THE UNIVERSITY OF CHICAGO

MULTIMODE CAVITIES FOR QUANTUM SCIENCE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
LUKAS PALM

CHICAGO, ILLINOIS
JUNE 2026

TABLE OF CONTENTS

LIST OF FIGURES v

LIST OF TABLES vii

ABSTRACT viii

1 INTRODUCTION 1
- 1.1 Quantum Simulation 1
- 1.2 Simulation of the Fractional Quantum Hall Effect 1
  - 1.2.1 Platforms 1
- 1.3 Quantum Computing and Networking 5
- 1.4 Thesis Overview 6

2 QUANTUM HALL EFFECT 7
- 2.1 The Integer Quantum Hall Effect 9

3 CAVITY QUANTUM ELECTRODYNAMICS 16
- 3.1 Cavity Finesse and Power Buildup 17
  - 3.1.1 Undercoupled and Overcoupled Regimes 19
- 3.2 Cavity Stability and Mode Spectrum 20
  - 3.2.1 Full solution to Helmholtz eq 21
  - 3.2.2 ABCD matrices 22
  - 3.2.3 operator formalism 24
  - 3.2.4 Hamiltonian Optics 29
  - 3.2.5 Perturbation theory 31
- 3.3 Ray tracing 34
- 3.4 Cavity QED 39
- 3.5 Atom coupled to a field 39
- 3.6 Light-matter coupling for a quantized field 40
- 3.7 Photon Extraction Efficiency and Cooperativity 40
  - 3.7.1 Summary of Cavity Design Principles 43
- 3.8 Jaynes-Cummings model 44
- 3.9 Non-Hermitian perturbation theory 45
  - 3.9.1 The Formalism 46
  - 3.9.2 VRS 49
  - 3.9.3 Collective States 49
- 3.10 Rydberg EIT 49
  - 3.10.1 Rydberg Polaritons 50
  - 3.10.2 Structured Hamiltonians 53
- 3.11 Rydberg interactions 55

4 MULTI-MODE CAVITY QED: METHODS . . . 62
4.1 Cavity Structure . . . 62
4.1.1 Laser System . . . 73
4.2 Cavity locking . . . 79
4.3 Light shaping with DMD . . . 82
4.4 Detection path . . . 86
4.5 Ultra-fast EMCCD . . . 87
4.5.1 Optical Design . . . 88
4.5.2 Noise characteristics . . . 90
4.6 Experimental Sequence . . . 90
4.6.1 Steps . . . 92
4.7 Characterization and Optimization of a Hybrid Cavity Rydberg Polariton experiment . . . 99
4.7.1 MOT . . . 101
4.7.2 PGC . . . 102
4.7.3 Transport . . . 102
4.7.4 dRSC . . . 103
4.7.5 Optical Pumping . . . 103
4.7.6 Trap Ramp . . . 105

5 MULTI-MODE CAVITY QED: RESULTS . . . 108
5.1 Cavity . . . 108
5.1.1 Cavity spectroscopy . . . 108
5.1.2 Landau level characterization . . . 113
5.1.3 Disorder measurement . . . 115
5.2 Single-mode VRS and EIT . . . 116
5.2.1 Single-mode Photon Blockade . . . 119
5.3 Multi-mode VRS and EIT . . . 119
5.3.1 Atomic lensing in VRS . . . 120
5.3.2 Spatially resolved cooperativity . . . 122
5.3.3 EIT for higher order modes . . . 122
5.3.4 Spatially resolved Polaritons . . . 124
5.4 Time-resolved dynamics in the LLL . . . 125
5.5 Flux lattice . . . 127
5.5.1 Background: mean field . . . 127
5.6 Förster resonance . . . 129

6 SMALL WAIST CAVITIES AND ARRAY . . . 136
6.1 Lens based degenerate cavities . . . 136
6.2 GRIN lens cavity for sub-micro waist . . . 138
6.2.1 Cavity Spectroscopy . . . 140
6.2.2 Scanning defect waist microscopy . . . 140
6.3 A first cavity array . . . 141
6.3.1 Mode picture for cavity array . . . 141

6.3.2 Aberrations . . . 146
6.3.3 Cavity resonance condition . . . 146

7 A CAVITY AFM . . . 148
7.1 Introduction . . . 148
7.2 Cantilever Lens Cavity . . . 152
7.3 Contact mode . . . 153
7.4 FM mode . . . 156

8 OUTLOOK . . . 160
8.1 Summary . . . 160
8.2 Laughlin state preparation . . . 160
8.2.1 Steady state driving . . . 160
8.3 Dissipative state preparation . . . 166
8.3.1 Flux threading and dynamical preparation . . . 169
8.3.2 Non reversible driving scheme . . . 171
8.4 Many-mode cavity QED . . . 172
8.5 Dynamical potentials using SLM . . . 174
8.6 Experiment Upgrades . . . 178
8.6.1 Magic lattice . . . 180
8.7 Spatio-temporal correlations . . . 182

A . . . 185
A.1 Laser Induced Contamination . . . 185
A.2 Motorized Mirror Model . . . 191
A.3 Magnetic Translation Operator . . . 202
A.3.1 The Magnetic Translation Operator . . . 202
A.4 Cavity mode shift and coupling from Astigmatism . . . 205
A.5 2. Effective Hamiltonian with Residual Detuning . . . 208
A.6 Second order astigmatism perturbation . . . 208
A.6.1 Coupling Calculation . . . 209
A.7 Analog Lockbox . . . 211
A.8 Piezo Hysteresis compensation . . . 213
A.9 Flexure design for precision optics alignment . . . 214
A.10 Spring washer design . . . 218
A.11 High bandwidth piezo driver . . . 219
A.12 Real-time feedback . . . 223

REFERENCES . . . 227

LIST OF FIGURES

2.1 Original measurement of the fractional quantum Hall effect 8
2.2 Landau level energy spectrum 12

3.1 Gaussian beam propagation parameters 21
3.2 Understanding aberrations in the mirror-based twisted resonator. 32
3.3 Aberrated mirror cavity mode images compared to theory 32
3.4 Mode dispersion of the lens cavity. 33
3.5 Ray tracing semiclassical orbits 36
3.6 Semi-classical ray orbits in lens cavities 39
3.7 Cavity optimization for power buildup and photon extraction 43
3.8 Jaynes-Cummings ladder 44
3.9 Cavity Rydberg EIT level scheme and transmission 52
3.10 Cavity Rydberg EIT window 54
3.11 Rydberg interactions and Förster defect 58
3.12 Cavity Rydberg polariton interaction potential 61

4.1 X-ray image of the cavity structure 63
4.2 EIT level scheme 66
4.3 X-ray image with in-vacuum SmarAct stage 67
4.4 X-ray image with electric field electrodes 72
4.5 Electronic level structure of ^{87}Rb 74
4.6 PDH locking with sideband on sideband technique 80
4.7 DMD blazed grating geometry and Lee holography 83
4.8 High level interleaved fast kinetics EMCCD idea and optical design. 88
4.9 Fast EMCCD Shack-Hartman alignment 89
4.10 All optical demonstration of the FAST EMCCD camera 91
4.11 Main optical setup 100
4.12 Degenerate Raman sideband cooling 104
4.13 Optical pumping optimization 106
4.14 ELAT absorption image 106

5.1 Cavity alignment and mode matching 110
5.2 Cavity transmission away from and at degeneracy 112
5.3 Landau Level characterization 114
5.4 Disorder tomography at degeneracy 116
5.5 Vacuum Rabi splitting transmission spectrum 117
5.6 EIT transmission spectrum 118
5.7 Cavity EIT spectrum vs control field 119
5.8 Single mode photon blockade 120
5.9 Intra-cavity lensing from atom misalignment 121
5.10 Spatially resolved VRS at degeneracy 123
5.11 Magnetic translation at degeneracy 124

