the first D_1 AOD is mounted on an adjustable axial stage to align the focal plane of the D_1 beams with the traps for colocalization. Both trap and D_1 beams are focused onto the atoms using a microscope objective. The output is collected by a second objective ($f = 4$ mm) mounted on a rotation mount, with the beam imaged by a CMOS camera. The camera and objective are mounted on a translation stage enabling 3D beam imaging. 93

4.5 **Fine-tuning strategy for colocalizing trapping beams and D_1 Raman beams for local single-qubit gates.** (a) A pictorial representation of the linear model relating the RF tones supplied to the AODs and the corresponding beam positions. The linear maps for both trap and D_1 beams are defined in Eqs. (4.4)–(4.5). These mappings allow for transforming coordinates between frequency space and position space in both directions, enabling frequency-space matching of the D_1 beams to the traps. (b) A visual depiction of the RTS decomposition of the 2×2 mapping matrix $\underline{\underline{\mathcal{M}}}$ into a rotation matrix $\underline{\underline{R}}$, a skew matrix $\underline{\underline{T}}$, and a scaling matrix $\underline{\underline{S}}$, as described in Eq. (4.12). The diagram illustrates the geometric meaning of the decomposition parameters: rotation angle θ, skew factor κ, and axis-specific scalings s_x and s_y (see Eq. (4.14)). 95

4.6 **Characterization of colocalization between AOD-generated trap arrays and D_1 Raman beams for local single-qubit gates.** (a) Trap array (gray circles) and D_1 Raman beam array (red circles) after coarse alignment, showing significant spatial mismatch. (b) The arrays after applying the fine-tuning colocalization procedure described in Sec. 4.2.1, showing close overlap. (c) Measured distances between matched sites in the xy plane. The average xy separation is 0.3 μm. (d) Full 3D spatial distances between matched sites, with an average of 0.6 μm. The increase in separation near the bottom-right corner is attributed to breakdown of the linear mapping model in Eqs. (4.4)–(4.5) for high-frequency AOD components far from the center of the array. 99

5.1 **Optimization of Imaging Frequency to Minimize Atom Loss and Classification Error.** (a) Experimental sequence showing three consecutive EMCCD triggers, each corresponding to a fluorescence image, taken at fixed optical power of 20 μW per beam: the first and third images are each 300 ms long, and the middle image is 100 ms long with variable detuning. (b) Fluorescence photon counts as a function of imaging frequency, increasing towards the AC Stark shifted resonance near $+1\,\Gamma$. (c) Atom survival probability decreases near resonance due to reduced polarization-gradient cooling efficiency. (d) Figure of merit (atom loss per detected fluorescence photon, see Eq. (5.9)) demonstrating a trade-off between fluorescence rate and survival probability, minimized near $-10\,\Gamma$, indicating the optimal imaging detuning. 106

5.2 **Imaging loss and classification fidelity.** (a) Photon-count thresholds are used to classify traps as occupied, vacant, or unknown, with representative point-spread functions shown. (b) Applying this heuristic across the array yields classification fidelities above 50%. (c) Signal and background distributions are extracted from consecutive images; their overlap sets the classification error, which averages 0.2% over 1305 traps. (d) Continuous imaging of a 45×29 array gives an atom lifetime of 66.7 ± 0.02 s and a loss probability of 0.12%. 107

List of Tables

2.1 **Transformation Rules for Generalized Gaussian Beams.** Rules governing the evolution of the $\mathbf{\Lambda}(z)$ matrix under free space propagation and arbitrarily oriented biaxial lenses. $\mathbf{\Lambda}(z)$ completely describes the intensity profile of generalized Gaussian beams, making these rules sufficient for determining the beam shape of a generalized Gaussian beam. 35

3.1 **Timing Analysis of the Uniformization and Stabilization Module.** Time taken for each step in the closed-loop feedback cycle for a 45×29 trap array. Measurements are averaged over 5000 feedback cycles. The total cycle duration is limited mainly by the CMOS camera readout time, which could be improved by using cameras with faster readout. 70

5.1 **Imaging Performance Metrics.** Summary of classification fidelity, atom loss probability, and imaging time for the optimized imaging configuration. 108

Chapter 1

Introduction

Quantum processors are emerging as powerful tools for solving computationally expensive problems that are intractable on classical computers, particularly those involving the dynamics of complex quantum systems. These platforms provide solutions for simulating strongly correlated materials, discovering new quantum phases of matter, and performing fault-tolerant quantum computation. Quantum simulation enables the emulation of many-body quantum systems (see Fig. 1.1), providing insight into challenging problems such as spin liquids and high-temperature superconductivity, which are difficult to study using classical methods due to the exponential scaling of the Hilbert space [2, 3]. Beyond simulation, scalable quantum processors also establish the foundation for universal quantum computation protected by error correction protocols [4]. A variety of competing hardware platforms exist for developing quantum processors. These include superconducting qubits [5], trapped ions [6], photonic systems [7], and neutral atom arrays [8], are being developed and refined to realize universal quantum computation.

Among these platforms, Rydberg atom array quantum processors have shown significant promise for enabling large-scale quantum information processing by demonstrating key properties such as scalability and controllability. Scalability has been illustrated by arrays exceeding 6100 qubits [9], underscoring the potential of the platform to reach system sizes relevant for practical quantum computing [10]. In terms of controllability, these systems allow for flexible, site-resolved addressing of individual qubits and have achieved high-fidelity two-qubit entangling gates operating in parallel [11]. Rydberg atom arrays have been applied to quantum simulation, facilitating the exploration of quantum phases and many-body dynamics including phase transitions, spin liquids [2, 3, 12], and spin transport phenomena [13, 14]. Beyond analog simulation, fault-tolerant quantum computation architectures have been proposed and experimentally realized on Rydberg platforms [8, 15],

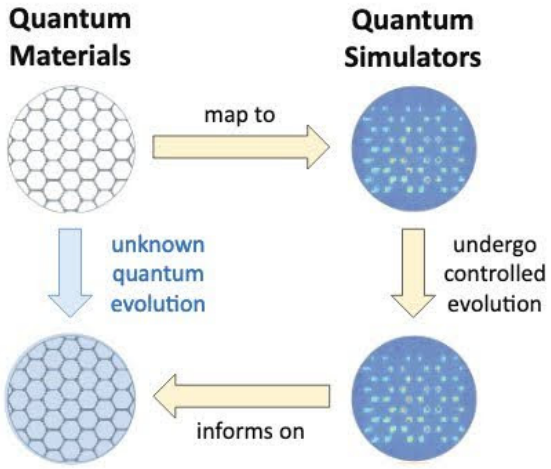


Figure 1.1: **Quantum Simulation for Material Discovery.** Quantum processors can serve as quantum simulators by encoding and controlling the dynamics of quantum many-body systems. This approach enables the study of strongly correlated systems and accelerates the discovery of new materials with desired properties.

including the preparation of logical Greenberger–Horne–Zeilinger (GHZ) states and teleportation protocols on encoded qubits [8]. Additionally, these arrays offer potential for quantum metrology and sensing applications [16]. For example, tweezer-trapped alkaline-earth atoms have demonstrated half-minute coherence times on optical clock transitions, enabling ultra-stable frequency references with implications for precision timekeeping [17]. Moreover, their sensitivity to electromagnetic fields, allow Rydberg atom arrays to be used for precise and high-resolution imaging of electromagnetic fields [18].

In this chapter, we first familiarize the reader by presenting an example of a ^{87}Rb-based Rydberg atom array quantum processor and describe how it enables applications in quantum simulation and fault-tolerant computation. This serves to familiarize the reader with the architecture and capabilities of Rydberg platforms. We then outline key challenges in building large-scale processors of this type, focusing on scalability, reliability, and controllability, and summarize how each is addressed in the subsequent chapters of this thesis.

1.1 Neutral Atom Arrays as a Quantum Platform

This section first outlines a representative ^{87}Rb-based Rydberg atom array quantum processor architecture and its physical properties (see Sec. 1.1.1) to familiarize the reader with this platform and how neutral atoms can be used as a quantum resource. Then, it discusses key applications in quantum simulation, computation, and sensing (see Sec. 1.1.2).

1.1.1 Architecture and Physical Properties

We now present an example architecture for a ^{87}Rb Rydberg atom array quantum processor, as illustrated in Fig. 1.3. The system begins with stochastic loading of single ^{87}Rb atoms into an array of optical tweezers formed by far-off-resonant traps (FORTs) at 813 nm. These traps are loaded from a magneto-optical trap (MOT) or a dipole trap. For appropriately chosen trap depths and atomic densities, the atoms enter the collisional blockade regime [19], in which light red-detuned from the D_2 transition (780 nm) preferentially ejects multiple-atom occupancies. As a result, each site is left with either zero or one atom. This probabilistic loading yields a partially filled array of traps containing single atoms.

Once loaded, atoms can be rearranged using dynamic tweezer control to form defect-free arrays with arbitrary geometries [8, 20, 21, 22]. The rich internal atomic structure of ^{87}Rb allows for versatile control over internal states and interactions. The D_1 line (795 nm) enables Raman transitions for sideband cooling and single qubit rotations (see Fig. 4.4a), while the D_2 line (780 nm) is used for polarization-gradient cooling, imaging (see Sec.5), and qubit state preparation and measurement (see Sec.4.1.1). Qubit states are encoded in long-lived hyperfine ground states, here defined as $|0\rangle = |F = 1, m_F = -1\rangle$ and $|1\rangle = |F = 2, m_F = -2\rangle$, as shown in Fig. 1.3(b). These states can be coherently coupled using microwave pulses near 6.8 GHz (see Sec. 4.1.2) or optically via Raman transitions (see Sec.4.2).

Spatially resolved single- and multi-qubit gates are achieved using optical beams [8] and microwave fields [23], with the experimental geometry of the beams and pulses illustrated in Fig. 1.3(c,d). High-fidelity single-qubit gates are implemented with globally or locally addressed beams, while two-qubit entangling gates rely on Rydberg interactions. These gates are realized by coherently exciting atoms to high principal quantum number states using a two-photon transition via 420 nm and 1013 nm lasers. In the Rydberg state, atoms interact via strong van der Waals or dipole-dipole forces, which can be exploited to implement multi-qubit gates [11, 12, 24].

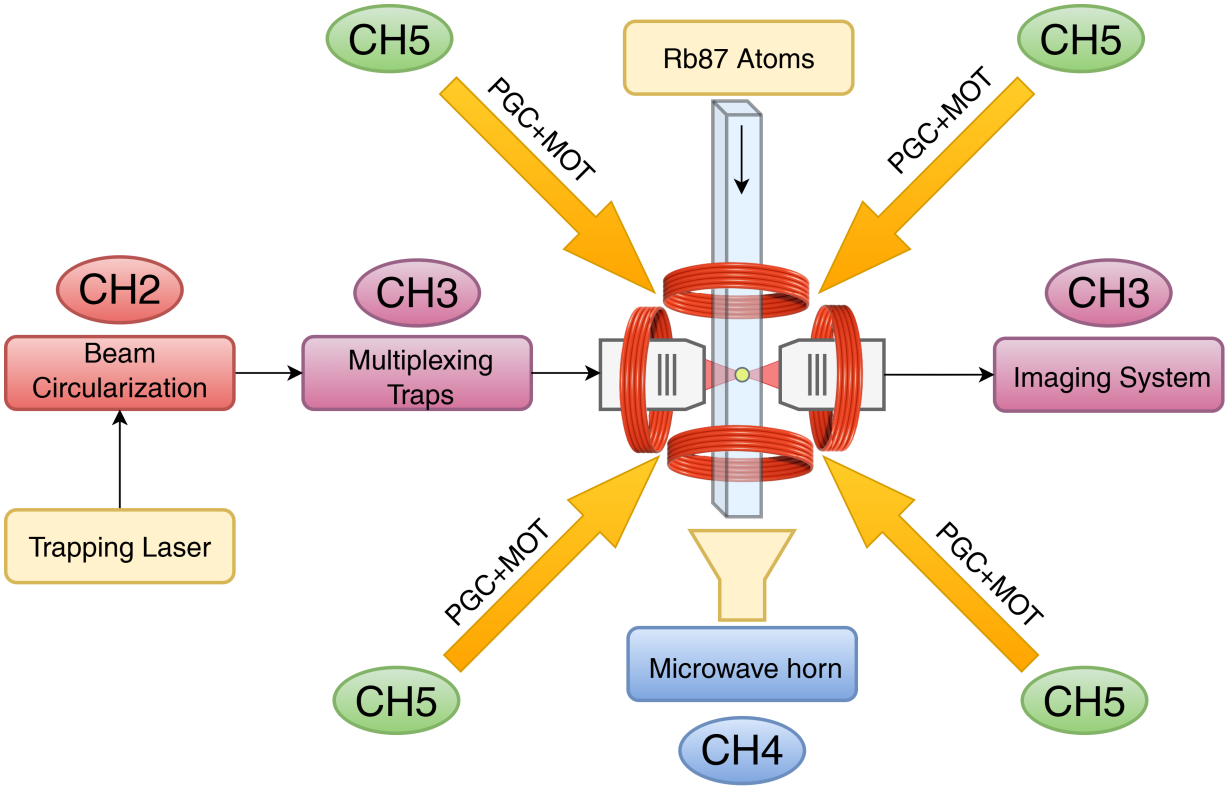


Figure 1.2: **Experimental Setup.** ^{87}Rb atoms are are dispensed from the atom source at the top of the vacuum cell. They are then trapped in a magneto-optical trap (MOT) generated by six MOT beams (orange arrows) and coils (red). After cooling, atoms are transferred into optical tweezers formed by 813 nm beams focused and collected by two microscope objectives (gray lenses). Chapters 2 and 3 describe trap engineering, including beam circularization (red) and 2D trap array multiplexing with power stabilization (purple boxes). Atoms are addressed with a microwave pulse from the microwave horn located beneath the vacuum cell (blue box). This setup is used for sensing electromagnetic fields as discussed in Chapter 4 across the trap array. Lastly, the six polarization-gradient cooling beams are used for imaging atoms with polarization-gradient cooling, optimized in Chapter 5 to maximize fidelity and minimize loss.

We have realized this architecture in our laboratory, following the design illustrated in Fig. 1.2. In our implementation, ^{87}Rb atoms are first loaded from the top atom dispenser into a magneto-optical trap (MOT) formed by six MOT beams (orange arrows) and surrounding magnetic coils (red). After cooling, single atoms are transferred into optical

tweezers created by 813 nm beams focused and collected by two microscope objectives (gray lenses). For tasks involving qubit control, microwave pulses are applied from below the vacuum cell (blue box), enabling coherent rotations and state manipulation. Imaging and feedback for trap stabilization are performed using the same six beams in a polarization-gradient cooling configuration, with parameters optimized to maximize detection fidelity and minimize atom loss, as detailed in Chapters 2–5.

An example of a trap array is shown in Fig. 1.4(a,b), where a 45×29 array of 1305 optical tweezers exhibits an average single-shot loading of 690 atoms on average, corresponding to a 53% filling fraction. Fluorescence imaging using the D_2 transition allows for reliable detection of individual atoms, with clear separation between signal and background levels in photon count histograms (Fig. 1.4(c)). Coherent control of qubit states is demonstrated through microwave Rabi spectroscopy, where coherent population oscillations between $|0\rangle$ and $|1\rangle$ are observed (Fig. 1.4(d)), indicating the ability to manipulate qubit populations with high fidelity.

1.1.2 Applications in Simulation, Computation, and Sensing

These atomic transitions, illustrated in Fig. 1.3(a,b), form the basis for coherent qubit state manipulation, which underlies both quantum simulation and general-purpose quantum computation. For quantum simulation, these platforms have enabled the exploration of rich many-body phenomena, including quantum spin liquids [3], topological spin models [24], and quantum phase transitions [2]. In the context of computation, the same gate toolbox can be extended to realize fault-tolerant architectures. In particular, high-fidelity multi-qubit gates and mid-circuit measurement capabilities have enabled the realizations of surface-code logical qubits [8]. Furthermore, the dynamic architecture of neutral atom arrays [21, 25], which allows atoms to be rearranged during operation, facilitates the implementation of more general low-density parity check (LDPC) codes that require non-local stabilizers [26], making Rydberg atom arrays a flexible platform for scalable error correction.

Importantly, the same transitions used for qubit operations are also sensitive to external fields via the Zeeman effect and Stark shifts, making them valuable probes for quantum sensing. Local variations in magnetic or electric fields shift the resonance frequencies of qubit transitions, enabling the use of Rydberg arrays as high-resolution precision sensors [18]. Moreover, this platform allows the development of highly stable atomic clocks based on tweezer-trapped alkaline-earth atoms, which achieve minute-scale coherence [17]. These capabilities demonstrate that Rydberg atom arrays are not only viable for computation and simulation, but also for advancing precision measurement and fundamental physics.

1.2 Challenges in Building Neutral Atom Array Quantum Processors

While Rydberg atom array quantum processors are promising platforms for computation, simulation, and sensing, they do not yet meet the performance and scale required for many industrially relevant applications [10]. Significant technical challenges remain in scaling up system size, improving operational fidelity, and ensuring precise control over qubits and environmental noise sources.

In this section, we outline several key challenges (see Fig. 1.5) in building large-scale Rydberg atom array quantum processors. These challenges fall into three main categories: scalability, reliability, and controllability. For each, we pose specific research questions

relevant to this thesis. We then describe how the remainder of this thesis addresses these questions through experimental and computational advances.

1.2.1 Scalability

Scalability refers to the challenge of increasing the number of qubits in a Rydberg atom array quantum processor, along with addressing the difficulties that arise in manipulating a larger number of qubits. Increasing the total number of qubits can be achieved either by increasing the number of optical traps or by improving the filling fraction of atoms within those traps [27]. However, the number of traps is limited by several practical constraints: the available laser power and the power required per trap [28], the field of view of the microscope objective that focuses the tweezers, and physical limits such as damage thresholds of optical elements and temperature-related issues.

One approach to reduce the required trapping power per trap is to improve the quality of the laser beam shape. Imperfections such as astigmatism and ellipticity distort the beam profile, reducing the effective trap depth after focusing. These distortions commonly occur in the outputs of solid-state lasers, such as Ti:sapphire lasers, which are used for far-off-resonance trapping of ^{87}Rb atoms. Existing solutions either rely on custom-manufactured aspheric lenses [29], which are difficult and expensive to produce, or on active optical elements like spatial light modulators [30], which typically have low power throughput. This raises the question of whether a passive optical solution using standard components with high power transmission can be developed to correct astigmatism and circularize the laser output efficiently (see Sec. 2.1).

As the number of qubits grows, maintaining consistently high gate fidelities and state preparation and measurement (SPAM) fidelities across the entire array becomes more difficult. Many quantum computations require every qubit to be correctly prepared and measured. As the system size increases, even small errors on individual qubits accumulate, making the overall error rate significantly higher. This means that the probability of encountering at least one error in the whole array rises sharply with more qubits, creating a strong demand for continuous improvements in SPAM performance as quantum processors scale up.

Fluorescence imaging is a critical tool for quantum state measurement, but it is subject to classification errors and atom loss. This motivates the question of how to optimize the beam properties of the imaging light to minimize these errors and maximize fidelity across large qubit arrays (see Sec. 5).

1.2.2 Reliability

Another key challenge in building practical Rydberg atom array quantum processors is reliability. Achieving high-fidelity qubit gates, state preparation, and imaging requires careful calibration of experimental parameters. For example, the frequency and polarization of optical pumping beams used for state preparation must be precisely tuned, and Rydberg excitation pulses must be fully overlapped with the traps. However, system properties can drift or degrade over time, necessitating repeated recalibrations. This instability limits the ability of the processor to reliably perform quantum computations continuously over long durations, which may be required for many industrially relevant algorithms that can run for weeks or even months [10].

To address this, one can either reduce the sources of drift and miscalibration or implement real-time feedback loops that continuously stabilize key experimental parameters. For example, PID controllers are commonly used to stabilize the power of cooling, trapping, and Rydberg lasers.

A particularly important parameter is the trap depth of the optical tweezers. Since the trapping potential causes AC Stark shifts on atomic transitions, calibration of frequencies for cooling, imaging, and state preparation beams must account for these shifts. If trap depths fluctuate or drift, many other parameters must be recalibrated as well.

However, these systems often contain many traps created by diffracting a single trapping beam through acousto-optic deflectors (AODs), as in several platforms [8, 11]. Although stabilizing the overall trapping laser power is necessary, it is insufficient because AOD diffraction efficiencies can drift, causing relative trap depths to fluctuate over time.

We therefore ask whether a real-time feedback system can be designed to continuously monitor and redistribute laser power among traps to maintain uniform, stable trap depths over extended periods. Furthermore, we explore whether such stabilization can be performed concurrently with quantum computations without disrupting operations (see Sec. 3).

1.2.3 Controllability

Controllability refers to the ability to perform precise, high-fidelity operations on individual and multiple qubits across the array. This capability is critical for reliable quantum information processing, accurate qubit readout, and applications in quantum sensing. In our platform, both global and site-resolved single-qubit gates are available via microwave

and optical Raman transitions. Given this level of control, we ask: how can we leverage these tools to enable high-resolution sensing of local electromagnetic fields, and what strategies are needed to suppress technical noise and achieve uniform qubit control across large arrays?

To address these questions, we first study spatial non-uniformities in the qubit transition frequency, which arise from position-dependent magnetic and AC Stark shifts. In Sec. 4.1.3, we use site-resolved microwave Ramsey spectroscopy to measure the local qubit resonance at each trap. This allows the mapping of spatial frequency offsets to background magnetic and electromagnetic fields, enabling high-resolution quantum imaging.

We then ask whether the D_1 Raman beams used for optical gates can be aligned with sufficient precision to ensure uniform, site-resolved control. Misalignment leads to spatial variations in gate fidelity due to beam inhomogeneity, reduced Rabi frequencies, and increased cross-talk between qubits. In Sec. 4.2, we develop and implement a procedure to precisely colocate the Raman beams with the atom array, improving the uniformity and fidelity of single-qubit gates across the system.

1.2.4 Outline of the Thesis

This thesis is structured around addressing the core challenges of scalability, reliability, and controllability in Rydberg atom array quantum processors. Each chapter targets one or more of these aspects:

- **Chapter 2:** The available laser power limits the maximum number of traps. Each trap must exceed a minimum depth to allow high-fidelity, low-loss imaging of atoms. To address this, we reduce the beam radius at the focus, which increases the intensity and lowers the power required per trap to maintain the same depth. We achieve this by implementing a passive, three-lens system that circularizes and de-astigmatizes the trapping beam. The design uses three cylindrical lenses, and relies entirely on standard, high-throughput optical components. This improves trap depth and uniformity without resorting to expensive custom optics or power-inefficient active elements such as spatial light modulators. As a result, more traps can be generated from the same laser source, advancing towards system sizes relevant for quantum advantage, with broader applicability to other applications such as laser machining and high-efficiency coupling of light to optical fibers.

- **Chapter 3:** High-fidelity quantum state preparation, imaging, and gate operations require careful calibration of experimental parameters, which are highly sensitive

to optical trap depth. Because atomic transitions experience AC Stark shifts from the tweezers, all calibrations must be matched to the specific trap depths used during computation. However, trap depths can drift over time—even when total laser power is stabilized—due to fluctuations in the relative diffraction efficiency of AOD-generated traps. To address this, we implement a real-time closed-loop feedback system that stabilizes the relative depths of individual traps using fast in-situ measurements. The noise in trap depths is decomposed into uncorrelated temporal modes, each regulated by an independent PID loop. This eliminates the need for frequent recalibration, maintaining consistent quantum control over long experiments. The system enables higher shot rates and supports extended quantum information processing without fidelity degradation.

- **Chapter 4:** We demonstrate controllability through two complementary developments. First, we perform magnetic field sensing via microwave Ramsey spectroscopy on a 45×29 trap array, mapping spatial variations in the qubit transition frequency due to magnetic field gradients and AC Stark shifts. This work lays the foundation for advanced quantum sensing protocols such as quantum lock-in amplification, enables in-situ electromagnetic field nulling using atoms as local sensors, and paves the way towards entanglement-based sensing of electromagnetic fields. Second, we develop an automated beam alignment algorithm that achieves near-perfect colocalization of the D_1 Raman beams with the trap array. This reduces gate cross-talk and increases the maximum achievable Rabi frequency, thereby enhancing single-qubit control fidelity across large arrays.

- **Chapter 5:** We perform a systematic optimization of the imaging beam parameters, specifically the power and detuning, to achieve a balance that minimizes both atom loss and classification errors during fluorescence imaging. By carefully tuning these parameters, we enhance the quantum state measurement fidelity, which is critical for accurate readout in large-scale neutral atom arrays. This improvement directly reduces measurement-induced errors, thereby enabling more reliable quantum information processing and state preparation verification across hundreds of qubits.

Collectively, these contributions provide experimental tools and protocols that improve the scalability, stability, and control required for advancing Rydberg atom arrays towards practical quantum information processing.

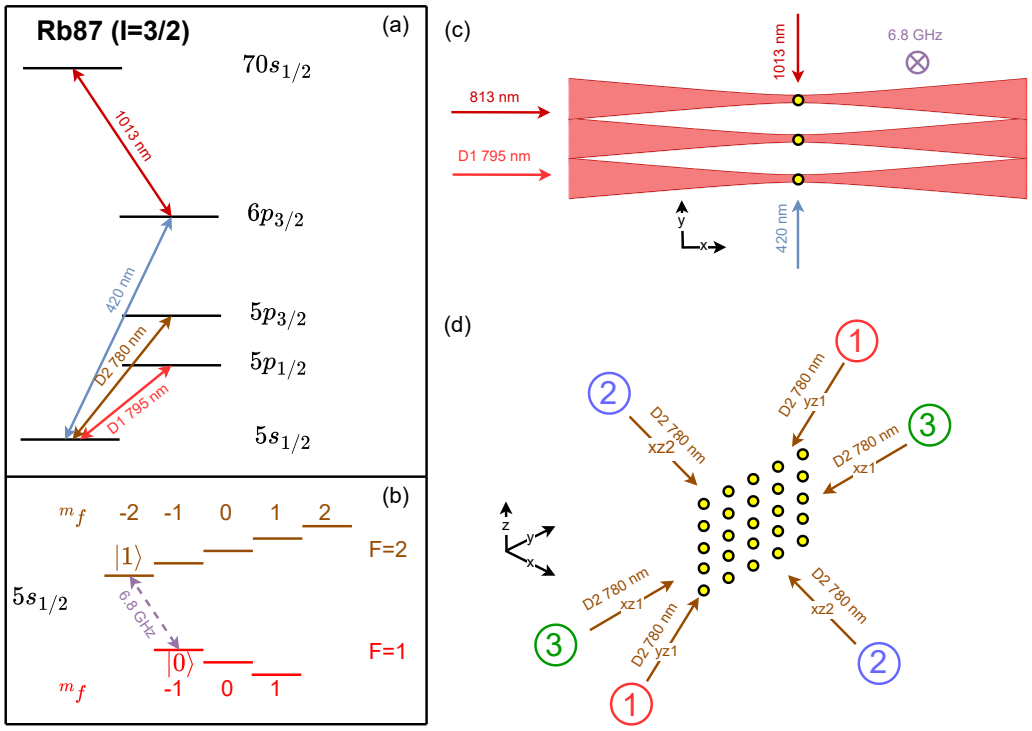


Figure 1.3: **Architecture of a ^{87}Rb Rydberg Atom Array Quantum Processor.** Overview of a ^{87}Rb-based Rydberg atom array quantum processor. (a) Atomic transitions used to control ^{87}Rb atoms. The D_1 line enables hyperfine transitions via Raman pulses, as well as enabling Raman sideband cooling. The D_2 line is used for polarization-gradient cooling, pushout, and fluorescence imaging. Two-photon Rydberg excitation is performed using 420 nm and 1013 nm beams. (b) Hyperfine structure of the ground state and definition of the qubit states as $|0\rangle = |F = 1, m_F = -1\rangle$ and $|1\rangle = |F = 2, m_F = -2\rangle$. Transitions between the qubit states are driven with a 6.834 GHz microwave pulse. (c) Spatial configuration of the atoms and beam paths used for implementing optical or microwave-based single- and multi-qubit gates. The atoms are trapped in arrays of optical tweezers at a wavelength of 813 nm. The D_1 beams are incident on the atoms along the same direction as the traps, while the Rydberg beams (420 nm and 1013 nm) are incident from opposite directions along the orthogonal axis. The microwave horn addresses the atoms with pulses traveling perpendicular to the page at 6.8 GHz. (d) Configuration of the D_2 beams used for cooling, atom ejection, and imaging. Two counter-propagating beam pairs lie along the diagonal and anti-diagonal of the xz plane. A third counter-propagating pair lies along the diagonal of the yz plane.

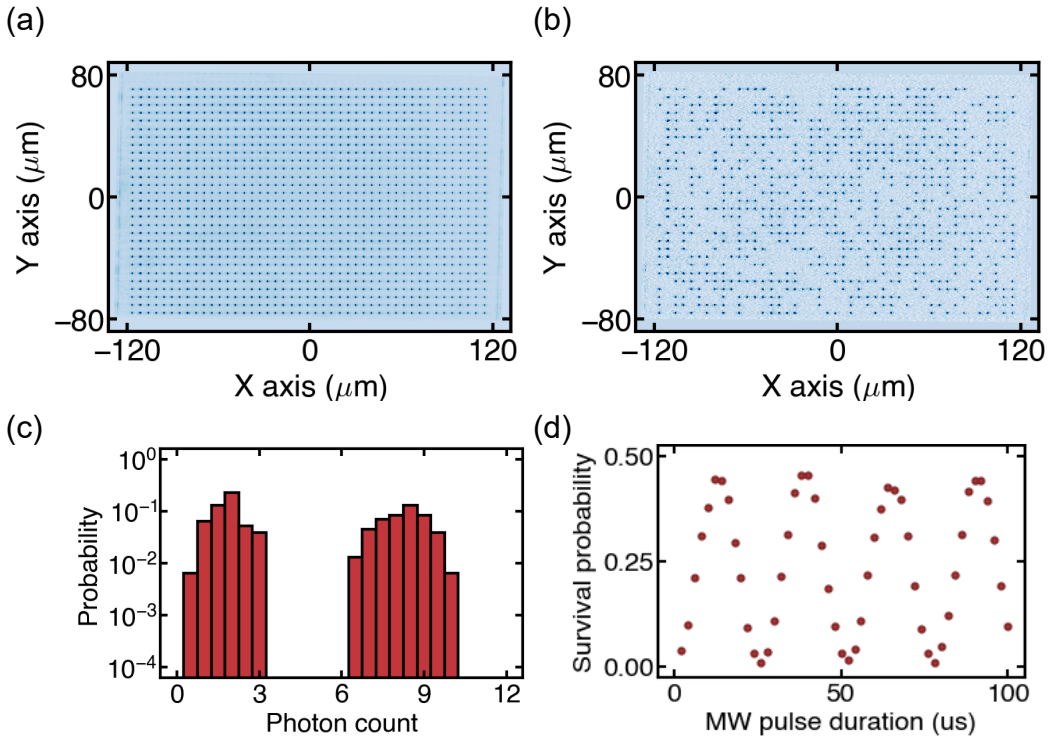


Figure 1.4: **Qubits in a ^{87}Rb Rydberg Atom Array Quantum Processor.** Development and characterization of a ^{87}Rb-based Rydberg atom array platform. (a) Average fluorescence image from a 45×29 trap array (1305 traps) with a loading probability of 53%. Each dark spot corresponds to fluorescence from a single atom at 780 nm. (b) Single-shot image showing approximately 690 occupied traps and 615 empty traps. (c) Photon count histogram from a single trap during imaging, showing distinct peaks for background (vacant) and atomic signal (occupied). (d) Microwave Rabi spectroscopy performed on the array, showing population oscillations between the qubit states. A global π-pulse occurs near 25 μs, corresponding to a full population transfer to the $|1\rangle$ state (see Sec. 4.1.2). The achieved population is limited by the qubit gate fidelity.

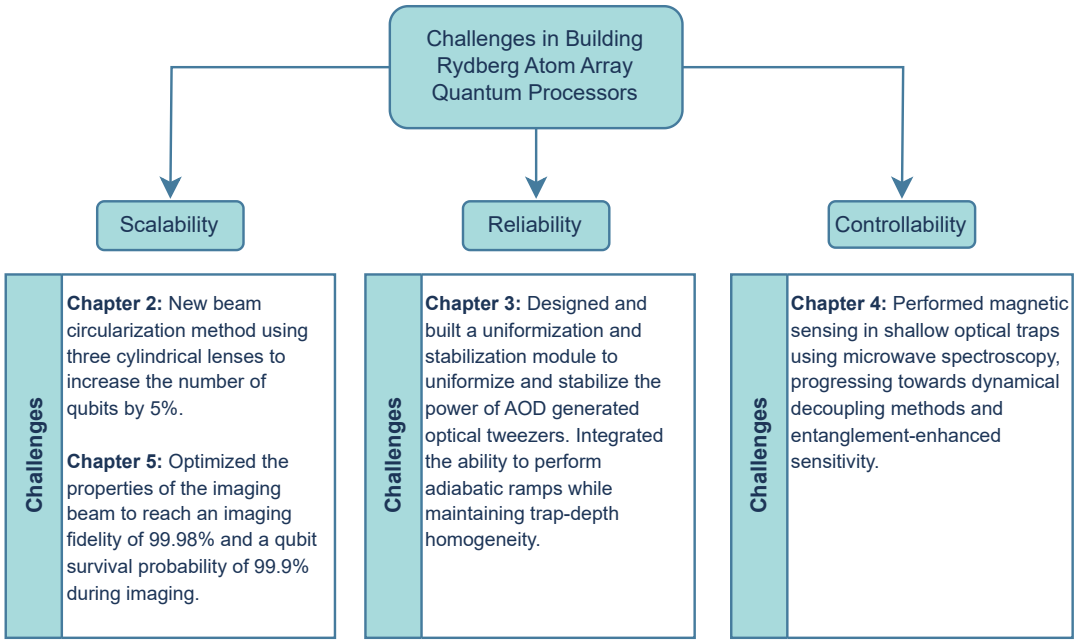


Figure 1.5: **Outline of the Thesis.** Overview of the contributions of this thesis to a set of key challenges in building Rydberg atom array quantum processors. **Scalability:** Chapter 2 introduces a new beam circularization method using three cylindrical lenses to increase the number of qubits by 5%, and Chapter 5 optimizes imaging to reach a fidelity of 99.98% with a qubit survival probability of 99.9%. **Reliability:** Chapter 3 develops a uniformization and stabilization module for AOD-generated tweezers, producing trap arrays with homogeneous and stabilized trap depths, and enabling adiabatic ramps while maintaining homogeneity. **Controllability:** Chapter 4 demonstrates magnetic sensing in shallow optical traps using microwave spectroscopy, advancing toward dynamical decoupling and entanglement-enhanced sensitivity.

Chapter 2

Correcting astigmatism and ellipticity in Gaussian beams using a cylindrical lens pair with tunable focal lengths

A major limiting factor in increasing the number of qubits in neutral atom array quantum processors is the limited optical power available for generating optical tweezers. Each trap must provide sufficient depth to enable high loading efficiencies and support lossless imaging and cooling of single atoms. Typically, the required trap depth is on the order of 1 mK [28].

We propose a method to reduce the optical power required per tweezer by improving the quality of the trapping beam. Specifically, we remove astigmatism and ellipticity, which are common features in the output of the solid-state lasers such as Ti:sapphire systems used for generating arrays of optical tweezers. This correction reduces the trap size at the focal plane of the microscope objective, increasing the beam intensity and, consequently, the trap depth for the same input power. As a result, a larger number of optical tweezers can be generated from a fixed total power (see Fig. 2.1).

Our solution consists of three standard cylindrical lenses, two of which act as an effective biaxial lens with tunable focal lengths. It provides higher power throughput compared to active beam-shaping techniques based on diffractive optical elements such as spatial light modulators [30], and avoids the need for custom-manufactured aspheric lenses tailored to the laser output [29]. The approach is cost-effective, compact, and easy to align, making

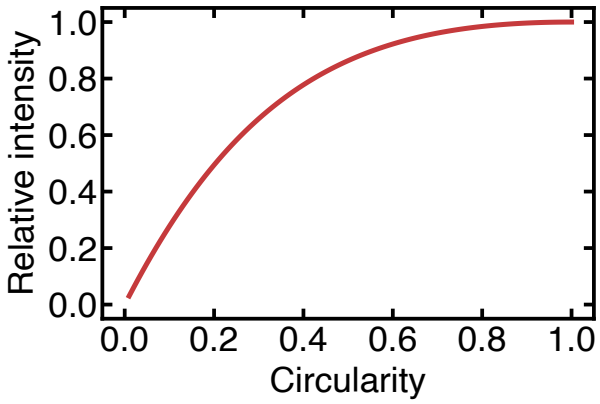


Figure 2.1: **Focusing elliptical Gaussian beams.** The figure shows the simulated intensity of an elliptical Gaussian beam focused with a Mitutoyo microscope objective of focal length $f = 4$ mm and aperture radius 4 mm, plotted as a function of the input beam circularity. For each circularity, the beam radii are chosen to maximize the intensity at the focal plane. The maximum focal intensity is achieved for a circular input beam (circularity = 1.0). As the input circularity decreases towards zero, the achievable focal intensity is reduced.

it suitable for a variety of applications where precise beam shaping is essential, including laser machining and high-efficiency fiber coupling [31].