5.12 Magnetic translation at degeneracy . . . 124
5.13 Magnetron motion of polaritons in the LLL . . . 126
5.14 Microwave polarization optimization . . . 130
5.15 AC Stark shift using MW drive . . . 131
5.16 Probe nonlinearity at resonance . . . 132
5.17 Correlation measurement at Förster resonance . . . 134

6.1 Mirror and lens cavity comparison . . . 137
6.2 GRIN lens cavity . . . 139
6.3 Scanning probe microscopy of diffraction limited waist . . . 141
6.4 CAM $8f$ cavity geometry . . . 143
6.5 Paraxial analysis of CAM local cavities . . . 145

7.1 Resonator design . . . 149
7.2 Contact mode AFM . . . 151
7.3 FM-AFM . . . 156

8.1 Quantum Hall Platform comparison . . . 162
8.2 Steady state Laughlin state preparation. . . . 163
8.3 Root partition Laughlin state preparation. . . . 165
8.4 Design of proposed dissipative state preparation scheme. . . . 167
8.5 Non-reversible driving schemes. . . . 168
8.6 Talbot effect of the CAM intra-cavity field . . . 173
8.7 Coupled cavities with SLM for local potentials . . . 175

A.1 Laser Induced Contamination . . . 187
A.2 LIC cleaning . . . 189
A.3 Alignment search space before and after model-based inversion . . . 195
A.4 Layered control architecture . . . 198
A.5 Analog lockbox schematic . . . 212
A.6 GPI model identification and correction . . . 215
A.7 Flexure design . . . 216
A.8 Spring washer design . . . 219
A.9 PA164 piezo driver schematic . . . 221
A.10 PA164 amplifier frequency response . . . 222
A.11 Real-time feedback system architecture . . . 224
A.12 Controller internal architecture . . . 225
A.13 Bayesian optimization workflow . . . 226

LIST OF TABLES

3.1 ABCD ray transfer matrices 23

8.1 Quantum Hall platform comparison 161

ABSTRACT

Over the last decade, quantum simulation and computation has progressed from a purely scientific testbed of quantum mechanics to a mature field that is starting to offer new insights into materials, quantum effects, and algorithms that are intractable with even the most state-of-the-art classical computation. Simulating certain quantum Hamiltonians that are theorized to be at the heart of high-T_c superconductivity, topological quantum systems such as the fractional quantum Hall effect or fundamental questions of thermalization quickly exceed the available memory of even the most advanced supercomputers – and in fact all memory in the world – due to the exponential growth of the Hilbert space with system size. As postulated by Feynman himself, the computational system best suited for solving quantum problems is to use a quantum system itself, as it can intrinsically represent the entanglement that is the key resource making quantum computing challenging but also promising. Apart from simulating problems in physics and chemistry, there is also tremendous excitement for quantum computers to provide speedups for general algorithms in optimization at a rate that is exponentially faster than existing classical programming. So far only a small class of algorithms has been actually proven to provide such a speedup, most prominently Shor's algorithm which can factor large prime numbers, and several other search and oracle algorithms.

In the last two decades, many experiments using ultracold atoms, ions, or superconducting circuits have focused on analog quantum simulation. Here, a model Hamiltonian that is believed to capture all the relevant physics of the system of study is experimentally realized using different building blocks and interactions but replicating condensed matter effects at an increased scale. This has led to a number of bespoke experiments such as BEC/Fermigass machines, quantum gas microscopes and more exotic setups that are purpose built to study a particular effect. More recently, there has been a shift to develop more universal simulation platforms and optical tweezer arrays have emerged as a leading solution due to their high

degree of reconfigurability. Through tremendous improvements in the degree of control and measurement of single qubits, error correction has become a viable route to build a true universal digital quantum computer. Even as we transition into the digital age of quantum simulation, certain classes of problems with additional physics constraints such as Fermionic statistics, gauge fields and lattice gauge theories still incur too large of an overhead to be encoded on these digital simulators. In this thesis, we use a highly degenerate twisted cavity to simulate the effect of a magnetic field on a two-dimensional system of interacting photons to enable realization of states of the quantum Hall effect. We design, build, and commission a new lens based cavity apparatus realizing a lowest Landau level for a mesoscopic number of modes. We demonstrate the integration of highly excited Rydberg atoms to mediate strong interactions and demonstrate photon blockade as well as mean field interactions. As a second application of cavities to quantum information leveraging our newfound comfort with intra-cavity optics, and experience with aberration modeling in resonators, we demonstrate a new class of small waist optical resonators enabled by lens-based cavities and then apply our understanding of optical aberrations in cavities to the development of a new class of cavity array. This *Cavity Array Microscope* (CAM) realizes a large number of cavities (up to currently ~ 600) using only a fixed number of free space optics and is both scalable with constant overhead as well as compatible with Rydberg atoms. This platform provides a route to highly parallel quantum repeater operation or entanglement generation between atom arrays in the future.

CHAPTER 1

INTRODUCTION

1.1 Quantum Simulation

The 21st century has seen a fundamental shift from understanding and observing the quantum world in hallmark experiments such as the observation of the first Bose Einstein condensate towards building and assembling quantum systems from fundamental building blocks to shape and control complex quantum systems on the single particle level using optical lattices, quantum gas microscopes and optical tweezers. Nowadays, these techniques enable microscopic control and imaging of increasingly large correlated states and combining them to realize an increasing range of physical effects including gauge fields, fermionic statistics or even full gauge theories with constraints. Quantum Simulation 2010s: Certain quantum experiments have focused on reducing errors for a subset of quantum operations to such an extreme extent that we are now entering the era of digital quantum simulation or quantum computing.

1.2 Simulation of the Fractional Quantum Hall Effect

1.2.1 Platforms

Several different experiments have succeeded in simulating quantum Hall states after their discovery in solid state semiconductor systems. Several different platforms have risen to the challenge that the quantum Hall effect presents, foremost realizing an artificial magnetic field for neutral particles that make up almost all quantum simulation systems, apart from charged ions. The neutral atoms and photons in these systems, be it Bosons or Fermions, do not experience a Lorentz force in a real magnetic field and also don't enjoy the strong and long-range interactions that electrons in condensed matter systems do. Many different

and creative ways to *mimic* the effect of a magnetic field have been invented over the years, and here a few of them are listed in approximately historical order:

- **Ultracold atoms in rotating traps [GSC10, CWG01]** The first attempts to realize a gauge field for ultracold atoms were based on the realization that particles in a trap that is rotating in the lab space will experience a Coriolis force in the rotating frame that is mathematically equivalent to an external magnetic field. In early BEC experiments, atoms were usually trapped in magnetic or optical dipole traps reaching macroscopically degenerate ground states by evaporative cooling. Optical lattices realized by interfering multiple laser beams were also available, however most of these experiments were bulk experiments due to the lack of high NA imaging systems able to resolve single particles. Depending on the atomic species, interactions between atoms happen mostly through scattering (elastic or also inelastic) where due to low temperatures the scattering amplitude can be decomposed in partial waves and ultracold atoms mostly experience s-wave or so-called "contact interaction." Depending on the detailed electronic structure of the valence electron in Alkali atoms, scattering resonances called *Feshbach* resonances can be used to tune the strength and also the sign of the interaction from attractive to repulsive. These experiments therefore already had all the necessary ingredients to simulate QHE states and pioneering experiments were performed [GSC10] using special optical lattices with elliptical sites that could be rotated optically, or using bulk BECs in rotating magnetic traps [CWG01, KTU03, ZAS$^+$05, Vie08, Fet09, MF07, RRD11]. While some signatures of QHE states like the Laughlin state could be observed, as the characteristic flat density profiles, they were unable to measure correlations that would have been a definitive signature due to the lack of imaging resolution at the time.

- **Quantum gas microscopes [BGP$^+$09, SWE$^+$10]** Optical lattices were first used in the bulk gas era, however the advent of quantum gas microscopes [BGP$^+$09, SWE$^+$10,

GB21] brought cold atom experiments to a new level of control and observability. In an optical lattice, single atoms are localized at the nodes of a periodic structure created by the interference of laser beams that can form many different geometries such as square, triangular [YLMS21], dimers, hexagonal, and kagome, which mimic different lattice structures found in the condensed matter. These systems are different from bulk gases as atoms no longer move in the continuum of space, but can only reside on discrete sites and tunnel between with a certain rate given by the lattice depth. Interactions between atoms are still enabled by the same scattering processes, however they take the much simpler form of an on-site non-linearity shifting the energy of a site depending on occupation. These simple ingredients already give rise to rich phase diagrams and physics such as insulators, magnetism and Bose- and Fermi-Hubbard model that can be intractable analytically or even for the best numerical simulations on classical computers. While the discrete lattice is fundamentally different from the QHE in the continuum it has equivalent states and effects such as Hofstadter's butterfly [Hof76]. Two major techniques have emerged to enable complex hopping amplitudes in these lattices that are equivalent to an Aharonov-Bohm phase from a magnetic flux piercing a plaquette: Raman transitions using two lasers incommensurate with the underlying lattice give rise to complex hopping, or periodic modulation or *shaking* of the lattice can achieve the same in a technique termed Floquet engineering.

- **Photons in engineered cavities and waveguides** Even for photons, which are massless particles that famously don't interact with each other nor can be trapped, artificial magnetic fields can be generated. While single photons can not be held in a trap like cold atoms due to their relativistic nature, their spatial wavefunction can be confined by placing them in optical fibers or optical resonators also known as *cavities*. Here, photons are repeatedly reflected by total internal reflection from a refractive index step or highly reflective coated mirrors. In both optical cavities or periodic waveguides

the photon returns to its initial position after one round trip. Depending on the geometrical structure of the phase space that the system undergoes during one cycle of Hamiltonian evolution, the photon can pick up a geometrical or Pancharatnam–Berry phase [Pan56]. A classical example of such a geometrical phase is Foucault's pendulum whose axis of oscillation precesses as it undergoes parallel transport by the earth's rotation. Similarly, transport of a photon around a closed cavity orbit or a helical waveguide also leads to a geometrical phase if an image rotation is incurred on a round trip.

- **Microwave circuit QED** Circuit QED has emerged as one of the most prominent quantum simulation and computing platforms: Instead of single electrons or atoms, the quantum states in superconducting circuits are formed by the macroscopic wavefunction that many electrons condense into at low temperatures, the BCS wavefunction. The fact that macroscopic structures made out of superconducting metal can behave in a quantum way opens them up to the whole toolbox of chip manufacturing techniques and complicated structures can be readily fabricated. Strong interactions are provided by special non-linear circuit elements called Josephson junctions that behave like a non-linear inductor and can be used to create anharmonic LC resonators. At extremely low temperatures of a few mK in special dilution refrigerators, these macroscopic wavefunctions can be coherent for several 10s of μs and since every site in a lattice needs to be fabricated it can also be controlled independently through dedicated control lines.

- **Few Fermions in rotating tweezers** The most recent platform to realize quantum Hall states is in many ways a return to the beginning: Equipped with much improved state preparation techniques that can generate a deterministic mesoscopic number of atoms [SZL$^+$11] and single particle resolved free space imaging [BKB$^+$18], the original idea of a rotating trap is again a viable platform. With modern optical techniques such as SLMs optical aberrations that could deform the trap and lead to imperfections that

prevent reaching the centrifugal limit can be compensated. Together with relatively strong s-wave interactions from Feshbach resonances long coherence times enable adiabatic preparation routines for both bosonic and fermionic species enabling unusual spinful QH effects [PGP20].

1.3 Quantum Computing and Networking

The core vision of quantum networking is to distribute entanglement, the unique resource quantum mechanics offers, across long distances to enable provably secure communication through quantum key distribution, enhance the sensitivity of distributed sensor arrays beyond the classical limit and allow different quantum nodes to communicate to exchange information between quantum computers or modules within one. For any real life distances outside of the lab, networking between cities or countries will be limited by photon loss in optical fibers and classical amplifiers must be replaced by quantum repeaters to maintain coherence. Realizing these repeaters and any quantum network requires high fidelity interfaces between flying (photons) and stationary (atoms, ions) qubits to exchange the quantum information. Optical cavities strongly coupled to single atoms have emerged as one of the most promising platforms [SSDRG11, AEE$^+$23] as the cavity provides Purcell-enhanced, efficient photon collection into a well-defined mode as well as strong qubit coupling from high cooperativity enabling dispersive readout schemes. Since the founding experiment [RNH$^+$12] demonstrating networking using single atoms, a variety of schemes [CFZ$^+$25, GGD$^+$25] using cavities have been demonstrated. With the emergence of optical tweezer arrays as one of the leading contenders for neutral atom quantum computing, a new challenge has emerged in how to parallelize cavities as an interface to couple multiple optical modes to many single atoms simultaneously as this will ultimately limit the bandwidth of quantum networking and distributed computing. A series of recent experiments [YHL$^+$23, HSW$^+$24, HSB$^+$25] have made first steps in this direction, however most current experiments utilize a single cavity to

interface with N atoms which incurs a serialization overhead that makes them not scalable.

Our work on novel cavity geometries has culminated in a new class of cavities, the *Cavity Array Microscope* (CAM) that realizes a large number of cavities (up to currently ~ 600) using only a fixed number of free space optics. This platform provides a route to highly parallel quantum repeater operation or entanglement generation between atom arrays.

1.4 Thesis Overview

In this thesis, we develop an understanding for multi-mode degenerate cavities both theoretically and experimentally and explore their applications in analog quantum simulation of quantum Hall physics as well as novel interfaces for quantum information. In chapter 3 we will give an overview over the different theoretical treatments of higher-order cavity modes and cavity QED when coupled to atoms. In chapter 4 we introduce the main experimental platform of our hybrid cavity Rydberg apparatus and the different cooling and preparation techniques. In chapter 5 we will explore a macroscopically degenerate lowest Landau level using non-interacting photons and interacting polaritons probing disorder, magnetron motion and the onset of organization in the mean field. In chapter 6 we will describe the evolution of a new kind of high-cooperativity cavity using diffraction limited waists and its application for single-atom and tweezer array interfacing. In chapter 7 we present an application of cavities to nano-sensing by combining cavities and atomic force microscopy. In chapter 8 we will summarize and discuss future directions for all these lines of cavity based science.