We use this solution on our trapping laser, which has a circularity of 0.69, to a beam with a circularity of 0.97 and almost no astigmatism, i.e. a beam waist plane separation of 0.8% of the beam Rayleigh range.

The content of this chapter is based on the recent paper by the author of this thesis, Anastasiia Mashko, and Prof. Alexandre Cooper-Roy [31].

Abstract. Correcting astigmatism and ellipticity in laser beams is critical for improving performance in many applications like microscopy, atomic physics, quantum information processing, and advanced manufacturing. Passive correction methods based on cylindrical lens telescopes require choosing lenses with precise focal lengths, effectively limiting the range of tunability when using standard catalog optics. Active solutions based on diffractive optical elements can achieve superior performance, but they are bulky, expensive, and suffer from finite diffraction efficiency and added complexity. Here, we introduce a simple method to convert astigmatic elliptical beams into circular Gaussian beams without astigmatism. Our method comprises three cylindrical lenses. The first lens focuses the beam along its

major axis to create a plane where the intensity profile is radially symmetric. The second and third lenses are placed one behind the other in that plane at a relative angle, acting as a biaxial lens pair with tunable focal lengths. By adjusting the relative angle of the lenses, the two separate beam waists of the astigmatic beam can be overlapped, resulting in a circular Gaussian beam without astigmatism. We theoretically validate our method, numerically quantify its robustness to experimental imperfections, and experimentally demonstrate its ability to circularize the output beam of a commercial laser source. Our method corrects astigmatism and ellipticity in laser beams without requiring precise focal length matching, offering greater flexibility than other passive solutions and greater cost-effectiveness than active methods. Its simple and compact design makes it well suited for integration into both tabletop optical setups and industrial systems.

2.1 Background

Correcting aberrations and imperfections in laser beams emitted by commercial laser sources is a common problem in laser optics [32, 33, 34, 35, 36, 37]. This problem is typically encountered when operating free-space-emitting diode lasers, which produce elliptical and astigmatic beams [35, 38]. It is also encountered when operating titanium:sapphire lasers, where astigmatism results from the Brewster-angle cut of the gain crystal and the potential alignment errors within the optical cavity [32, 39, 40]. Addressing this problem is critical in many applications, such as those involving high-power optical systems where fiber coupling and spatial filtering are not viable solutions.

A common approach to circularizing an elliptical beam (see Fig. 2.2) involves two cylindrical lenses oriented along the principal axes of the beam (see Fig. 2.3). This approach benefits from low cost and simplicity; however, it requires precisely choosing the focal lengths of the lenses according to a precise ratio, which is not always possible when sourcing lenses from standard optics catalogs. These lenses may also deviate from being perfectly uniaxial, leading to unwanted changes in the beam divergence along the weak axis of the cylindrical lens, thus preventing circularization. Another approach to circularizing the beam emitted by diode lasers involves using anamorphic prism pairs combined with a circular clipping aperture. This method benefits from a compact and tunable design, low optical loss, high-power handling, and intrinsic correction of chromatic dispersion. However, these prism pairs correct only for the ellipticity of the beam far away from the beam waist by expanding or compressing the beam in one direction; they do not affect the relative separation between the beam waists, and thus do not correct astigmatism [41, 37]. Other passive circularization methods exist and have demonstrated high circularity and

astigmatism correction (see e.g., Refs. [42, 43, 36, 29]); however, these approaches provide limited tunability and depend on specialized, custom-made components tailored to the properties of the input beam.

Active methods based on active optical diffracting elements, such as spatial-light-modulators, benefit from a high level of tunability, enabling the active suppression of high-order aberration terms [44, 45]; however, these elements are bulky, expensive, and rely on complex software optimization to achieve optimal performance. They also suffer from power losses due to their finite diffraction efficiency, making them unsuitable for compact, energy-efficient, portable systems. Their cost and complexity are typically not justified for simple, compact applications such as maximizing the coupling efficiency of light in an optical fiber [46].

In this section, we introduce a simple method to correct astigmatism and ellipticity in laser beams. Our solution employs three standard cylindrical lenses mounted on rotation mounts (see Fig. 2.4). The first cylindrical lens enables the creation of a plane in which the intensity profile is circular. The other two cylindrical lenses are placed one behind the other in that plane at a relative angle, acting as an effective biaxial lens with tunable focal lengths that can match the divergence angles along the two principal axes of the beam. Our method reduces alignment constraints by eliminating the need for precise selection of the focal lengths of these lenses. It is both compact and flexible, allowing for the use of off-the-shelf components with a broad range of focal lengths. Our method remains effective even when the cylindrical lenses are not perfectly uniaxial, accommodating manufacturing variations without significant loss of performance. Our solution can be used to circularize laser beams in various settings such as neutral-atom quantum computing, where high-power, circular Gaussian beams are critical. For example, these circular beams enable trapping individual neutral atoms in large arrays of optical traps [45, 30, 47], displacing atoms to correct for lost atoms in large arrays of optical tweezers [20, 21, 22, 48], and performing site-selective quantum gates with minimal crosstalk [49, 50]. A similar solution has been used to construct an anamorphic beam expander for adjusting the magnification and aspect ratio of a laser beam at a fixed plane [51]. It is also employed in the Jackson Cross Cylinder, which uses two cylindrical lenses with equal but opposite focal lengths mounted one behind the other to correct astigmatism in ophthalmology [52].

The organization of this section is as follows. We first describe the problem of correcting astigmatism and ellipticity in Gaussian beams. We then describe our three-lens solution using a simple analytical model that describes the propagation of Gaussian beams through optical components. We show that two counter-rotated cylindrical lenses act as an effective biaxial lens with tunable focal lengths that can correct astigmatism and ellipticity using a preceding cylindrical lens. We then show that the minimum circularity of the output

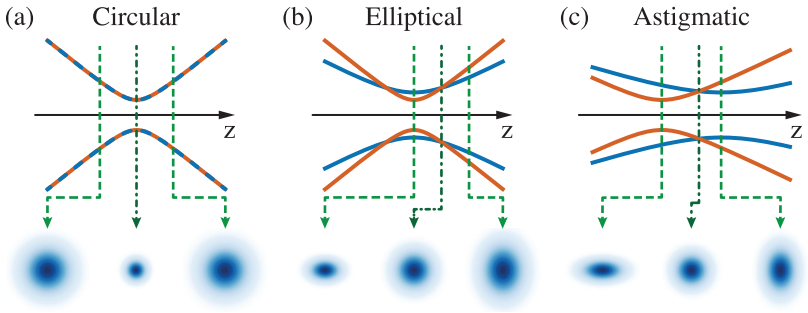


Figure 2.2: **Gaussian beams.** (a) A circular Gaussian beam has a circular intensity profile with radial symmetry. (b) An elliptical Gaussian beam has distinct divergence angles along its major (blue) and minor (red) axes. The beam radius along each axis is minimized in the same focal plane, where the intensity profile is elliptical. The intensity profile becomes circular in two planes symmetrically located on either side of the beam waist. (c) An astigmatic Gaussian beam is an elliptical beam whose beam waists do not overlap. There exists at least one plane where the intensity profile is circular.

beam is robust against variations in the separation distance between the two lenses of the pair (see Sec. 2.2.2). We finally validate our model by mapping the effect of a tunable biaxial lens on a circular beam and by circularizing the output beam of a commercial laser source (see Sec. 2.3.1).

2.2 Theory – Gaussian Beam Circularization

We seek a simple solution to the problem of converting an astigmatic elliptical beam into a circular Gaussian beam (see Fig. 2.2). We refer to a circular Gaussian beam (see Fig. 2.2a) as the fundamental TEM00 mode described by the Laguerre-Gaussian polynomials, which are solutions to the Helmholtz equation under the paraxial approximation [53]. The electric field vector of a circular Gaussian beam is given by

$$\boldsymbol{E}(x,y,z) = \boldsymbol{E}_0 \frac{w_0}{w(z)} \exp\left[-\frac{x^2+y^2}{w^2(z)}\right] \exp\left[-jk\frac{x^2+y^2}{2R(z)}\right] \exp\left[jkz - j\arctan\left(\frac{z}{z_R}\right)\right], \tag{2.1}$$

where $\boldsymbol{E}_0$ is the field normalization vector, $w(z) = w_0\sqrt{1+(z-z_0)^2/z_R^2}$ is the position-dependent beam radius at position z along the propagation axis, $w_0 = w(z_0)$ is the beam radius at the beam waist $z = z_0$, $z_R = \pi w_0^2/\lambda$ is the Rayleigh range, λ is the wavelength,

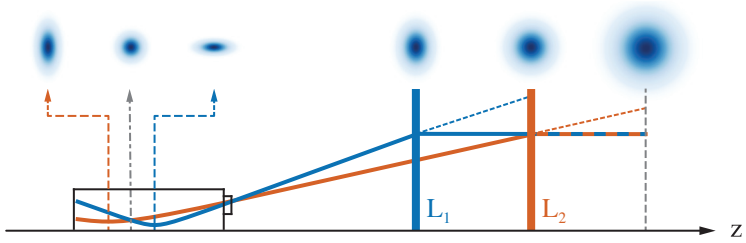


Figure 2.3: **Typical Two-Cylindrical-Lens Solution.** An astigmatic elliptical Gaussian beam can be converted into a circular Gaussian beam by placing two cylindrical lenses, L_1 and L_2, each oriented along one of the principal axes of the diverging beam. Each cylindrical lens is placed approximately one focal length away from its corresponding beam waist.

and

$$R(z) = (z - z_0)\left[1 + \left(\frac{z_R}{z - z_0}\right)^2\right] \tag{2.2}$$

is the radius of curvature of the phase front.

The transverse intensity profile of a circular Gaussian beam in any z plane is given by a two-dimensional Gaussian distribution with radial symmetry,

$$I(x, y, z) = \frac{1}{2}\epsilon c|\boldsymbol{E}(x, y, z)|^2 = \frac{1}{2}\epsilon c|\boldsymbol{E}_0|^2 \frac{w_0^2}{w^2(z)} \exp\left[-\frac{2\,(x^2 + y^2)}{w^2(z)}\right], \tag{2.3}$$

where ϵ is the electric permittivity, and c is the speed of light.

The circular Gaussian beam is an idealization that often fails to describe the output beam of a commercial laser source. For example, as mentioned in the introduction, cavity-based lasers often contain an active medium crystal typically oriented at the Brewster angle to minimize reflection loss. This crystal breaks axial symmetry, resulting in beams whose intensity profiles are not radially symmetric [32, 39]. To describe Gaussian beams with ellipticity and astigmatism (see Fig. 2.2b-c), we instead rely on a more general solution to the paraxial Helmholtz equation,

$$\boldsymbol{E}(x, y, z) = \boldsymbol{E}_0 \sqrt{\frac{w_{0x} w_{0y}}{w_x(z) w_y(z)}} \exp\left[-\left(\frac{x^2}{w_x^2(z)} + \frac{y^2}{w_y^2(z)}\right)\right] \exp\left[-jk\left(\frac{x^2}{2R_x(z)} + \frac{y^2}{2R_y(z)}\right)\right] \exp\left[jkz - \frac{j}{2}\left(\arctan\left(\frac{z - z_{0x}}{z_{Rx}}\right) + \arctan\left(\frac{z - z_{0y}}{z_{Ry}}\right)\right)\right], \tag{2.4}$$

where, without loss of generality, the axes of the beam are chosen to be oriented along the x and y axes. Here, $w_\mu = w_{0\mu}\sqrt{1 + (z - z_{0\mu})^2/z_{R\mu}^2}$ is the position-dependent beam radius for $\mu \in \{x, y\}$, $z_{R\mu} = \pi w_{0\mu}^2/\lambda$ is the Rayleigh range, and

$$R_\mu(z) = (z - z_{0\mu})\left[1 + \left(\frac{z_{R\mu}}{z - z_{0\mu}}\right)^2\right] \tag{2.5}$$

is the radius of curvature of the phase front of the beam. The beam waist $z_{0\mu}$ denote the plane in which the beam radius achieves its minimal value, such that $w_\mu(z_{0\mu}) = w_{0\mu}$. In the presence of astigmatism, the beam waists z_{0x} and z_{0y} do not overlap (see Fig. 2.2c).

To further characterize the Gaussian beam, we define its divergence angle, α_μ, as the asymptotic rate of change of the beam radius,

$$\tan(\alpha_\mu) = \lim_{z\to\infty} \frac{dw_\mu}{dz} = \frac{w_{0\mu}}{z_{R\mu}}, \ \mu \in \{x, y\}. \tag{2.6}$$

This equation relies on the fact that, far from the beam waists, the beam radii increase linearly with distance. Equations (2.1) and (2.4) coincide for circular Gaussian beams with equal beam waist radii, $w_{0x} = w_{0y}$, and overlapping beam waists, $z_{0x} = z_{0y}$.

2.2.1 Three-Lens Solution to Gaussian Beam Circularization

Converting an elliptical beam with astigmatism into a circular beam without astigmatism requires suppressing both astigmatism and ellipticity. Suppressing astigmatism requires overlapping the beam waists z_{0x} and z_{0y} (see Fig. 2.2b). In the absence of astigmatism, the resulting beam may still exhibit ellipticity if the beam waist radii differ, $w_{0x} \neq w_{0y}$. Suppressing ellipticity requires equalizing the beam waist radii to achieve $w_{0x} = w_{0y}$. We refer to *beam circularization* as the process of correcting astigmatism and ellipticity. Mathematically, beam circularization involves converting a Gaussian beam described by Eq. (2.4) into a circular Gaussian beam whose intensity profile is radially symmetric in all planes.

A typical approach to solving the beam circularization problem is to use two uniaxial cylindrical lenses, L_1 and L_2 (represented as thick blue and red-orange vertical lines in Fig. 2.3, respectively). Each lens is aligned with either the major or minor axis of the beam and positioned approximately one focal length away from its corresponding beam waist. The ratio of the focal lengths must be chosen to satisfy

$$f_x/f_y \approx w_{0x}/w_{0y}, \tag{2.7}$$

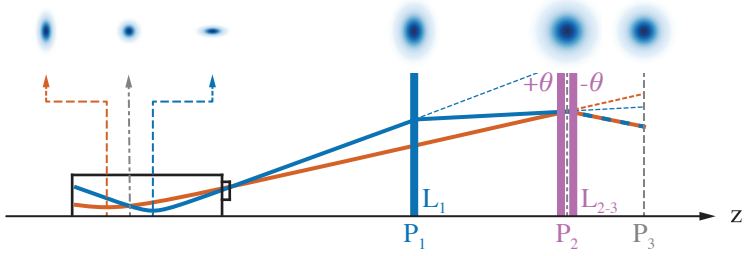


Figure 2.4: **Three-Lens Solution.** An astigmatic elliptical Gaussian beam is converted into a circular Gaussian beam using three cylindrical lenses. The first lens (L_1) is oriented along the major axis of the beam and placed in an arbitrary plane, not necessarily located a focal length away from the beam waist. The second and third lenses (L_{2-3}) are mounted one behind the other at an angle $\pm\theta$ in the plane where the intensity profile of the beam is circular. The relative angle is chosen to equalize the divergence angles of the beam along the major and minor axes. The solution is robust against small changes in the relative distance between the two lenses. The resulting circular beam can readily be collimated using a spherical lens.

where the exact ratio for Gaussian beams has a complex dependence on w_{0x}, w_{0y}, $\Delta z = z_{0y} - z_{0x}$, and the wavelength λ. Achieving this exact ratio is typically challenging in practice when using lenses from standard optics catalogs. These lenses might also deviate from being perfectly uniaxial, introducing unwanted focusing along the weak axis of the lens.

Our approach to solving the circularization problem involves three cylindrical lenses (see Fig. 2.4). The first cylindrical lens (L_1, represented as a thick blue line in Fig. 2.4) is placed in an arbitrary plane, P_1, and oriented along the major axis of the beam (the axis along which the beam divergence is the largest). The lens reduces the divergence of the beam along the major axis so as to create a plane, P_2, in which the beam radii along the major and minor axes are equal, i.e., the intensity profile of the beam is circular. The second and third lenses (L_{2-3}, represented by thick purples lines in Fig. 2.4) are mounted one behind the other in the plane P_2. Both lenses are aligned along the same axis, then counter-rotated by $\pm\theta$. The relative angle between the two lenses is thus 2θ. The relative angle is increased until the beam waists overlap, thereby correcting astigmatism. The resulting beam is a circular Gaussian beam, free of astigmatism and ellipticity. Its intensity profile is radially symmetric in any subsequent z plane. It is readily collimated using standard spherical lenses.

We now demonstrate the theoretical validity of our approach in three steps. We first

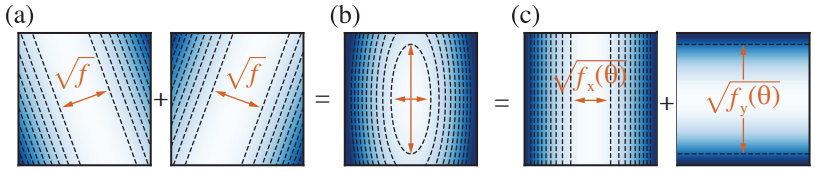


Figure 2.5: **Equivalent Phase Profiles of a Cylindrical Lens Pair.** (a) The cumulative phase profile of two uniaxial cylindrical lenses of focal length f, counter-rotated by $\pm\theta$, is equivalent to the phase profile of a (b) biaxial lens, which is itself equivalent to the cumulative profile of (c) two uniaxial cylindrical lenses with tunable focal lengths $f_x(\theta) = f\sec^2(\theta)/2$ and $f_y(\theta) = f\csc^2(\theta)/2$.

show that the cylindrical lens pair acts as an effective biaxial lens with tunable focal lengths (Sec. 2.2.1). We then show that this tunable biaxial lens can be used to correct astigmatism (Sec. 2.2.1). We finally show that placing the biaxial lens in a plane where the beam intensity is radially symmetric corrects both ellipticity and astigmatism simultaneously (Sec. 2.2.1).

Tunable Biaxial Lens from Two Uniaxial Cylindrical Lenses

We show that two cylindrical lenses of focal length f counter-rotated by an angle $\pm\theta$ about the x axis act as an effective biaxial lens with tunable focal lengths given by

$$f_x(\theta) = f\sec^2(\theta)/2, \tag{2.8}$$
$$f_y(\theta) = f\csc^2(\theta)/2. \tag{2.9}$$

This result can be obtained by considering the transmittance function of a uniaxial cylindrical lens, which is given by

$$\mathcal{T}_x(f) = \exp\left[-\frac{i\pi}{\lambda f}x^2\right] \tag{2.10}$$

when its major axis is oriented along the x axis, and by

$$\mathcal{T}_{\pm\theta}(f) = \exp\left[-\frac{i\pi}{\lambda f}\left[x\cos(\theta) \pm y\sin(\theta)\right]^2\right] \tag{2.11}$$

when its major axis is rotated by $\pm\theta$ from the x axis. Neglecting the propagation of the beam between the two lenses or, equivalently, assuming a zero inter-lens distance, the

cumulative transmittance function of the two cylindrical lenses is given by the product of their individual transmittance functions,

$$\begin{aligned}
\mathcal{T}_{\text{pair}}(f,\theta) &= \mathcal{T}_{+\theta}(f)\,\mathcal{T}_{-\theta}(f) && (2.12)\\
&= \exp\left[-\frac{i\pi}{\lambda f}\left[x\cos(\theta)+y\sin(\theta)\right]^2\right]\exp\left[-\frac{i\pi}{\lambda f}\left[x\cos(\theta)-y\sin(\theta)\right]^2\right] && (2.13)\\
&= \exp\left[-\frac{i\pi}{\lambda(f/2)}x^2\cos^2(\theta)\right]\exp\left[-\frac{i\pi}{\lambda(f/2)}y^2\sin^2(\theta)\right] && (2.14)\\
&= \mathcal{T}_x\left[\frac{f}{2}\sec^2(\theta)\right]\;\mathcal{T}_y\left[\frac{f}{2}\csc^2(\theta)\right]. && (2.15)
\end{aligned}$$

The transmittance function of the cylindrical lens pair is equal to the product of the transmittance functions of two uniaxial lenses oriented along the x and y axes with focal lengths $f_x = f\sec^2(\theta)/2$ and $f_y = f\csc^2(\theta)/2$, respectively. The cylindrical lens pair thus acts as an effective biaxial lens with tunable focal lengths. This result can be understood visually by considering the spatial representation of the phase profiles of the cylindrical lens pair (see Fig. 2.5).

Astigmatism Correction

The biaxial lens can suppress astigmatism, converting an elliptical beam with astigmatism into an elliptical beam without astigmatism (see Fig. 2.6). By rotating the relative angle between the two cylindrical lenses, the focal lengths of the effective biaxial lens is chosen to minimize the relative separation between the beam waists, denoted Δz (see Fig. 2.6b). The relative separation between the beam waists is defined as

$$\Delta z(\theta) = z_{0y}(\theta) - z_{0x}(\theta), \tag{2.16}$$

where the dependence of the beam waist on θ can be calculated analytically (see Eq. (2.43) in App. 2.5.1). For sufficiently small values of the focal length $|f|$, there exists at least one angle θ^* for which the two beam waists coincide, i.e., $\Delta z(\theta^*) = 0$ (see App. 2.5.3 for a proof of existence of θ^*). This statement means that the lens pair can always be used to correct astigmatism in any input Gaussian beam.

Even after overlapping the beam waists, the waist radii may still differ, and some ellipticity may remain (see Fig. 2.6c). This limitation arises because a thin lens can only modulate the phase of the input beam without directly modifying its amplitude profile. Although phase modulation can influence the intensity profile of the beam upon propagation, it cannot, in general, eliminate radial asymmetry in the immediate vicinity of the

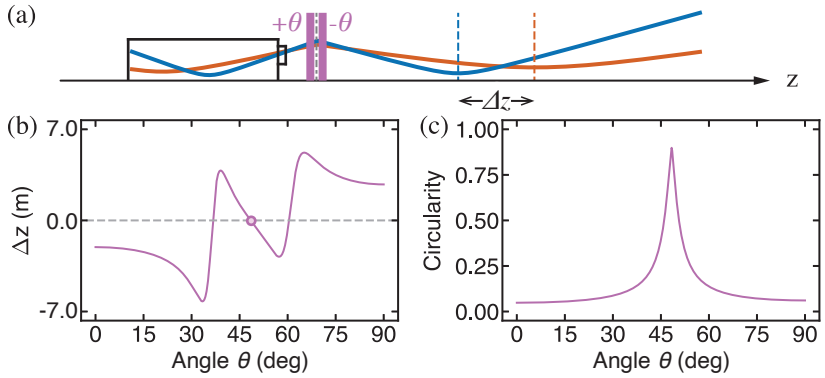


Figure 2.6: **Astigmatism Correction.** (a) An elliptical beam is focused through a pair of cylindrical lenses with a tunable relative angle 2θ. The major and minor axes focus in different planes (blue and orange dashed lines). (b) The relative distance between the two focal planes (see Eq. (2.16)) changes with the lens angle and vanishes at three specific values. (c) Beam circularity (see Eq. (2.22)) is maximized close to one of the angles where the focal planes overlap; however, the biaxial lens pair alone is not sufficient to achieve perfect circularity. Note that the angle for which the circularity is maximized does not necessarily coincide with the angle for which astigmatism is suppressed, and the plots are not necessarily symmetric.

lenses. As a result, the output beam remains an elliptical Gaussian beam, unless the beam has a circular intensity profile at the input of the lens pair. Making the intensity profile of the input beam circular requires the use of another cylindrical lens before the lens pair.

Ellipticity Correction

Suppressing ellipticity requires placing the cylindrical lens pair in a plane P_2 where the intensity profile of the beam is circular. This plane can be found by placing a uniaxial cylindrical lens oriented along the major axis of the beam in a plane P_1 preceding P_2. This lens reduces the divergence of the beam along the major axis of the beam such that, for sufficiently small focal lengths $f_1 > 0$, there exists a plane P_2 in which the beam radii along the minor and the major axes overlap (see App. 2.5.4 for a proof of existence for f_1).

Given that the focal lengths of the cylindrical lenses of the lens pair are equal, i.e., $f_2 = f_3 = f$, the optimal rotation angle at which the beam is circular is given by:

$$\theta^* = \frac{1}{2}\arccos\left[\frac{f}{2}\left(\frac{\lambda}{\pi w_r}\right)^2\left\{\frac{z_{0y}}{w_{0y}^2} - \frac{z_{0x}}{w_{0x}^2}\right\}\right], \tag{2.17}$$

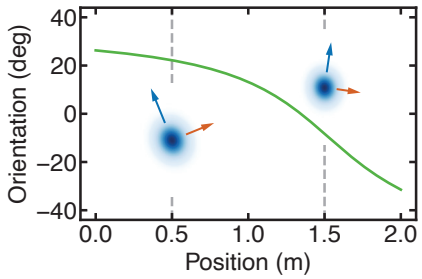


Figure 2.7: **Generalized Gaussian Beam Twisting.** The orientation of a generalized Gaussian beam rotates as it propagates through free space, while its transverse intensity profile in each plane is elliptical, similar to the intensity profile of an elliptical Gaussian beam (see insets).

where $w_{0\mu}$ is the beam waist radius of the input beam for $\mu \in \{x, y\}$, $z_{0\mu}$ is the beam waist, and $w_r = w_x(z_{P_2}) = w_y(z_{P_2})$ is the beam radius at P_2 (see App. 2.5.2 for a full derivation). This equation is valid as long as

$$|f| \leq f_{\max} = \left| \frac{2\left(\frac{\pi w_r}{\lambda}\right)^2}{\frac{z_{0y}}{w_{0y}^2} - \frac{z_{0x}}{w_{0x}^2}} \right|. \tag{2.18}$$

When f satisfies this inequality, an astigmatic elliptical beam is converted into a circular beam with overlapping beam waists, $z_{0x}(\theta^*) = z_{0y}(\theta^*)$, and beam waist radii, $w_{0x}(\theta^*) = w_{0y}(\theta^*)$. When f violates this inequality, no real-valued solution for θ^* exists, and circularization cannot be achieved.

The above results are valid when the two lenses of the cylindrical lens pair are perfectly uniaxial. In practice, these lenses may have a finite focal length $f_\perp$ along their weak axis due to manufacturing imperfections. We can prove that for a finite $f_\perp$, Eqs. (2.17) and (2.18) are still valid when replacing the nominal focal length f by the effective focal length (see App. 2.5.2)

$$f' = \left(\frac{1}{f} - \frac{1}{f_\perp}\right)^{-1}. \tag{2.19}$$

This result is valid when both lenses have the same transverse focal lengths, $f_\perp = f_{2\perp} = f_{3\perp}$. However, when the two focal lengths are not equal, $f_{2\perp} \neq f_{3\perp}$ the beam at the output of the lens pair exhibits residual ellipticity and astigmatism for all angles θ, though these effects are suppressed compared to the typical two-cylindrical-lens solution (see App. 2.5.2 for details).

2.2.2 Robustness Against Finite Inter-Lens Distance

Our calculations have so far assumed a negligible distance between the two lenses of the cylindrical lens pair. In practice, this distance is lower-bounded by the thickness of the lenses and the rotation mounts to which they are attached. A nonzero distance between the two cylindrical lenses results in a breakdown of the transmittance results presented in Sec. 2.2.1. In this case, the lens pair no longer acts as an effective biaxial lens, and the beam after the lens pair can no longer be accurately modeled by a standard Gaussian beam (see Eq. (2.4)).

To study the effects of the nonzero distance between the two cylindrical lenses, a generalized model for beam propagation is required. We now introduce this model and use it to quantify the degradation in beam circularity resulting from variations in the inter-lens distance. We show that the imperfections introduced by the nonzero inter-lens distance can be mitigated by using a large focal length f for the two cylindrical lenses.

Generalized Gaussian Beams

A sufficient model to describe Gaussian beams passing through cylindrical lenses with arbitrary orientations is given by

$$\boldsymbol{E}(x, y, z) = \boldsymbol{E_0}(z) \exp\left(-\boldsymbol{r}^T \boldsymbol{\Lambda}(z) \boldsymbol{r}\right), \tag{2.20}$$

where $\boldsymbol{E_0}(z) \in \mathbb{C}^2$ is the complex-valued electric field amplitude, $\boldsymbol{r} = (x, y)^T$ is the transverse coordinate vector, and $\boldsymbol{\Lambda}(z) \in \mathbb{C}^{2\times 2}$ is a complex-valued 2×2 matrix (see App. 2.5.1 for a detailed presentation). We refer to beams described by Eq. (2.20) as *generalized Gaussian beams*.

Gaussian beams with ellipticity and astigmatism are a special case of generalized Gaussian beams for which the commutator of $\Re\boldsymbol{\Lambda}(z)$ and $\Im\boldsymbol{\Lambda}(z)$ is zero, i.e., $[\Re\boldsymbol{\Lambda}(z), \Im\boldsymbol{\Lambda}(z)] = 0$. When Gaussian beams pass through a cylindrical lens oriented along an arbitrary axis, they become non-Gaussian beams with $[\Re\boldsymbol{\Lambda}(z), \Im\boldsymbol{\Lambda}(z)] \neq 0$. Despite maintaining an intensity profile identical to that of a Gaussian beam, the nonzero commutator gives rise to a "twisting" effect, wherein the principal axes of the beam rotate during free-space propagation (see Fig. 2.7).

Dependence of Beam Circularity on Inter-Lens Distance

The generalized Gaussian beam model provides a means to study changes in the intensity profile of the beam as a function of the inter-lens distance. By computing the evolution

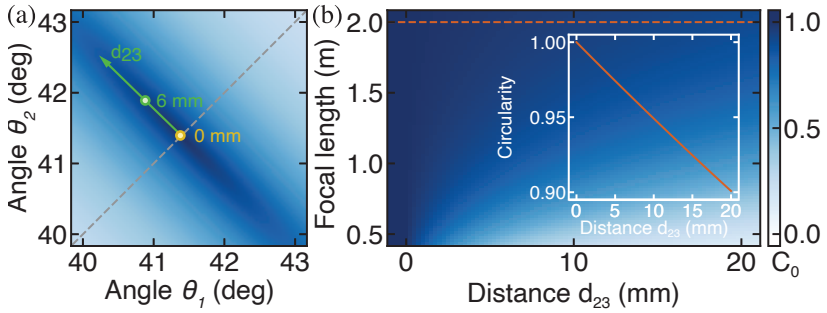


Figure 2.8: **Dependence of Minimum Circularity on Nonzero Inter-Lens Spacing.** (a) The minimum circularity $\mathcal{C}_0$ (white-to-blue surface) changes with the rotation angle of the cylindrical lenses, θ_2 and θ_3. At $d_{23} = 0$ (yellow circle), the optimal minimum circularity is achieved for $\theta_2 = \theta_3 = \theta^*$ (gray dashed line), where $\theta^* = 41.4°$. For $d_{23} \neq 0$, the optimal angles deviates from the optimal predicted values (green line). The green disk indicates the optimal angles obtained for $d_{23} = 6$ mm, which approximately corresponds to the distance between the lenses when mounted in rotation mounts. (b) The minimum circularity of the beam depends on both the focal length of the lenses and their inter-lens distance. The circularity decreases with inter-lens spacing. For a fixed inter-lens spacing, increasing the focal length improves the minimum achievable circularity. These results were obtained for a Gaussian beam with a wavelength of $\lambda = 415$ nm whose beam waist radii are $w_{0x} = 260$ μm and $w_{0y} = 1.015$ mm at the beam waists located $z_{0x} = -6$ m and $z_{0y} = -22$ m, respectively, where $z = 0$ denotes the plane where the first lens of the lens pair is located.

of the matrix $\mathbf{\Lambda}(z)$ with respect to z (see Table 2.1), we can determine the beam radii and orientation in all planes z by diagonalizing $\Re\mathbf{\Lambda}(z)$: the principal axes of the beam are oriented along the eigenvectors of $\Re\mathbf{\Lambda}(z)$, whereas the beam radii are given by $w_j = 1/\sqrt{\lambda_j}$, where λ_j are the associated eigenvalues for $j = 1, 2$. We note that the eigenvectors are not necessarily aligned with the lab-frame axes, x and y, and may vary with propagation distance due to beam twisting.

We define the *circularity* of a generalized Gaussian beam in the plane z as the ratio of the smaller radius to the larger radius of the beam in that plane,

$$\mathcal{C}(z) = \frac{\min(w_1(z), w_2(z))}{\max(w_1(z), w_2(z))}. \tag{2.21}$$

We then define the *minimum circularity* of a beam as

$$\mathcal{C}_0 = \inf_{z \geq z_3} \mathcal{C}(z). \tag{2.22}$$

This quantity represents the infimum of $\mathcal{C}(z)$ over all z beyond the position of the third lens of the three-lens system, located at $z_3 = z_{P_2} + d_{23}/2$, assuming the beam propagates indefinitely in free space.

We numerically compute the minimum circularity of a generalized Gaussian beam for different values of inter-lens distance, d_{23}, as a function of the rotation angles θ_2 and θ_3 of the two cylindrical lenses with respect to the x axis (see Fig. 2.8a). The input beam is chosen to be an elliptical Gaussian beam with astigmatism that is circular in the plane P_2 where the lens pair is placed. The focal length of the two cylindrical lenses is chosen to be $f = 2000$ mm.

When the inter-lens distance is zero (see yellow disk in Fig. 2.8a for $d_{23} = 0$), the optimal minimum circularity is obtained for a specific value of $\theta^* = 41.4°$ located along the main diagonal defined by $\theta = \theta_2 = \theta_3$ (see gray dashed line). Deviation away from this optimal angle leads to degradation in the minimum circularity (blue-to-white shaded surface). As the inter-lens distance increases (see green line in Fig. 2.8a with the green marker corresponding to $d_{23} = 6$ mm), the optimal angles for reaching the optimal minimum circularity are no longer equal $\theta_2^* \neq \theta_3^*$.

We further compute the change in minimum circularity for different focal lengths and inter-lens distances, d_{23} (see Fig. 2.8b). The minimum circularity is computed for the case where $\theta_2 = \theta_3$. For $d_{23} = 0$, the circularity of the beam is equalt for all focal lengths. As d_{23} departs from 0, the minimum circularity decreases approximately linearly (see inset in Fig. 2.8b). The loss of circularity with increasing inter-lens distance is more pronounced for smaller focal lengths. We conclude that, for the specific test beam used in our simulations, mitigating the detrimental effects of inter-lens distance requires choosing a large focal length f, ideally approaching its upper bound $f_{\max}$ defined in the right hand side of Eq. (2.18). This conclusion is generally valid for any input beam with a sufficiently small divergence angle. Indeed, larger focal lengths correspond to weaker lenses, which cause a slower evolution of the beam's transverse profile across the inter-lens distance. As a result, the beam properties remain approximately unchanged between the lenses. In contrast, smaller focal lengths induce rapid variations in the beam radii between the two lenses, increasing the sensitivity to inter-lens spacing and ultimately degrading the circularization performance. However, the optimal choice of f should generally be checked for any specific beam, especially for beams with large divergence angles.

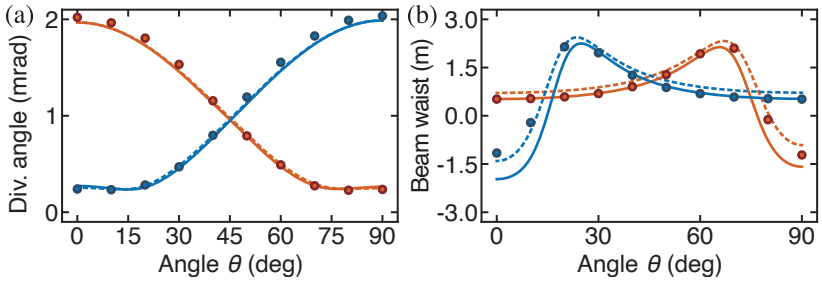


Figure 2.9: **Model Validation.** (a) Divergence angles and (b) beam waists of an astigmatic elliptical Gaussian beam measured along its principal axes, which are aligned with the x (red) and y axes (blue). The beam is generated by transmitting a nearly circular Gaussian beam at 780 nm through a pair of cylindrical lenses with a tunable relative angle 2θ, oriented asymmetrically about the x axis. Each data point represents a fit of the longitudinal intensity profile to a Gaussian beam model. Solid lines indicate theoretical predictions based on the input beam and the geometry of the lens system. The cylindrical lenses have a focal length of $f = 1000$ mm. The dashed lines show the expected behavior obtained when choosing the lenses to have an effective focal length along the transverse direction $f_\perp = 51.4$ m.

2.3 Experiment – Astigmatism and Ellipticity Suppression

Having established the validity and robustness of our method, we now demonstrate its applicability in an experimental setting. We begin by confirming that a pair of cylindrical lenses acts as a biaxial lens with tunable focal lengths. We then apply our method to circularize the beam emitted by a commercial titanium:sapphire laser source.

Biaxial Lens with Tunable Focal Lengths

To experimentally confirm that a pair of counter-rotated cylindrical lenses acts as an effective biaxial lens with tunable focal lengths, we measured the changes in the properties of a circular Gaussian beam and compared them against those computed using the transmittance function of Eqs. (2.12)–(2.15).

We created a collimated circular Gaussian beam using a fiber collimator at 780 nm. The beam had a minimum circularity of $\mathcal{C}_0 \geq 0.95$ with a collimated beam radius of 1.1 mm and a beam divergence of 0.1 mrad. We then placed the cylindrical lens pair

at an arbitrary distance of 240 mm away from the fiber collimator. The two cylindrical lenses, each with a focal length of $f = 1000$ mm (Thorlabs, LJ4530RM), were mounted in a rotation mount (Thorlabs, CRM1T) and oriented parallel to the optical table along the x axis. We counter-rotated each lens by an angle $\pm\theta$ and measured the intensity profile of the focusing beam at various longitudinal planes near its beam waists (see App. 2.5.5 for details on the image acquisition and processing procedure). We then fit each transverse intensity profile to an elliptical Gaussian beam with beam radii

$$w_\mu(z) = w'_{0\mu}\sqrt{1 + \left((z - z'_{0\mu})\frac{\lambda}{\pi w'^2_{0\mu}}\right)^2}, \ \mu \in \{x, y\} \tag{2.23}$$

to obtain the beam waists $z_{0\mu}$ along each axis, as well as the beam divergence, α (see Eq. (2.6)).

The measurements of the divergence angles and beam waists are in good agreement with theoretical predictions (see solid lines on Fig. 2.9), except for small deviations in the beam waist locations near $\theta = 0°$ and $\theta = 90°$. At $\theta = 0°$, the lens pair should ideally have no effect on the beam along the y axis, as both lenses are nominally aligned along the x axis. The observed deviation from theory suggests a shift in the beam waist along the y axis at $\theta = 0°$. This discrepancy cannot be attributed to inaccurate lens orientation alone. Our calculations show that an angular misalignment of up to 5° would be required to produce such an effect, which is substantially greater than the alignment precision of less than 1° for the our rotation mounts. We attribute the observed discrepancy to the cylindrical lenses not being perfectly uniaxial, exhibiting a small residual focusing power along their transverse axis. The discrepancy between theory and experiment is minimized when we assume that the two lenses of the lens pair have a residual transverse focal length of $f_\perp = 51.4$ m along their weak axis (see dashed lines in Fig. 2.9). A similar explanation accounts for the deviation observed near $\theta = 90°$. These experimental data confirm that the lens pair acts as an effective biaxial lens system with tunable focal lengths.

2.3.1 Beam Circularization of a Commercial Laser Source

We now show that the three-lens system can be used to circularize the output of a commercial titanium:sapphire laser source at 813 nm. The laser source emits a diverging beam with an elliptical intensity profile whose principal axes align with the x and y directions of the lab frame, defined as parallel and perpendicular to the optical table. The circularity of the beam in the far-field is $\mathcal{C}_\infty = 0.69$ (see Fig. 2.10a).

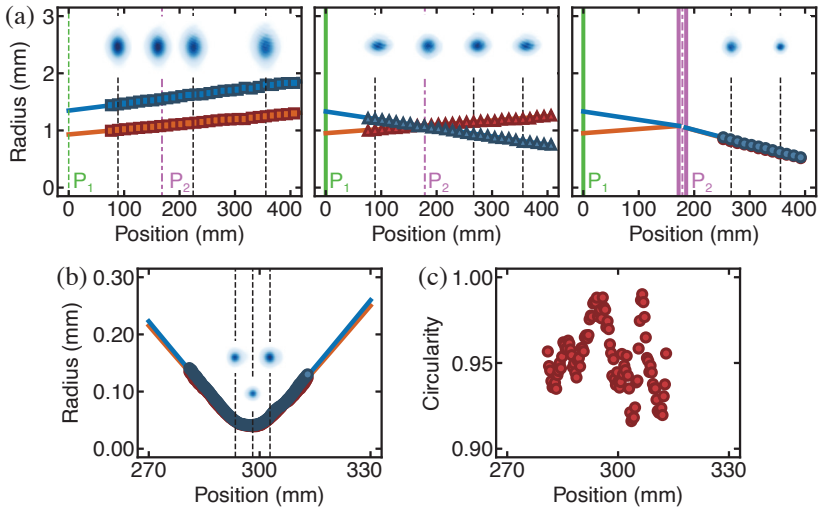


Figure 2.10: **Experimental Demonstration of Beam Circularization.** (a-left) A commercial laser source emits a diverging beam with an elliptical intensity profiles (insets). The beam radii (squares) are measured along the x (red) and y (blue) axes and fitted to Gaussian beam profiles (solid lines). (a-middle) Focusing the beam along its major axis using a cylindrical lens ($f_1 = 500$ mm) placed at P_1 (thick green line) creates a plane P_2 at $z = 176$ mm where the intensity profile becomes circular. (a-right) Focusing the beam along both its major and minor axes using a cylindrical lens pair ($f_2 = f_3 = 500$ mm) with an optimized relative angle suppresses astigmatism and restores circularity. (b) Focusing the beam after P_2 with an achromatic doublet ($f = 150$ mm) further demonstrates the suppression of astigmatism, with the beam waists overlapping at $z = 298$ mm. (c) The minimum beam circularity near the beam waist averages to 0.95(1), bounded between 0.92(1) and 0.99(1).

We placed the first cylindrical lens with focal length $f_1 = 500$ mm (Thorlabs, LJ4147RM) at a distance of 1300 mm away from the output window of the laser. This location was chosen arbitrarily to satisfy the constraints of our scientific apparatus while providing sufficient space for beam characterization. The lens was mounted in a rotation mount (Thorlabs, CRM1PT) and oriented along the major axis of the elliptical beam. Because of the properties of generalized Gaussian beams, whose principal axes might undergo an anti-crossing effect (see App. 2.5.4),any relative orientation misalignment between the lens and the major axis results in the beam circularity to be strictly less than 1 in all subsequent planes after the lens. To minimize misalignment errors, we measured the beam profile across a range surrounding the expected location of P_2 and adjusted the orientation

to maximize the circularity within this range. After optimization, we measured a circularity of $\mathcal{C}_{P_2} = 0.9996^{+0.0004}_{-0.0050}$ in the plane P_2 located $z = 186$ mm away from the lens (see Fig. 2.10b).

We then placed two plano-convex cylindrical lenses with equal focal length $f = 500$ mm (Thorlabs, LJ4147RM) symmetrically around the plane P_2, each offset by ± 3 mm along the optical axis ($d_{23} = 6$ mm). The lenses were mounted in rotation mounts (Thorlabs, CRM1PT) with their convex surfaces facing the incoming beam. Their major axes were aligned with the beam's minor x-axis and counter-rotated by an angle $\pm\theta$ to correct astigmatism. After optimization, the circularity of the beam in the far field was measured to be $\mathcal{C}_\infty = 0.97 \pm 0.01$ at the optimal lens angle of $\theta^* = 28.7°$. We confirmed that the beam maintained a circular intensity profile throughout the entire propagation range beyond P_2 (see Fig. 2.10c).

To confirm that the lens pair corrected not only ellipticity far from the beam waists but also astigmatism, we focused the beam using an achromatic spherical lens with focal length $f = 150$ mm (Thorlabs, AC254-150) at a distance of 30 mm away from P_2. We measured the intensity profile near the beam waist to extract the relative distance between the beam waists, Δz. We observed that this relative distance is highly sensitive to the orientation angle of the cylindrical lenses, with angular deviations as small as 0.05° resulting in measurable changes. After optimization, we achieved $|\Delta z| = 0.008 z_R$, less than 0.8% of the Rayleigh range.

These results demonstrate that our three-lens solution can be used to circularize the elliptical laser beam of a commercial laser source, transforming a highly elliptical Gaussian beam into a nearly circular one. We attribute the residual imperfections to imperfect lens orientation, nonzero inter-lens spacing, and other minor unaccounted factors.

2.4 Conclusion

In conclusion, we have introduced a simple three-lens solution to convert astigmatic, elliptical Gaussian beams into circular Gaussian beams without astigmatism. We have validated our solution using theoretical results, numerical simulations, and experimental demonstrations. We have shown that its optimality and robustness to imperfections can be understood through the theory of generalized Gaussian beams, which explains phenomena such as beam twisting and axis anti-crossings. Our three-lens solution can be readily used to suppress astigmatism and ellipticity in laser beams emitted from commercial laser sources, as well as those degraded by dispersive or imperfect optical elements. Its key advantage

over the typical two-lens solution is that it does not impose strict constraints on the choice of lenses. It readily supports many applications that require tunable beam shaping, such as creating large arrays of optical traps with spatially homogeneous properties [54, 30], minimizing cross-talk in site-selective quantum gate operations [49, 50], and coupling laser beams into optical fibers with high efficiency [46, 55]. More broadly, any spherical lens in an optical system can be replaced with a tunable cylindrical lens pair to correct residual ellipticity and astigmatism.

2.5 Appendices

2.5.1 Matrix Formalism for Generalized Gaussian Beams

Our beam circularization method uses cylindrical lenses with arbitrary orientations to shape the spatial profile of Gaussian beams. After passing through a cylindrical lens with an arbitrary orientation, the beam is no longer Gaussian and cannot be described by Eq. (2.4). To describe the propagation of beams through optical systems composed of cylindrical lenses with arbitrary orientations, we introduce generalized Gaussian beams, also called Gaussian beams with general astigmatism [56].

Using a matrix representation, we show that the computation of the beam shape reduces to the diagonalization of a 2×2 complex-valued and symmetric matrix $\mathbf{\Lambda}(z)$. This matrix depends on the position z along the propagation axis of the beam. We derive a set of rules that determine how $\mathbf{\Lambda}(z)$ changes with z as the beam propagates in free-space and transmits as it passes through cylindrical lenses with arbitrary orientation. We use this matrix formalism to analyze the shape of the beam through various optical systems.

Generalized Gaussian beams

A generalized Gaussian beam can be thought of as a Gaussian beam with an orientation angle ϕ that is complex-valued [56]. Its complex electric field is given by

$$\boldsymbol{E}(x,y,z) = \boldsymbol{E}_0 \sqrt{\frac{w_{0x} w_{0y}}{w_x(z) w_y(z)}} \exp\left[-\left(\frac{\mu(\phi)^2}{w_x^2(z)} + \frac{\nu(\phi)^2}{w_y^2(z)}\right)\right] \exp\left[-jk\left(\frac{\mu(\phi)^2}{2R_x(z)} + \frac{\nu(\phi)^2}{2R_y(z)}\right)\right] \exp\left[jkz - \frac{j}{2}\left(\arctan\left(\frac{z-z_{0x}}{z_{Rx}}\right) + \arctan\left(\frac{z-z_{0y}}{z_{Ry}}\right)\right)\right], \tag{2.24}$$

This equation is a valid solution to Maxwell's equations even for complex values of ϕ. It accurately describes the beam transmitted through an arbitrarily oriented cylindrical lens, where the orientation of the beam can transform from a real-valued angle before the lens to a complex-valued angle afterwards. The complex-valued orientation results in non-Gaussian features in the beam. For instance, while the intensity profile takes the form of a radially asymmetric Gaussian distribution in all planes, the principal axes of the beam can rotate as the beam propagates in free space (see Fig. 2.7). We refer to this phenomenon as *twisting*.

The complex electric field given in Eq. (2.24) can be rewritten in the following compact form:

$$\boldsymbol{E}(\boldsymbol{r}, z) = \boldsymbol{A}(z) \exp\left(-\boldsymbol{r}^T \boldsymbol{\Lambda}(z) \boldsymbol{r}\right), \tag{2.25}$$

where $\boldsymbol{r} = (x, y)^T$ is the transverse position vector, $\boldsymbol{\Lambda}(z) \in \mathbb{C}^{2\times 2}$ is a symmetric rank-2 tensor represented as a two-dimensional matrix with complex coefficients, and $\boldsymbol{A}(z)$ is a position-dependent complex-valued envelope function independent of the transverse coordinates x and y. The associated intensity profile of the generalized Gaussian beam in each plane z is given by

$$I(\boldsymbol{r}, z) = \frac{1}{2}\epsilon c |\boldsymbol{A}(z)|^2 \exp\left(-2\boldsymbol{r}^T \, \Re\boldsymbol{\Lambda}(z) \, \boldsymbol{r}\right). \tag{2.26}$$

The $\boldsymbol{\Lambda}(z)$ matrix contains sufficient information to determine the beam radii and orientation of the generalized Gaussian beam in all planes z. The principal axes of the beam, $\boldsymbol{v}_1(z)$ and $\boldsymbol{v}_2(z)$, are given by the eigenvectors of $\Re\boldsymbol{\Lambda}(z)$, and the beam radii are given by $w_\mu(z) = 1/\sqrt{\lambda_\mu(z)}$, where $\lambda_\mu(z)$ are the eigenvalues of $\Re\boldsymbol{\Lambda}(z)$, for $\mu \in 1, 2$.

Matrix transformations for computing beam propagation through optical components

We now show that the computation of $\boldsymbol{\Lambda}(z)$ can be performed independently of $\boldsymbol{A}(z)$ using a simple set of rules involving only inversion and addition of 2×2 matrices (see App. 2.5.1). This set of rules describes how $\boldsymbol{\Lambda}(z)$ evolves under free-space propagation and transmission through cylindrical lenses, assuming the thin-lens approximation (see Table 2.1 for a summary of these rules). Each rule is derived using a matrix formalism analogous to the ABCD matrix method [57]. This formalism eliminates the need to compute the convolution of the electric field of the beam with the free-space propagation kernel from Fresnel diffraction theory [58], thereby reducing both the computational cost and the analytical complexity of computing the propagation of generalized Gaussian beams through optical systems comprising multiple cylindrical lenses.

Evolution	**Rule**
Free-space propagation by Δz.	$\boldsymbol{\Lambda}(z+\Delta z) = \left(\boldsymbol{\Lambda}(z)^{-1} + \frac{i\lambda\Delta z}{\pi}\mathbb{I}_{2\times 2}\right)^{-1}$
Transmission through a spherical lens with focal length f.	$\boldsymbol{\Lambda}(z^+) = \boldsymbol{\Lambda}(z^-) + \frac{i\pi}{\lambda f}\mathbb{I}_{2\times 2}$
Transmission through a cylindrical lens with focal length f rotate, with the major axis rotated at an angle θ from the x axis.	$\boldsymbol{\Lambda}(z^+) = \boldsymbol{\Lambda}(z^-) + \frac{i\pi}{\lambda f}\begin{bmatrix} \cos^2(\theta) & -\frac{1}{2}\sin(2\theta) \\ -\frac{1}{2}\sin(2\theta) & \sin^2(\theta) \end{bmatrix}$
Transmission through a biaxial lens with major focal length f_1 and minor focal length f_2, with the major axis rotated at an angle θ from the x axis.	$\boldsymbol{\Lambda}(z^+) = \boldsymbol{\Lambda}(z^-) + \frac{i\pi}{\lambda f_1}\begin{bmatrix} \cos^2(\theta) & -\frac{1}{2}\sin(2\theta) \\ -\frac{1}{2}\sin(2\theta) & \sin^2(\theta) \end{bmatrix} + \frac{i\pi}{\lambda f_2}\begin{bmatrix} \sin^2(\theta) & +\frac{1}{2}\sin(2\theta) \\ +\frac{1}{2}\sin(2\theta) & \cos^2(\theta) \end{bmatrix}$

Table 2.1: **Transformation Rules for Generalized Gaussian Beams.** Rules governing the evolution of the $\boldsymbol{\Lambda}(z)$ matrix under free space propagation and arbitrarily oriented biaxial lenses. $\boldsymbol{\Lambda}(z)$ completely describes the intensity profile of generalized Gaussian beams, making these rules sufficient for determining the beam shape of a generalized Gaussian beam.

Propagation of Gaussian Beams in Free Space. The free-space propagation of a Gaussian beam can be computed using Fresnel diffraction theory [58]. Given $\boldsymbol{E}(x, y, z = 0)$,

the complex electric field at $z = \Delta z$ is given by

$$\boldsymbol{E}(x, y, z = \Delta z) = \mathcal{F}^{-1}\left[\mathcal{F}\left(\boldsymbol{E}(x, y, z = 0)\right) \cdot H_{\Delta z}(x, y)\right], \tag{2.27}$$

where

$$H_{\Delta z}(f_x, f_y) = e^{ikz} \exp\left(-i\pi\lambda\Delta z(f_x^2 + f_y^2)\right) \tag{2.28}$$

is the Fourier transform of the free space evolution kernel [58], and

$$\mathcal{F}\left[\boldsymbol{E}(x, y, z0)\right](f_x, f_y) = \iint \boldsymbol{E}(x, y, z = 0) \exp\left(-2\pi i f_x x - 2\pi i f_y y\right)\ dx\ dy, \tag{2.29}$$

is the Fourier transform with respect to the x and y degrees of freedom.

Given the Fourier transform of a generalized Gaussian beam,

$$\mathcal{F}\left[E(\boldsymbol{r}, z = 0)\right](f_x, f_y) = \frac{\pi}{\sqrt{\det(\boldsymbol{\Lambda}(0))}} \boldsymbol{A}(0) \exp\left(-\pi^2 \boldsymbol{f}^T \boldsymbol{\Lambda}^{-1}(0) \boldsymbol{f}\right), \tag{2.30}$$

where $\boldsymbol{f} = (f_x, f_y)^T$, we can compute the free-space propagation of a generalized Gaussian beam,

$$E(\boldsymbol{r}, z = \Delta z) = e^{ik\Delta z} \boldsymbol{A}(0) \cdot \sqrt{\frac{\det(\boldsymbol{\Lambda}(\Delta z))}{\det(\boldsymbol{\Lambda}(0))}} \exp\left(-\boldsymbol{r}^T \boldsymbol{\Lambda}(\Delta z) \boldsymbol{r}\right), \tag{2.31}$$

where

$$\boldsymbol{\Lambda}(\Delta z) = \left(\boldsymbol{\Lambda}(0)^{-1} + \frac{i\lambda\Delta z}{\pi} \mathbb{I}_{2\times 2}\right)^{-1}, \tag{2.32}$$

and $\mathbb{I}_{2\times 2}$ is the 2-by-2 identity matrix. In the case of a Gaussian beam, the coefficient $\sqrt{\det(\boldsymbol{\Lambda}(\Delta z))/\det(\boldsymbol{\Lambda}(0))}$ accounts for both the Gouy phase shift and a corrective factor to ensure power conservation. The intensity profile and phase curvature of the beam are fully described by the $\exp(-\boldsymbol{r}^T \boldsymbol{\Lambda}(\Delta z) \boldsymbol{r})$ term.

Transmission of Gaussian Beams through Cylindrical Lenses. We compute the properties of a generalized Gaussian beam at the output of a biaxial lens oriented along an arbitrary direction. The transmittance function of an elliptical lens with focal lengths f_1 and f_2, oriented at an angle θ with respect to the x axis, is given by

$$\mathcal{T}_\theta(f_1, f_2) = \exp\left(-i\frac{\pi}{\lambda f_1}(x\cos(\theta) - y\sin(\theta))^2 - i\frac{\pi}{\lambda f_2}(x\sin(\theta) + y\cos(\theta))^2\right). \tag{2.33}$$

The complex electric field of the output beam is computed as the product of the complex electric field of the input beam and the transmittance function. This product effectively maps $\mathbf{\Lambda}(z)$ to $\mathbf{\Lambda}'(z,\theta)$, where

$$\mathbf{\Lambda}'(z,\theta) = \mathbf{\Lambda}(z) + \frac{i\pi}{\lambda f_1}\begin{bmatrix} \cos^2(\theta) & -\sin(2\theta)/2 \\ -\sin(2\theta)/2 & \sin^2(\theta) \end{bmatrix} + \frac{i\pi}{\lambda f_2}\begin{bmatrix} \sin^2(\theta) & \sin(2\theta)/2 \\ \sin(2\theta)/2 & \cos^2(\theta) \end{bmatrix}. \quad (2.34)$$

This mapping can be used to describe the effect of both spherical and uniaxial cylindrical lenses. For spherical lenses with $f \equiv f_1 = f_2$ that are invariant under rotation, we obtain

$$\mathbf{\Lambda}'(z,\theta) = \mathbf{\Lambda}(z) + \frac{i\pi}{\lambda f}\mathbb{I}_{2\times 2}, \quad (2.35)$$

whereas, for uniaxial cylindrical lenses with $f_2 \to \infty$, we obtain

$$\mathbf{\Lambda}'(z,\theta) = \mathbf{\Lambda}(z) + \frac{i\pi}{\lambda f_1}\begin{bmatrix} \cos^2(\theta) & -\sin(2\theta)/2 \\ -\sin(2\theta)/2 & \sin^2(\theta) \end{bmatrix}. \quad (2.36)$$

This mapping can be used to demonstrate the equivalence between a biaxial lens and a pair of uniaxial cylindrical lenses (see Sec. 2.2.1). Specifically, two cylindrical lenses, each with focal length f and counter-rotated by the angle $\pm\theta$, transform the $\mathbf{\Lambda}(z)$ matrix in the same way as two cylindrical lenses with focal lengths $f_x = f\sec^2(\theta)/2$ and $f_y = f\csc^2(\theta)/2$, aligned along the x and y axes, respectively. This transformation is given by

$$\mathbf{\Lambda}'(z,\theta) = \mathbf{\Lambda}(z) + \frac{i\pi}{\lambda\frac{f\sec^2(\theta)}{2}}\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + \frac{i\pi}{\lambda\frac{f\csc^2(\theta)}{2}}\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}. \quad (2.37)$$

Properties of Gaussian Beams transmitted through cylindrical lenses

Having established the matrix formalism, we now use it to compute the properties of a Gaussian beam transmitted through a biaxial lens pair. We specifically seek to compute the output beam waists, $z'_{0\mu}$, and beam waist radii, $w'_{0\mu}$, as a function of the input beam properties for $\mu \in \{x, y\}$. We choose the principal axes of the biaxial lens pair to be aligned with the principal axes of the beam chosen to be oriented along the x and y axes. Since the beam propagates independently along both axes, we only compute these quantities for the x axis; the corresponding properties of the beam along the orthogonal y axis can readily obtained.

We consider a diverging beam transmitted through a lens located at $z = 0$. The input beam is described by its $\mathbf{\Lambda}(z)$ matrix,

$$\mathbf{\Lambda}(z_{0x}) = \begin{bmatrix} \dfrac{1}{w_{0x}^2} & 0 \\ 0 & * \end{bmatrix}, \tag{2.38}$$

where w_{0x} is the beam waist radius measured at its beam waist $z = z_{0x}$.

The $\mathbf{\Lambda}(z)$ matrix before the lens, $\mathbf{\Lambda}(0^-)$, is obtained by propagating $\mathbf{\Lambda}(-z_{0x})$ in free space over a distance $\Delta z = z_{0x}$. Using the transformation rule for free-space propagation, we obtain

$$\mathbf{\Lambda}(0^-) = \left(\mathbf{\Lambda}(z_{0x})^{-1} - \frac{i\lambda z_{0x}}{\pi}\mathbb{I}_{2\times 2}\right)^{-1} = \begin{bmatrix} w_{0x}^2 - \dfrac{i\lambda z_{0x}}{\pi} & 0 \\ 0 & * \end{bmatrix}^{-1}, \tag{2.39}$$

where $*$ denotes some non-zero arbitrary value. The $\mathbf{\Lambda}(z)$ matrix after the lens, $\mathbf{\Lambda}(0^+)$, is obtained by updating $\mathbf{\Lambda}(0^-)$ using the transformation rule for a uniaxial cylindrical lenses oriented along $\theta = 0$,

$$\mathbf{\Lambda}(0^+) = \left(\mathbf{\Lambda}(0^-) + \frac{i\pi}{\lambda f}\right)\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \left(w_{0x}^2 - \dfrac{i\lambda z_{0x}}{\pi}\right)^{-1} + \dfrac{i\pi}{\lambda f} & 0 \\ 0 & 0 \end{bmatrix}. \tag{2.40}$$

We seek to compute the beam radii, w'_{0x}, and beam waist, z'_{0x} such that

$$\mathbf{\Lambda}(z'_{0x}) = \left(\mathbf{\Lambda}(0^+)^{-1} + \frac{i\lambda z'_{0x}}{\pi}\mathbb{I}_{2\times 2}\right)^{-1} = \begin{bmatrix} \dfrac{1}{w_{0x}'^2} & 0 \\ 0 & * \end{bmatrix}. \tag{2.41}$$

This equation is equivalent to

$$\mathbf{\Lambda}(0^+) = \left(\mathbf{\Lambda}(z'_{0x})^{-1} - \frac{i\lambda z'_{0x}}{\pi}\mathbb{I}_{2\times 2}\right)^{-1} = \begin{bmatrix} w_{0x}'^2 - \dfrac{i\lambda z'_{0x}}{\pi} & 0 \\ 0 & * \end{bmatrix}. \tag{2.42}$$

By comparing Eq. (2.40) and Eq. (2.42), we can solve for the beam waist location

$$z'_{0x} = \frac{\left(1 + \dfrac{z_{0x}}{f}\right) z_{0x} + \left(\dfrac{z_{Rx}}{f}\right)^2 f}{\left(1 + \dfrac{z_{0x}}{f}\right)^2 + \left(\dfrac{z_{Rx}}{f}\right)^2}, \tag{2.43}$$

and beam waist radius

$$w'_{0x} = \frac{w_{0x}}{\sqrt{\left(1 + \frac{z_{0x}}{f}\right)^2 + \left(\frac{z_{Rx}}{f}\right)^2}}. \tag{2.44}$$

Assuming $|(f + z_{0x})z_{0x}| \gg z_{Rx}^2$, we can approximate Eq. (2.43) as

$$z'_{0x} = \frac{z_{0x}}{1 + \frac{z_{0x}}{f}} + \mathcal{O}\left(\frac{z_{Rx}^2}{(f + z_{0x})z_{0x}}\right) \approx \frac{f z_{0x}}{f + z_{0x}}, \tag{2.45}$$

which can be rearranged into the familiar thin-lens equation from ray optics

$$\frac{1}{z'_{0x}} + \frac{1}{-z_{0x}} = \frac{1}{f}. \tag{2.46}$$

Similarly, the beam waist radius equation can be obtained from Eq. (2.44) by choosing $f = -z_{0x}$,

$$w'_{0x} = \frac{w_{0x}}{z_{Rx}/|f|} = \frac{\lambda |f|}{\pi w_{0x}}. \tag{2.47}$$

2.5.2 Beam Circularization Using a Cylindrical Lens Pair

We now derive the conditions on the rotation angle θ and focal length f required to circularize an astigmatic, elliptical Gaussian beam using two perfect cylindrical lenses. We then show that, even if the lenses are not perfectly uniaxial, there exists an angle θ^* at which the beam becomes circular.

Beam Circularization with Two Perfect Cylindrical Lenses

If the intensity profile of the beam is not radially symmetric in the plane where the lens pair is placed, then circularization is not possible. The reason is that the lenses only modulate the phase of the beam and do not immediate change its intensity profile. For the beam to become circular after the lens pair, its intensity profile must be radially symmetric at the position of lens pair, i.e., in the plane where the major and minor radii of the Gaussian beam are equal. We therefore assume that the intensity profile of the beam is circular in the plane of the lens pair. We now seek the conditions on the focal length f and rotation angle θ required to achieve a circular intensity profile in all subsequent planes following the cylindrical lens pair.

We consider an elliptical Gaussian beam whose principal axes are oriented along the x and y axes. The beam radii along each axis can be described by

$$w_\mu(z) = w_{0\mu}\sqrt{1 + \frac{(z - z_{0\mu})^2}{z_{R\mu}^2}} = w_{0\mu}\sqrt{1 + \frac{\lambda^2(z - z_{0\mu})^2}{\pi^2 w_{0\mu}^4}}, \tag{2.48}$$

where $w_{0\mu}$ is the beam waist radius at the beam waist $z = z_{0\mu}$ for $\mu \in \{x, y\}$.

We assume that there exists a plane, which we choose as the origin at $z = 0$, where the intensity profile is radially symmetric, i.e., $r \equiv r_x(0) = r_y(0)$. Then, the $\mathbf{\Lambda}(z)$ matrix at $z = 0$ is given by

$$\mathbf{\Lambda}(0) = \begin{bmatrix} \dfrac{1}{w_{0x}^2 - \dfrac{j\lambda}{\pi} z_{0x}} & 0 \\ 0 & \dfrac{1}{w_{0y}^2 - \dfrac{j\lambda}{\pi} z_{0y}}. \end{bmatrix} \tag{2.49}$$

Given that $r_x(0) = r_y(0)$, and recalling that the eigenvalues of $\Re\mathbf{\Lambda}(z)$ are given by $r_\mu(z)^{-2}$ for $\mu \in \{x, y\}$, we can decompose $\mathbf{\Lambda}(0)$ into its real and imaginary parts

$$\mathbf{\Lambda}(0) = \begin{bmatrix} \dfrac{1}{r^2} & 0 \\ 0 & \dfrac{1}{r^2} \end{bmatrix} + j \begin{bmatrix} \dfrac{\dfrac{\lambda z_{0x}}{\pi}}{w_{0x}^4 + \dfrac{\lambda^2 z_{0x}^2}{\pi^2}} & 0 \\ 0 & \dfrac{\dfrac{\lambda z_{0y}}{\pi}}{w_{0y}^2 + \dfrac{\lambda^2 z_{0y}^2}{\pi^2}} \end{bmatrix}, \tag{2.50}$$

where

$$r = r_x(0) = r_y(0) = w_{0x}\sqrt{1 + \frac{z_{0x}^2}{z_{Rx}^2}} = w_{0y}\sqrt{1 + \frac{z_{0y}^2}{z_{Ry}^2}}. \tag{2.51}$$

The real part of $\mathbf{\Lambda}(0)$ has equal eigenvalues, resulting in a radially symmetric intensity profile in the plane $z = 0$. However, the imaginary part of $\mathbf{\Lambda}(0)$ has unequal eigenvalues, causing the intensity profile to become radially non-symmetric after $z = 0$.

The addition of a cylindrical lens pair at $z = 0$, where both lenses have a focal length

f and are counter-rotated by an angle $\pm\theta$, transforms the $\mathbf{\Lambda}(0)$ matrix to

$$\mathbf{\Lambda}'(0^+) = \begin{bmatrix} \frac{1}{r^2} & 0 \\ 0 & \frac{1}{r^2} \end{bmatrix} + j \begin{bmatrix} \frac{\frac{\lambda z_{0x}}{\pi}}{w_{0x}^4 + \frac{\lambda^2 z_{0x}^2}{\pi^2}} & 0 \\ 0 & \frac{\frac{\lambda z_{0y}}{\pi}}{w_{0y}^4 + \frac{\lambda^2 z_{0y}^2}{\pi^2}} \end{bmatrix} + j\frac{2\pi}{f\lambda} \begin{bmatrix} \cos^2(\theta) & 0 \\ 0 & \sin^2(\theta) \end{bmatrix}. \quad (2.52)$$

If there exists a choice of f and θ^* such that $\Im\mathbf{\Lambda}'(0^+)$ has two equal eigenvalues, then the beam is circularized. Circularization is thus achieved when

$$\frac{\frac{\lambda z_{0x}}{\pi}}{w_{0x}^4 + \frac{\lambda^2 z_{0x}^2}{\pi^2}} + \frac{2\pi}{\lambda f}\cos^2(\theta^*) = \frac{\frac{\lambda z_{0y}}{\pi}}{w_{0y}^4 + \frac{\lambda^2 z_{0y}^2}{\pi^2}} + \frac{2\pi}{\lambda f}\sin^2(\theta^*), \quad (2.53)$$

which can be further simplified to

$$\frac{2}{f}\cos^2(\theta^*) - \frac{2}{f}\sin^2(\theta^*) = \left(\frac{\lambda}{\pi r}\right)^2 \left\{\frac{z_{0y}}{w_{0y}^2} - \frac{z_{0y}}{w_{0x}^2}\right\}. \quad (2.54)$$

Solving for θ^*, we find the required angle for circularization to be

$$\theta^* = \frac{1}{2}\arccos\left[\frac{f}{2}\left(\frac{\lambda}{\pi r}\right)^2 \left\{\frac{z_{0y}}{w_{0w}^2} - \frac{z_{0x}}{w_{0x}^2}\right\}\right]. \quad (2.55)$$

Since θ^* is real-valued, the upper bound on $|f|$ for which circularization is possible is given by

$$|f| \leq f_{\max} = \left|\frac{2\left(\frac{\pi r}{\lambda}\right)^2}{\frac{z_{0x}}{w_{0x}^2} - \frac{z_{0y}}{w_{0y}^2}}\right|. \quad (2.56)$$

We note that this upper bound is always well-defined for a non-circular input Gaussian beam. If the denominator in Eq. (2.56) vanishes, then the beam is already circular. Therefore, circularization conditions are satisfied for any focal length f provided that $\theta = \pi/4$, i.e., $f_{\max} \to \infty$. The reason is that, at $\theta = \pi/4$, the cylindrical lens pair behaves as a spherical lens, which preserves the circular symmetry of the Gaussian beam.

Beam Circularization with Two Imperfect Cylindrical Lenses

Having found the parameters required to achieve beam circularization using two perfectly uniaxial lenses, we now derive similar conditions for the case where the two cylindrical lenses forming the lens pair are not perfectly uniaxial. We choose each cylindrical lens, denoted as L_2 and L_3, to have a focal length f along its strong focusing axis, while also possessing a finite, nonzero focal length, $f_{2\perp}$ for L_2 and $f_{3\perp}$ for L_3, along the weak focusing axis, which is perpendicular to the strong axis..

Two Identical Lenses. In the symmetric case where $f_{2\perp} = f_{3\perp} = f_\perp$, each of the two imperfect cylindrical lenses can be modeled as two perfectly-uniaxial cylindrical lenses. The four equivalent lenses can then be paired into two effective biaxial lenses. The first two equivalent lenses, $L_{2,1}$ and $L_{3,1}$, are aligned with the strong focusing axis of L_2 and L_3, respectively. Together, both lenses form an effective biaxial lens pair comprising two uniaxial lenses with focal length f and a relative angle of 2θ. The two other equivalent lenses, $L_{2,2}$ and $L_{3,2}$, are oriented perpendicular to the weak focusing axis of L_2 and L_3, respectively, with a focal length $f_\perp$. Together, both lenses form a biaxial lens pair with focal length of $f_\perp$, each counter-rotated by an angle $\pi/2 + \theta$.

Recalling that two uniaxial cylindrical lenses with focal length f mounted one behind the other at a relative angle 2θ formed an effective biaxial lens with (see Eq. (2.15)), we find

$$f_x^A = \frac{f}{2}\sec^2(\theta) \tag{2.57}$$

$$f_y^A = \frac{f}{2}\sec^2(\theta) \tag{2.58}$$

for the first effective biaxial lens, and

$$f_x^B(\theta) = \frac{f_\perp}{2}\sec^2(\pi/2+\theta) = \frac{f_\perp}{2}\csc^2(\theta) \tag{2.59}$$

$$f_y^B(\theta) = \frac{f_\perp}{2}\csc^2(\pi/2+\theta) = \frac{f_\perp}{2}\sec^2(\theta) \tag{2.60}$$

for the second effective biaxial lens. Then, considering that two cylindrical lenses, both oriented in the same μ direction, with focal lengths f_μ^A and f_μ^B, act as a single cylindrical lens with focal length

$$\frac{1}{f_\mu} = \frac{1}{f_\mu^A} + \frac{1}{f_\mu^B}, \tag{2.61}$$

we find

$$\frac{1}{f_x(\theta)} = 2\left(\frac{\cos^2(\theta)}{f} + \frac{\sin^2(\theta)}{f_\perp}\right) \tag{2.62}$$

$$\frac{1}{f_y(\theta)} = 2\left(\frac{\sin^2(\theta)}{f} + \frac{\cos^2(\theta)}{f_\perp}\right) \tag{2.63}$$

for the effective focal lengths of the cylindrical lens pair. The circularization condition for this pair is almost identical to the one derived in Eq. (2.54),

$$\frac{1}{f_x(\theta^*)} - \frac{1}{f_y(\theta^*)} = \left(\frac{\lambda}{\pi r}\right)^2 \left\{\frac{z_{0y}}{w_{0y}^2} - \frac{z_{0x}}{w_{0x}^2}\right\}, \tag{2.64}$$

from which we find

$$\theta^* = \frac{1}{2}\arccos\left[\frac{f'}{2}\left(\frac{\lambda}{\pi r}\right)^2 \left\{\frac{z_{0yy}}{w_{0y}^2} - \frac{z_{0x}}{w_{0x}^2}\right\}\right], \tag{2.65}$$

where

$$f' = \left(\frac{1}{f} - \frac{1}{f_\perp}\right)^{-1}. \tag{2.66}$$

The condition for the existence of an angle θ^* for which circularization is possible is given by

$$|f'| \leq f_{\max} = \left|\frac{2\left(\frac{\pi r}{\lambda}\right)^2}{\frac{z_{0y}}{w_{0y}^2} - \frac{z_{0x}}{w_{0x}^2}}\right|. \tag{2.67}$$

These results are identical to those derived for a perfectly uniaxial cylindrical lens pair (see Eqs. (2.55)–(2.56)), except that the nominal focal length f is replaced by the equivalent focal length f'.