CHAPTER 2

QUANTUM HALL EFFECT

When electrons are confined to two spatial dimensions and placed in a strong magnetic field at low temperatures, they produce a strongly correlated quantum liquid. This gives rise to a number of phenomena summarized as the quantum Hall effect (QHE). Originally, the QHE was discovered in semiconductor samples that were cooled to low temperatures and subjected to a strong perpendicular magnetic field. The Hall conductance σ, defined as the ratio of the driving current to the resulting Hall voltage, was found to take quantized values, opposed to the classical effect. In the first discovery by Klaus von Klitzing, the Hall conductance showed only integer multiples of the so-called von Klitzing constant. This effect can be understood through the theory of Landau levels, the quantized orbitals of cyclotron motion of electrons in a perpendicular magnetic field. A few years later, additional plateaus of the Hall conductivity were observed [TSG82] at lower temperatures and higher magnetic fields, that did not correspond to integer, but fractional multiples of the quantization constant. This measurement is shown in Figure 2.1 and plateaus of the Hall resistance R_H can be observed at specific magnetic field strengths. The filling fraction ν is directly related to the magnetic field as it is defined as $\nu = p/q$ the ratio of the number of electrons (p) to magnetic flux quanta q. A number of integer as well as fractional filling fractions could be observed in Figure 2.1. The existence of these fractional values requires a theory involving electron interactions, whereas the integer effect could be explained on a single particle level. This was first achieved by Robert Laughlin [Lau83], who described the fractional quantum Hall states as an incompressible electron liquid and received a Nobel prize for this work together with the experimentalists Horst Störmer and Daniel Tsui [TSG82] in 1998. However, Laughlin's theory only applied to filling fractions $\nu = 1/m$ (with m an odd integer) and was based on a specific trial wavefunction. Therefore, the discovery of the fractional quantum Hall effect stimulated a multitude of theoretical work to explain the microscopic origin of the different

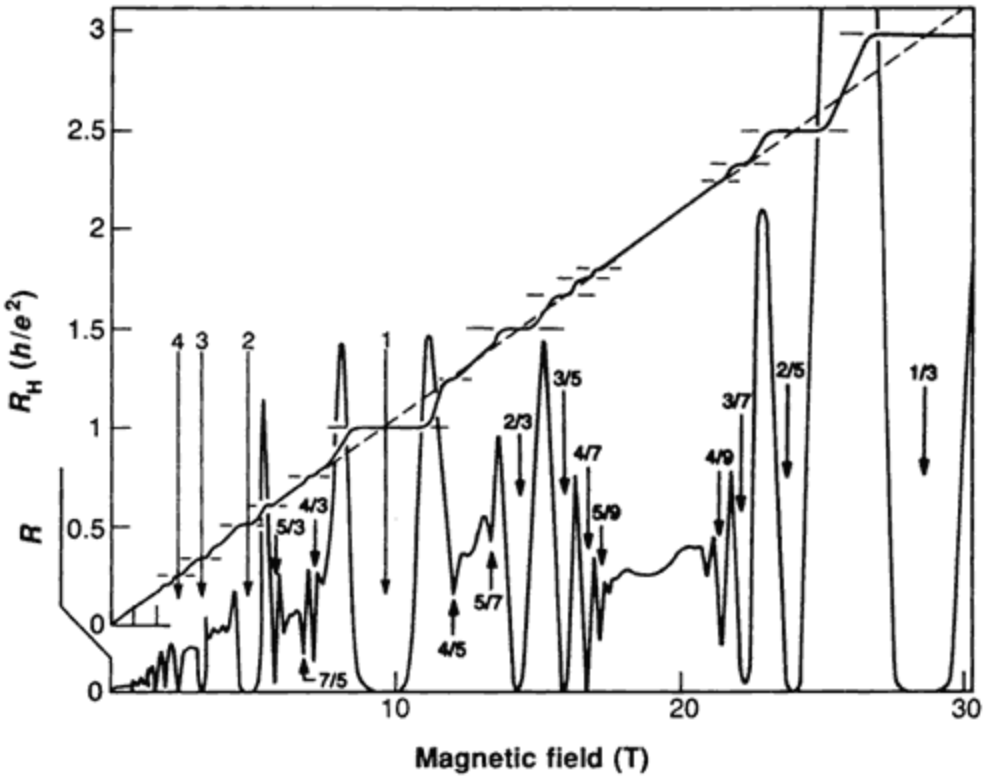


Figure 2.1: Original measurement of the fractional quantum Hall effect by [TSG82]. Upper curve: the Hall resistance R_H perpendicular to the driving current shows plateaus at certain magnetic fields that correspond to filling fractions ν indicated by the arrows. Lower curve: the longitudinal resistance R vanishes at these fractions.

fractions. A great advance was achieved through the theory of composite fermions [Jai89]. Its central idea is that electrons can capture magnetic flux quanta and the corresponding vortices in order to form quasi-particles minimizing their interaction energy. Consequently these composite fermions experience a reduced effective magnetic field. By this mechanism, a whole sequence of states with filling fractions

$$\nu = \frac{\nu^*}{2p\nu^* \pm 1} \tag{2.1}$$

in terms of the composite filling factor ν^* could be explained. However, the exact mechanism of the experimentally observed $\nu = 5/2$ state still remains an open question until today.

What makes the QHE so special is the precise quantisation of the Hall conductance into

multiples of the fundamental constant $e^2/h = 1/(25812.807572\Omega)$ (von Klitzing constant). This value is independent of any geometrical details of the sample or imperfections of the material, making it suited as a standard of resistance in metrology. The phases of matter present in the FQHE can only be described through non-local, so-called topological order. This is in contrast to the Landau theory of e.g. BCS superconductivity, where the different phases are classified via a local order parameter and their related symmetries. States described by topological order only depend on the global properties of the system and are therefore insensitive to local disorder and defects. The energy gap above the ground-state of FQHE systems arises through strong correlations induced by the Coulomb interaction among electrons. In this case macroscopic properties like the Hall conductance can be related to a topological invariant called the Chern number. Also a number of excited states exist that contain quasi-particles and -holes displaying unusual properties such as fractional charge and exchange statistics. They are called anyons because they neither obey fermionic nor bosonic properties.

In order to understand the fractional quantum Hall effect, which arises through interactions between electrons, we first have to look at the single particle picture. The derivations in this chapter closely follow [Ton16].

2.1 The Integer Quantum Hall Effect

The Integer Quantum Hall Effect can be understood in terms of the electron motion in a two dimensional system subject to a strong magnetic field. Therefore we look at the Hamiltonian of a free particle in a vector potential

$$\mathcal{H} = \frac{1}{2m}\left(\vec{p} + e\vec{A}\right)^2 = \frac{1}{2m}\vec{\pi}^2 \tag{2.2}$$

with the kinetic momentum $\vec{\pi} = \vec{p} + e\vec{A}$. The canonical commutation relations are

$$\left[x_i, p_j\right] = i\hbar\delta_{ij}, \quad \left[x_i, x_j\right] = 0 = \left[p_i, p_j\right] \tag{2.3}$$

It is important to note that two kinetic momenta do not commute:

$$\left[\vec{\pi}_x, \vec{\pi}_y\right] = -ie\hbar B \tag{2.4}$$

To solve the Hamiltonian Equation (2.2) we introduce ladder operators similar to the quantum harmonic oscillator by

$$a = \frac{1}{\sqrt{2e\hbar B}}\left(\vec{\pi}_x - i\vec{\pi}_y\right) \text{ and } a^\dagger = \frac{1}{\sqrt{2e\hbar B}}\left(\vec{\pi}_x + i\vec{\pi}_y\right) \tag{2.5}$$

which obey the commutation relation $\left[a, a^\dagger\right] = 1$. If we write the Hamiltonian in terms of these operators, we obtain the same form as in the harmonic oscillator case:

$$\mathcal{H} = \frac{1}{2m}\vec{\pi}^2 = \hbar\omega_B\left(a^\dagger a + \frac{1}{2}\right) \tag{2.6}$$

with the cyclotron frequency ω_B. The Hilbert space of this Hamiltonian is a Fock space containing states $|n\rangle$ with an energy of

$$E_n = \hbar\omega_B\left(n + \frac{1}{2}\right). \tag{2.7}$$

These energy levels are called Landau levels and will be of particular importance later. There are several possible choices for the vector potentials gauge describing the same magnetic field $\vec{B} = \nabla \times \vec{A} = B\vec{e}_z$. The symmetric gauge reflects rotational invariance

$$\vec{A} = (A_x, A_y) = \frac{B}{2}(-y, x), \tag{2.8}$$

while two other (called Landau gauge) preserve the translational invariance of the system

$$\vec{A} = B(-y, 0), \quad \vec{A} = B(0, x). \tag{2.9}$$

Even when the magnetic field $\vec{B}$ is invariant under rotation and translation in the plane, the choice of $\vec{A}$ is not. Choosing the symmetric gauge breaks translation invariance in the x- and y-direction, but it preserves the rotational symmetry and the angular momentum is a good quantum number.