Two Non-Identical Lenses. In the asymmetric case where $f_{2\perp} \neq f_{3\perp}$, perfect circularization is no longer achievable, because for any angle θ, the lens pair introduces a residual phase curvature to the beam that makes the beam elliptical and astigmatic. This statement can be proved mathematically. Suppose $f_{3\perp} = f_{2\perp} + \delta$, where $\delta \neq 0$. The lens with focal length $f_{2\perp}$ is equivalent to a pair of two uniaxial cylindrical lenses with focal lengths $f_{3\perp}$ and f_δ, where

$$f_\delta = \left(\frac{1}{f_{2\perp}} - \frac{1}{f_{3\perp}}\right)^{-1} = f_{2\perp}\left(1 - \frac{1}{1 + \delta/f_{2\perp}}\right)^{-1}. \tag{2.68}$$

This result is a direct consequence of the fact that two cylindrical lenses with focal lengths f_1 and f_2 mounted one behind the other along the same orientation are equivalent to a uniaxial cylindrical lens with focal length

$$f = \left(\frac{1}{f_1} + \frac{1}{f_2}\right)^{-1}. \tag{2.69}$$

An important consequence of Eq. (2.68) is that when $f_{2\perp}$ and $f_{3\perp}$ have the same sign, then $|f_\delta| > \min(|f_{2\perp}|, |f_{3\perp}|)$. This statement indicates that the equivalent lens with focal length f_δ introduces less focusing power than the more defective of the two cylindrical lenses in the direction perpendicular to their nominal focusing axes.

If the effect of f_δ is neglected, then the system effectively reduces to a lens pair with $f_{2\perp} = f_{3\perp}$, for which perfect circularization is possible. Therefore, the lens pair first circularizes the input beam and subsequently acts on the circularized beam with a uniaxial cylindrical lens of focal length f_δ, oriented at an angle $\theta^* + \pi/2$ with respect to the x axis. The resulting beam is an elliptical Gaussian beam with astigmatism oriented at $\theta^* + \pi/2$ relative to the x axis.

Note that this effect is expected to be negligible in practice for many tabletop optics applications. For instance, from our experimental measurements (see Sec. 2.3 in the main text), we estimate a transverse focal length of $f_\perp = 51.4$ m, where $f_{2\perp}$ and $f_{3\perp}$ are equal to within measurement precision, i.e. $\delta \ll f_{2\perp}$. These results imply

$$f_\delta = \frac{f_{2\perp}^2}{\delta} + \mathcal{O}(\delta^2) \gg f_{2\perp}, \tag{2.70}$$

making the residual astigmatism and ellipticity caused by the difference between the lenses negligible.

2.5.3 Astigmatism Correction Using a Cylindrical Lens Pair

We now prove that under appropriate conditions, a pair of cylindrical lenses with focal lengths f can transform an elliptical Gaussian beam with astigmatism into an elliptical Gaussian beam without astigmatism. We then propose a method to compute the set of viable focal lengths for which correcting astigmatism is possible.

Proof of the Existence of an Angle for Astigmatism Correction

We first show that, for sufficiently small values of $|f|$, a cylindrical lens pair can correct astigmatism for any elliptical Gaussian beam. We assume that the lens pair is located in

the plane $z = 0$. The input beam has a beam waist radii of w_{0x} and w_{0y} at the beam waists located at $z = z_{0x}$ and $z = z_{0y}$, respectively, where z_{0x}, z_{0y} may be positive or negative. After passing through the lens pair, the output beam has beam waist radii w'_{0x} and w'_{0y} at locations z'_{0x} and z'_{0y}, respectively (see Eqs. (2.43)–(2.44) in Sec. 2.5.1 for explicit definitions derived in terms of the input beam parameters). The lens pair acts as an effective biaxial lens with focal lengths $f_x(\theta) = f\sec^2(\theta)/2$ and $f_y(\theta) = f\csc^2(\theta)/2$ along the x and y axes, respectively. We define the distance between the beam waist locations

$$\Delta z(\theta) = z'_{0y}(\theta) - z'_{0x}(\theta). \tag{2.71}$$

Correcting astigmatism requires showing that there exists values of f and θ for which the beam waists overlap, i.e., $\Delta(z)(\theta) = 0$. We solve this problem for three possible cases that cover all possible configurations of z_{0x} and z_{0y}.

Case 1: $z_{0x} \leq 0$ **and** $z_{0y} \leq 0$ We prove that, for sufficiently small and positive values of f, there always exists at least one angle θ^* such that $\Delta(z)(\theta^*) = 0$. We choose f such that $0 < f \ll 1$. When $\theta = 0$, it follows from Eq. (2.43) that $z'_{0x}(0) = f/2 + \mathcal{O}(f^2)$ and $z'_{0y}(0) = z_{0y}$. For small enough values of f where the quadratic term is smaller than $f/2$, we have $z'_{0x}(0) > 0$ and $z'_{0y}(0) \leq 0$.

Similarly, when $\theta = \pi/2$, we have $z'_{0x}(\pi/2) = z_{0x}$, and $z'_{0y}(\pi/2) = f/2 + \mathcal{O}(f^2)$, such that, for small values of f, we have $z'_{0x}\left(\frac{\pi}{2}\right) \leq 0$ and $z'_{0y}\left(\frac{\pi}{2}\right) > 0$. As a result of these inequalities, we have $\Delta z(0) < 0$ and $\Delta z\left(\frac{\pi}{2}\right) > 0$, so that by the Intermediate Value Theorem for continuous functions, there exists at least one angle $\theta \in [0, \pi/2]$ such that $\Delta z(\theta) = 0$.

Case 2: $z_{0x} \geq 0$ **and** $z_{0y} \geq 0$ The proof for Case 2 is similar to Case 1. We choose f such that $f < 0$ and close to 0. We can then show that $\Delta z(0) > 0$ and $\Delta z(\pi/2) < 0$ for sufficiently small values of $|f|$. Using the Intermediate Value Theorem, we can show that there exists some $\theta \in [0, \pi/2]$ for which $\Delta z(\theta) = 0$.

Case 3: $z_{0x} \cdot z_{0y} < 0$ In this case, the cylindrical lens pair is placed in between the planes $z = z_{0x}$ and $z = z_{0y}$, making the proof slightly different from the previous two cases. We choose f such that $f < 0$ and close to 0. Without loss of generality, assuming that $z_{0x} < 0$ and $z_{0y} > 0$, we find $z'_{0x}(\pi/2) = z_{0x}$, and $z'_{0y}(\pi/2) = f/2 + \mathcal{O}(f^2)$. Therefore, for sufficiently small values of $|f|$

$$\Delta z(\pi/2) = z'_{0y}(\pi/2) - z'_{0x}(\pi/2) > 0. \tag{2.72}$$

Now we prove that there exists some angle θ close to 0 such that $\Delta z(\theta) < 0$ when choosing a small enough value of $|f|$. We choose ϕ such that $f_y(\phi) = f\csc^2(\phi)/2 = -z_{0y}$, when

$$\phi = \sin^{-1}\left(\sqrt{\frac{-f}{2z_{0y}}}\right). \tag{2.73}$$

Since we have chosen f to satisfy the condition $-1 \ll f < 0$, ϕ can be made arbitrarily close to 0. As a result, $f_x(\phi)$ will be close to $f_x(0) = f/2$. More precisely

$$f_x(\phi) = \frac{f}{2}\sec^2(\phi) = \frac{f}{2} + \mathcal{O}(f^2). \tag{2.74}$$

For this choice of ϕ and following Eq. (2.43), we conclude that $z'_{0x}(\phi) = \dfrac{f}{2} + \mathcal{O}(f^2)$, and $z'_{0y}(\phi) = -z_{0y}$. Therefore, $z'_{0x}(\phi) > 0$, and $z'_{0y}(\phi) < 0$. As a result of these inequalities we have $\Delta z(\phi) < 0$. From these inequalities and the Intermediate Value Theorem, we conclude that there exists some $\theta \in [\phi, \pi/2] \subset [0, \pi/2]$ where $\Delta z(\theta) = 0$.

Computing the Bounds on the Lens Pair Focal Length for De-Astigmatization

We seek the necessary conditions on the focal length f to correct astigmatism in elliptical Gaussian beams. Our approach involves finding the optimal angle $\theta^*(f)$ for different values of f. The domain of this function corresponds to the range of valid choices of f that allow for correcting astigmatism. It follows from Eq. (2.43) that given an angle θ, a necessary condition to correct astigmatism in an elliptical Gaussian beam is

$$z'_{0x}(f_x) = \frac{z_{0x}f_x(\theta)^2 + (z_{0x}^2 + z_{Rx}^2)f_x(\theta)}{(z_{0x} + f_x(\theta))^2 + z_{Rx}^2} = \frac{z_{0y}f_y(\theta)^2 + (z_{0y}^2 + z_{Ry}^2)f_y(\theta)}{(z_{0y} + f_y(\theta))^2 + z_{Ry}^2} = z'_{0y}(f_y(\theta)), \tag{2.75}$$

where $f_x = f_x(\theta) = f\sec^2(\theta)/2$ and $f_y = f_y(\theta) = f\csc^2(\theta)/2$. Solving this equation for θ results in a complicated equation involving multiple powers of $\sin(\theta)$ and $\cos(\theta)$ without a simple representation.

Alternatively, we can fix f_y and search for the f_x required to the correct astigmatism. Having both f_y and $f_x(f_y)$, we can infer the required angle θ and focal length f of the lens pair. Solving for f_x in Eq. (2.75), we find that

$$f_x(f_y) = \frac{A(f_y) \pm \sqrt{A(f_y)^2 - 4z'_{0y}(f_y)(z'_{0y}(f_y) - z_{0x})(z_{0x}^2 + z_{Rx}^2)}}{2\left(z'_{0y}(f_y) - z_{0x}\right)}, \tag{2.76}$$

where $A(f_y) = z_{0x}^2 + z_{Rx}^2 - 2z'_{0y}(f_y)z_{0x}$. For every f_y, we find the values of f and θ that correct astigmatism, namely

$$\frac{2}{f(f_y)} = \frac{1}{f_y} + \frac{1}{f_x(f_y)}, \tag{2.77}$$

and

$$\theta(f_y) = \arctan\left(\sqrt{\frac{f_x(f_y)}{f_y}}\right). \tag{2.78}$$

These equations produce valid choices of f and θ as long as the conditions $|f_y| > \frac{|f|}{2}$ and $f \cdot f(f_y) > 0$ are satisfied. These conditions guarantee that $f_x(f_y)$ and f_y have the same sign and are both larger than $f/2$. Consequently, there always exists a valid choice of f and θ that can be used to correct astigmatism in elliptical Gaussian beams.

2.5.4 Ellipticity and Astigmatism Correction with Three Cylindrical Lenses

In this section, we demonstrate that a cylindrical lens oriented along the major axis of an elliptical Gaussian beam, with an appropriate choice of focal length, guarantees the existence of a plane after the lens where the intensity profile becomes radially symmetric. We first prove that such a plane exists, provided the lens is aligned precisely along the major axis. We then study the resulting intensity profile when the orientation of the cylindrical lens deviates from the major axis.

Circularization with Perfectly Oriented Cylindrical Lenses

To circularize an elliptical Gaussian beam with astigmatism using a pair of cylindrical lenses, the intensity profile of the beam must be made radially symmetric in the plane where the lens pair is located. The first cylindrical lens in our three-lens solution, placed in the plane P_1 at $z = z_{P_1}$, fulfills this role by generating a radially symmetric intensity profile in a plane P_2 at $z = z_{P_2}$, where the lens pair is located (see Sec. 2.2.1).

We now prove that, for any input beam, there always exists a suitable choice of focal length f_1 such that the plane P_2 exists. We assume that at $z = z_{P_1}$, the radius of the beam along the y axis is larger than the radius along the x axis, i.e., $w_{0y}(z_{P_1}) > w_{0x}(z_{P_1})$,

such that $\Delta w_{xy}(z_{P_1}) = w_y(z_{P_1}) - w_x(z_{P_1}) > 0$. We choose the focal length of the first lens, oriented along y, to be $f_1 > 0$ and $f_1 \ll z_{R\mu}$ for both $\mu \in \{x, y\}$.

Using the transformation properties of Gaussian beams under lens action (see App. 2.5.1), we can show that the beam radius along the y axis focuses at

$$z_{0y} = z_{P_1} + f_1 + \mathcal{O}(f_1^2), \tag{2.79}$$

with a beam waist radius of

$$w_{0y} = \mathcal{O}(f_1). \tag{2.80}$$

For sufficiently small f_1, we have $w_{0y} < w_{0x}$ and $z_{0y} > z_{P_1}$. The beam focuses along the y axis after $z = z_{P_1}$ and achieves a smaller waist than along the x axis. As a result, the beam radius difference becomes negative at $z = z_{0y}$, i.e., $\Delta w(z_{0y}) = w_y(z_{0y}) - w_x(z_{0y}) < 0$. Since $\Delta w(z)$ is a continuous function of z, and $\Delta w(z_{P_1}) > 0$ whereas $\Delta w(z_{0y}) < 0$, it follows from the Intermediate Value Theorem that there exists some $z_{P_2} \in [z_{P_1}, z_{0y}]$ such that $\Delta w(z_{P_2}) = 0$. The plane P_2 is the plane at which the beam achieves a radially symmetric intensity profile, i.e., $w_x(z_{P_2}) = w_y(z_{P_2})$. Hence, for any elliptical Gaussian beam, we can always choose a sufficiently small focal length f_1 such that a symmetry plane z_{P_2} exists.

Limitations of Imperfectly Oriented Cylindrical Lenses

We have so far assumed that the first cylindrical lens is perfectly aligned with the major axis of the elliptical beam. In that case, given an appropriate choice of focal length f_1, there exists a plane P_2 at which the beam acquires a circular intensity profile. We now show that, if the lens is rotated with respect to the the elliptical beam, then perfect circularization cannot be realized. We then analyze how the circularity of the beam in the plane P_2 degrades as a function of the misalignment angle (see Fig. 2.11).

Let the location of the first cylindrical lens to be $z = 0$. When the lens is not properly aligned, the output is a generalized Gaussian beam with a $\mathbf{\Lambda}$ matrix whose real and imaginary parts do not commute: $[\Re\mathbf{\Lambda}(0^+), \Im\mathbf{\Lambda}(0^+)] \neq 0$. This nonzero commutator implies that the beam remains asymmetric at all planes beyond $z > 0$. Since the beam propagates in free space for $z > 0$, the commutator $[\Re\mathbf{\Lambda}(z), \Im\mathbf{\Lambda}(z)]$ remains nonzero at all such positions due to the transformation properties of the $\mathbf{\Lambda}$ matrix (see Table 2.1). However, for the beam to exhibit a radially symmetric intensity profile at some plane $z = z_0$, the commutator at that plane must vanish. This contradiction implies that no such plane can exist.

To clarify this point, recall that for a generalized Gaussian beam (see Eq. (2.25)), the real part of $\mathbf{\Lambda}(z)$ can be expressed in terms of the beam radii r_μ with $\mu \in \{0, 1\}$ and the

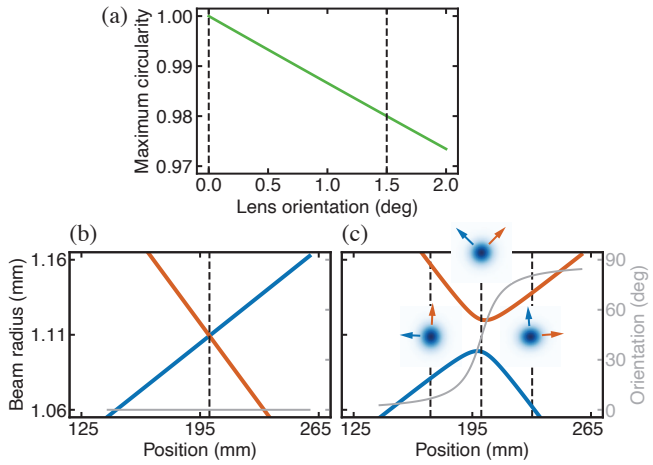


Figure 2.11: **Imperfect Lens Orientation.** A cylindrical lens of focal length $f = 500$ mm is placed at $z = 0$ mm at an angle θ with respect to the major axis of an incoming elliptical beam. The maximum circularity of the beam decreases away from 1.0 when the cylindrical lens is rotated away from $\theta = 0°$. At $\theta = 0°$ the beam radii intersect at $z = 200$ mm and the beam orientation remains constant. At $\theta = 1.5°$ the beam radii of the beam do not intersect, and the beam rotates by almost 90°. The beam reaches a maximum circularity of 0.98 at $z = 200$ mm.

beam orientation angle θ:

$$\Re\{\mathbf{\Lambda}(z)\} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} \frac{1}{r_0^2} & 0 \\ 0 & \frac{1}{r_1^2} \end{bmatrix} \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}. \tag{2.81}$$

If the intensity profile of the beam is radially symmetric, the beam radii must be equal, $r = r_0 = r_1$, simplifying the previous expression to

$$\Re\mathbf{\Lambda}(z) = \frac{1}{r^2}\mathbb{I}_2. \tag{2.82}$$

Since the $\mathbb{I}_2$ identity matrix commutes with any other matrix, we have $[\Re\mathbf{\Lambda}, \Im\mathbf{\Lambda}] = 0$. Hence, the beam is a Gaussian beam. This derivation proves that achieving a radially symmetric intensity profile is only possible if the beam is not a generalized Gaussian beam.

We now study through numerical simulations how misalignment of the first cylindrical lens affects the maximum circularity of the beam. We consider a perfectly uniaxial cylindrical lens with a focal length of $f = 500$ mm. We find a 2° misalignment of between the

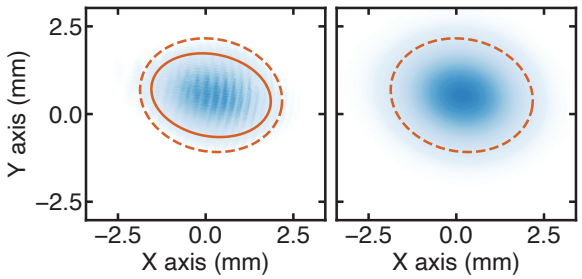


Figure 2.12: **Image Processing.** (left) The raw image captured by the camera exhibits interference effects caused by its protective glass cover. The beam radii along the principal axes are extracted from a rescaled best-fit ellipse (solid red) applied to the filtered image (dashed red). (right) The filtered image is obtained by applying a Gaussian filter with $\sigma = 0.55$ mm to the raw image. The principal axes are extracted by fitting the filtered image to the best-fit ellipse (dashed red).

focusing axis of the lens and the major axis of the beam and the reduces the maximum achievable circularity from 1.0 when perfectly aligned to 0.97 (see Fig. 2.11).

We further examine how the beam profile obtained for $\theta = 1.5°$ compares against the one obtained for $\theta = 0°$ (see insets in Fig. 2.11). When $\theta = 0°$, the beam orientation remains fixed at 0°, and the beam radii intersect at $z = 200$ mm. In contrast, when $\theta = 1.5°$, the beam orientation undergoes a rotation of nearly 90°, and the radii never intersect. As a result, the circularity peaks at 0.98 near $z = 200$ mm, placing a practical limit on the degree of circularization achievable using our three-lens solution.

2.5.5 Image Processing Method for Extracting Beam Properties

We now describe the image processing algorithm used to extract properties of the beam from intensity profiles captured at various positions along the propagation axis. From each image, we determine the orientation of the beam and extract the beam radii along the principal axes. For the analysis, we assume that the beam can be modeled by a generalized Gaussian beam (see App. 2.5.1) whose intensity profile is given by Eq. (2.26). Given a measurement of the intensity profile on a camera, the image processing task consists of extracting and diagonalizing the 2×2 real-valued and symmetric matrix $\Re\mathbf{\Lambda}(z)$.

Filtering of Interference Patterns

We first describe the method used to digitally filter the interference patterns observed when imaging the intensity profile of the beam on a CMOS camera (Basler ace2 a2A2840-48umBAS). The sensor of this camera has a protective glass cover that results in multiple reflections of the beam between the sensor and the glass cover. These reflections result in Fabry–Pérot fringes on the image (see Fig. 2.12a), which increase the complexity of the image processing algorithm and can lead to inaccurate estimates of the beam properties.

To remove these interference patterns, we first apply a Gaussian blur using a Gaussian kernel with a suitably chosen standard deviation, σ. The Gaussian blur acts as a two-dimensional low-pass filter, suppressing high-frequency interference fringes while slightly broadening the apparent size of the beam. The apparent increase in the beam size can be quantified using the properties of Gaussian functions. The convolution of two Gaussian functions with zero mean and standard deviations σ_1 and σ_2 results in another Gaussian function with zero mean and standard deviation $\sqrt{\sigma_1^2 + \sigma_2^2}$. The radius r of a generalized Gaussian beam is related to the standard deviation σ of the Gaussian function describing its intensity profile by the relation $r = 2\sigma$. Therefore, given the standard deviation σ of the Gaussian kernel used in the blur and the apparent radius of the filtered image r_{apparent}, the true radius of the beam r_{true} is given by

$$r_{\text{true}} = \sqrt{r_{\text{apparent}}^2 - 4\sigma^2}. \tag{2.83}$$

Extraction of Beam Parameters

We now describe the method used for extracting the orientation and radii of the beam from the filtered images. We denote the measured pixel values at the i-th row and j-th column of the camera sensor by $p_{ij} \in [0, 255]$. We first apply a binary mask to the filtered image, setting all pixels with values less than $\max_{ij} p_{ij}/e^2$ to zero. We then extract the coordinates of the pixels forming the boundary of the island of non-zero pixels remaining after the binary mask is applied using the *findContours* function of the OpenCv2 image analysis library [59]. We finally extract the ellipse of best fit to the pixel coordinates of the boundary using the *fitEllipse* function of OpenCv2, which uses a direct least square method [60]. This fit returns the orientation and radii of the best-fit ellipse, from which the beam orientation and radii are obtained after applying a Gaussian blur correction.

Chapter 3

A Feedback System for the Uniformization and Stabilization of Trap Depths in Optical Trap Arrays

Performing quantum computations with scientific or industrial applications requires a large number of logical and physical qubits, along with the execution of many-qubit gates. These computations typically run over extended durations and often must be repeated many times (see Sec. 1.2.2). For example, estimating the properties of the ground state of the unit cell of Li_2FeSiO_4, which is an industrially relevant problem for the design and manufacturing of rechargeable batteries, requires approximately 6652 logical qubits and an estimated run time of 6 days on a neutral atom array quantum processor, assuming each two-qubit gate takes 200 ns [10]. This highlights the need for quantum processors with a large number of physical qubits and reliable operation, where gate fidelities and SPAM errors remain stable over time. Even for more modest computational tasks, such as simulating quantum many-body dynamics of a strongly interacting spin system with fewer than 100 qubits and thousands of shots, a neutral atom array quantum processor may need to perform computations over several hours, assuming a cycle time of 1 second [12].

To support such computations, the quantum processor must operate over extended periods of time using a large number of qubits, without compromising performance. In the context of Rydberg atom array quantum processors, this requires satisfying two key conditions:

1. **The qubits must form an indistinguishable spatio-temporal ensemble:** The properties of the atoms, such as transition frequencies, coherence times, and fluo-

rescence rates, must remain stable throughout the computation and uniform across all traps in the array. Temporal drifts can degrade gate and measurement fidelities over time, while spatial inhomogeneities introduce systematic errors and prevent the simultaneous optimization of global gates and measurement fidelity for all atoms in the array.

2. **SPAM errors and qubit gate fidelity must remain stable over time and consistent across the array:** To ensure uniform qubit behavior, the control parameters, such as the intensity and frequency of the optical beams and the amplitude and phase of microwave pulses, must remain stable throughout the experiment and uniform across all sites. Shot-to-shot fluctuations or slow drifts in these parameters lead to variations in state preparation, measurement fidelity, and gate performance. Spatial inhomogeneities cause qubits to respond differently to global operations, reducing the overall fidelity of gates and readout. Stability and uniformity of the control fields are therefore essential for reliable and scalable quantum computation.

Achieving stable and uniform qubit properties and gate fidelities requires precise control over many aspects of the experimental environment. This includes stabilizing the intensity and frequency of the optical and microwave fields used for cooling, imaging, and gate operations, as well as maintaining a stable magnetic field across the array.

A key requirement for consistent qubit performance is the uniformity and stability of the optical tweezer trap depths, which are determined by the optical power delivered to each trap. The optical transitions of the atoms experience AC Stark shifts that scale with the local trap depth. Since atomic cooling, imaging, and qubit gate operations are sensitive to these shifts, shot-to-shot fluctuations and slow drifts in trap depth cause changes in atomic temperatures, loading probabilities, and gate fidelities across the array over time. Such fluctuations and drifts undermine the conditions necessary for long quantum computations in practical applications.

Arrays of optical traps in Rydberg atom array quantum processors can be generated using passive optical elements, such as microlens array [61], or active elements, including spatial light modulators [30, 62, 8], acousto-optical deflectors [12, 21], and digital micromirror devices [63].

In this chapter, we address the problem of trap depth non-uniformity and introduce a method for real-time stabilization in large arrays of optical tweezers. We focus specifically on trap arrays generated using one or two acousto-optical deflectors (AODs). While stabilization of the total laser power is commonly implemented in the literature [64, 54, 30, 12, 65], the individual depths of the traps generated by AODs can still fluctuate and

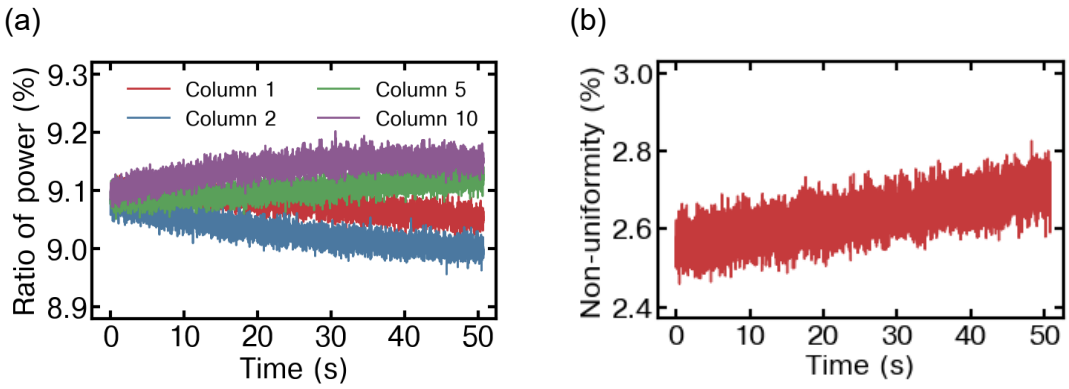


Figure 3.1: **Drift in the Relative Power of Optical Traps in an 11×11 Tweezer Array.** (a) Fraction of the total power contained in four selected columns of an initially uniformized 11×11 trap array over a duration of 50 seconds. Each column exhibits independent drift and fluctuations. For example, column 2 shows a steady increase in relative power while column 10 decreases. (b) Temporal evolution of the non-uniformity metric σ/μ, where σ is the standard deviation and μ is the mean of the trap depths. The metric increases over time, indicating a loss of uniformity in trap depths due to drift in the relative optical power distribution.

drift relative to one another, even when the total optical power is stabilized and the array is initially uniformized at the beginning of computation (see Fig. 3.1). This degradation in trap depth uniformity can undermine the reliability requirements outlined above. We propose a solution to this problem based on closed-loop control.

First, we present the design and evaluate the performance of a closed-loop feedback system for real-time uniformization and stabilization of trap depths in one- or two-dimensional arrays (see Sec. 3.1). This system continuously updates the control parameters in the RF signals supplied to the AODs to suppress fluctuations and drifts in trap depths. We begin by describing the hardware implementation of the feedback system (see Sec. 3.1.1), followed by a mathematical formulation of the trap depth uniformization problem (see Sec. 3.1.2). We then describe how correlated fluctuations in trap depth across the array can be decomposed into a set of uncorrelated temporal modes and efficiently suppressed through appropriate modifications to the control parameters (see Sec. 3.1.3). Lastly, we outline the feedback logic used in the feedback loop (see Sec. 3.1.4) and analyze its runtime performance (see Sec. 3.1.5).

Building on the stabilization of trap depths during steady-state operation, we extend our approach to maintain uniformity throughout dynamic processes such as adiabatic ramping. Specifically, we describe a method for ramping down the trap depths while preserving their uniformity across the array using acousto-optical deflectors (see Sec. 3.2). This is particularly important for applications such as microwave spectroscopy in shallow optical traps, where spatial variations in trap depth introduce position-dependent AC Stark shift and reduce the effective coherence time T_2^* [66]. We first explain how the adiabatic ramps are implemented using the AODs (see Sec. 3.2.1), and then describe how these ramps enable temperature measurements of the atoms in the tweezers (see Sec. 3.2.2).

3.1 Trap Depth Uniformization and Stabilization Module

In AOD-generated optical tweezer arrays, even when the total trap power is actively stabilized, slow drifts and fluctuations persist in the relative depths of individual traps (see Fig. 3.1). These relative variations originate from imperfections and time-dependent changes in optical and electronic components such as acousto-optical deflectors (AODs), amplifiers, and beam shaping optics.

This section addresses the problem of measuring, controlling, and uniformizing the trap depths across the entire array in a closed-loop feedback system. We develop a linear model relating changes in control parameters to variations in trap depths, describe the hardware setup and calibration procedures, and introduce an update rule for adjusting the control parameters to efficiently correct errors. Additionally, we present a temporal mode decomposition technique that isolates distinct noise sources, which are then addressed using independent digital PID loops to improve feedback performance.

3.1.1 Hardware Implementation

This section describes the hardware setup used to implement the closed-loop feedback system for real-time stabilization and uniformization of trap depths across a two-dimensional array of optical tweezers (see Fig. 3.2).

An Arbitrary Waveform Generator (AWG, Spectrum M4i.6621-x8) generates two poly-

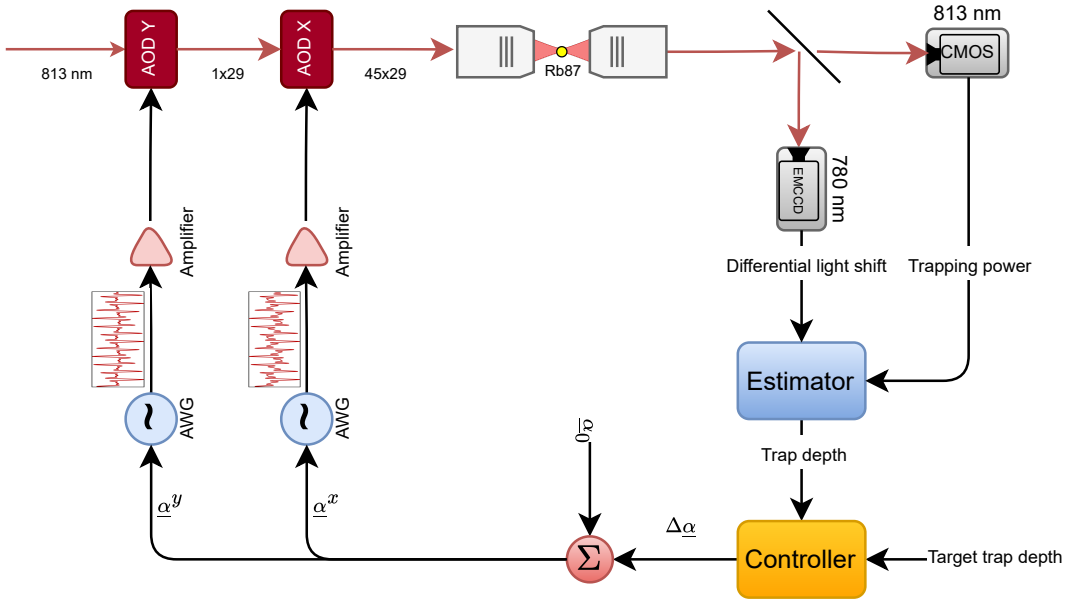


Figure 3.2: **Hardware Architecture for Closed-Loop Stabilization and Uniformization.** Schematic of the closed-loop feedback system used to stabilize and uniformize trap depths in a 45×29 ^{87}Rb atom array. Polychromatic RF waveforms are generated by a pair of AWG channels, amplified, and sent to orthogonal AODs to create a multiplexed tweezer array. The laser beam is focused by a microscope objective to trap single atoms and is collected by a second objective. The outgoing light is split: trapping light at 813 nm is detected by a CMOS camera to estimate relative tweezer powers, while fluorescence at 780 nm is imaged by an EMCCD camera to extract atom-dependent light shifts. These measurements feed into an estimator module (see Fig. 3.3) that reconstructs the trap depth vector $\underline{D}^{\text{estimate}}$. A controller module (see Fig. 3.6) computes a correction vector $\Delta\underline{\alpha}$ to update the AWG waveforms and minimize the deviation $\|\underline{D}^{\text{estimate}} - \underline{D}^{\text{target}}\|$ (see Sec. 3.1.2).

chromatic RF waveforms:

$$V_x(t) = \sum_{i=0}^{M-1} \alpha_i^x \cos(2\pi\nu_i^x t + \phi_i^x), \tag{3.1}$$

$$V_y(t) = \sum_{i=0}^{N-1} \alpha_i^y \cos(2\pi\nu_i^y t + \phi_i^y), \tag{3.2}$$

where ν_i^x and ν_i^y are the frequency components that determine the horizontal and vertical positions of the traps, respectively. The amplitudes $\alpha_i^x, \alpha_i^y \geq 0$ control the optical power delivered to each column or row in the array, and the phases ϕ_i^x, ϕ_i^y are free parameters chosen to minimize the crest factor of the waveforms to suppress undesired non-linearities in the system [67].

These RF waveforms are amplified using high-power linear RF amplifiers (Mini-Circuits ZHL-50W-52-S+) before being sent to a pair of orthogonally oriented acousto-optical deflectors (AA Optoelectronics DTSX-400-800.850). The AODs diffract an input trapping laser beam to generate an $N \times M$ two-dimensional array of optical tweezers, with each site corresponding to a unique pair of frequency components (ν_i^x, ν_j^y).

The traps are focused using a high-NA microscope objective (Mitutoyo G PLAN APO 50x) lens inside a vacuum cell and are used to trap single atoms in the $N \times M$ array. The trapping beam and the atomic fluorescence signal are collected by another high-NA microscope objective. A dichroic mirror separates the two paths. The trapping beam is monitored using a CMOS camera (Basler a2A2840-48umBAS) and the atomic signal is captured using an EMCCD camera (Andor iXon Ultra 888).

The data captured from the trapping power and atomic signal is used in real-time to estimate the trap depths in the array. The estimated trap depths are compared to the target values, and an error signal is calculated. This error signal is processed to update the waveforms generated by the AWG in Eqs. (3.1)–(3.2) to correct the trap depth deviations.

3.1.2 Formulation of the Uniformization Problem

In this section, we define the variables and formulate the trap depth uniformization and stabilization problem mathematically.

The uniformization problem is formulated as an optimization task to adjust the waveforms supplied to the AODs in order to minimize the discrepancy between the desired and actual trap depths experienced by the atoms. This is mathematically expressed as

$$\min_{V_x(t), V_y(t)} \left\| \underline{D}^{\text{traps}}(t) - \underline{D}^{\text{target}} \right\|_2 , \tag{3.3}$$

where $V_x(t)$ and $V_y(t)$ are the RF waveforms supplied to the AODs (see Eqs. (3.1)–(3.2)). The vectors $\underline{D}^{\text{traps}}(t) \in \mathbb{R}^{NM \times 1}$ and $\underline{D}^{\text{target}} \in \mathbb{R}^{NM \times 1}$ denote the measured and target trap depths, respectively, arranged in row-major order across the $N \times M$ trap array. Since the goal is to uniformize the array, the entries of $\underline{D}^{\text{target}}$ are equal. However, we do not assume this in our formulation to keep the treatment general.

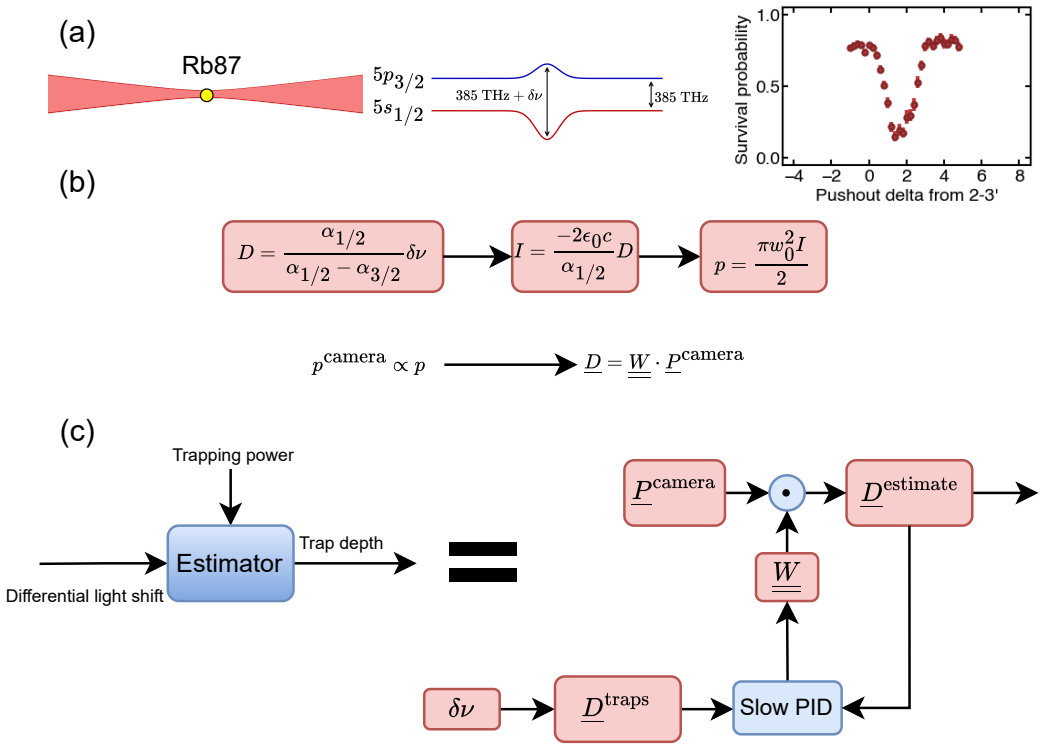


Figure 3.3: **Trap Depth Estimator for Closed-Loop Stabilization and Uniformization.** Implementation of the estimator module within the stabilization and uniformization feedback system (see Fig. 3.2). (a) The differential light shift $\delta\nu$ of the D_2 transition is measured using pushout spectroscopy: a near-resonant beam is swept across frequencies, and the frequency that maximizes atom loss is taken as the light-shifted resonance. (b) The measured $\delta\nu$ is related to the trap intensity via atomic polarizabilities of the $5S_{1/2}$ and $5P_{3/2}$ states, allowing extraction of trap depth and peak intensity. Assuming Gaussian beam profiles, the trap power can be inferred. A diagonal weight matrix $\underline{\underline{W}}$ is introduced to map CMOS-measured power $\underline{P}^{\text{camera}}$ to trap depth $\underline{D}$. (c) The estimator combines fast CMOS measurements of $\underline{P}^{\text{camera}}$ with slowly refreshed $\delta\underline{\nu}$ estimates from the EMCCD to compute $\underline{D}^{\text{estimate}} = \underline{\underline{W}} \cdot \underline{P}^{\text{camera}}$. A slow PID loop updates $\underline{\underline{W}}$ to reduce error between $\underline{D}^{\text{estimate}}$ and the EMCCD-inferred $\underline{D}^{\text{traps}}$ (see Sec. 3.1.2).