To see the degeneracy of the Landau levels, we have to introduce a different momentum operator $\tilde{\vec{\pi}} = \vec{p} - e\vec{A}$ similar to the kinetic momentum, but with a minus sign. It fulfils the commutation relation

$$\left[\tilde{\vec{\pi}}_x, \tilde{\vec{\pi}}_y\right] = ie\hbar B \tag{2.10}$$

This momentum is not gauge invariant and depends on the specific gauge chosen, as we see from its commutators with the kinetic momentum:

$$\left[\vec{\pi}_x, \tilde{\vec{\pi}}_x\right] = 2ie\hbar \frac{\partial A_x}{\partial x}, \ \left[\vec{\pi}_y, \tilde{\vec{\pi}}_y\right] = 2ie\hbar \frac{\partial A_y}{\partial y} \tag{2.11}$$

$$\left[\vec{\pi}_x, \tilde{\vec{\pi}}_y\right] = \left[\vec{\pi}_y, \tilde{\vec{\pi}}_x\right] = ie\hbar \left(\frac{\partial A_x}{\partial y} + \frac{\partial A_y}{\partial x}\right) \tag{2.12}$$

As we can see, all of these commutators only vanish if we choose the symmetric gauge. Then, $\tilde{\vec{\pi}}$ commutes with the Hamiltonian and we can diagonalize them simultaneously, resulting in a new quantum number. We will later identify this quantity as the angular momentum, supporting our previous argument on rotational symmetry. So far, the calculation of the energy spectrum was independent of a specific gauge. But for the reasons given above, we will continue the remaining calculations in the symmetric gauge, as this is also the natural choice for a rotating system. After fixing the gauge we have $\left[\vec{\pi}_i, \tilde{\vec{\pi}}_j\right] = 0$ and can define two

new ladder operators

$$b = \frac{1}{\sqrt{2e\hbar B}}\left(\tilde{\vec{\pi}}_x + i\tilde{\vec{\pi}}_y\right) \text{ and } b^\dagger = \frac{1}{\sqrt{2e\hbar B}}\left(\tilde{\vec{\pi}}_x - i\tilde{\vec{\pi}}_y\right) \tag{2.13}$$

which also obey the canonical commutation relation $\left[b, b^\dagger\right] = 1$. The general Hilbert space is then constructed from the vacuum $|0,0\rangle$ as

$$|n,m\rangle = \frac{\left(a^\dagger\right)^n \left(b^\dagger\right)^m}{\sqrt{n!m!}}|0,0\rangle \tag{2.14}$$

where n is a label for the Landau level and $L = m - n$ denotes the total angular momentum. In the lowest Landau level (LLL), L and m coincide. The energy spectrum of this Hamiltonian and the different ladder operators actions are shown in Figure 2.2. To construct the

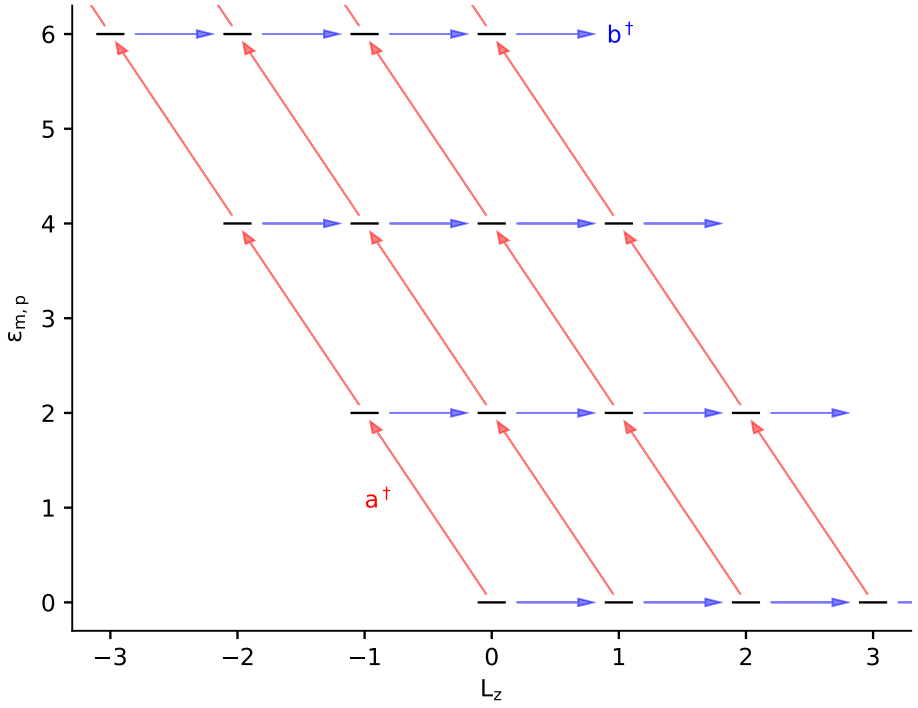


Figure 2.2: Landau levels (energy eigenstates of the Hamiltonian Equation (2.2)) as a function of angular momentum L_z. The arrows indicate the action of the creation operators $a^\dagger$ (blue, increase Landau level) and $b^\dagger$ (red, increase angular momentum).

real-space wave-functions, we express the creation and annihilation operators in terms of a complex coordinate z, because this fits the description of a plane.

$$a = \frac{1}{\sqrt{2e\hbar B}} \left(p_x - ip_y + e\left(A_x - iA_y\right)\right) \tag{2.15}$$

$$= \frac{1}{\sqrt{2e\hbar B}} \left(-i\hbar \left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right) + \frac{eB}{2}\left(-y - ix\right)\right) \tag{2.16}$$

$$= -i\sqrt{2}\left(l_B\bar{\partial} + \frac{z}{4l_B}\right) \tag{2.17}$$

Where we introduced the complex coordinate and derivatives

$$z = x - iy \text{ and } \bar{z} = x + iy \tag{2.18}$$

$$\partial = \frac{\partial}{\partial z} = \frac{1}{2}\left(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\right) \text{ and } \bar{\partial} = \frac{\partial}{\partial \bar{z}} = \frac{1}{2}\left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right) \tag{2.19}$$

which fulfil $\partial z = \bar{\partial}\bar{z} = 1$ and $\partial\bar{z} = \bar{\partial}z = 0$ and where we introduced the magnetic length $l_B = \sqrt{\frac{\hbar}{eB}}$ that sets the length scale of the state. To construct the *lowest Landau level (LLL)* $n = 0$, we make use of the property $a\,|n = 0, m\rangle = 0$ and solve the resulting differential equation to obtain

$$\psi_{LLL}(z, \bar{z}) = f(z)e^{-|z|^2/4l_B} \tag{2.20}$$

where f is an arbitrary function. To fix this, we employ that the ground state of the LLL $|0, m\rangle$ is also annihilated by $b\,|0, m = 0\rangle = 0$ giving

$$\psi_{LLL,m=0}(z, \bar{z}) \propto e^{-|z|^2/4l_B} \tag{2.21}$$

up to normalization. The excited states in the LLL can then be obtained by acting with $b^\dagger$, yielding

$$\psi_{LLL,m}(z, \bar{z}) \propto \left(\frac{z}{l_B}\right)^m e^{-|z|^2/4l_B}. \tag{2.22}$$