However, the formulation presented in Eq. (3.3) is inadequate for real-time uniformization and stabilization. This is because the trap depths $\underline{D}^{\text{traps}}$ are measured via in-situ atomic measurements using pushout spectroscopy, which can take on the order of minutes to perform. In-situ measurements can be accelerated using alternative techniques such as fluorescence spectroscopy or by using the loading efficiency as a proxy for trap depth [64, 54], but the update rate of these methods is ultimately limited by the cycling rate of the neutral atom array quantum processor, typically around 1 Hz.

As an alternative, the trap depths can be estimated more rapidly by measuring the optical power at each trap site using a CMOS camera (see Sec. 3.1.1), while the CMOS camera is continuously calibrated to the signal measured by the EMCCD camera (see Fig. 3.3). These measurements are represented by $\underline{P}(t) \in \mathbb{R}^{NM\times 1}$ in row-major order. The trap depths are then estimated as

$$\underline{D}^{\text{estimate}}(t) = \underline{\underline{W}}[t_0(t)] \cdot \underline{P}(t), \tag{3.4}$$

where $\underline{\underline{W}}[t_0(t)] \in \mathbb{R}^{NM\times NM}$ is a diagonal weight matrix calibrated at time $t_0(t)$ via atomic spectroscopy. Here, $t_0(t)$ denotes the most recent calibration time prior to t. The matrix $\underline{\underline{W}}[t_0(t)]$ maps the measured optical powers to the trap depths experienced by the atoms.

While $\underline{\underline{W}}[t]$ may vary slowly over time due to factors such as changes in trap polarization or drift in the sensitivity of the pixels of the camera, we assume that these variations are slow enough for $\underline{D}^{\text{estimate}}(t)$ to remain a reliable proxy for $\underline{D}^{\text{traps}}$ over timescales much longer than the interval between calibrations. As a result, $\underline{D}^{\text{estimate}}(t)$ can be used in the feedback loop, while $\underline{\underline{W}}$ is recalibrated periodically during extended computations.

Lastly, we simplify the control parameters by representing the waveforms supplied to the AODs as a single vector in $\mathbb{R}^{(N+M)\times 1}$. We assume that in the waveforms generated by the AWG (see Eqs. (3.1)–(3.2)), the phases ϕ_i^x, ϕ_i^y and frequencies ν_i^x, ν_i^y remain constant. Uniformization and stabilization are performed solely by updating the amplitudes $\alpha_i^x(t)$ and $\alpha_i^y(t)$, which we refer to as the control parameters of the optimization problem. These are represented vectorially as

$$\underline{\alpha}^x(t) = \begin{pmatrix} \alpha_0^x(t) \\ \alpha_1^x(t) \\ \vdots \\ \alpha_{M-1}^x(t) \end{pmatrix} \in \mathbb{R}^{M\times 1}, \qquad \underline{\alpha}^y(t) = \begin{pmatrix} \alpha_0^y(t) \\ \alpha_1^y(t) \\ \vdots \\ \alpha_{N-1}^y(t) \end{pmatrix} \in \mathbb{R}^{N\times 1} \tag{3.5}$$

and

$$\underline{\alpha}(t) = \begin{pmatrix} \underline{\alpha}^x(t) \\ \underline{\alpha}^y(t) \end{pmatrix} \in \mathbb{R}^{(N+M)\times 1}. \tag{3.6}$$

Therefore, the optimization problem presented in Eq. (3.3) can be reformulated as

$$\min_{\underline{\alpha}(t)} \left\| \underline{D}^{\text{estimate}}(t) - \underline{D}^{\text{target}} \right\|_2, \tag{3.7}$$

where $\underline{D}^{\text{estimate}}$ is computed using a weight matrix $\underline{\underline{W}}$ that is calibrated periodically (see Eq. (3.4)).

3.1.3 Control Updates and Temporal Modes

In this section, we introduce the notions of control updates and temporal modes, which are used to improve the performance of the feedback loop.

We begin by describing how the control parameters $\underline{\alpha}$ are updated to reduce the error signal $\underline{D}^{\text{estimate}} - \underline{D}^{\text{target}}$ (see Sec. 3.1.3). Due to the frequency-dependent diffraction efficiency of the AODs and the nonlinear response of the electronic amplifiers preceding them (see Fig. 3.2), changes in $\underline{\alpha}$ can affect the trapping powers in ways that deviate from idealized behavior. By accurately modeling these effects, we can update $\underline{\alpha}$ more intelligently, effectively increasing the speed of the PID feedback loop. Specifically, this approach moves the control parameters towards the solution manifold along the shortest path, enabling faster and more precise stabilization of the trap depths.

Afterward, we describe the concept of temporal modes, which are used to address uncorrelated noise components using distinct control parameters (see Sec. 3.1.3). The error signal $\underline{D}^{\text{target}} - \underline{D}^{\text{estimate}}(t)$ is an NM-dimensional vector whose elements are typically correlated in time. We show that this signal can be decomposed into a set of NM uncorrelated modes, each with a distinct power spectrum. These modes can be suppressed independently, for example, using separate PID controllers with mode-specific parameters.

Control Updates

In this section we explain how to determine the required changes $\Delta\underline{\alpha}$ to the control parameters $\underline{\alpha}$ such that a measured error signal $\left\| \underline{D}^{\text{estimate}} - \underline{D}^{\text{target}} \right\|_2$ is suppressed. We first

approximate the system using a linear model. Then, we show that in general the $\Delta\underline{\alpha}$ required to minimize the error is not unique, and describe how to find the solution with the least norm $\|\Delta\underline{\alpha}\|_2$.

Consider the case where the control parameters are initially set to $\underline{\alpha}^0$ at $t = 0$, for which $\left\|\underline{D}^{\text{estimate}}(0) - \underline{D}^{\text{target}}\right\|_2$ is minimized, and only a small change $\Delta\underline{\alpha}$ is applied to stabilize the trap depths at $t > 0$, i.e. $\underline{\alpha}(t) = \underline{\alpha}^0 + \Delta\underline{\alpha}(t)$. We approximate the trap powers by

$$\underline{P}(t) = \underline{P}(0) + \underline{\underline{J}} \cdot \Delta\underline{\alpha}(t) + \delta\underline{P}(t), \tag{3.8}$$

where $\delta\underline{P}(t)$ represents time-dependent fluctuations in the trap powers, and $\underline{\underline{J}} \in \mathbb{R}^{NM\times(N+M)}$ is defined as

$$\underline{\underline{J}} = \left\{ \left.\frac{\partial \underline{P}_i}{\partial \underline{\alpha}_j}\right|_{\underline{\alpha}^0} \right\}_{ij} \quad (t = 0), \tag{3.9}$$

i.e., the Jacobian of the power vector with respect to the control parameters, evaluated at $\underline{\alpha} = \underline{\alpha}^0$. We assume that $\underline{\underline{J}}$ varies slowly with time and requires recalibration at a rate much lower than the feedback loop. The matrix $\underline{\underline{J}}$ can be experimentally determined by perturbing each component of $\underline{\alpha}$ individually and measuring the corresponding changes in the trap powers, starting from the initial value $\underline{\alpha}^0$.

In an ideal system, such as one with frequency-independent AOD diffraction efficiency and perfectly linear electronic amplifiers, each control parameter α_i^x or α_i^y would only affect the optical power of traps in the corresponding column or row, respectively. However, in practice, the frequency-dependent diffraction efficiency of the AODs and nonlinearities in the amplifiers lead to intermodulation, causing each control parameter to influence multiple rows or columns across the array (see Fig. 3.8). Moreover, thermalization effects can alter how the trap powers depend on the control parameters over time. These factors make $\underline{\underline{J}}$ generally dense and time dependent. The time dependence necessitates periodic recalibration of $\underline{\underline{J}}$. We will later introduce a method to quantify this time dependence, allowing us to estimate an approximate recalibration interval for $\underline{\underline{J}}$, which in our system is approximately 52 min (see Sec. 3.1.6).

Having introduced the model described in Eq. (3.8), we now seek a value of $\Delta\underline{\alpha}$ that minimizes the error $\left\|\underline{D}^{\text{estimate}}(t) - \underline{D}^{\text{target}}\right\|_2$. This is equivalent to solving the optimization problem

$$\min_{\Delta\underline{\alpha}} \left\| \left[\underline{\underline{W}}^{-1} \cdot \underline{D}^{\text{target}} - (\underline{P}(0) + \delta\underline{P}(t))\right] - \underline{\underline{J}} \cdot \Delta\underline{\alpha} \right\|_2, \tag{3.10}$$

where $\underline{\underline{W}}$ is the weight matrix (see Eq. (3.4)).

In general, the matrix $\underline{\underline{J}}$ is neither injective (one to one) nor surjective (onto), which affects how the optimization problem can be solved. If $\underline{\underline{J}}$ were injective, every change in

trap power $\Delta\underline{P}$ would correspond to a unique change in the control parameters $\Delta\underline{\alpha}$. This is not the case. For example, increasing all components of $\underline{\alpha}^x$ or all components of $\underline{\alpha}^y$ produces the same uniform increase in trap power across the array, demonstrating that $\underline{\underline{J}}$ is not injective. Moreover, $\underline{\underline{J}}$ is not surjective when $N + M < NM$ because of the rank-nullity theorem. This inequality holds for any two-dimensional trap array except the 2×2 case, meaning the dimension of the control parameter space is smaller than the number of generated traps.

As a result, the optimization problem in Eq. (3.10) may not admit a solution that drives the error to zero. Moreover, the choice of $\Delta\underline{\alpha}$ which minimizes the error is not unique.

The solution $\Delta\underline{\alpha}^*$ to the optimization problem with the smallest norm $\|\Delta\underline{\alpha}\|_2$ is calculated using the Moore–Penrose pseudoinverse of $\underline{\underline{J}}$, using the relation

$$\Delta\underline{\alpha}^* = \underline{\underline{J}}^+ \cdot \left(\underline{\underline{W}}^{-1} \cdot \underline{D}^{\text{target}} - (\underline{P}(0) + \delta\underline{P}(t))\right), \tag{3.11}$$

where $\underline{\underline{J}}^+$ denotes the pseudoinverse of $\underline{\underline{J}}$. This equation determines both the magnitude and direction of the change in control parameters $\underline{\alpha}$ needed to reduce the error in trap depths. In practice, the magnitude from this equation may be disregarded, using only its direction to guide the adjustment. The actual magnitude of the update may then be set independently by a PID feedback loop, which controls the strength and speed of the parameter changes based on the error signal.

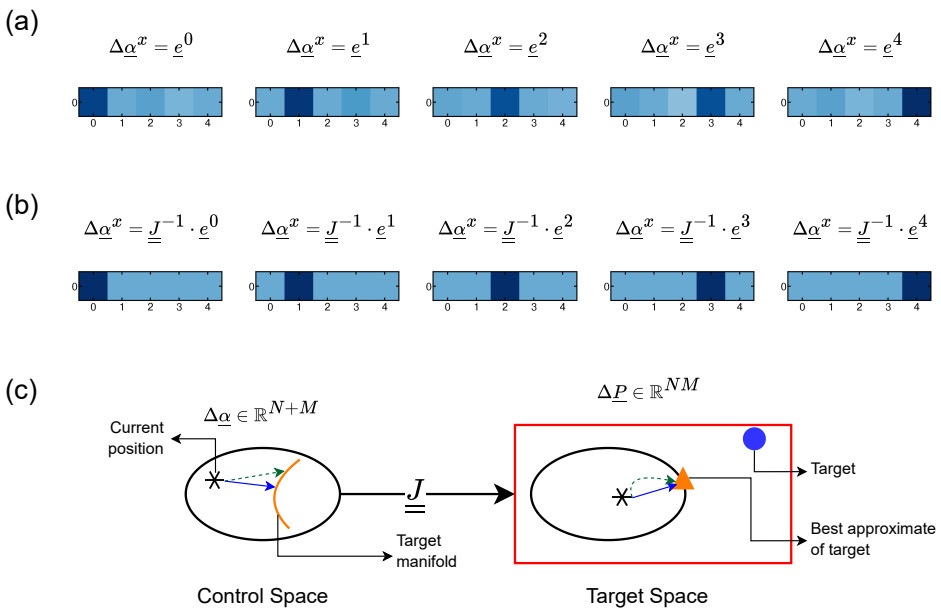


Figure 3.4: **Update Rule for Control Parameters to Correct for Trap Depth Noise.** (a) Effect of perturbing the control parameters $\underline{\alpha}^x$ along individual basis directions $\underline{e}^i$ in a 5×1 trap array. Each change primarily affects the corresponding column, but also induces residual changes in other traps due to nonlinear intermodulation from the RF amplifier. The varying color intensity on the ith trap indicates non-uniform sensitivity to each control parameter. (b) Applying the Jacobian-based update rule (see Eq. (3.11)) using the system's Jacobian matrix $\underline{\underline{J}}$ (see Eq. (3.8)) compensates for intermodulation and non-uniform response. This enables targeted, uniform updates of individual columns. In this 1D example, $\underline{\underline{J}}$ is invertible. (c) Conceptual illustration for larger arrays where $\underline{\underline{J}}$ is not invertible. The control space (left) and the higher-dimensional trap depth space (right) are connected via $\underline{\underline{J}}$. The accessible set of trap depths (black ellipse on the right) is a submanifold of the full target space. The current trap depths (black asterisk) and desired trap depths (blue circle on the right) are generally not directly connected via the control parameters. The best approximation of the target (orange triangle on the right) lies on the accessible manifold. Its pre-image under $\underline{\underline{J}}$ is a submanifold in control space (orange curve on the left). The update rule computes the smallest-norm change $\Delta\underline{\alpha}$ that brings the current control parameters to the closest point on this pre-image, moving along the shortest path (solid blue arrow), while naive rules may follow suboptimal directions (dashed green arrow).

Temporal Modes

This section presents a method to decompose the measured trap depth error signal into a set of uncorrelated components called temporal modes. By isolating distinct sources of noise that affect the trap depths differently over time, we can apply mode-specific feedback controls to each mode, improving overall stabilization performance.

Each atomic spectroscopy measurement or image of the trapping beams produces a vector of NM trap depths. The error in each trap depth, $|D_i^{\text{target}} - D_i^{\text{estimate}}(t)|$, can be modeled as a time-dependent random variable. These errors are generally correlated across traps due to shared optical and electronic components. For instance, all traps are illuminated by the same laser, so fluctuations in laser power induce correlated errors throughout the array. Additionally, since the traps are multiplexed by the same AODs and focused by a common lens, thermal fluctuations affecting the AODs and the lens can cause correlated variations in trap powers among traps in the same row, column, or local neighborhood.

We represent the error signal

$$\underline{E}(t) = \underline{D}^{\text{target}} - \underline{D}^{\text{estimate}}(t), \tag{3.12}$$

in a rotated basis whose components are temporally uncorrelated (see Fig. 3.5). The rotated basis vectors are defined to be the temporal modes. This decomposition separates distinct noise sources, each characterized by its own frequency spectrum. Feedback loops with mode-specific parameters, such as independent PID controllers, can then be applied to suppress fluctuations in each temporal mode selectively. Modes with negligible impact on system performance may be discarded. Additionally, the mode corresponding to uniform changes across all traps, likely due to laser power fluctuations, can be handled separately by a PID controller acting directly on the laser power rather than through the control parameters $\underline{\alpha}$ driving the AODs.

We now describe how the rotated basis vectors can be computed. Two random variables $X(t)$ and $Y(t)$ are said to be correlated if and only if

$$\left\langle \left[X - \langle X \rangle\right] \left[Y - \langle Y \rangle\right] \right\rangle_t \neq 0. \tag{3.13}$$

We expect the components of the error signal $\underline{E}$ to be correlated (see Eq. (3.12)). To find the temporal modes, we compute the correlation matrix of $\underline{E}$, defined as

$$\underline{\underline{\mathcal{C}}}\left(\underline{E}\right) = \left\{ \frac{\left\langle \left[E_i - \langle E_i \rangle\right] \left[E_j - \langle E_j \rangle\right] \right\rangle_t}{\sqrt{\left\langle \left[E_i - \langle E_i \rangle\right]^2 \right\rangle_t \left\langle \left[E_j - \langle E_j \rangle_t\right]^2 \right\rangle_t}} \right\}_{ij}, \tag{3.14}$$

where $\underline{\underline{\mathcal{C}}}(\underline{E}) \in \mathbb{R}^{NM\times NM}$ is the correlation matrix. Correlation between different components of $\underline{E}$ implies that the off-diagonal elements of $\mathcal{C}$ are nonzero. However, the correlation matrix is diagonal in its eigenbasis by construction. Therefore, the temporal modes are simply the normalized eigenvectors of the correlation matrix.

Let the orthonormal eigenvectors of the correlation matrix, that is, the temporal modes, be denoted by $\underline{\tilde{\mathcal{C}}}^i$ for $i \in \{0, 1, \ldots, NM-1\}$. Then the error signal can be expressed in the temporal mode basis as

$$\underline{E} = \sum_{i=0}^{NM-1} \tilde{E}^i \, \underline{\tilde{\mathcal{C}}}^i, \tag{3.15}$$

where the mode amplitudes are given by

$$\tilde{E}^i = \langle \underline{\tilde{\mathcal{C}}}^i, \underline{E} \rangle. \tag{3.16}$$

The control parameters can now be updated to selectively address the desired temporal mode components of the error signal. These updates may be mode specific, allowing different PID controllers to be tuned individually for each mode. This is useful because different modes often exhibit distinct frequency spectra, so different PID parameters may be optimal for suppressing fluctuations in each component.

In practice, estimating the correlation matrix requires acquiring a time series of the trap depth errors. Once this time trace is obtained, $\underline{\underline{\mathcal{C}}}(\underline{E})$ can be computed and diagonalized to yield the temporal modes.

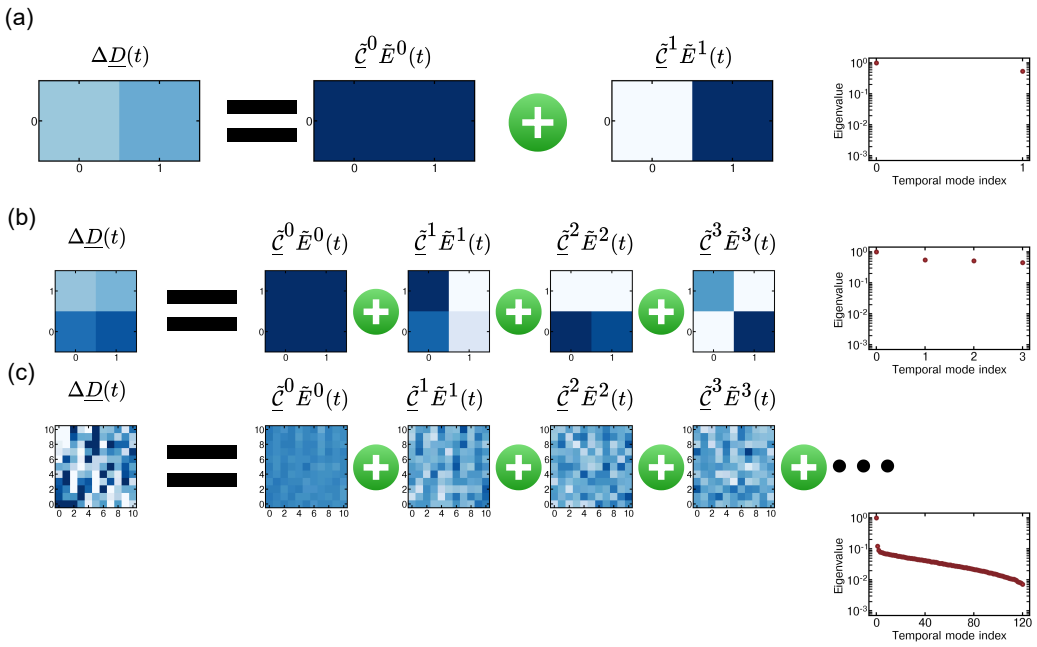


Figure 3.5: **Decomposition of Trap Depth Noise into Uncorrelated Temporal Modes.** (a) Noise in a 2×1 trap array is decomposed into uncorrelated temporal modes by diagonalizing the correlation matrix of the estimated trap depths over a 4 s period (see Eq. (3.14)). The first mode corresponds approximately to a global fluctuation $(1\ 1)^T/\sqrt{2}$, and the second to a differential fluctuation $(1\ -1)^T/\sqrt{2}$, indicating two independent sources of noise: fluctuations in total laser power and in AOD diffraction efficiency, which affects the power balance between traps. The two eigenvalues of the correlation matrix $\underline{\underline{\mathcal{C}}}$ are shown to the right. (b) Decomposition of trap depth fluctuations in a 2×2 array into four orthogonal modes. The first mode is nearly global, while the second and third correspond to row-wise and column-wise variations. The column mode has a larger eigenvalue, indicating more significant column fluctuations. The four eigenvalues of $\underline{\underline{\mathcal{C}}}$ are shown in descending order. (c) Temporal modes for an 11×11 array, illustrating spatial structure in the trap depth noise and the relative importance of collective fluctuation modes.

3.1.4 Stabilization Protocol

In this section, we describe the protocol used to uniformize and stabilize the trap depths in optical trap arrays. We use the notation introduced in the problem formulation (see Sec. 3.1.2) and combine the methods for control parameter updates (see Sec. 3.1.3) and temporal mode decomposition (see Sec. 3.1.3) to achieve this goal.

First, we describe the initialization process of the feedback loop. Then, we describe the feedback loop itself.

Initialization: The weight matrix $\underline{\underline{W}}$ is calibrated by comparing the power measurements from the CMOS camera and the trap-depth spectroscopy. Then, heuristic updates are applied to the control parameters $\underline{\alpha}$ until the estimated trap depths approximately match the target depths. The control parameters after this step are labeled $\underline{\alpha}^0$. The Jacobian matrix $\underline{\underline{J}}$ and the temporal modes $\left\{\underline{\tilde{\mathcal{C}}}^i\right\}$ are then computed (see Sec. 3.1.3).

After initialization, the following four steps are repeated continuously during the stabilization process (see Fig. 3.2 and Fig. 3.6):

1. **Measure the error signal:** The trap powers $\underline{P}$ are measured using the CMOS camera and used to estimate the trap depths via $\underline{D}^{\text{estimate}} = \underline{\underline{W}} \cdot \underline{P}$. The error signal is calculated as $\underline{E} = \underline{D}^{\text{target}} - \underline{D}^{\text{estimate}}$.

2. **Decompose the error signal into its temporal modes:** The error signal is represented in the temporal mode basis, and the component $\tilde{E}^i$ along each mode vector $\underline{\tilde{\mathcal{C}}}^i$ is extracted (see Eq. (3.16)).

3. **Calculate the control parameter changes required to correct each temporal mode:** A digital PID loop is created for each mode amplitude $\tilde{E}^i$. Each PID loop determines the required correction magnitude $|\Delta\underline{\alpha}^i|$, while the correction direction $\underline{\hat{\alpha}}^i$ is computed using the control parameter update rule described in Eq. (3.11).

4. **Update the control parameters** The net change in control parameters is computed by summing the contributions from all PID loops

$$\Delta\underline{\alpha} = \sum_{i=0}^{NM-1} \Delta\underline{\alpha}^i. \tag{3.17}$$

The control parameters are then updated using

$$\underline{\alpha} = \underline{\alpha}^0 + \Delta\underline{\alpha}. \tag{3.18}$$

Lastly, the weight matrix $\underline{\underline{W}}$, the Jacobian matrix $\underline{\underline{J}}$, and the temporal modes $\left\{\underline{\tilde{C}}^i\right\}$ are periodically recalibrated, depending on how quickly they drift over time in the system. This prevents the performance of the uniformization and stabilization loop from degrading.

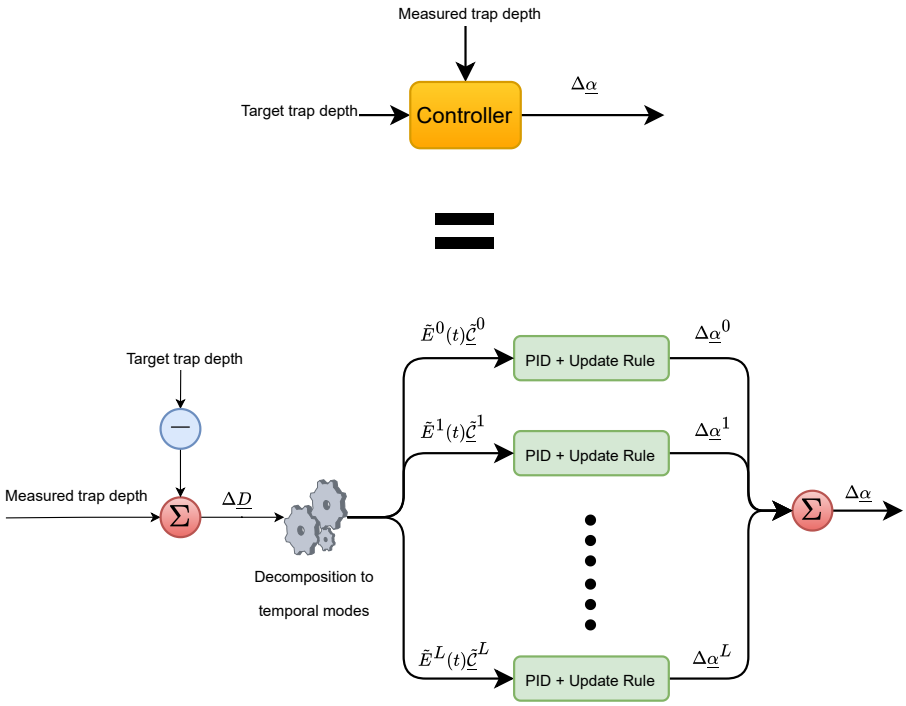


Figure 3.6: **Controller Module Used in the PID Protocol to Update Control Parameters.** The controller module used in the uniformization and stabilization feedback system for an $N \times M$ trap array (see Fig. 3.2). The measured trap depth $\underline{D}^{\text{estimate}}$ is subtracted from the target depth $\underline{D}^{\text{target}}$ to compute the trap depth error $\Delta\underline{D}$. This error is decomposed into a truncated basis of $L+1$ uncorrelated temporal modes $\{\tilde{\underline{C}}^i\}_{i=0}^{L}$, where $L+1 \leq NM$ and $L+1$ is the number of independent modes the controller is designed to correct. Each mode amplitude $\tilde{E}^i(t)$ is extracted via projection and sent to a dedicated PID loop, which computes the magnitude of the corrective response. The update direction is given by the update rule (see Eq. (3.11)) and the temporal mode vector $\tilde{\underline{C}}^i$, resulting in a control update $\Delta\underline{\alpha}^i$. The total update to the control parameters is $\Delta\underline{\alpha} = \sum_{i=0}^{L} \Delta\underline{\alpha}^i$.

Table 3.1: **Timing Analysis of the Uniformization and Stabilization Module.** Time taken for each step in the closed-loop feedback cycle for a 45×29 trap array. Measurements are averaged over 5000 feedback cycles. The total cycle duration is limited mainly by the CMOS camera readout time, which could be improved by using cameras with faster readout.

Step	**Duration (ms)**
CMOS imaging and readout	44.3
Image analysis	11.9
Upading control parameters	2.2
Waveform synthesis and streaming	6.1
Total	**64.6**

3.1.5 Quantifying Trap Stability and Uniformity

In this section, we characterize the performance of the stabilization and uniformization system described using a 21×21 trap array as an example. The trap depth non-uniformity is quantified by the ratio σ/μ, where σ is the standard deviation and μ is the mean of the trap depths taken over all traps. Prior to activating the feedback system, when identical control parameters are supplied to all RF tones, the non-uniformity is approximately 17%. After applying the control protocol (see Sec. 3.1.4), the system reduces the non-uniformity to 1.5% (see Fig. 3.7a and Fig. 3.7b), and this level of uniformity is maintained over extended timescales (see Fig. 3.7c).

This indicates that the mean trap power across tweezers remains close to the target value over time. To assess whether the feedback loop also suppresses temporal fluctuations, we examine the time-averaged relative fluctuations, computed as the standard deviation of each trap depth over time divided by its mean. After tuning the PID parameters, the system achieves a substantial reduction in temporal fluctuations compared to the case with no feedback control (see Fig. 3.7d).

Lastly, we report a timing analysis of the uniformization and stabilization module. As shown in Table 3.1, the feedback protocol significantly improves both spatial uniformity and temporal stability of the trap depths. The performance is currently limited by the frame rate of the camera used to acquire measurements. Further improvements are possible using imaging systems with faster readout times.

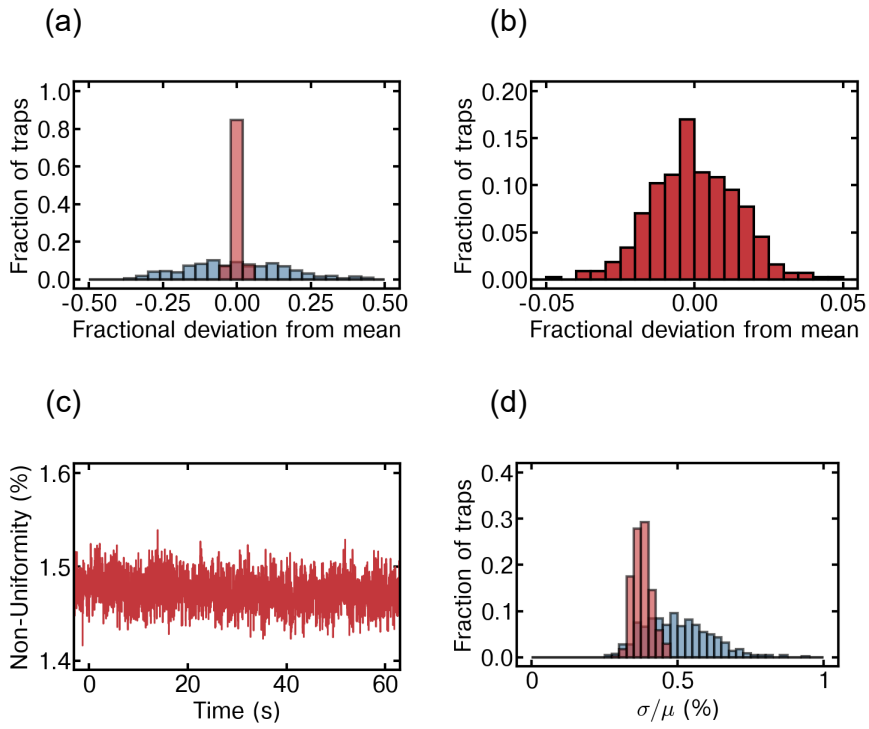


Figure 3.7: **Performance of the Stabilization and Uniformization System.** (a) Histogram of the relative deviation of time-averaged trap depths from the target trap depth for a 21×21 array. The red distribution (with feedback) is significantly narrower and centered around zero compared to the blue distribution (without feedback), indicating substantial improvement in time-averaged uniformity. (b) A high-resolution view of the red distribution in (a), showing its sharp peak around zero. (c) Evolution of the non-uniformity metric (σ/μ) over a 60 s period, showing long-term stability of the uniformization system (c.f. Fig. 3.1b). (d) Histogram of the relative temporal fluctuation σ/μ, where σ is the standard deviation of the trap depth over time and μ is the mean. Data is measured over 4 s for each trap. The red distribution (with feedback) is more compressed and closer to zero compared to the blue distribution (without feedback), indicating significant suppression of temporal fluctuations by the feedback system.

3.1.6 Time Dependence of the Jacobian Matrix

In this section, we propose a method for the characterization the time dependence of the Jacobian matrix to estimate the frequency with which recalibration is required, and use this method for our trapping system (see Fig. 3.8). The measurements are performed on an 11×11 trap array, corresponding to 22 control parameters: 11 for $\underline{\alpha}^y$ and 11 for $\underline{\alpha}^x$. Each control parameter is perturbed by approximately 8% in both directions. For each perturbation, trap powers are measured at 21 equidistant alpha values. By sweeping individual control parameters and observing the resulting changes in trap powers, we extract the slope of the dependence of each trap's power on the corresponding alpha. From these slopes, we reconstruct the Jacobian matrix, as illustrated in Fig. 3.8(a).

This measurement is repeated over a period of 12000 s with data collected at intervals of approximately 23 s. To quantify the variation of the Jacobian over time, we compute the Frobenius norm of its deviation from the mean, which is defined by

$$\text{variation from the mean}(t) = \|\underline{\underline{J}}(t) - \langle \underline{\underline{J}} \rangle\|_{\mathrm{F}}, \tag{3.19}$$

where the Frobenius norm is defined as

$$\|\underline{\underline{J}}\|_{\mathrm{F}} = \sqrt{\sum_{i,j} |J_{ij}|^2}. \tag{3.20}$$

The relative variation reaches a minimum after approximately 1500 s and remains stable at 30×10^{-3}, with periodic peaks every 1–2 minutes (see Fig. 3.8(b)). A power spectral density analysis of the variations (Fig. 3.8(c)) shows that significant frequency components are below 10 mHz, consistent with the periodic peaks observed.

Finally, we calculate the Allan variance for each element of the Jacobian matrix as a function of the bin size τ, constructing the Allan variance vector. The norm of this vector decreases with increasing τ, reaches a minimum at approximately 3132 s, and then increases again (Fig. 3.8(d)). This behavior indicates slow drift in the Jacobian and suggests a recalibration period of roughly 3132 s, or 52 min.

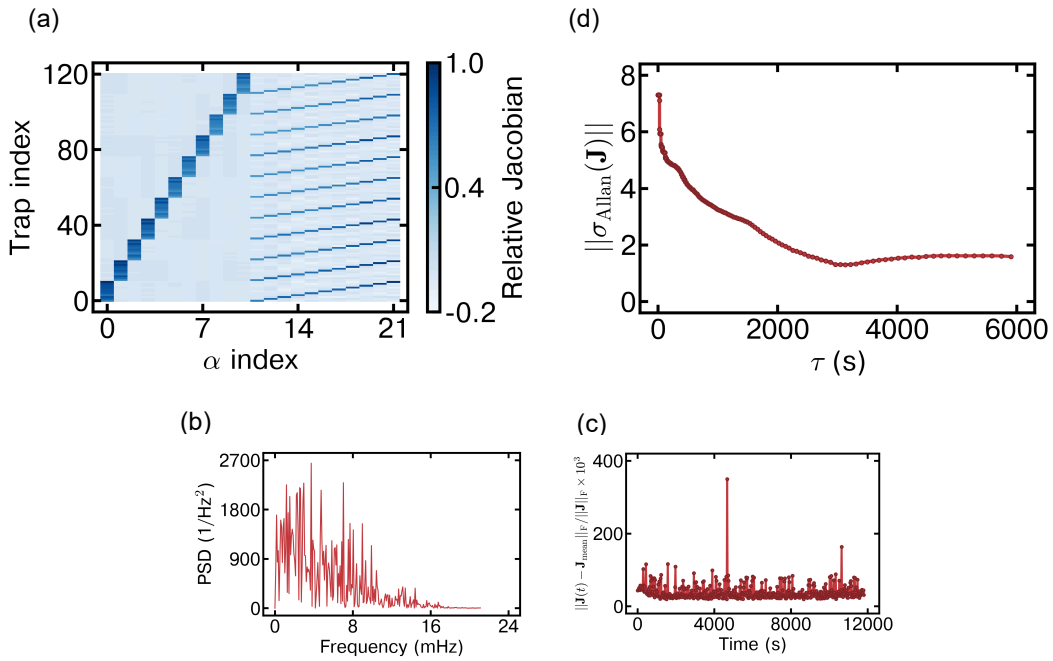


Figure 3.8: **Time Dependence of the Jacobian Matrix.** (a) Pictorial representation of the measured relative Jacobian matrix, $\underline{\underline{J}}_{\text{relative}} = \underline{\underline{J}}/\|\underline{\underline{J}}\|_\infty$, over ≈ 12000 s for an 11×11 trap array (121 traps). Traps are indexed in row-major order. The first 11 control parameters correspond to $\underline{\alpha}^y$ and predominantly affect their corresponding rows (dark blue blocks in the first 11 columns), while the second 11 control parameters correspond to $\underline{\alpha}^x$ and predominantly affect their corresponding columns. Notably, for the second set of parameters, there are non-zero correlations with trap powers outside the corresponding columns due to intermodulations between frequency tones applied to the AODs. (b) Relative variations of the Jacobian matrix as a function of time over a 12000 s period, with data points spaced by 23 s. The relative variation stabilizes after approximately 1500 s and remains roughly constant around 30×10^{-3}, with periodic peaks every 1–2 minutes. (c) Power spectral density (PSD) of the variations. Major frequency components are below 10 mHz, corresponding to periods of ~ 100 s, likely related to the periodic peaks observed in (b). (d) Norm of the Allan variance vector of the Jacobian as a function of the time bin size τ. The norm reaches a minimum at ≈ 3132 s, indicating drift in the Jacobian and suggesting a recalibration period of approximately 3132 s, i.e., 52 min.

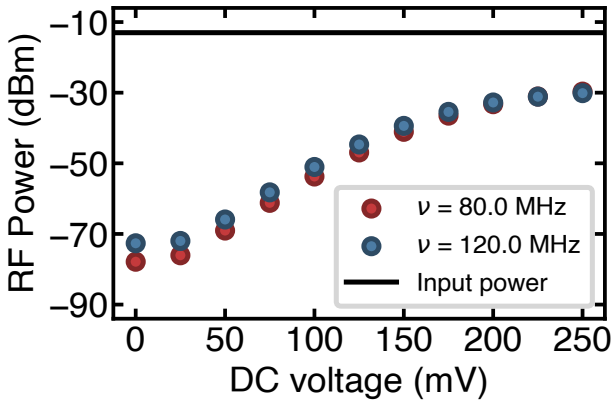


Figure 3.9: **Non-Uniformity in Output RF Power from a Mixer at Different Frequencies.** Output RF power of a double-balanced mixer (Mini-Circuits, ZFM-2+) as a function of DC voltage applied to the I input, with a fixed RF input power of approximately −13 dBm supplied to the L input. The output is shown for two RF frequencies: 80 MHz (red) and 120 MHz (blue). As the DC voltage approaches 0, a significant frequency-dependent discrepancy in output power emerges, with a maximum difference of up to 5.2 dBm. This non-uniformity implies that using a single mixer for global control of multiple traps can introduce large variations in trap depth during ramp-down.

3.2 Uniform Adiabatic Control of Trap Depths

A common requirement in applications such as magnetometry [18], electromagnetic field imaging [1, 68], and atomic thermometry [69] is the ability to adiabatically reduce the trap depth from typical values of about 1 mK to depths on the order of 1 μK or less. Here, adiabatic refers to a process in which the motional quantum state of the trapped atoms remains unchanged as the trap depth evolves. Operating in shallow traps suppresses temperature-induced broadening of the Ramsey fringe spectrum, enabling narrower resonance features and improved frequency resolution [66].

A standard method for reducing trap depth is to modulate the total laser power before the AODs, for example using an acousto-optic modulator (AOM) [70]. However, AOMs have limited diffraction efficiency, so this approach reduces the total optical power delivered to the traps. To avoid this limitation, we instead modulate the RF signals driving the AODs themselves, allowing the same devices responsible for multiplexing the trapping beam to also perform adiabatic ramps.

One method for rapid RF power modulation is to pass the waveform through a double-

balanced mixer, which adjusts the amplitude of the polychromatic RF signal. However, mixers exhibit strongly frequency-dependent transmission characteristics (see Fig. 3.9). As a result, this approach leads to non-uniform scaling of the RF tones, and therefore the trap depths after and during the ramp. This non-uniformity results in spatial inhomogeneities in observables such as atomic temperature, atomic survival probability after the ramp, and AC Stark shifts across the array. In particular, variations in the AC Stark shifts across different traps mean that a global microwave pulse will rotate each qubit differently, significantly reducing the fidelity of global single-qubit gates. This imposes a severe limitation on protocols that rely on coherent, homogeneous control, such as those used in electrometry, magnetometry, and atomic clock operations.

To address this issue, we extend the stabilization and uniformization system described in Sec. 3.1 by integrating it with a double-balanced mixer for global amplitude control. The mixer enables high-bandwidth analog regulation of the total RF power, while digital feedback compensates for frequency-dependent imbalances and slow drift. This hybrid scheme ensures both global power stabilization and spatial uniformity during the ramp. The resulting architecture is well suited for protocols requiring precise, homogeneous control over trap depth throughout the array.

We first describe the control architecture used to implement uniform adiabatic ramps, including waveform modulation and mixer-based amplitude control. We then demonstrate how this system enables thermometry measurements by extracting atomic temperatures from survival probabilities after the ramp.

3.2.1 Adiabatic Trap Depth Control Using Waveform Modulation and Mixer-Based Power Stabilization

In this section, we describe how adiabatic trap depth ramps are implemented using a combination of waveform modulation and mixer-based power control. The goal is to reduce the trap depth while maintaining uniformity across the array. This is achieved by simultaneously adjusting the control waveforms streamed by the AWG and modulating the overall RF power delivered to the AODs using an analog PID controller. The RF waveform sent to the second AOD passes through a double-balanced mixer (Mini-circuits, ZFM-2+), with its in-phase (I) input connected to the analog output of the PID controller. This setup enables global amplitude control over the RF tones, while the waveform provides spatially resolved control over individual traps (see Fig. 3.10).

This hybrid architecture combines high-bandwidth global power stabilization, handled by the analog PID controller (Moku:Lab), with lower-bandwidth digital feedback for trap

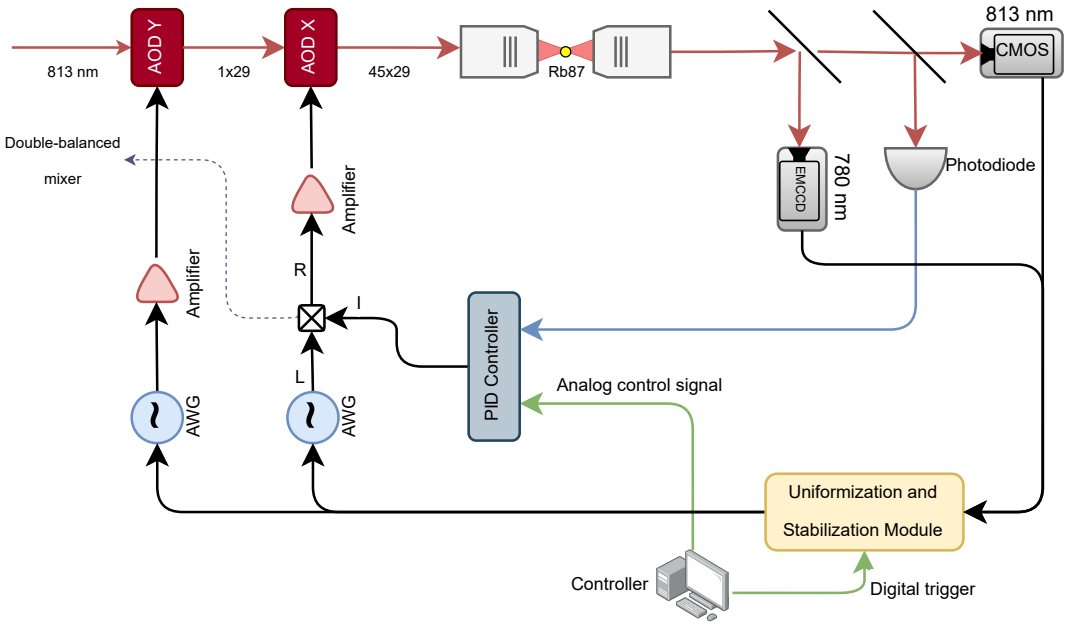


Figure 3.10: **Hardware Setup for Adiabatic Ramps with Uniformized Trap Depths.** An updated hardware setup extending the uniformization and stabilization system previously introduced in Fig. 3.2 to enable adiabatic ramping while maintaining uniform trap depths. In addition to the CMOS and EMCCD cameras used to measure trapping power and trap depth, respectively, a photodiode is added to monitor the total trapping power. A double-balanced mixer is inserted between the second AWG channel and AOD X to allow analog control over the global RF power via its I input. The uniformization and stabilization module continues to update the AWG waveforms using camera measurements, but it leaves the global power mode unchanged, i.e. $\underline{\tilde{C}}^0$. Instead, the photodiode signal is sent to an analog PID controller (Moku:Lab), which compares it to an external analog input that defines the desired adiabatic power ramp. The controller outputs an analog signal that adjusts the mixer voltage through the analog PID controller, scaling the RF drive to achieve the desired global power profile. Simultaneously, a digital trigger initiates a coordinated sweep of the control parameters $\underline{\alpha}$ within the uniformization module, ensuring trap depth uniformity is preserved throughout the ramp.

depth uniformity. The digital system, described in Sec. 3.1, operates through a set of PID loops targeting different temporal modes of the error signal and is limited by the CMOS camera frame rate (see Sec. 3.1.5). To avoid redundancy, the digital PID loop corresponding to the global temporal mode is disabled (see Table 3.1), so that only relative fluctuations are corrected digitally.

During steady operation, the digital feedback loop compensates for non-uniformities, while the analog PID controller stabilizes the total laser power using photodiode measurements. To initiate an adiabatic ramp, a digital trigger is sent to the AWG to switch to transition to a new set of control parameters $\underline{\alpha}$, which are smoothly varied using a pre-calibrated lookup table. At the same time, the analog control signal to the PID controller is ramped down. This results in a coordinated ramp, in which the analog loop controls and stabilizes the global power, while the digital loop maintains uniformity across traps.

We demonstrate the improvement in trap depth homogeneity by measuring the atomic survival probability after the ramp (see Fig. 3.11). If trap depth homogeneity is preserved during the ramp, and assuming the atoms have approximately uniform initial temperatures in the tweezers, the post-ramp survival probability should also be more homogeneous. We compare our method against a control case where the ramp is performed using only a mixer without waveform compensation. In the control case, spatial non-uniformities increase during the ramp. In contrast, our method significantly reduces survival probability non-uniformity, defined as

$$\text{Non-Uniformity} = \frac{\text{Standard deviation of survival probabilities}}{\text{Mean survival probability}}, \tag{3.21}$$

from 1.0 to 0.5 across a 45×29 trap array, which is a significant improvement. It is observed that the column dependence in the non-uniform case is no longer visible when using the uniformization and stabilization module in the post-ramp survival probability (see Fig. 3.11).

The strong column dependence when using the mixer alone, arises from the frequency-dependent properties of the double-balanced mixer. In contrast, this column dependence is absent when the control parameters $\underline{\alpha}$ are updated using the uniformization and stabilization module, which compensate for the frequency dependence of the mixer.

Ultimately, the homogeneity of the post-ramp survival probability is limited by two factors: inhomogeneity in the initial atomic temperature across the optical tweezers and residual non-uniformity in the trap depths during the ramp. The latter is a consequence of the lower dimensionality of the control space relative to the target space in the two-AOD system (see Fig. 3.4c). Using trap-depth spectroscopy, we confirm that the observed

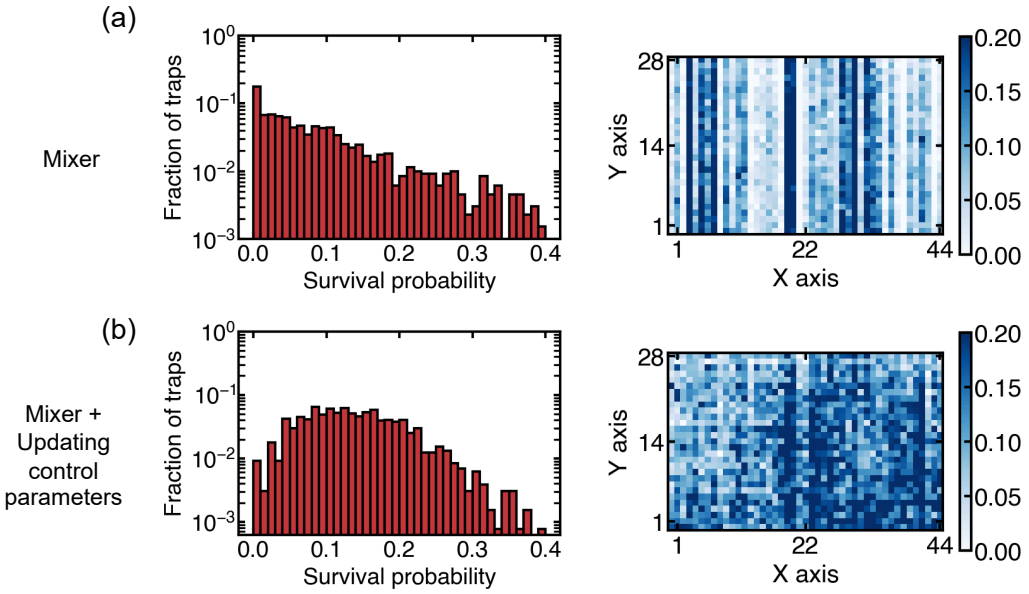


Figure 3.11: **Characterizing Post-Ramp Survival Probability with and without Uniformization of Trap Depths.** Comparison of atom survival probability distributions after a trap depth ramp in a 45×29 array. (a) When using only the mixer to perform the ramp, a strong column dependence is observed in the survival probability, with a high non-uniformity of $\sigma/\mu = 1.0$. (b) When using the mixer together with the uniformization and stabilization module to update the control parameters $\underline{\alpha}$ during the ramp, the survival probability becomes more homogeneous, with non-uniformity reduced to 0.5 and no significant column dependence. Residual spatial variation is attributed to initial non-uniformity in trap depths prior to the ramp, a consequence of the reduced control space dimensionality in a two-AOD system (see Fig. 3.4c).

inhomogeneity in survival probability is strongly correlated with the trap depth distribution prior to the ramp.

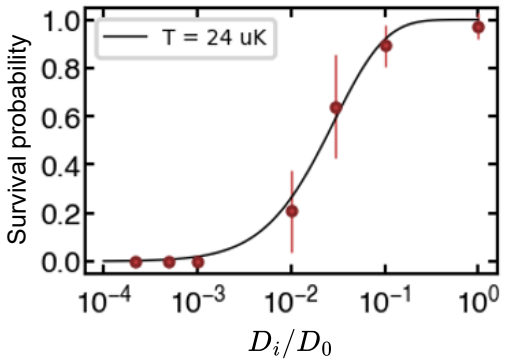


Figure 3.12: **Thermometry of Single Atoms in Optical Tweezer Arrays.** Atoms are loaded into traps with initial depth U_0 and adiabatically ramped down to a final depth U_i. The survival probability of atoms is measured as a function of U_i and fitted using the model provided in Eq. (3.25). The fit yields an estimated atomic temperature of 24 μK.

3.2.2 Thermometry with Adiabatic Ramps

In this section, we use the adiabatic ramping system developed in Sec. 3.2.1 to perform atomic thermometry. The technique relies on measuring the survival probability of atoms after an adiabatic ramp as a function of the initial and final trap depths. By varying these parameters and recording the corresponding survival probability, we obtain a characteristic curve that depends primarily on the ratio of the final to initial trap depth.

This survival curve can be fitted using a model that accounts for the thermal energy distribution of atoms in a finite-depth potential [69]. From the fit, we extract the atomic temperature in the trap prior to the ramp. Using this method, we estimate a temperature of $T = 24$ μK (see Fig. 3.12).

We now describe the equation used for the fit to extract the temperature (see Fig. 3.12). Given an atomic temperature T, the energy density of the atoms in the trap is given by the Maxwell-Boltzmann distribution,

$$P(E) = \frac{1}{2k_B^3 T^3} E^2 \exp\left(-\frac{E}{k_B T}\right), \tag{3.22}$$

where k_B is the Boltzmann constant. When the trap depth is adiabatically ramped down, the energy of the atoms scales with the square root of the trap depth. This follows from the fact that the energy of the nth motional state is given by $\hbar\omega(n + \frac{1}{2})$, where ω is the trapping frequency. Since $\omega \propto \sqrt{D}$, where D is the trap depth, it follows that the energy scales with $\sqrt{D}$.

Given initial and final trap depths D_0 and D_i, the final energy of an atom initially at energy E_0 is

$$E_i = E_0 \cdot \sqrt{\frac{D_i}{D_0}}. \tag{3.23}$$

Atoms are lost from the trap if $E_i > D_i$, or equivalently if $E_0 > \sqrt{D_0 D_i}$. The fraction of lost atoms is then

$$P_{\text{lost}}(T) = \int_{\sqrt{D_0 D_i}}^{\infty} \frac{1}{2k_B^3 T^3} E_0^2 \exp\left(-\frac{E_0}{k_B T}\right) \cdot dE_0, \tag{3.24}$$

and the survival probability is

$$P_{\text{survive}}(T) = 1 - P_{\text{lost}}(T). \tag{3.25}$$

By fitting the measured survival probabilities as a function of D_0 and D_i to this model, we extract the temperature T.

Chapter 4

Magnetic Field Sensing with Arrays of Neutral Atom Atoms

Neutral atom array quantum processors have emerged as a promising platform for both digital quantum computation and quantum sensing (see Sec.1.1). The intrinsic sensitivity of neutral atoms such as ^{87}Rb to external electromagnetic fields has enabled applications in electromagnetic field imaging and magnetometry [1, 18, 68].

For example, Rydberg atom arrays have been used to measure DC magnetic fields with high spatial resolution [18]. Moreover, trapped ion based platforms have been used to probe near-infrared electromagnetic fields using quantum lock-in detection protocols [1]. These applications rely critically on the ability to perform high-fidelity single-qubit gates, either globally across the array or locally at selected sites.

These quantum sensing applications have broad relevance. Magnetic field sensing with neutral atoms enables spatially resolved magnetometry and noise spectroscopy, providing a means to characterize magnetic field fluctuations in the environment. This is crucial for diagnosing decoherence sources, as such fluctuations can shift atomic transition frequencies via the Zeeman effect. Near-infrared electromagnetic field imaging allows direct visualization of optical fields, such as the 1013 nm Rydberg excitation beam used in qubit-qubit gates. This capability is especially valuable for stabilizing the spatial alignment of the beam with the atom array over long experimental runs, ensuring that the atoms experience a consistent Rabi frequency across repeated shots and thereby improving the fidelity and reproducibility of multi-qubit gates.

A critical requirement for quantum sensing protocols such as Rabi and Ramsey spectroscopy and quantum quantum lock-in detection is high-fidelity single-qubit gates. Im-

plementing both global and local gate operations is therefore a necessary step towards realizing practical neutral-atom-based quantum sensors.

Global single-qubit gates can be implemented using microwave radiation that drives the hyperfine transition in ^{87}Rb between the $5S_{1/2}, F = 1$ and $5S_{1/2}, F = 2$ levels, simultaneously addressing all atoms in the array. Alternatively, site-resolved single-qubit control can be achieved via two-photon Raman transitions mediated by the D_1 line, enabling selective single-qubit rotations on individual atoms [8].

In addition, performing single-qubit gates using the optical D_1 line enables an alternative form of dynamic decoupling in quantum sensing, known as *quantum double lock-in amplification* [68], which enhances the sensitivity of the quantum sensor by a factor of two.

In this chapter, we begin by specifying what we mean by a *qubit* in the context of neutral atom arrays. We encode qubits in the hyperfine ground states of ^{87}Rb atoms by restricting the internal state of each atom to a two-dimensional Hilbert space spanned by the states $|0\rangle \equiv |F = 1, m_F = -1\rangle$ and $|1\rangle \equiv |F = 1, m_F = -1\rangle$. We then demonstrate the ability to coherently manipulate population and phase within this subspace, thereby realizing well-defined qubits suitable for quantum information processing and quantum sensing.

We present two complementary approaches for implementing high-fidelity single-qubit control. First, we demonstrate global single-qubit gates using microwave radiation and their application in quantum sensing protocols such as Rabi and Ramsey spectroscopy (see Sec. 4.1), enabling magnetometry and electromagnetic imaging with neutral atom arrays. Next, we describe progress towards local single-qubit control using Raman transitions via the D_1 line. In particular, we show how the addressing beam is multiplexed using a pair of orthogonal AODs and aligned in three dimensions with the optical traps (see Sec. 4.2).

4.1 Magnetic Field Sensing using Microwave Spectroscopy

In this section, we describe the steps taken towardsquantum sensing using single ^{87}Rb atoms in arrays of optical traps. We outline the procedures for initializing and measuring atoms in a qubit state (see Sec. 4.1.1), performing microwave spectroscopy (see Sec. 4.1.2), and applying these techniques to quantum sensing tasks such as magnetometry and electromagnetic field imaging (see Sec. 4.1.3). We demonstrate the ability to prepare atoms in the qubit state and measure them with high fidelity: atoms are optically pumped into

the qubit state $|1\rangle$, and during measurement, atoms in the $|1\rangle$ state are ejected from the traps via a state-selective pushout beam. We also perform single-qubit gates with a π-pulse fidelity of 98.8%. Finally, we describe how these capabilities enable applications such as magnetometry and near-infrared electromagnetic field imaging.

Throughout this section, we define the qubit states using hyperfine sublevels of the $5s_{1/2}$ ground state (see Fig. 1.3b):

$$|0\rangle = |F = 1, m_F = -1\rangle, \tag{4.1}$$

$$|1\rangle = |F = 2, m_F = -2\rangle. \tag{4.2}$$

These hyperfine states were chosen to maximize sensitivity to external magnetic fields via the Zeeman effect, which is essential for magnetometry. This is because the differential shift of the transition frequency between two hyperfine states $|F = 1, m_F = m_{F1}\rangle$ and $|F = 2, m_F = m_{F2}\rangle$ approximately scales with $|m_{F1} + m_{F2}|$. Selecting $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = -2\rangle$ therefore yields the largest magnetic sensitivity possibles. Additionally, the state preparation process for these levels is relatively simple and can be implemented with high fidelity.

4.1.1 State Preparation and Measurement

Reliable state preparation and high-fidelity readout are essential prerequisites for quantum sensing. In our setup, ^{87}Rb atoms are trapped in optical tweezers and initialized in the ground hyperfine state $|F = 2, m_F = -2\rangle$ using optical pumping, which we define as the qubit state $|1\rangle$. To perform qubit state readout, atoms in the $|1\rangle$ state are ejected from the traps using a state-selective pushout beam. The probability of an atom being ejected from the trap is proportional to the population of its wavefunction in the $|1\rangle$ state.

State Initialization. State preparation is performed in two stages: coarse pumping followed by fine pumping. The goal is to initialize all atoms in the qubit state $|1\rangle = |F = 2, m_F = -1\rangle$.

In the coarse pumping stage, atoms are first transferred from the $F = 1$ to the $F = 2$ hyperfine manifold. This is achieved by illuminating the atoms with σ^+ polarized light resonant with the $F = 1 \rightarrow F' = 2$ transition of the D_2 line. The light is applied from three different directions to ensure uniform coupling regardless of the initial state of the atom or its quantization axis (see Fig. 1.3d for the beam geometry). This process efficiently depopulates the $F = 1$ manifold, but does not resolve the magnetic sublevels of $F = 2$.

In the subsequent fine pumping stage, we apply an additional σ^- polarized beam resonant with the $F = 2 \rightarrow F' = 3$ transition. This drives atoms towardslower m_F states, ultimately accumulating population in $|F = 2, m_F = -2\rangle$. Simultaneous illumination on the $F = 1 \rightarrow F' = 2$ transition is maintained during this stage to repump any population that falls back into the $F = 1$ manifold back to $F = 2$.

At the end of this process, atoms are initialized in the $|1\rangle$ qubit state. The fidelity of the state initialization process theoretically approaches unity as a function of the time spent in the fine pumping stage.

State-Selective Pushout. As discussed in Sec. 5, atoms in optical traps can be imaged using polarization gradient cooling (PGC) beams. To perform qubit state readout, we selectively remove atoms in the $F = 2$ hyperfine manifold using a resonant pushout beam. After this state-selective ejection, the presence or absence of an atom in each trap is detected via fluorescence imaging. Atoms that remain trapped are inferred to have been in the $F = 1$ manifold, while those that are lost were in the $F = 2$ state.

To eject atoms in the $F = 2$ manifold, we illuminate them with resonant light on the $F = 2 \rightarrow F' = 3$ transition of the D_2 line. Atoms absorb photons from this beam and gain momentum. After sufficient scattering events, the atoms are heated out of the trap. Atoms that remain after the pushout are inferred to be in the $|0\rangle$ state. The fidelity of the pushout is theoretically limited by the probability of off-resonant excitation to the $F' = 2$ level, from which atoms can decay to the $F = 1$ manifold. If this occurs, the atom no longer interacts with the pushout beam and remains trapped, resulting in a state measurement error.

Characterization. The overall state preparation and measurement (SPAM) error is quantified by preparing atoms in the $|1\rangle$ state and comparing the measured outcomes to the expected result. We obtain a total SPAM error of 5%, which is significantly larger than the contributions from atom loss or classification errors in the imaging process (see Sec. 5). We attribute this residual error to the finite fidelity of the pushout process. Increasing the duration of both fine optical pumping and the pushout does not lead to further reduction in the SPAM error, suggesting that atoms in the $F = 2$ manifold decay to the $F = 1$ manifold before acquiring enough energy to be ejected from the trap.

In principle, this measurement error could be further reduced by adopting alternative pushout schemes. One approach is to transfer atoms in the $F = 2$ manifold to long-lived Rydberg states, which experience an anti-trapping potential in optical tweezers and

naturally escape from the traps without requiring photon scattering [4]. Another possibility is to optimize the state-selective pushout protocol itself, for example, by minimizing the number of required scattering events. This would reduce the likelihood of atoms decaying into the dark $F = 1$ manifold before being successfully ejected.

Another way of directly observing the state preparation and measurement fidelity is through microwave spectroscopy. In this protocol, we initialize the atoms in a known hyperfine state, apply a microwave π-pulse at a specific frequency, and measure the survival probability using state-selective pushout. The π-pulse drives transitions between sublevels of the $F = 1$ and $F = 2$ hyperfine manifolds. By scanning the frequency of the microwave pulse and recording the resulting survival probability, we can infer the population distribution across magnetic sublevels. This technique allows us to directly verify the effectiveness of both coarse and fine optical pumping schemes. Figure 4.1 shows representative spectra: in the case of coarse pumping, multiple peaks are visible, consistent with partial population in several magnetic sublevels. With fine pumping, however, only a single transition remains, confirming that the atoms are predominantly prepared in the desired $|F = 2, m_F = -2\rangle$ state.

4.1.2 Microwave Spectroscopy in Shallow Traps

In this section, we characterize coherent control over hyperfine qubit states in ^{87}Rb atoms trapped in optical tweezers. Accurate manipulation of these states is essential for quantum sensing applications such as magnetometry, where the transition frequency encodes external field fluctuations. We begin by describing the microwave setup used to drive qubit transitions, followed by Rabi spectroscopy to identify the relevant spectral features and calibrate gate durations. Finally, we perform Ramsey spectroscopy to assess the coherence properties of the system and identify the dominant sources of decoherence.

The atoms are addressed using a microwave horn positioned directly below the vacuum chamber (see Fig. 4.1). The horn emits microwave radiation with both σ^+ and σ^- polarization components relative to the quantization axis defined by a bias magnetic field. The microwave field drives transitions between the $F = 1$ and $F = 2$ hyperfine manifolds of the $5S_{1/2}$ ground state, with a measured Rabi frequency of 42 kHz on the qubit transition. The microwave signal is generated by mixing a local oscillator at 6.83 GHz with two radio-frequency (RF) tones using an IQ mixer, enabling full control over the amplitude, frequency, and phase of the output.

To identify the relevant transition and calibrate the π-pulse duration, we first perform Rabi spectroscopy. A sequence of pulses with varying durations is applied to atoms pre-

pared in the $|1\rangle$ state, and the resulting population dynamics are used to extract the Rabi frequency. Once the π-pulse duration is determined, we fix the pulse length and sweep the carrier frequency to obtain the hyperfine spectrum. When only coarse optical pumping is applied before the pulse (see Sec. 4.1.1), the spectrum exhibits multiple peaks, each corresponding to transitions between different Zeeman sublevels of the $F = 1$ and $F = 2$ manifolds. After including the fine pumping stage, all but a single peak disappear. The remaining resonance corresponds to the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = -2\rangle$ transition that defines our qubit see Fig. 4.1.

To reduce spectral broadening caused by thermal motion and motional-state-dependent AC Stark shifts, we apply the adiabatic ramp protocol described in Sec. 3.2 prior to the spectroscopy sequence. This ramp lowers the trap depth, thereby reducing the atomic temperature. In deeper traps, atoms in different motional states experience varying AC Stark shifts, leading to energy level variations on the order of $\Delta\alpha \cdot \hbar\omega$, where $\Delta\alpha$ is the differential polarizability between the $F = 1$ and $F = 2$ states, and ω is the trapping frequency. The total AC Stark shift measured in deep traps is approximately 50 kHz before the ramp, while in the ramped-down traps, it is reduced to about 500 Hz.

Finally, we perform Ramsey spectroscopy to extract the qubit coherence time T_2^*. In this protocol, two $\pi/2$ pulses are applied with a variable delay t between them, during which the relative phase of the second pulse is modulated as $\nu\tau$, with $\nu = 15$ MHz. The resulting Ramsey fringes are used to extract the coherence time from the decay envelope. We measure a coherence time of $T_2^* = 896$ μs, currently limited by temporal magnetic field fluctuations and spatial inhomogeneities arising from current noise in the coil drivers (see Fig. 4.3). In addition, fluctuations in the microwave pulse power and frequency contribute to decoherence. Future improvements to T_2^* are expected through stabilization of both the magnetic field environment and the microwave control system.

From this, the fidelity of a global π pulse applied to the entire array can be inferred to be 98.8% over a 45×29 trap array. This fidelity is limited by shot-to-shot fluctuations in the transition frequency and non-uniform responses of atoms across the array, including:

1. **Magnetic Field Inhomogeneities:** Spatial variations or gradients in the magnetic field cause position-dependent Zeeman shifts, resulting in different transition frequencies for atoms at different sites.

2. **Magnetic Field Fluctuations and Noise:** Temporal fluctuations in the magnetic field lead to shot-to-shot variations in the transition frequency, broadening the spectral line and reducing the coherence time T_2^*.

3. **Trap Depth Variations and AC Stark Shifts:** Differences in trap depths across the array cause inhomogeneous AC Stark shifts. Since the traps are ramped down prior to spectroscopy, we estimate these shifts to be below 1 kHz, making this effect comparatively minor.

4. **Microwave Field Non-uniformity and Fluctuations:** Imperfections in the spatial profile of the microwave pulse—caused by reflections and standing waves within the vacuum chamber, lead to spatial variations in the Rabi frequency experienced by the atoms [71]. Additionally, temporal fluctuations in the microwave pulse amplitude and frequency contribute to decoherence.

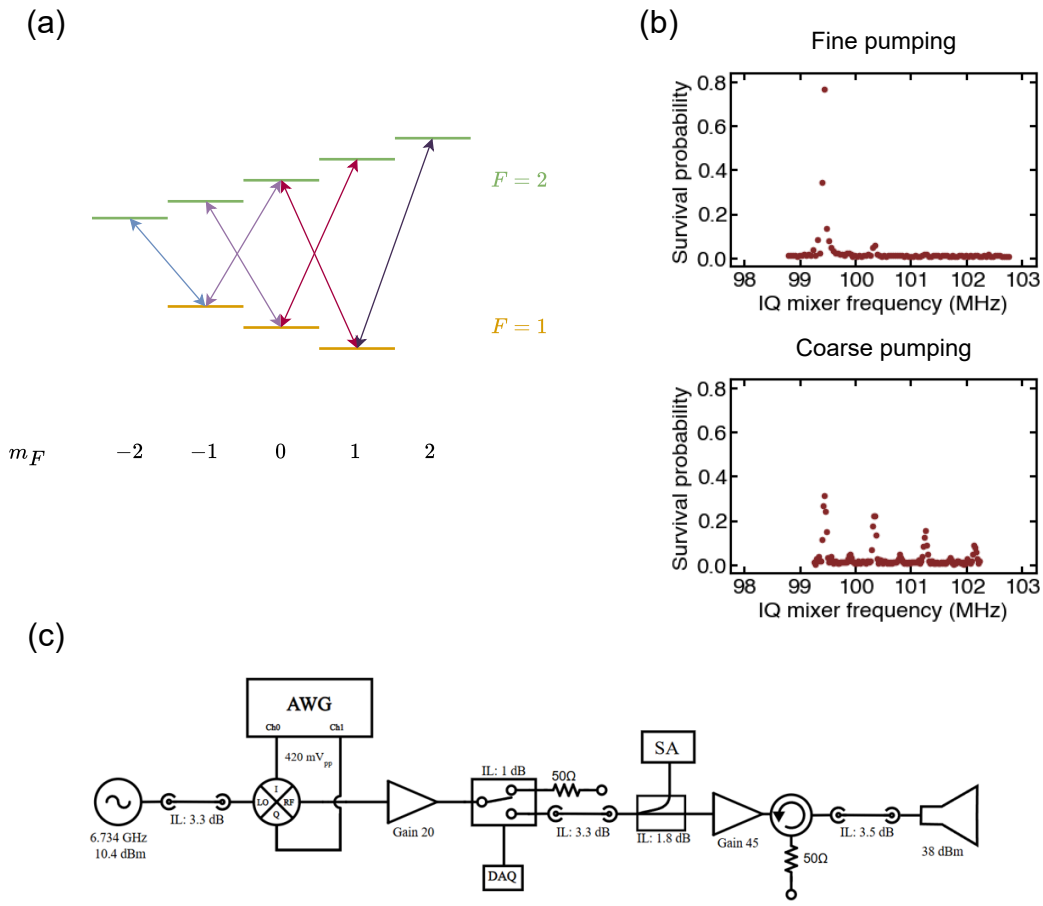


Figure 4.1: **Microwave addressing of hyperfine transitions.** (a) Microwave radiation from a horn antenna drives hyperfine transitions between the $F = 1$ and $F = 2$ manifolds of the $5s_{1/2}$ ground state. In an external magnetic field, $\sigma_\pm$ components couple to seven allowed transitions, grouped into four distinct frequencies (colored). (b) Microwave spectroscopy is performed by sweeping the IQ mixer frequency and applying a π-pulse. With coarse optical pumping, four peaks corresponding to the transitions are observed, while fine pumping prepares atoms in $|F = 2, m_F = -2\rangle$, leaving only a single peak. (c) Experimental setup: a local oscillator at ~ 6.734 GHz is IQ-modulated by an AWG, amplified, and delivered to the atoms with a horn antenna.

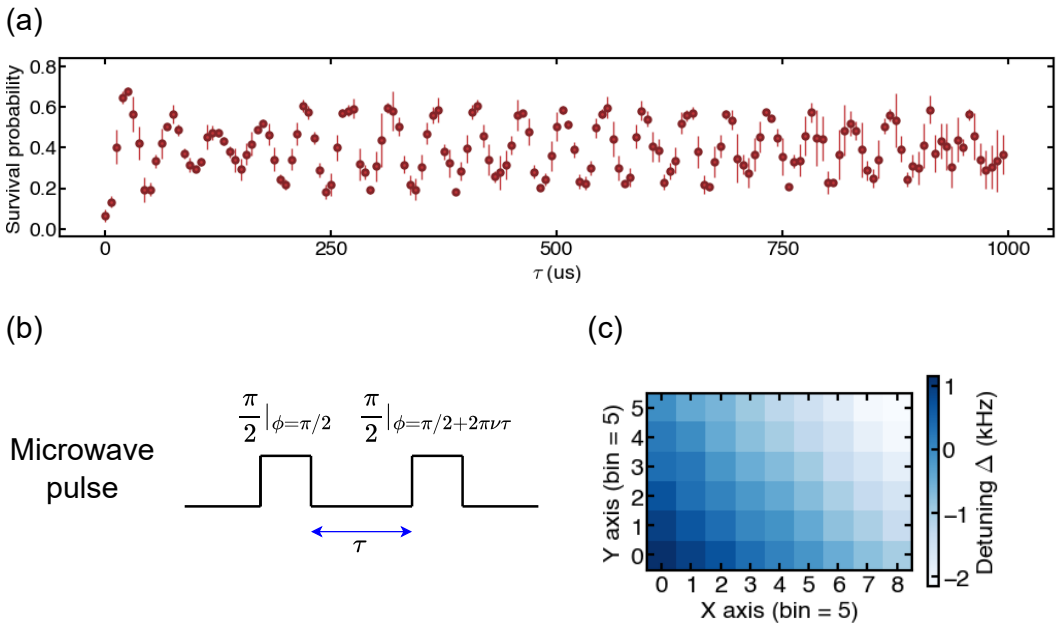


Figure 4.2: **Magnetometry using Ramsey spectroscopy.** (a) Ramsey signal measured across a 45×29 trap array. The observed oscillations correspond to a Rabi frequency of 21 kHz, and the fitted decay envelope yields a coherence time of $T_2^* = 896$ μs. (b) Pulse sequence: a $\pi/2$ pulse is applied along the y axis, i.e. $\phi = \pi/2$, followed by a variable delay τ, and a second $\pi/2$ pulse with phase $\phi = \pi/2 + \nu\tau$, where $\nu = 15$ kHz. (c) Spatially resolved magnetometry. Atoms are grouped into 5×5 bins (with the final row containing 4 traps). For each bin, the Ramsey signal is averaged and the oscillation frequency is extracted. In the absence of a magnetic field, the frequency should match the reference $\nu = 15$ kHz. Deviations from this value reveal spatial variations in the magnetic field introduced by external coils. The observed detuning pattern indicates the presence of a magnetic field gradient along both the x and y axes, which we have intentionally introduced using magnetic coils.