The function f is then a monomial in z, making the state $\psi_{LLL,m}$ an eigenstate of the

angular momentum operator

$$L_z = i\hbar\left(x\frac{\partial}{\partial y} - y\frac{\partial}{\partial x}\right) = \hbar\left(z\partial - \bar{z}\bar{\partial}\right) \tag{2.23}$$

$$L_z\psi_{LLL,m} = \hbar m\psi_{LLL,m} \tag{2.24}$$

with eigenvalue $\hbar m$. The real-space density of these wave-functions looks like a ring around the origin with radius $r = \sqrt{2m}l_B$ for an angular momentum m. In the same way, one could also construct the wave-functions of higher Landau levels by acting with $a^\dagger$, but the subsequent discussion will only focus on the LLL. As we can see from Equation (2.7) and also Figure 2.2, the energy of a state only depends on the Landau level index n, but not m. This means that the ground state (and every Landau level) is highly degenerate. We can estimate the number of states $\mathcal{N}$ in a disk shaped region of area $A = \pi R^2$ as

$$\mathcal{N} = \frac{R^2}{2l_B^2} = \frac{A}{2\pi l_B^2} = \frac{eBA}{2\pi\hbar} = \frac{BA}{\Phi_0}, \text{ with } \Phi_0 = \frac{2\pi\hbar}{e}, \tag{2.25}$$

where Φ_0 is the flux quantum. In the presence of an electric field from a voltage applied to the sample, the Hall conductivity can be calculated [Ton16] as

$$\rho_{xy} = \frac{2\pi\hbar}{e^2}\frac{1}{\nu}, \quad B = \frac{n}{\nu}\Phi_0 \tag{2.26}$$

where B is the magnetic field at which plateaus of ρ_{xy} occur and n denotes the electron density. In semiconductor systems, a number of additional effects comes into play. The finite size of the sample gives rise to chiral edge modes, whose direction of motion is fixed for each edge. Furthermore, no reason was given so far on why the plateaus in conductivity (see Figure 2.1) extend around the exact magnetic fields given in Equation (2.26). This effect is a result of disorder caused by impurities in the sample, that give rise to localised states. Both of these effects are not significant for a cold atoms system, which is why they are not

described in detail here.



CHAPTER 3

CAVITY QUANTUM ELECTRODYNAMICS

Photons are an amazing platform for quantum physics as we have all the tools of optics, such as lasers, beam splitters, and wave plates at our disposal to create and manipulate individual particles of light. By choosing the right source of light, such as parametric down conversion, entangled pairs of photons can be generated. These have provided some of the cleanest demonstrations of quantum physics, as the violations of Bell's inequalities, to date. However, photons are ill suited for realizing correlated many-body states, as optical photons do not interact with each other in vacuum due to the superposition principle of the linear Maxwell equations. In order to study any kind of interacting physics using photons, we need to use a non-linear medium whose properties change in the electric field of light and therefore appear different for additional photons. While any material can be non-linear at a high enough field strength, there are specific classes of crystal with an internal electronic structure that provide a strong non-linearity even at intermediate optical powers. Such crystals are used in non-linear optics to convert photons between different wavelengths in second harmonic generation or sum- and difference-frequency generation, modulate temporal dynamics in mode-locked lasers or solitons or octave spanning frequency combs. In all of these systems, non-linearity can be described as an optical susceptibility that is proportional to an intensity of a classical electric field, which means that all these effects require many photons to observe any non-linearity. Although a lot of recent progress has been made to enhance the non-linearity in on-chip devices down to the level of a few photons, no current material can provide a non-linear response for single photons. While some schemes like difference-frequency generation can be used for the interconversion of single photons, there is always a strong classical pump field required that will add noise to the output field due to the shot noise of photons in the pump beam.

3.1 Cavity Finesse and Power Buildup

We consider a Fabry-Perot cavity formed by two mirrors, M_1 (input/output coupler) and M_2 (end mirror), separated by a length ℓ. Each mirror j is characterized by its power transmission T_j and parasitic loss L_j (due to scattering and absorption), with the power reflectivity following from energy conservation as $R_j = 1 - T_j - L_j$. For high-finesse cavities it is more intuitive to specify both transmission and loss in parts per million (ppm) rather than quoting reflectivities with many nines.

From these power coefficients we obtain the field amplitude coefficients $t_j = \sqrt{T_j}$ and $r_j = \sqrt{R_j}$. We calculate all fields inside and outside the cavity using the self-consistent technique, assuming a steady-state intra-cavity field. Let E_{in} be the incident field on M_1, E_{circ} the circulating field inside the cavity (just after reflection from M_1), and ϕ the round-trip phase accumulation with $\phi = 2k\ell$. The self-consistent equation for the circulating field is formed by the sum of the transmitted input field and the field that has completed a round trip:

$$E_{\text{circ}} = it_1 E_{\text{in}} + r_1 r_2 e^{i\phi} E_{\text{circ}} \tag{3.1}$$

where the factor i accounts for the $\pi/2$ phase shift associated with transmission to satisfy the Stokes relations. Solving for E_{circ} yields

$$E_{\text{circ}} = \frac{it_1 E_{\text{in}}}{1 - r_1 r_2 e^{i\phi}}. \tag{3.2}$$

For many applications such as trapping in cavity-enhanced optical lattices or dipole traps, or for Rydberg excitation, a cavity can be ideal for amplifying available laser power. As the atoms absorb a negligible amount of power, the field inside the cavity is increased through constructive interference of successive round trips, enhancing trap depths or Rabi frequencies.

The circulating power is proportional to $|E_{\text{circ}}|^2$, and with $P_{\text{in}} \propto |E_{\text{in}}|^2$ we obtain

$$P_{\text{circ}} = \frac{T_1}{|1 - \sqrt{R_1 R_2} e^{i\phi}|^2} P_{\text{in}} \overset{\phi=0}{=} \frac{T_1}{(1 - \sqrt{R_1 R_2})^2} P_{\text{in}} \tag{3.3}$$

where the second equality holds on resonance with $e^{i\phi} = 1$. The power buildup factor β is defined as the ratio of circulating power to input power on resonance:

$$\beta = \frac{P_{\text{circ}}^{\text{res}}}{P_{\text{in}}} = \frac{T_1}{(1 - \sqrt{R_1 R_2})^2}. \tag{3.4}$$

The finesse $\mathcal{F}$ is defined as the ratio of the free spectral range (FSR) to the cavity linewidth (FWHM). For a single-ended cavity with $T_2 = 0$ (high reflector), the total internal round-trip power loss is $L_{\text{tot}} = T_2 + L_1 + L_2$. In the high-finesse limit where $T_1, L_{\text{tot}} \ll 1$, we may use the approximation $\sqrt{R_1 R_2} \approx 1 - (T_1 + L_{\text{tot}})/2$, yielding

$$\mathcal{F} \approx \frac{2\pi}{T_1 + L_{\text{tot}}}. \tag{3.5}$$

With this approximation, the buildup factor simplifies to

$$\beta = \frac{4T_1}{(T_1 + L_{\text{tot}})^2}. \tag{3.6}$$

The on-resonance amplitude reflection coefficient of the cavity, measuring the interference between directly reflected light and light leaking back out through M_1, is

$$r_{\text{cav}} = \frac{T_1 - L_{\text{tot}}}{T_1 + L_{\text{tot}}}. \tag{3.7}$$

Setting $r_{\text{cav}} = 0$ yields the impedance matching condition $T_1 = L_{\text{tot}}$, where all incident light is coupled into the cavity and the buildup reaches its maximum value of $\beta_{\text{max}} = 1/L_{\text{tot}}$. This condition defines the critical coupling point, relative to which we distinguish

two regimes: the cavity is overcoupled when $T_1 > L_{\text{tot}}$ and undercoupled when $T_1 < L_{\text{tot}}$.