4.1.3 Future Steps – Quantum Sensing Applications

In this section, we describe how the protocols developed for state preparation and measurement (see Sec. 4.1.1) and microwave control (see Sec. 4.1.2) enable applications in quantum sensing. These include magnetometry using large arrays of neutral atoms [18] and imaging of near-infrared (NIR) electromagnetic fields via dynamically decoupled quantum protocols [1].

We begin by describing how our current system enables high-resolution magnetic field imaging using Ramsey spectroscopy. We then outline how dynamic decoupling techniques can extend this capability to sensing electromagnetic fields in the NIR regime.

Magnetometry with Ramsey Spectroscopy. Magnetometry is performed using Ramsey interferometry, as discussed in Sec. 4.1.2. The relative phase between the two $\pi/2$ pulses is modulated as a function of the free evolution time t, with the phase chosen to be $\nu\tau$, where ν is a known modulation frequency. The experiment is repeated for a range of values of t.

For each atom at site (i, j), the probability $p_{ij}(t)$ of being measured in the $|1\rangle$ state is recorded as a function of t. By taking the Fourier transform of $p_{ij}(t)$, we obtain the spectral response for each atom. In the absence of environmental magnetic fields, the central frequency of the spectrum for all atoms is ν. However, external magnetic fields along the quantization axis introduce site-dependent Zeeman shifts, causing deviations in the spectral peak locations. The shift in the resonance frequency corresponds directly to the local magnetic field at each site via the linear Zeeman effect.

Using this method, we can reconstruct the spatial variation of the magnetic field across the trap array. We demonstrate this capability by detecting a magnetic field gradient intentionally applied using external coils in the vacuum chamber (see Fig. 4.3c).

Electromagnetic Sensing using Quantum Lock-In Amplification. The Ramsey-based sensing method described above for magnetometry can also be adapted to measure the intensity profile of an oscillating electromagnetic field. Specifically, the local AC Stark shift induced by the field modifies the energy splitting between qubit states, resulting in a measurable shift in the spectrum of the survival probability $p_{ij}(t)$. This shift is proportional to the local intensity of the field and can thus be used to spatially map the intensity of the beam across the atomic array.

However, the measurements can be perturbed by noise sources such as magnetic field fluctuations, which can obscure the signal. To overcome this, we use *quantum lock-in amplification* [1], a dynamical decoupling technique that enhances sensitivity to a narrowband signal while rejecting broadband noise. This method has been previously demonstrated in trapped ion platforms and can be readily implemented in neutral atom array quantum processors. In this protocol, a train of equidistant π pulses is inserted between the initial and final $\pi/2$ pulses in the Ramsey sequence. These π pulses refocus dephasing from low-frequency noise sources while constructively interfering with signal components oscillating at half the frequency of the π-pulse repetition rate (see Fig. 4.3).

This modulation effectively filters out environmental noise, making the system selectively sensitive to oscillating electromagnetic fields at a frequency determined by the spacing between the inserted π pulses in the Ramsey sequence. If the electromagnetic field is modulated such that its power oscillates at this same frequency, e.g. using an AOM, then the atoms become sensitive to the field intensity at that frequency while rejecting noise at other frequencies. As a result, the spatial intensity profile of the field can be measured across the atomic array with significantly reduced contamination from spectrally separated noise sources.

We can use this technique to stabilize the spatial alignment of Rydberg excitation beams, such as those at 1013 nm. Precise alignment ensures consistent overlap between the trap array and the excitation beam, which is essential for maximizing the effective Rabi frequency of the Rydberg transition. This, in turn, improves the fidelity and repeatability of Rydberg gates, especially important for quantum computations performed over extended durations.

4.2 Colocated Optical Beams for Site-Selective Single-Qubit Control

In this section, we describe the progress made towards implementing site-selective single-qubit gates in our large-scale atom array quantum processor. Site-selective control over individual atoms is achieved by spatially multiplexing a Raman beam that drives a hyperfine transition using the D_1 atomic line (see Fig. 4.4a). This is accomplished by passing the Raman beam through a pair of orthogonal acousto-optic deflectors (AODs) to generate a 2D array of beams (see Fig. 4.4b).

A critical requirement for this approach is precise spatial overlap between the trap array and the D_1 beam array used to address individual atoms. This overlap is essential

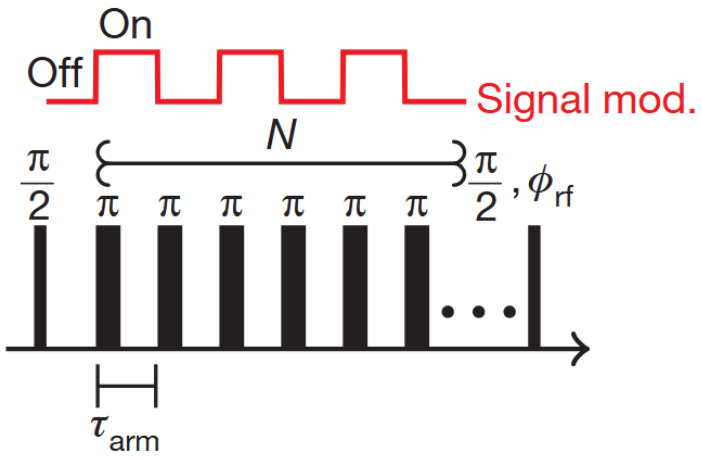


Figure 4.3: **Quantum lock-in amplification sequence for electromagnetic field imaging.** A Ramsey spectroscopy sequence incorporating a train of equidistant π pulses between two $\pi/2$ pulses. The signal to be measured, in our case a near-infrared (NIR) electromagnetic field, is modulated at a frequency equal to half the repetition rate of the π pulses. This modulation ensures that the signal constructively interferes with the Ramsey sequence, while broadband noise is dynamically decoupled and suppressed. The technique acts as a quantum lock-in amplifier, enabling sensitive detection of weak, narrowband signals in the presence of dominant noise. Adapted from Ref. [1], Fig. 2b.

for achieving site-resolved qubit control, which in turn enables local single-qubit gates for digital quantum computation, as well as advanced sensing protocols such as quantum double lock-in amplification [68].

Achieving this colocalization requires careful calibration and alignment due to several factors. First, the tweezer beams (813 nm) and the D_1 Raman beams (795 nm) have different wavelengths, which leads to chromatic shifts when passing through the same microscope objective. This causes a mismatch in their focal planes along the optical axis, necessitating axial adjustment to bring both beams into focus at the same plane. Second, the AODs generating the D_1 beam array must be rotated to match the angular orientation of the tweezer array; any misalignment results in systematic offsets across the array. Finally, the diffraction angle from an AOD depends on the wavelength of the incident beam, so the frequency-to-position mapping differs between the D_1 and 813 nm trapping beams. Even if the central beams are aligned and the AODs are correctly oriented, the mapping must still be recalibrated to ensure uniform overlap across the entire field of view.

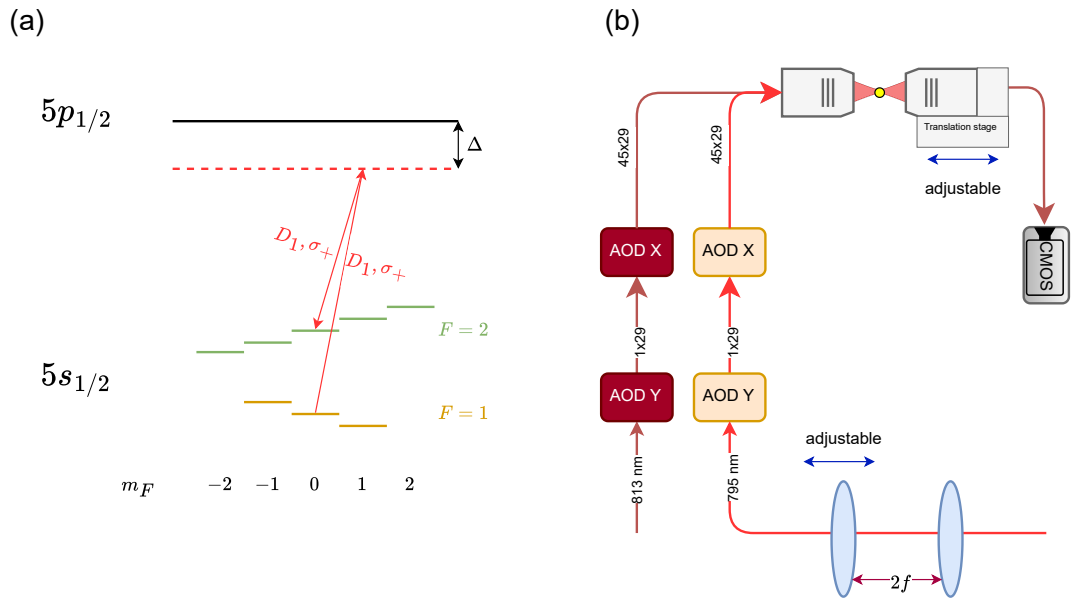


Figure 4.4: **Hardware setup for implementing local single-qubit gates using the Raman D_1 transition.** (a) Atomic transitions used for coherent control between $|F = 1, m_F = 0\rangle$ and $|F = 2, m_F = 1\rangle$ via a Raman D_1 beam composed of two σ_+ polarized frequency components addressing the $F = 1 \rightarrow F' = 2$ and $F = 2 \rightarrow F' = 2$ transitions, both red-detuned from resonance by Δ. (b) Experimental setup for generating the optical tweezers and multiplexed D_1 Raman beams. The traps and Raman beams are each created using pairs of orthogonal AODs. The D_1 beam passes through a $4f$ imaging system before the AODs; the lens before the first D_1 AOD is mounted on an adjustable axial stage to align the focal plane of the D_1 beams with the traps for colocalization. Both trap and D_1 beams are focused onto the atoms using a microscope objective. The output is collected by a second objective ($f = 4$ mm) mounted on a rotation mount, with the beam imaged by a CMOS camera. The camera and objective are mounted on a translation stage enabling 3D beam imaging.

Our objective is to generate a two-dimensional array of D_1 beams that is geometrically matched to and spatially colocalized with the optical tweezer trap array. To quantify the

quality of overlap, we define the following figure of merit:

$$\text{Figure of Merit} = \sum_{ij} \|\underline{r}_{ij}^{\text{traps}} - \underline{r}_{ij}^{D_1}\|^2, \tag{4.3}$$

where $\underline{r}_{ij}^{\text{traps}}$ and $\underline{r}_{ij}^{D_1}$ denote the (x, y) positions of the (i, j) site in the trap and D_1 beam arrays, respectively.

In the following sections, we first describe the procedure used to align the D_1 beams with the optical tweezers (see Sec. 4.2.1), and then characterize the alignment quality using the figure of merit (See Eq. (4.3)) and discuss the limiting factors (see Sec. 4.2.2).

4.2.1 Colocalization Procedure

The goal of the colocalization procedure is to ensure that each D_1 Raman beam spatially overlaps with its corresponding optical tweezer site, such that both are focused at the same location in three dimensions. This alignment is essential for achieving high-fidelity, site-resolved single-qubit control. This is because any misalignment between the tweezer and Raman beams could result in reduced Rabi frequencies, or undesired cross-talk between qubits.

To achieve alignment, we first coarsely align the central D_1 beam with the central optical trap. The central beam is defined as the one generated by the RF tone at the center of the AOD bandwidth. This step ensures approximate spatial overlap at the center of the field of view. We then proceed with the following optimization steps to fine-tune the alignment across the entire array:

1. **Angular Matching:** We first rotate the AODs that generate the D_1 beams so that the angular orientation of the resulting D_1 array matches that of the optical tweezer array.

2. **2D Lateral Alignment:** Both the trap and D_1 beam arrays are imaged using a front-facing CMOS camera (Basler, ace2 a2A2840-48umBAS). Two-dimensional alignment is then performed (neglecting axial displacement) by tuning the RF tones driving the D_1 AODs to match the beam positions to those of the traps in the camera plane.

3. **Axial (z-Axis) Colocalization:** To ensure the D_1 beams are focused in the same plane as the traps along the propagation (z) axis, we finely adjust the position of a

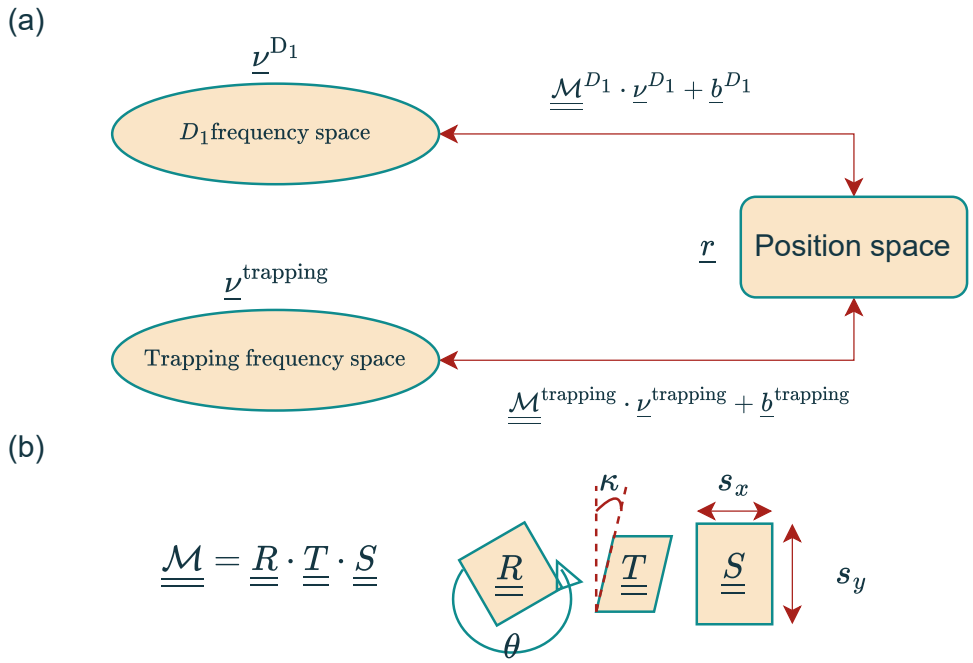


Figure 4.5: **Fine-tuning strategy for colocalizing trapping beams and D_1 Raman beams for local single-qubit gates.** (a) A pictorial representation of the linear model relating the RF tones supplied to the AODs and the corresponding beam positions. The linear maps for both trap and D_1 beams are defined in Eqs. (4.4)–(4.5). These mappings allow for transforming coordinates between frequency space and position space in both directions, enabling frequency-space matching of the D_1 beams to the traps. (b) A visual depiction of the RTS decomposition of the 2×2 mapping matrix $\underline{\underline{\mathcal{M}}}$ into a rotation matrix $\underline{\underline{R}}$, a skew matrix $\underline{\underline{T}}$, and a scaling matrix $\underline{\underline{S}}$, as described in Eq. (4.12). The diagram illustrates the geometric meaning of the decomposition parameters: rotation angle θ, skew factor κ, and axis-specific scalings s_x and s_y (see Eq. (4.14)).

lens placed before the first AOD in the D_1 beam path. This ensures the D_1 beams and traps are simultaneously focused on the atom plane, enabling high-fidelity site-resolved addressing.

The third step, axial colocalization, is the simplest. It involves measuring the point-

spread function (PSF) of the D_1 beams to determine the location of their focal plane. Assuming each beam is approximately Gaussian, the PSF, as well as the focal plane of the beams, can be extracted using the method described in Sec. 2.1. This allows for precise axial alignment of the D_1 beams with the focal plane of the trap array.

To perform the first two steps, angular matching and lateral alignment, we will make use of a linear model that maps the frequency tones supplied to the AODs to positions on the camera. In particular, for an $N_x \times N_y$ trap array, we define two linear maps $\underline{\underline{\mathcal{M}}}^{\text{traps}} \in \mathbb{R}^{2\times 2}$ and $\underline{\underline{\mathcal{M}}}^{D_1} \in \mathbb{R}^{2\times 2}$, and two offset vectors $\underline{b}^{\text{traps}} \in \mathbb{R}^{2\times 1}$ and $\underline{b}^{D_1} \in \mathbb{R}^{2\times 1}$, such that

$$\underline{r}^{\text{traps}}(\underline{\nu}) = \underline{\underline{\mathcal{M}}}^{\text{traps}} \cdot \underline{\nu} + \underline{b}^{\text{traps}}, \tag{4.4}$$

and

$$\underline{r}^{D_1}(\underline{\nu}) = \underline{\underline{\mathcal{M}}}^{D_1} \cdot \underline{\nu} + \underline{b}^{D_1}, \tag{4.5}$$

where $\underline{\nu} = (\nu_x, \nu_y)^T \in \mathbb{R}^{2\times 1}$ is the vector describing the RF tones supplied to the x and y AODs generating either the tweezer array or the D_1 beam array.

Here, $\underline{\underline{\mathcal{M}}}^{\text{traps}}$ and $\underline{b}^{\text{traps}}$ are fixed constants determined by the trapping AODs and optical setup, while $\underline{\underline{\mathcal{M}}}^{D_1}$ and $\underline{b}^{D_1}$ vary depending on the orientation of the D_1 AODs, which we adjust during the initial fine-tuning step.

First, we describe how to estimate the matrices $\underline{\underline{\mathcal{M}}}^{\text{traps}}, \underline{\underline{\mathcal{M}}}^{D_1}$ and vectors $\underline{b}^{\text{traps}}, \underline{b}^{D_1}$ by imaging the arrays generated by the respective AODs and extracting $\underline{r}_{ij}$ for both the traps and the D_1 beams. We then explain how these estimates are used to align the orientation of the D_1 AODs and determine the RF frequencies needed to overlap the D_1 beams with the trap array (see Fig. 4.5).

Estimation. Given any set of position vectors $\underline{r}_{ij}$ describing the positions of optical beams in the xy plane, generated by the frequency components $\underline{\nu}_{ij} = (\nu_{0x} + i\Delta\nu_x,\ \nu_{0y} + j\Delta\nu_y)^T$, where $i \in \{0, 1, \ldots, N_x - 1\}$ and $j \in \{0, 1, \ldots, N_y - 1\}$, we aim to find the optimal matrix $\underline{\underline{\mathcal{M}}}$ and vector $\underline{b}$ such that the linear model in Eqs. (4.4)–(4.5) accurately reproduces the beam positions.

To do so, we first minimize the following error function:

$$E(\underline{\underline{\mathcal{K}}}, \underline{c}) = \sum_{ij} \left\| \underline{r}_{ij} - \underline{\underline{\mathcal{K}}} \cdot \underline{e}_{ij} - \underline{c} \right\|^2, \tag{4.6}$$

where $\underline{e}_{ij} = (i, j)^T$ is the grid index vector.

The optimal parameters $\underline{\underline{\mathcal{K}}}$ and $\underline{c}$ are then used to compute the matrix $\underline{\underline{\mathcal{M}}}$ and offset vector $\underline{b}$ as follows:

$$\underline{\underline{\mathcal{M}}} = \underline{\underline{\mathcal{K}}} \cdot \begin{bmatrix} 1/\Delta\nu_x & 0 \\ 0 & 1/\Delta\nu_y \end{bmatrix}, \tag{4.7}$$

$$\underline{b} = \underline{c} - \underline{\underline{\mathcal{K}}} \cdot \begin{bmatrix} 1/\Delta\nu_x & 0 \\ 0 & \Delta\nu_y \end{bmatrix} \cdot \begin{bmatrix} \nu_{0x} \\ \nu_{0y} \end{bmatrix} \tag{4.8}$$

The parameters $\underline{\underline{K}}$ and $\underline{c}$ can be found by minimizing the error function in Eq. (4.6), whose solution is given by

$$\begin{bmatrix} \mathcal{K}_{00} \\ \mathcal{K}_{01} \\ \mathcal{K}_{10} \\ \mathcal{K}_{11} \\ c_0 \\ c_1 \end{bmatrix} = \underline{\underline{\mathcal{L}}}^{-1} \cdot \underline{d}, \tag{4.9}$$

where

$$\mathcal{L} = \begin{bmatrix} \frac{N_x(N_x-1)(2N_x-1)N_y}{6} & \frac{N_x(N_x-1)N_y(N_y-1)}{4} & 0 & 0 & \frac{N_x(N_x-1)N_y}{2} & 0 \\ \frac{N_x(N_x-1)N_y(N_y-1)}{4} & \frac{N_xN_y(N_y-1)(2N_y-1)}{6} & 0 & 0 & \frac{N_xN_y(N_y-1)}{2} & 0 \\ 0 & 0 & \frac{N_x(N_x-1)(2N_x-1)N_y}{6} & \frac{N_x(N_x-1)N_y(N_y-1)}{4} & 0 & \frac{N_x(N_x-1)N_y}{2} \\ 0 & 0 & \frac{N_x(N_x-1)N_y(N_y-1)}{4} & \frac{N_xN_y(N_y-1)(2N_y-1)}{6} & 0 & \frac{N_xN_y(N_y-1)}{2} \\ \frac{N_x(N_x-1)N_y}{2} & \frac{N_xN_y(N_y-1)}{2} & 0 & 0 & N_xN_y & 0 \\ 0 & 0 & \frac{N_x(N_x-1)N_y}{2} & \frac{N_xN_y(N_y-1)}{2} & 0 & N_xN_y \end{bmatrix}, \tag{4.10}$$

and

$$\underline{d} = \begin{bmatrix} \sum_{ij} i \left(\underline{r}_{ij}\right)_x \\ \sum_{ij} j \left(\underline{r}_{ij}\right)_x \\ \sum_{ij} i \left(\underline{r}_{ij}\right)_y \\ \sum_{ij} j \left(\underline{r}_{ij}\right)_y \\ \sum_{ij} \left(\underline{r}_{ij}\right)_x \\ \sum_{ij} \left(\underline{r}_{ij}\right)_y \end{bmatrix}. \tag{4.11}$$

Overlapping Procedure. The estimates of $\underline{\underline{\mathcal{M}}}^{\text{traps}}, \underline{\underline{\mathcal{M}}}^{D_1}, \underline{b}^{\text{traps}}, \underline{b}^{D_1}$ can be used to determine the appropriate angular alignment for the D_1 AODs, as well as the RF frequencies required to overlap each D_1 beam with its corresponding optical tweezer.

First, we decompose the matrices $\underline{\underline{\mathcal{M}}}^{\text{traps}}$ and $\underline{\underline{\mathcal{M}}}^{D_1}$ into the product of a rotation matrix $\underline{\underline{R}}$, a *skewing matrix* $\underline{\underline{T}}$, and a scaling matrix $\underline{\underline{S}}$:

$$\underline{\underline{\mathcal{M}}} = \underline{\underline{R}} \cdot \underline{\underline{T}} \cdot \underline{\underline{S}}, \tag{4.12}$$

where

$$\underline{\underline{R}} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}, \qquad \underline{\underline{T}} = \begin{bmatrix} 1 & \kappa \\ 0 & 1. \end{bmatrix}, \qquad \underline{\underline{S}} = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix}, \quad s_x, s_y \geq 0. \tag{4.13}$$

The decomposition parameters θ, κ, s_x, and s_y can be extracted from the matrix elements $\mathcal{M}_{ij}$ via:

$$s_x = \sqrt{\mathcal{M}_{00}^2 + \mathcal{M}_{10}^2}, \tag{4.14}$$

$$\theta = \arctan\left(\frac{\mathcal{M}_{10}}{\mathcal{M}_{00}}\right) + \frac{\pi}{2} \pm \frac{\pi}{2}, \tag{4.15}$$

$$s_y = \mathcal{M}_{11}\cos\theta - \mathcal{M}_{01}\sin\theta, \tag{4.16}$$

$$\kappa = \frac{\mathcal{M}_{01}\cos\theta + \mathcal{M}_{11}\sin\theta}{\mathcal{M}_{11}\cos\theta - \mathcal{M}_{01}\sin\theta}. \tag{4.17}$$

The sign in the expression for θ is chosen such that $s_y \geq 0$. Here, θ denotes the orientation of the array, κ is the skew factor (with $\kappa = 0$ indicating a perfectly rectangular lattice), and s_x, s_y denote the inter-site spacings along the principal axes of the array per unit of frequency .

The parameters θ and κ are functions of the orientation of the D_1 AODs. We fine tune the orientation of the D_1 AODs such that these parameters are the same for both the D_1 array and the optical tweezers array. This concludes the first fine-tuning step.

For the second step, we use the relations in Eqs. (4.4)–(4.5) to find $\underline{\nu}^{D_1}$ as a function of $\underline{\nu}^{\text{traps}}$. In particular, using the relation

$$\underline{\nu}^{D_1} = \left(\underline{\underline{\mathcal{M}}}^{D_1}\right)^{-1} \cdot \left(\underline{\underline{\mathcal{M}}}^{\text{traps}} \cdot \underline{\nu}^{\text{traps}} + \underline{b}^{\text{traps}} - \underline{b}^{D_1}\right). \tag{4.18}$$

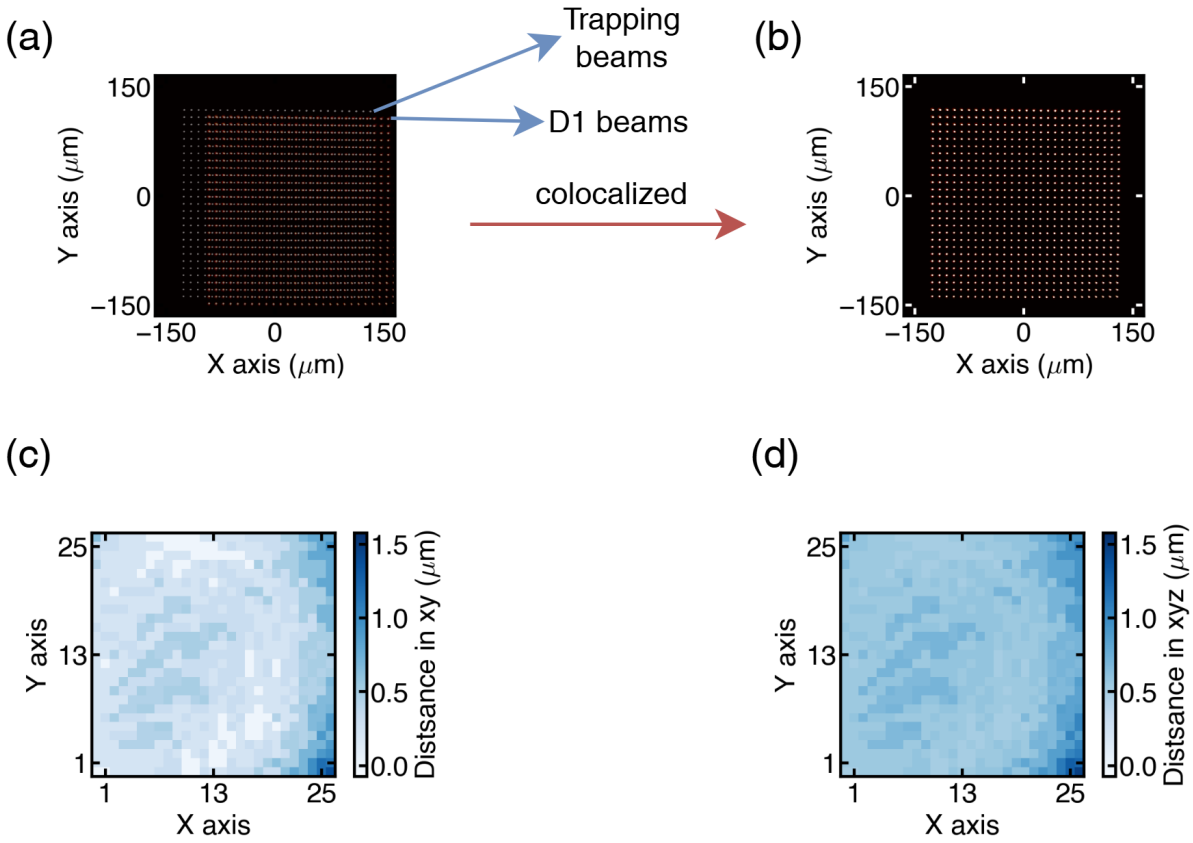


Figure 4.6: **Characterization of colocalization between AOD-generated trap arrays and D_1 Raman beams for local single-qubit gates.** (a) Trap array (gray circles) and D_1 Raman beam array (red circles) after coarse alignment, showing significant spatial mismatch. (b) The arrays after applying the fine-tuning colocalization procedure described in Sec. 4.2.1, showing close overlap. (c) Measured distances between matched sites in the xy plane. The average xy separation is 0.3 μm. (d) Full 3D spatial distances between matched sites, with an average of 0.6 μm. The increase in separation near the bottom-right corner is attributed to breakdown of the linear mapping model in Eqs. (4.4)–(4.5) for high-frequency AOD components far from the center of the array.

4.2.2 Assessment of Performance

To quantify the quality of colocalization between the traps and the D_1 beams, we use the residual error metric defined in Eq. (4.6), which captures the spatial misalignment

between the two arrays. The spatial profiles of the traps and D_1 beams are characterized by imaging their point-spread functions (PSFs) while translating the imaging lens along the axial (z) direction. The imaging system uses the second microscope objective (focal length $f = 4$ mm), which is mounted on a precision translation stage (Newport CONEX-SAG) to allow controlled axial displacement (see Fig. 4.4).

From these measurements, we extract an average spatial separation of 0.3 μm in the xy plane between the trap centers and the D_1 beams, for traps with radial dimensions of ≈ 0.8 μm and axial extent of ≈ 5 μm (see Fig. 4.6c). The full 3D spatial distance between colocalized sites is larger, approximately 0.6 μm (see Fig. 4.6d), primarily due to limited axial resolution. Since this 3D distance is comparable to the measurement precision in the z direction (≈ 0.5 μm), the axial resolution is the main limiting factor. Improved axial alignment may be achievable through in-situ Rabi frequency measurements to fine-tune the overlap.

In the xy plane, colocalization is near perfect at the center of the array but gradually worsens towardsthe edges (see Fig. 4.6c). This limitation arises from the linear frequency-to-position model described in Eqs. (4.4)–(4.5), which does not capture higher-order dependencies. As these non-linear effects differ for the trap and D_1 beam paths, further improvement in full-array colocalization is currently limited.

The practical impact of this level of alignment is that we can now perform single-qubit gates simultaneously across all traps with minimal cross-talk. Improved colocalization in the xy plane could be attained when addressing only a subset of traps at a time, allowing for individual frequency tuning and localized optimization. This is because simultaneous addressing of all traps constrains us to use equi-spaced frequency components to drive the D_1 AODs.

Chapter 5

Optimizing Fluorescence Imaging for Quantum State Readout

A key step in projective qubit state measurement in neutral atom array quantum processors is fluorescence imaging of atoms in optical tweezers. During this process, atoms can be unintentionally ejected from their traps due to heating and photon recoil, leading to qubit loss. Additionally, the occupation state of traps may be misidentified due to the finite number of detected photons, which causes the atomic signal distribution to partially overlap with the background noise distribution. These issues become more pronounced as the array size grows, increasing the probability of at least one error. It is therefore essential to optimize the imaging process to minimize both atom loss and state misclassification.

Fluorescence imaging is performed by illuminating the optical traps with three pairs of counter-propagating polarization-gradient cooling beams, which in our platform are near-resonant with the atomic D_2 transition (see Fig. 1.3). The imaging light consists of two components: a cooler beam addressing the $F = 2 \rightarrow F' = 3$ transition and a repumper beam addressing the $F = 1 \rightarrow F' = 2$ transition. The repumper prevents atoms from becoming trapped in the $F = 1$ dark state, which can occur because the cooler beam occasionally drives the $F = 2 \rightarrow F' = 2$ transition with low probability, from which atoms may decay into the $F = 1$ manifold. By returning these atoms to the $F = 2 \rightarrow F' = 3$ cycling transition, the repumper ensures continuous scattering and fluorescence during imaging.

During imaging, atoms are repeatedly excited from the ground state $5s_{1/2}$ to the excited state $5p_{3/2}$ and subsequently decay by spontaneously emitting photons at 780 nm. These photons are collected and detected using an EMCCD camera. The imaging performance

depends on several factors, including the properties of the imaging beams (frequency and power of the cooler beam, duration of imaging), the magnetic field, and the polarization of the polarization-gradient cooling beams. In this work, we focus primarily on the properties of the cooler beam, specifically its frequency and power, as well as the total imaging time. Optimizing these parameters involves a trade-off between the number of photons collected per atom and atom loss: higher imaging powers increase the scattering rate, which raises the number of photons collected by the camera from each atom and improves classification accuracy within a fixed imaging time, but also increases heating and recoil, raising the probability of atom loss during imaging.

In the following sections, we propose a method for optimizing the imaging baem parameters in neutral atom arrays to achieve low-loss and high-fidelity imaging of single atoms in optical traps. We begin by deriving an expression for a figure of merit that captures the trade-off between atom loss and detection fidelity. Then, we describe how this figure of merit can be measured experimentally (see Sec. 5.1). Afterwards, we use our method to identify optimal imaging parameters, specifically the imaging beam frequency and power, that minimize atom loss per detected photon (see Sec. 5.2). We then extract the resulting classification error and atom loss probabilities to characterize the fidelity of our imaging sequence (see Sec. 5.3). This method allows a higher fidelity for quantum state readout in neutral atom array, thus reducing the number of repetitions and increasing the accuracy for measurements in applications such as quantum sensing of electromagnetic fields or quantum simulations.

5.1 Derivation of the Figure of Merit

In this section, we introduce a figure of merit (FOM) to characterize the efficiency of the fluorescence imaging process in terms of atom survival and photon collection [72]. We first define this FOM based on the atom loss rate per photon fluoresced, which captures the trade-off between minimizing atom loss and maximizing photon detection. Then, we propose a practical method to experimentally measure this FOM by monitoring atom survival probability and fluorescence rate under varying imaging conditions.

Definition. To quantify the performance of the imaging process, we model the survival probability of an atom during fluorescence imaging as an exponential decay:

$$P_{\text{surv}}(t) = e^{-t/T}, \tag{5.1}$$

where T is the characteristic lifetime of the atom in the trap under the given imaging conditions.

Simultaneously, the atom fluoresces photons at a rate ν, which depends on the imaging beam parameters such as intensity and detuning.

The loss rate of atoms during imaging is therefore given by $1/T$, while the total number of photons emitted over the imaging time $\Delta t_{\text{imaging}}$ is approximately $\nu \Delta t_{\text{imaging}}$. To balance the trade-off between imaging speed and atom loss, we define the loss rate per photon fluoresced as

$$\text{FOM} = \frac{1/T}{\nu} = \frac{1}{T\nu}. \tag{5.2}$$

Minimizing this quantity ensures that the atom remains trapped while emitting the maximum number of photons, thus improving detection fidelity with minimal loss. This alternate figure of merit provides a direct metric for optimizing imaging parameters to achieve the best compromise between atom survival and measurement quality.

In practice, directly measuring the figure of merit defined above can be time-consuming, as it requires simultaneous tracking of atom survival and photon emission over time. Instead, we describe an indirect method for estimating this quantity under the assumptions that the atomic fluorescence rate remains constant and the atom's survival probability decays exponentially with time. These assumptions are valid in the regime where imaging parameters (beam power and detuning) remain fixed, and allow the figure of merit to be extracted efficiently from a series of measurements taken at different imaging durations.

Measurement. We now describe a method to measure the FOM defined in Eq. (5.2). This procedure assumes that atoms can be imaged with a classification fidelity greater than 50% using a heuristic method. We will later describe how to systematically classify atoms as present or absent using an optimized thresholding procedure (see Sec. 5.3). However, in order to optimize the imaging parameters and determine this threshold, we require a means of performing atom imaging in the first place. For this purpose, a heuristic classification method suffices, as long as the atom survival probability is nonzero and the classification fidelity exceeds 50%. Under these conditions, the occupancy of each trap can be reliably inferred by repeating the same imaging parameters and threshold multiple times, and using the resulting statistics to estimate the true occupancy state.

The experimental sequence consists of three steps (see Fig. 5.1a). In the first step, we verify the presence of an atom in a trap using a preliminary imaging step with heuristic parameters. In the second step, we apply the candidate imaging parameters, whose figure

of merit we aim to evaluate, for a fixed duration $\Delta t_{\text{imaging}}$. During this period, fluorescence photons emitted by the atom are collected on the EMCCD camera, and the corresponding photon count is recorded as N. In the third step, we repeat the same heuristic imaging method to determine whether the atom remains trapped. This three-step sequence is repeated over M experimental shots, indexed by $i \in \{0, 1, \ldots, M-1\}$, with M chosen to be large enough to obtain statistically meaningful averages of the recorded photon counts N_i.

Let $L \subseteq \{0, 1, \ldots, M-1\}$ be the subset of shots in which an atom was detected in the first step, and $S \subseteq L$ be the subset of those shots where the atom also survived the third step. Then, the average photon count collected during the imaging period is given by

$$N_{\text{avg}} = \frac{\sum_{i \in S} N_i}{\|S\|}, \tag{5.3}$$

and the corresponding survival probability is

$$p_{\text{surv}} = \frac{\|S\|}{\|L\|}, \tag{5.4}$$

where $\|L\|$ and $\|S\|$ denote the cardinalities of the sets L and S, respectively.

The quantities p_{surv} and N_{avg} can be experimentally measured using the procedure described above, and the figure of merit can be estimated from them through a simple relation. First, we model the survival probability as

$$p_{\text{surv}} = \exp\left(-\frac{\Delta t_{\text{imaging}}}{T}\right), \tag{5.5}$$

where T is the atom lifetime under the given imaging conditions. Additionally, assuming the fluorescence rate ν is constant, the average photon count is given by

$$N_{\text{avg}} = \nu \Delta t_{\text{imaging}}. \tag{5.6}$$

Using these expressions, the figure of merit defined as the atom loss rate per emitted photon,

$$\text{FOM} = \frac{1/T}{\nu}, \tag{5.7}$$

can be rewritten as

$$\text{FOM} = \frac{\Delta t_{\text{imaging}}}{N_{\text{avg}}} \cdot \frac{-\ln(p_{\text{surv}})}{\Delta t_{\text{imaging}}}, \tag{5.8}$$

which simplifies to

$$\mathrm{FOM} = \frac{-\ln(p_{\mathrm{surv}})}{N_{\mathrm{avg}}}. \tag{5.9}$$

This final expression provides a convenient way to compute the figure of merit from measured survival probabilities and average photon counts. The optimal imaging parameters are those that minimize this quantity.

5.2 Optimization of Imaging Parameters

To minimize atom loss during quantum state readout while maintaining high-fidelity classification, we optimize the imaging frequency and power by minimizing the figure of merit defined in Eq. (5.9), which penalizes atom loss per detected fluorescence photon. Since photons are only captured during the 100 ms middle image within a total 250 ms interval between the first and third images, the counted photon number is scaled by a factor of $\frac{250}{100}$ to account for the full interval.

The optimization protocol consists of three consecutive fluorescence images. The first and third images are each 300 ms long and use a fixed imaging configuration: 20 μW per beam and a detuning of -9 Γ from the $F = 2 \rightarrow F' = 3$ transition of the D_2 line, where $\Gamma = 6.06$ MHz is the natural linewidth. The middle image is 100 ms long and separated from the others by 75 ms. It is acquired at the same power (~ 20 μW), while the frequency is varied to evaluate performance.

The presence or absence of an atom in the first and third images is determined by a thresholding algorithm (see Sec. 5.3). If the atom is detected in both the first and third images, it is considered to have survived the middle imaging pulse. If it is present in the first but absent in the third, it is assumed to have been lost during the middle exposure. Simultaneously, the scaled number of detected photons in the second image provides a measure of the fluorescence signal at each imaging frequency.

As the imaging frequency approaches the resonance of the D_2 transition, the fluorescence rate increases, consistent with the standard scattering rate expression given by

$$R_{\mathrm{sc}} = \frac{\Gamma}{2} \cdot \frac{I/I_{\mathrm{sat}}}{1 + I/I_{\mathrm{sat}} + 4(\Delta/\Gamma)^2}, \tag{5.10}$$

where I is the imaging beam intensity, I_{sat} is the saturation intensity, and Δ is the detuning [73, 74]. Due to the AC Stark shift from the optical tweezer, the effective resonance is shifted by approximately $+1$ Γ from the free-space value (see Sec. 3.1).

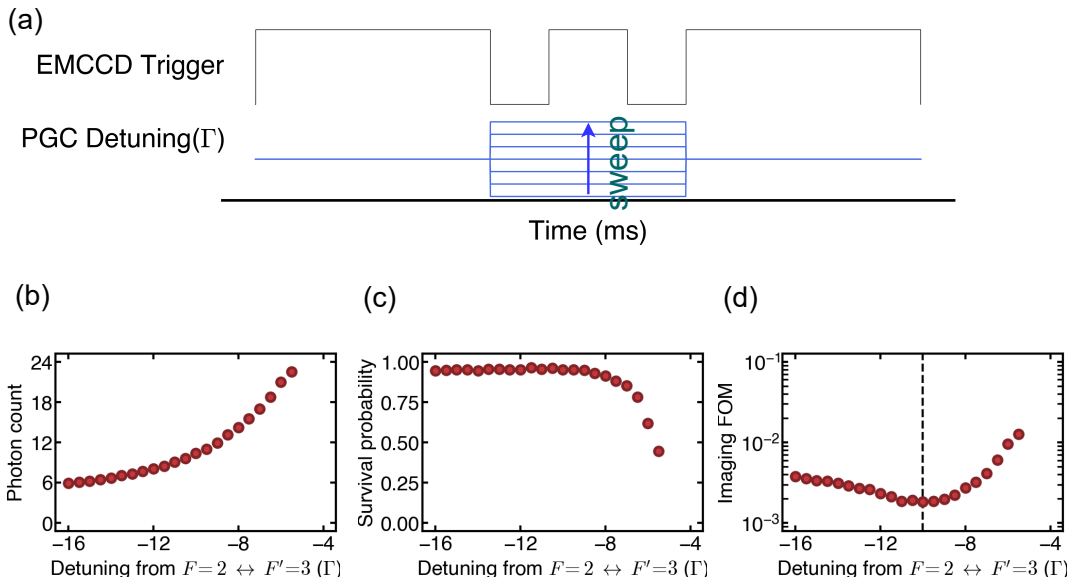


Figure 5.1: **Optimization of Imaging Frequency to Minimize Atom Loss and Classification Error.** (a) Experimental sequence showing three consecutive EMCCD triggers, each corresponding to a fluorescence image, taken at fixed optical power of 20 μW per beam: the first and third images are each 300 ms long, and the middle image is 100 ms long with variable detuning. (b) Fluorescence photon counts as a function of imaging frequency, increasing towards the AC Stark shifted resonance near +1 Γ. (c) Atom survival probability decreases near resonance due to reduced polarization-gradient cooling efficiency. (d) Figure of merit (atom loss per detected fluorescence photon, see Eq. (5.9)) demonstrating a trade-off between fluorescence rate and survival probability, minimized near −10 Γ, indicating the optimal imaging detuning.

However, as the detuning approaches resonance, atom survival probability decreases. polarization-gradient cooling is most effective for red-detuned beams and fails when the frequency becomes too close to resonance [75]. This trade-off results in a nontrivial minimum in the figure of merit. We find the optimal imaging frequency to be −10 Γ at a power of 20 μW.

Repeating the measurement for different powers, we observe that this optimum remains near 20 μW. With both frequency and power fixed, the remaining parameter, the imaging duration, is chosen to ensure the target classification fidelity is achieved. This choice

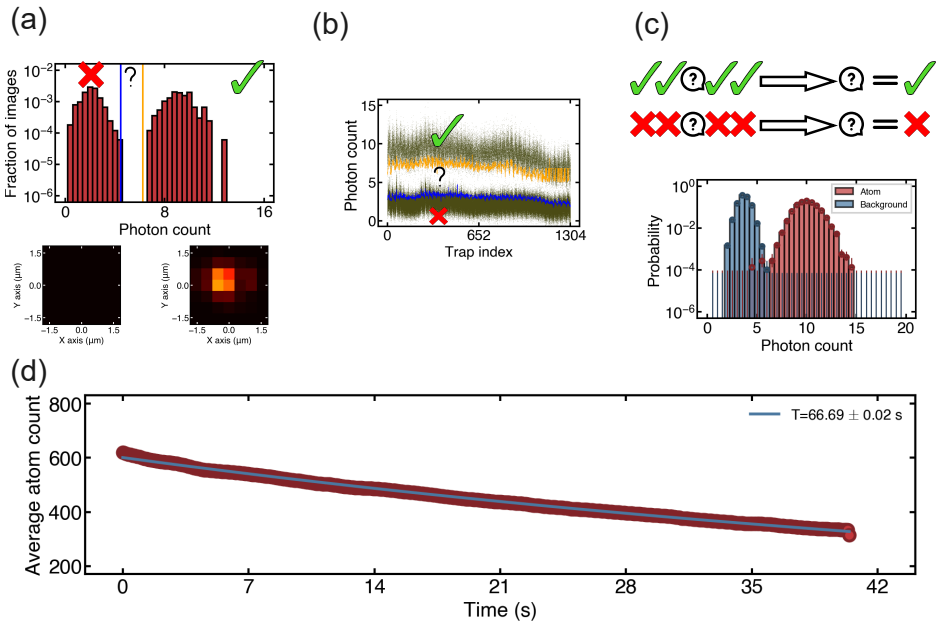


Figure 5.2: **Imaging loss and classification fidelity.** (a) Photon-count thresholds are used to classify traps as occupied, vacant, or unknown, with representative point-spread functions shown. (b) Applying this heuristic across the array yields classification fidelities above 50%. (c) Signal and background distributions are extracted from consecutive images; their overlap sets the classification error, which averages 0.2% over 1305 traps. (d) Continuous imaging of a 45×29 array gives an atom lifetime of 66.7 ± 0.02 s and a loss probability of 0.12%.

yields a state readout configuration with minimal qubit loss while maintaining the desired classification fidelity.

Table 5.1: **Imaging Performance Metrics.** Summary of classification fidelity, atom loss probability, and imaging time for the optimized imaging configuration.

Parameter	**Value**
Imaging time	75 ms
Classification error	0.2×10^{-3}
Atom loss probability	1.2×10^{-3}

5.3 Characterization of the Imaging Classification Fidelity and Loss Probability

Having found the optimal imaging parameters, we now characterize the atom loss probability and the classification fidelity of the imaging procedure (see Tab. 5.1). For this analysis, we fix the imaging duration to an arbitrary value of $\Delta t_{\text{imaging}} = 75$ ms. In general, increasing $\Delta t_{\text{imaging}}$ leads to improved classification fidelity due to greater photon collection but also results in higher atom loss. Conversely, reducing $\Delta t_{\text{imaging}}$ decreases atom loss at the cost of reduced classification fidelity. The choice of 75 ms here serves as a representative example to illustrate the characterization procedure.

Loss Probability. The loss probability can be inferred from the optimization procedure in Sec. 5.3, since to find the FOM we calculated $p_{\text{sur}} = 1 - p_{\text{loss}}$. Alternatively, we measure the average number of atoms in the trap array when illuminated by the imaging beam. We then fit the function

$$A \exp(-t/T) \tag{5.11}$$

to the number of atoms, with the fit parameters A and T. The loss probability during imaging is then given by

$$p_{\text{loss}} = 1 - \exp(-\Delta t_{\text{imaging}}/T). \tag{5.12}$$

From the results, it follows that the loss probability of atoms in the imaging process is 0.12% (see Fig. 5.2).

Classification Error. Here, we characterize the classification fidelity of the imaging process. To do this, we first describe the image processing algorithm responsible for classification. We use a simple method that integrates the number of detected photons over

a 5×5 pixel region centered at the expected trap location on the camera, and compares the result against a fixed threshold. We first describe how an appropriate threshold can be chosen, and then explain how the classification fidelity can be estimated.

To estimate the classification fidelity, we analyze the distribution of integrated photon counts over many imaging shots. Due to the probabilistic nature of atom loading, some traps will be occupied while others will be vacant. This leads to a bimodal distribution in the detected photon counts, with one peak corresponding to background noise and the other to atomic fluorescence. However, significant overlap can occur between the two distributions due to various noise sources.

Photon count fluctuations arise from the Poissonian statistics of both background noise and atomic fluorescence. Additional variability is introduced by the analog-to-digital converter (ADC) of the camera, whose bias voltage fluctuations constitute the readout noise. The EMCCD sensor contributes further noise through clock-induced charges and the probabilistic nature of the electron multiplication process. These combined effects broaden the photon count distributions for both occupied and vacant traps, leading to non-zero overlap, thereby limiting the classification fidelity.

To extract these distributions separately, we first perform a rough classification using a pair of thresholds $t_1 < t_2$. These serve a role analogous to digital signal discrimination: photon counts above t_2 are classified as "occupied," below t_1 as "vacant," and counts in between are treated as ambiguous. The values of t_1 and t_2 are determined heuristically by fitting a bimodal Gaussian to the overall histogram of photon counts. Specifically, each threshold is chosen to lie two standard deviations away from the mean of its corresponding peak (see Fig. 5.2).

We then take five consecutive images using the optimal imaging parameters (see Sec. 5.2). We assign the third image in the sequence to be occupied if and only if the photon counts in all four surrounding images exceed t_2. Similarly, the third image is labeled vacant if all four surrounding images have photon counts below t_1. This voting scheme provides high-confidence samples for both occupied and vacant states, which enables the construction of separate distributions $d_{\text{atom}}(t')$ and $d_{\text{vacant}}(t')$ for the two cases.

With these distributions, we compute the optimal threshold t that minimizes the classification error. The classification error is the sum of the false negative and false positive

probabilities, defined by

$$p_{\text{false negative}}(t) = \int_0^t d_{\text{atom}}(t')\,dt', \tag{5.13}$$

$$p_{\text{false positive}}(t) = \int_t^\infty d_{\text{vacant}}(t')\,dt'. \tag{5.14}$$

We assume that the a priori probability of a trap being occupied is η, and define the total classification error as

$$p_{\text{error}}(t) = \eta\, p_{\text{false negative}}(t) + (1-\eta)\, p_{\text{false positive}}(t). \tag{5.15}$$

In the absence of any prior knowledge of the trap state, we take $\eta = \frac{1}{2}$, in which case the classification error simplifies to

$$p_{\text{error}}(t) = \frac{1}{2}\left(p_{\text{false negative}}(t) + p_{\text{false positive}}(t)\right). \tag{5.16}$$

For each individual atom, the classification threshold t is chosen to minimize $p_{\text{error}}(t)$.

From this, we obtain both the optimal threshold t and the corresponding classification error $p_{\text{error}}(t)$. The average classification error is found to be 0.02%.

Chapter 6

Conclusion and Outlook

In this thesis, we have developed and demonstrated experimental methods that enhance the scalability, stability, and performance of neutral atom array quantum processors. (see Fig. 1.5)

First, we reduced the per-trap laser power requirement by circularizing and de-astigmatizing the trapping beam with a compact three-lens system, improving beam circularity from 0.69 to 0.97 and suppressing astigmatism to within 0.8% of the Rayleigh range (see Sec. 2). This modification enables more efficient use of the trapping laser power and supports the scaling number of qubits. Applied to an M2 Ti:sapphire laser in our laboratory, the method increases the number of traps generated by 5% for the same available optical power.

Second, we implemented a real-time closed-loop feedback system that stabilizes trap depths using both fast CMOS and EMCCD measurement (see Sec. 3). By decomposing noise into uncorrelated temporal modes and regulating each with independent PID controllers, we achieved long-term trap-depth stability exceeding one minute and suppressed high-frequency fluctuations. This approach eliminates the need for frequent recalibration and ensures stable operation during extended experimental runs. Furthermore, we integrated this module with analog PID controllers to perform adiabatic ramps of the traps, enabling manipulation of trap depths while maintaining spatial homogeneity across the array. This capability serves as a critical building block for quantum sensing protocols in neutral atom array processors, which often require traps to be ramped down to very low depths.

Third, using the developments presented in this thesis, we demonstrate how neutral atom arrays can be employed to perform measurements of magnetic fields across the trap array (see Sec. 4). We illustrate this capability by measuring a magnetic field gradient

across the array. This is significant, since the measurement of electromagnetic fields on the trap array allows the characterization and nulling of environmental magnetic fields, as well as the stabilization of the alignment of Rydberg excitation beams used for multi-qubit gates on this platform. Additionally, we show that the platform can serve as a building block for entanglement-enhanced quantum sensing and for implementing dynamic decoupling techniques to sense electromagnetic fields with improved precision.

Finally, we proposed methods to optimize the properties of the global imaging beams used for probing single atoms in optical tweezers, thereby maximizing atom survival and state-detection fidelity 5. These improvements directly reduce measurement errors in large-scale experiments and enhance the overall reliability of neutral atom array quantum processors.

Together, these advances demonstrate a hardware-aware framework for improving both the reliability and efficiency of neutral atom quantum processors. The techniques developed here provide a foundation for scaling up system sizes, maintaining high-fidelity operation over long experimental timescales, and enabling precise measurements of electromagnetic fields using atom arrays. Looking ahead, these results can be extended to support fault-tolerant gate operations [62], entanglement-enhanced metrology [76], and quantum simulations of classically intractable many-body systems, thereby contributing to the realization of quantum advantage with neutral atoms array quantum processors.

References

[1] Shlomi Kotler, Nitzan Akerman, Yinnon Glickman, Anna Keselman, and Roee Ozeri. Single-ion quantum lock-in amplifier. *Nature*, 473(7345):61–65, 2011.

[2] Sepehr Ebadi, Tout T. Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595:227–232, 2021.

[3] Giulia Semeghini, Harry Levine, Alexander Keesling, Sepehr Ebadi, Tout T. Wang, Daniel Bluvstein, Ruben Verresen, Hannes Pichler, M. Kalinowski, Ahmed Omran, Sylvain Schwartz, Sean M. Sommer, Deborah S. Weiss, Mikhail Greiner, Vladimir Vuletić, Ashvin Vishwanath, and Markus D. Lukin. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.

[4] Iris Cong, Harry Levine, Alexander Keesling, Dolev Bluvstein, Sheng-Tao Wang, and Mikhail D. Lukin. Hardware-efficient, fault-tolerant quantum computation with rydberg atoms. *Phys. Rev. X*, 12(2):021049, 2022.

[5] 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. *Annual Review of Condensed Matter Physics*, 11:369–395, 2020.

[6] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 2019.

[7] Pieter Kok, W. J. Munro, Kae Nemoto, T. C. Ralph, Jonathan P. Dowling, and G. J. Milburn. Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics*, 79:135–174, 2007.

[8] Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Logical quantum processor based on reconfigurable atom arrays. *Nature*, 624:54–61, 2023.

[9] Hannah J. Manetsch, Gyohei Nomura, Elie Bataille, Kon H. Leung, Xudong Lv, and Manuel Endres. A tweezer array with 6100 highly coherent atomic qubits. *arXiv*, 2403.12021, 2024.

[10] Alain Delgado, Pablo A. M. Casares, Roberto dos Reis, Modjtaba Shokrian Zini, Roberto Campos, Norge Cruz-Hernández, Arne-Christian Voigt, Angus Lowe, Soran Jahangiri, M. A. Martin-Delgado, Jonathan E. Mueller, and Juan Miguel Arrazola. Simulating key properties of lithium-ion batteries with a fault-tolerant quantum computer. *Physical Review A*, 106(3):032428, 2022.

[11] Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature*, 622:287–292, 2023.

[12] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D Lukin. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 2017.

[13] Kent Ueno and Alexandre Cooper. Distributing entanglement at the quantum speed limit in rydberg chains. *arXiv*, quant-ph/2506.19228, June 2025.

[14] Christopher Wyenberg, Kent Ueno, and Alexandre Cooper. Quantum channel for modeling spin-motion dephasing in rydberg chains. *arXiv*, quant-ph/2506.24082, June 2025.

[15] Dolev Bluvstein, Alexandra A. Geim, Sophie H. Li, Simon J. Evered, J. Pablo Bonilla Ataides, Gefen Baranes, Andi Gu, Tom Manovitz, Muqing Xu, Marcin Kalinowski, Shayan Majidy, Christian Kokail, Nishad Maskara, Elias C. Trapp, Luke M. Stewart, Simon Hollerith, Hengyun Zhou, Michael J. Gullans, Susanne F. Yelin, Markus

Greiner, Vladan Vuletić, Madelyn Cain, and Mikhail D. Lukin. Architectural mechanisms of a universal fault-tolerant quantum computer. *arXiv*, quant-ph/2506.20661, 2025.

[16] Tommaso Roscilde, Meenu Kumari, Alexandre Cooper, and Fabio Mezzacapo. Exponential onset of scalable entanglement via twist-and-turn dynamics in xy models. *arXiv*, quant-ph/2507.08206, July 2025.

[17] Aaron W. Young, William J. Eckner, William R. Milner, Dhruv Kedar, Matthew A. Norcia, Eric Oelker, Nathan Schine, Jun Ye, and Adam M. Kaufman. A tweezer clock with half-minute atomic coherence at optical frequencies and high relative stability. *Nature*, 619:697–703, 2023.

[18] Dominik Schäffner, Tobias Schreiber, Fabian Lenz, Malte Schlosser, and Gerhard Birkl. Quantum sensing in tweezer arrays: Optical magnetometry on an individual-atom sensor grid. *arXiv*, 2307.08055, 2023.

[19] N. Schlosser, G. Reymond, and P. Grangier. Collisional blockade in microscopic optical dipole traps. *Physical Review Letters*, 89(2):023005, 2002.

[20] Barry Cimring, Remy El Sabeh, Marc Bacvanski, Stephanie Maaz, Izzat El Hajj, Naomi Nishimura, Amer E. Mouawad, and Alexandre Cooper. Efficient algorithms to solve atom reconfiguration problems. I. redistribution-reconfiguration algorithm. *Phys. Rev. A*, 108(2):023107, 2023.

[21] Amir H. Dadpour, Timur Khayrullin, Fouad Afiouni, Remy El Sabeh, Amer E. Mouawad, Izzat El Hajj, and Alexandre Cooper. Low-latency control system for feedback experiments with optical tweezer arrays. *arXiv*, 2504.06528, 2025.

[22] Fouad Afiouni, Remy El Sabeh, Naomi Nishimura, Izzat El Hajj, Amer E. Mouawad, and Alexandre Cooper. Efficient algorithms to solve atom reconfiguration problems. iii. the bird and batching algorithms and other parallel implementations on GPUs. *Physical Review A*, 112(2):023109, August 2025.

[23] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman. Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits. *Physical Review Letters*, 114(10):100503, 2015.

[24] Pascal Scholl, Michael Schuler, Hannah J. Williams, Alexander A. Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C. Lang, Thierry Lahaye, Andreas M. Läuchli, and Antoine Browaeys. Programmable quantum

simulation of 2d antiferromagnets with hundreds of rydberg atoms. *Nature*, 595:233–238, 2021.

[25] Rui Lin, Han-Sen Zhong, You Li, Zhang-Rui Zhao, Le-Tian Zheng, Tai-Ran Hu, Hong-Ming Wu, Zhan Wu, Wei-Jie Ma, Yan Gao, Yi-Kang Zhu, Zhao-Feng Su, Wan-Li Ouyang, Yu-Chen Zhang, Jun Rui, Ming-Cheng Chen, Chao-Yang Lu, and Jian-Wei Pan. AI-Enabled rapid assembly of thousands of defect-free neutral atom arrays with constant-time-overhead. *arXiv*, 2412.14647, 2024.

[26] Omprakash Chandra, Gopikrishnan Muraleedharan, and Gavin K. Brennen. Non-local resources for error correction in quantum ldpc codes. *arXiv*, 2409.05818, September 2024.

[27] Adam L. Shaw, Pascal Scholl, Ran Finklestein, Ivaylo S. Madjarov, Brandon Grinkemeyer, and Manuel Endres. Dark-state enhanced loading of an optical tweezer array. *Physical Review Letters*, 131(18):183001, May 2023.

[28] Kai-Niklas Schymik. *Scaling-up the Tweezer Platform — Trapping Arrays of Single Atoms in a Cryogenic Environment*. Phd thesis, Université Paris-Saclay, 2022.

[29] Mert Serkan and Hulya Kirkici. Reshaping of a divergent elliptical gaussian laser beam into a circular, collimated, and uniform beam with aspherical lens design. *IEEE Sensors Journal*, 9(1):36–43, 2009.

[30] Y. T. Chew, M. Poitrinal, T. Tomita, S. Kitade, J. Mauricio, K. Ohmori, and S. de Léséleuc. Ultraprecise holographic optical tweezer array. *Phys. Rev. A*, 110:053518, Nov 2024.

[31] Soroush Khoubyarian, Anastasiia Mashko, and Alexandre Cooper. Correcting astigmatism and ellipticity in gaussian beams using a cylindrical lens pair with tunable focal lengths. *arXiv*, 2506.22308, 2025.

[32] D. Hanna. Astigmatic gaussian beams produced by axially asymmetric laser cavities. *IEEE Journal of Quantum Electronics*, 5(10):483–488, 1969.

[33] W. A. Clarkson and D. C. Hanna. Two-mirror beam-shaping technique for high-power diode bars. *Optics Letters*, 21(6):375–377, 1996.

[34] Matthew Knecht. Optimization of optical tweezers. Southern Scholars Senior Research Project, Southern Adventist University, 2007.

[35] Mert Serkan and Hulya Kirkici. Optical beam-shaping design based on aspherical lenses for circularization, collimation, and expansion of elliptical laser beams. *Applied Optics*, 47(2):230–241, 2008.

[36] Zhen-Nan Tian, Li-Jie Wang, Qi-Dai Chen, Tong Jiang, Li Qin, Li-Jun Wang, and Hong-Bo Sun. Beam shaping of edge-emitting diode lasers using a single double-axial hyperboloidal micro-lens. *Optics Letters*, 38(24):5414–5417, 2013.

[37] Haiyin Sun. *Laser Diode Beam Manipulations*, pages 53–104. Springer Netherlands, Dordrecht, 2015.

[38] Md Nazmul Hasan, Muttahid-Ull Haque, and Yung Chun Lee. Deastigmatism, circularization, and focusing of a laser diode beam using a single biconvex microlens. *Optical Engineering*, 55(9):095107, 2016.

[39] Catalina Ramírez-Guerra, José Agustín Moreno-Larios, Martha Rosete-Aguilar, and Jesús Garduño-Mejía. Mode-coupling enhancement by pump astigmatism correction in a ti:sapphire femtosecond laser. *Applied Optics*, 55(34):9889–9894, 2016.

[40] Spectra-Physics. *Model 3900S CW Ti:Sapphire Laser User's Manual*. The Solid-State Laser Company, 1335 Terra Bella Avenue, Mountain View, CA 94043, 2002.

[41] M. V. R. K. Murty and Jesus R. Moya-Cessa. Magnification of prisms: an apparent anomaly and its explanation. *Optical Engineering*, 23(1):83–87, 1984.

[42] Zhou Xiao-Qun, Bryan Ngoi Kok Ann, and Koh Soon Seong. Single aspherical lens for deastigmatism, collimation, and circularization of a laser beam. *Applied Optics*, 39(7):1148–1151, 2000.

[43] Shan Yang, Ming Bai, and Jungang Miao. Deastigmatism and circularization of an elliptical gaussian beam by off-axis ellipsoid reflector based off-focus configuration. *Progress In Electromagnetics Research B*, 10:91–103, 2008.

[44] He Cheng, Chun Xia, Stephen M. Kuebler, and Xiaoming Yu. Aberration correction for slm-generated bessel beams propagating through tilted interfaces. *Optics Communications*, 475:126213, 2020.

[45] Richard W. Bowman, Amanda J. Wright, and Miles J. Padgett. An slm-based shack–hartmann wavefront sensor for aberration correction in optical tweezers. *Journal of Optics*, 12(12):124004, 2010.

[46] R. E. Wagner and W. J. Tomlinson. Coupling efficiency of optics in single-mode fiber components. *Applied Optics*, 21(15):2671–2688, 1982.

[47] Kurt D. Wulff, Daniel G. Cole, Robert L. Clark, Roberto DiLeonardo, Jonathan Leach, Jon Cooper, Graham Gibson, and Miles J. Padgett. Aberration correction in holographic optical tweezers. *Optics Express*, 14(9):4169–4174, 2006.

[48] Remy El Sabeh, Jessica Bohm, Zhiqian Ding, Stephanie Maaz, Naomi Nishimura, Izzat El Hajj, Amer E. Mouawad, and Alexandre Cooper. Efficient algorithms to solve atom reconfiguration problems. II. assignment-rerouting-ordering algorithm. *Phys. Rev. A*, 108(2):023108, 2023.

[49] Ali Binai-Motlagh, Matthew L Day, Nikolay Videnov, Noah Greenberg, Crystal Senko, and Rajibul Islam. A guided light system for agile individual addressing of ba^+ qubits with 10^{-4} level intensity crosstalk. *Quantum Science and Technology*, 8(4):045012, 2023.

[50] Sungjoo Lim, Seunghyun Baek, Jacob Whitlow, Marissa D'Onofrio, Tianyi Chen, Samuel Phiri, Stephen Crain, Kenneth R. Brown, and Junki Kim. Design and characterization of individual addressing optics based on multi-channel acousto-optic modulator for $^{171}Yb^+$ qubits. *Optics and Laser Technology*, 180:111436, 2025.

[51] E. A. DeHoog. Anamorphic zoom lens based on rotating cylindrical lenses. *Optics Express*, 29(8):12206–12215, 2021.

[52] D. B. Elliott. *Clinical Procedures in Primary Eye Care*. Butterworth-Heinemann, Oxford, 3rd edition, 2007.

[53] Joseph T. Verdeyen. *Laser Electronics*. Prentice Hall Series in Solid State Physical Electronics. Prentice Hall, Upper Saddle River, NJ, third edition, 1995.

[54] 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. *Phys. Rev. A*, 106(2):022611, 2022.

[55] D. V. Gris-Sánchez and T. A. Birks. The airy fiber: an optical fiber that guides light diffracted by a circular aperture. *Optica*, 3(3):270–276, 2016.

[56] J. A. Arnaud and H. Kogelnik. Gaussian light beams with general astigmatism. *Appl. Opt.*, 8(8):1687–1693, Aug 1969.

[57] A. E. Siegman. *Lasers*. University Science Books, Mill Valley, California, 1986.

[58] Joseph W. Goodman. *Introduction to Fourier Optics*. Roberts and Company Publishers, 3rd edition, 2005.

[59] OpenCV Development Team. *The OpenCV Reference Manual*. OpenCV, release 2.3 edition, 2011.

[60] Andrew William Fitzgibbon and Robert B. Fisher. A buyer's guide to conic fitting. In *British Machine Vision Conference*, 1995.

[61] Dominik Schäffner, Tilman Preuschoff, Simon Ristok, Lukas Brozio, Malte Schlosser, Harald Giessen, and Gerhard Birkl. Arrays of individually controllable optical tweezers based on 3d-printed microlens arrays. *Optics Express*, 28(6):8640–8645, 2020.

[62] Dolev Bluvstein, Harry Levine, Giulia Semeghini, Adam M. Kaufman, Alexander Keesling, Thomas Thiele, Soonwon Choi, Rhine Samajdar, Hannes Pichler, Subir Sachdev, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906):451–456, 2022.

[63] Dustin Stuart and Axel Kuhn. Single-atom trapping and transport in dmd-controlled optical tweezers. *New Journal of Physics*, 20(2):023013, 2018.

[64] Kai-Niklas Schymik, Vincent Lienhard, Daniel Barredo, Pascal Scholl, Hannah Williams, Antoine Browaeys, and Thierry Lahaye. Enhanced atom-by-atom assembly of arbitrary tweezer arrays. *Phys. Rev. A*, 102(6):063107, 2020.

[65] Florence Nogrette, Henning Labuhn, Sylvain Ravets, Daniel Barredo, Lucas Béguin, Aline Vernier, Thierry Lahaye, and Antoine Browaeys. Singleatom trapping in holographic 2d arrays of microtraps with arbitrary geometries. *Phys. Rev. X*, 4(2):021034, 2014.

[66] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. Analysis of dephasing mechanisms in a standing-wave dipole trap. *Physical Review A*, 72:023406, August 2005.

[67] Manuel Endres, Hannes Bernien, Alexander Keesling, Harry Levine, Eric R Anschuetz, Alexandre Krajenbrink, Crystal Senko, Vladan Vuletić, Markus Greiner, and Mikhail D Lukin. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315):1024–1027, 2016.

[68] Sijie Chen, Min Zhuang, Ruihuang Fang, Yun Chen, Chengyin Han, Bo Lu, Jiahao Huang, and Chaohong Lee. Quantum iuble lock-in amplifier. *Communications Physics*, 7:189, 2024.

[69] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier. Energy distribution and cooling of a single atom in an optical tweezer. *Physical Review A*, 78(3):033425, 2008.

[70] B. Fröhlich, T. Lahaye, B. Kaltenhäuser, H. Kübler, S. Müller, T. Koch, M. Fattori, and T. Pfau. A two-frequency acousto-optic modulator driver to improve the beam pointing stability during intensity ramps, 2007.

[71] B. Nikolov, E. Diamond-Hitchcock, J. Bass, N. L. R. Spong, and J. D. Pritchard. Randomized benchmarking using nondestructive readout in a two-dimensional atom array. *Phys. Rev. Lett.*, 131:030602, 2023.

[72] Alexandre Cooper, Jacob P. Covey, Ivaylo S. Madjarov, Sergey G. Porsev, Marianna S. Safronova, and Manuel Endres. Alkaline-earth atoms in optical tweezers. *Phys. Rev. X*, 8:041055, 2018.

[73] C. J. Foot. *Atomic Physics*. Oxford University Press, Oxford, UK, 2005.

[74] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Graduate Texts in Contemporary Physics. Springer-Verlag, New York, 1999.

[75] Jean Dalibard and Claude Cohen-Tannoudji. Laser cooling below the doppler limit by polarization gradients: simple theoretical models. *Journal of the Optical Society of America B*, 6(11):2023–2045, 1989.

[76] William J. Eckner, Nelson Darkwah Oppong, Alec Cao, Aaron W. Young, William R. Milner, John M. Robinson, Jun Ye, and Adam M. Kaufman. Realizing spin squeezing with rydberg interactions in an optical clock. *Nature*, 625:120–125, 2023.

[77] Kevin D. Whitley, Matthew J. Comstock, and Yann R. Chemla. *High-resolution "fleezers": Dual-Trap Optical Tweezers Combined with Single-Molecule Fluorescence Detection*, volume 1486, pages 183–256. Humana Press, New York, 2017.

Glossary

AC Stark shift A shift in the energy levels of a quantum system induced by an oscillating electromagnetic field. In atomic systems, it arises from off-resonant light and can be used to implement state-dependent potentials or control qubit frequencies. xxvi, 8–10, 53, 55, 75, 86, 87, 90, 105, 106

beam circularity A measure of how close a laser beam's cross-section is to a perfect circle, defined as the minimum ratio of the beam's minor axis to its major axis achieved as it propagates in free space. A circularity of 1 indicates a perfectly circular beam, while lower values indicate elliptical distortion. xiv, xv, xvii, 15, 24, 26, 31

cylindrical lens A lens with different focal lengths along two orthogonal axes. A uni-axial cylindrical lens has one infinite focal length, meaning it does not focus along that axis. A bi-axial cylindrical lens has finite focal lengths along both axes. xiv–xvii, 9, 14, 16, 17, 19–35, 37–45, 47–49

dynamic architecture A neutral atom quantum processor design in which the positions of atoms in the array can be reconfigured during operation. This design provides flexible connectivity between qubits, adaptive circuit implementation, and optimization of gate operations for specific algorithms or protocols. 6

Jacobian matrix A matrix of all first-order partial derivatives of a vector-valued function. xix, xxi, 63, 67, 68, 72, 73

logical qubit A qubit encoded using multiple physical qubits to protect quantum information from errors. Logical qubits enable fault-tolerant quantum computation by allowing error detection. 6

PID controller A feedback control loop that continuously calculates an error signal as the difference between a desired setpoint and a measured variable, and applies a correction to the control parameters based on proportional, integral, and derivative terms to minimize the error. xxii, 60, 64, 65, 75–77, 111

quantum processor A physical device that encodes and manipulates quantum information using qubits. It can implement any quantum gate on its qubits and perform measurements, enabling computation, simulation, or sensing tasks. xii, xiii, 1–3, 6–9, 11

quantum simulation The use of a controllable quantum system to mimic the dynamics of another quantum system that is difficult to study directly. The Hamiltonian describing the target system is mapped to a Hamiltonian on the controllable system, which is then engineered and evolved to reproduce the dynamics of the target system. 1–3, 6

qubit A quantum mechanical system whose wavefunction is constrained to a two-dimensional subspace of the total Hilbert space of the system. Simply having qubits does not imply that one can control their states or induce interactions between them. xi, xiii, xiv, xxv, xxvi, 1–3, 5–14, 52, 53, 75, 81–86, 90–95, 99–101, 107, 112

Rabi spectroscopy A spectroscopic technique in which a two-level quantum system is driven by an oscillating field. By varying the drive frequency, one observes a resonance when it matches the transition frequency. Alternatively, by varying the pulse duration at a fixed frequency, one observes Rabi oscillations in the state populations, from which the transition frequency and Rabi frequency of the system can be extracted. xiii, 5, 12, 85

Ramsey spectroscopy A spectroscopic technique in which a two-level quantum system is subjected to two short coherent pulses separated by a free-evolution period. By scanning the phase or frequency of the second pulse relative to the first, one observes interference fringes whose pattern reveals the transition frequency and coherence properties of the system. xxiv, 9, 10, 81, 82, 85, 86, 90, 92

spherical lens An optical lens whose surfaces are portions of a sphere, focusing light symmetrically in all directions. xv, 21, 32, 33, 35, 37, 41

Zeeman shift The splitting or shifting of atomic or molecular energy levels in the presence of an external magnetic field. The shift is proportional to the magnetic field strength

and the magnetic moment of the state, and it is used to manipulate and address specific quantum states. 86, 90