3.1.1 Undercoupled and Overcoupled Regimes

The distinction between undercoupled and overcoupled operation has important practical consequences for power buildup and sensitivity to loss variations.

In the undercoupled regime where $T_1 \ll L_{\text{tot}}$, the cavity is loss-dominated. Expanding eq. (3.6) in this limit gives

$$\beta_{\text{uc}} = \frac{4T_1}{L_{\text{tot}}^2(1 + T_1/L_{\text{tot}})^2} \approx \frac{4T_1}{L_{\text{tot}}^2} \qquad (T_1 \ll L_{\text{tot}}). \tag{3.8}$$

The buildup is linear in T_1 and drops quadratically with internal loss, making the undercoupled regime extremely sensitive to contamination or degradation of the mirrors.

In the overcoupled regime where $T_1 \gg L_{\text{tot}}$, the input coupler transmission dominates the total cavity decay. Expanding eq. (3.6) yields

$$\beta_{\text{oc}} = \frac{4}{T_1} \cdot \frac{1}{(1 + L_{\text{tot}}/T_1)^2} \approx \frac{4}{T_1}\left(1 - \frac{2L_{\text{tot}}}{T_1}\right) \qquad (T_1 \gg L_{\text{tot}}). \tag{3.9}$$

Here the buildup scales inversely with the design parameter T_1 and is only weakly sensitive to the internal loss L_{tot}.

The contrast between these regimes is most clearly seen in the fractional change of buildup for a given loss increment δL. In the undercoupled case, the fractional degradation is $\delta\beta/\beta \approx -2\delta L/L_{\text{tot}}$, which diverges as $L_{\text{tot}} \to 0$. In the overcoupled case, $\delta\beta/\beta \approx -2\delta L/T_1$, which is suppressed by a factor that we control through our choice of input coupler.

The buildup is maximized at critical coupling $T_1 = L_{\text{tot}}$, but this represents a knife-edge optimum. If one designs for critical coupling at a nominal loss L_0 but the actual loss drifts to $L_{\text{tot}} = L_0 + \delta L$ with $\delta L > 0$, the cavity enters the undercoupled regime where buildup degrades quadratically with the total loss. A cavity designed to be moderately overcoupled

that experiences the same loss drift instead merely approaches critical coupling from above; the buildup initially increases before eventually degrading.

To hedge against unforeseen loss degradation, one should design the input coupler to be critically coupled at the worst-case expected loss $T_1^{\text{hedge}} = L_{\text{tot}}^{\text{worst}}$, accepting reduced peak buildup in exchange for robustness.

3.2 Cavity Stability and Mode Spectrum

Apart from the resonance condition discussed previously defining the longitudinal modes of a cavity, the field inside the resonator also has a transverse profile. In the simplest case, this is a Gaussian profile in the transverse x and y coordinates forming the *fundamental* mode and the size of the Gaussian is given by the cavity waist ω_0 in a reference plane. In this section, we will develop the theoretical formalism to calculate these fundamental properties such as stability, mode waist and spectrum of higher order modes. Starting from a simple ray picture in the regime of paraxial optics, we will determine what makes a cavity stable and we can use simple transfer matrices (ABCD matrices) to calculate its modes and spectrum. More formally, we will explore the connection to diffraction and full solutions of the wave equation.

To rigorously compute the transverse mode profiles of a cavity, we need to solve Maxwell's equations of free space with boundary conditions defined on the mirrors that confine the light field. Ignoring polarization and only considering a scalar field, we can use the wave equation $(\nabla^2 + k^2)\mathbf{E}(x, y, z) = 0$ to solve for the transverse field profiles. As for macroscopic Fabry-Perot type cavities the mirror spacing L is far greater than the optical wavelength λ, we can factor out the fast oscillations e^{ikz} and apply the *slowly-varying envelope* approximation given that the transverse size changes much slower than λ. These two approximations lead to the *paraxial* wave equation

$$2ik\partial_z u(\mathbf{x}) = \nabla_\perp^2 u(\mathbf{x}) \tag{3.10}$$

where $\nabla_\perp = (\partial_x, \partial_y)$ is the transverse derivative. For a stable cavity, the local slope of the field has to be perpendicular to the mirror surface. In the paraxial limit, this means the curvature of the Gaussian beam wavefront needs to match the mirror's curvature.

3.2.1 Full solution to Helmholtz eq

Gaussian beam

$$\mathbf{E}(r,z) = E_0 \times \frac{w_0}{w(z)} \exp\left(\frac{-r^2}{w(z)^2}\right) \exp\left(-i\left(kz + k\frac{r^2}{2R(z)} - \psi(z)\right)\right) \tag{3.11}$$

with waist size $w(z) = w_0\sqrt{1+\left(\frac{z}{z_R}\right)^2}$, Rayleigh range $z_R = \frac{\pi\, w_0^2}{\lambda}$, radius of curvature $R(z) = z\left[1+\left(\frac{z_R}{z}\right)^2\right]$, Gouy phase $\psi(z) = \arctan\left(\frac{z}{z_R}\right)$. All of these properties can be combined into a single complex number, the so-called q parameter $q(z) = z + i\, z_R$ or inversely $\frac{1}{q(z)} = \frac{1}{R(z)} - i\,\frac{\lambda}{\pi\, w(z)^2}$ in terms of physical observables.

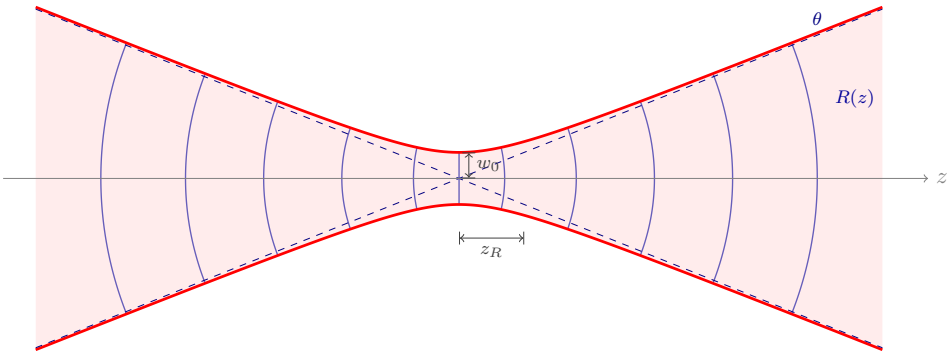


Figure 3.1: Gaussian beam propagation showing the beam envelope $w(z)$, wavefronts with radius of curvature $R(z)$, beam waist w_0, Rayleigh range z_R, and far-field divergence angle θ.

The set of transverse modes form an orthonormal basis of the transverse 2d space, and depending on the symmetry of the cavity two bases are natural choices. In cartesian coordinates, eq. (3.10) can be solved through separation of variables leading to the so-called

Hermite-Gaussian modes

$$E_{l,m}(x,y,z) = E_0 \frac{w_0}{w(z)} H_l\left(\frac{\sqrt{2}\,x}{w(z)}\right) H_m\left(\frac{\sqrt{2}\,y}{w(z)}\right) \exp\left(-\frac{x^2+y^2}{w^2(z)}\right) \times \exp\left(-i\,\frac{k(x^2+y^2)}{2R(z)}\right) \exp\big(i\psi(z)\big) \exp(-ikz) \quad (3.12)$$

in terms of Hermite polynomials $H_n(x)$.

LG

$$u(r,\phi,z) = C_{lp}^{LG} \frac{1}{w(z)} \left(\frac{r\sqrt{2}}{w(z)}\right)^{|l|} \exp\left(-\frac{r^2}{w^2(z)}\right) L_p^{|l|}\left(\frac{2r^2}{w^2(z)}\right) \times \exp\left(-ik\frac{r^2}{2R(z)}\right) \exp(-il\phi)\ \exp(i\psi(z)), \quad (3.13)$$

3.2.2 ABCD matrices

We begin with an overview of the ABCD matrix formalism, which is perfectly suited to describe cavities comprised of flat or curved mirrors, lenses and propagation in between them. These ABCD or ray transfer matrices describe linear transformations of the rays position and slope caused by paraxial (i.e. quadratic in x or p) optics and this linear representation enables us to use all the tools of linear algebra. As ABCD matrices always describe paraxial optics, one of the fundamental limitations is that we can only describe rays close to the optical axis where $\sin\theta \approx \theta$. By convention, we usually denote the z-coordinate as the optical axis and describe the dynamics in the two transverse coordinates x and y. Consider a ray $\boldsymbol{v} = \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{p} \end{pmatrix}$ where $\boldsymbol{x}$ is the ray's position and $\boldsymbol{s}$ is the slope, where either can be taken to be either one- or two-dimensional, leading to 2×2 or 4×4 ABCD matrices.

For simplicity or symmetric cavities that are identical in the sagittal and tangential plane, we start our discussion with simple 2×2 matrices. As long as there is no rotation

or Astigmatism in the system that couples the two axes, the full 4×4 matrices have block form and the x and y axis can be treated independently. The ABCD matrices for common cavity optics are and they transform the input ray to an output ray $\begin{pmatrix} \boldsymbol{x'} \\ \boldsymbol{p'} \end{pmatrix} = M \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{p} \end{pmatrix}$

Table 3.1: Common ABCD ray transfer matrices for paraxial optics.

Element	Matrix	Remark
Free propagation	$\begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix}$	d: propagation distance
Thin lens	$\begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}$	f: focal length ($f > 0$ converging)
Curved mirror	$\begin{pmatrix} 1 & 0 \\ -2/R & 1 \end{pmatrix}$	R: radius of curvature ($R > 0$ concave)
Dielectric interface	$\begin{pmatrix} 1 & 0 \\ 0 & n_1/n_2 \end{pmatrix}$	n_1, n_2: refractive indices before/after
Curved interface	$\begin{pmatrix} 1 & 0 \\ \frac{n_1-n_2}{n_2 R} & \frac{n_1}{n_2} \end{pmatrix}$	R: radius of curvature, $n_1 \to n_2$

through a matrix vector product. Optical systems are then comprised by a series of matrices, representing the optical element and the free propagation in between, where the matrices are applied right to left. Closed systems such as cavities can be described in the same way by unrolling a single round trip through the cavity. The only important choice is the choice of a reference plane in the cavity that the rays will return back to after a round trip, and while arbitrary, denotes the location where waist etc. are calculated. The first question that the round-trip ABCD matrix M_{rt} of a cavity allows us to answer is whether the cavity is *stable*: While for a simple Fabry-Perot two-mirror cavity the stability can be deduced from matching the curvature of the mirrors to the curvature of the Gaussian beam wavefront, in more complicated cavities stability can sensitively depend on any distance.

For a cavity to be stable, it needs to support a ray that returns back to itself after a

round-trip, a so called *eigen-ray* $\boldsymbol{\mu}_j^{\pm}$ where

$$M_{rt}\boldsymbol{\mu}_j^{\pm} = \nu_j^{\pm}\boldsymbol{\mu}_j^{\pm} \tag{3.14}$$

and $\nu_j^{\pm}$ are the corresponding eigenvalues. An arbitrary input ray can then be decomposed into these eigenvectors $\boldsymbol{r}_0 = a_1\boldsymbol{\mu}_1 + a_2\boldsymbol{\mu}_2$ and after n roundtrips is transformed to $\boldsymbol{r}_n = M_{rt}^n\boldsymbol{r}_0 = a_1\nu_1^n\boldsymbol{\mu}_1 + a_2\nu_2^n\boldsymbol{\mu}_2$. For the ray to not exponentially grow to infinity and leave the cavity, or decay to the trivial null solution, the absolute value of the complex eigenvalues has to be 1. Because the matrix M_{rt} is real, the two eigenvectors and eigenvalues must be each others complex conjugate and we can write $\nu_j^{\pm} = exp(\pm i\chi_j)$. The resulting phase χ_j is the Gouy phase of the cavity for the transverse dimension j and determines the whole spectrum of the cavity in the paraxial limit. For a mode u_{qmn} with longitudinal mode index q and transverse mode indices n, m, the resonant frequency is

$$\nu_{qmn} = \nu_{FSR} \times \left(q + m\frac{\chi 1}{2\pi} + n\frac{\chi 2}{2\pi} + \frac{1}{\pi}\right) \tag{3.15}$$

with ν_{FSR} the free spectral range of the resonator. This is what we call the spectrum of a cavity, and the higher order modes are linearly spaced with the separation determined by the Gouy phases. We can see that for a symmetric cavity $\chi_1 = \chi_2$ and all modes with a total mode number $N = n + m$ will be degenerate such as the $HG_{0,2}$, $HG_{1,1}$ and $HG_{2,0}$ mode.

3.2.3 operator formalism

So far we have treated the ABCD matrices for one and two spatial dimensions on equal footing, but now we'll focus on 4×4 ABCD matrices that will be necessary to model twisted cavities and reveal more about the underlying structure. As we will later see through the connection to Hamiltonian optics in much greater detail, the underlying mathematical

structure of the ABCD transfer matrices is the *symplectic group* $\mathrm{Sp}(2n, \mathbb{R})$ for n transverse dimensions. The 4×4 matrices for two transverse dimensions are therefore elements of $M \in \mathrm{Sp}(4, \mathbb{R})$ which encodes a special structure

$$M^T G M = G \tag{3.16}$$

where G is the antisymmetric *symplectic matrix* $G = \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix}$. All matrices M satisfying this condition form a group that describes all possible lossless, paraxial systems. Also, in two transverse dimensions the determinant of M is always equal to one which will become useful later. By inserting the eigenvectors $\boldsymbol{\mu}_j^{\pm}$ of equation eq. (3.14) into the symplectic condition eq. (3.16) we get

$$\nu_i \nu_j \boldsymbol{\mu}_i G \boldsymbol{\mu}_j = \boldsymbol{\mu}_i G \boldsymbol{\mu}_j \tag{3.17}$$

and can observe two important facts: For $i = j$ the value of $\boldsymbol{\mu}_i G \boldsymbol{\mu}_i = 0$ exactly vanishes by construction, however this is not true for $\bar{\boldsymbol{\mu}}_i G \boldsymbol{\mu}_i$ which can have an imaginary part which we use to normalize the eigenvectors such that $\bar{\boldsymbol{\mu}}_1 G \boldsymbol{\mu}_1 = \bar{\boldsymbol{\mu}}_2 G \boldsymbol{\mu}_2 = 2i$. With this, we see that the condition eq. (3.17) can only hold for $\bar{\nu}_i \nu_i = 1$ which means the complex eigenvalues must have an absolute value $|\nu_i| = 1$ of one, and the eigenvalues come in conjugate pairs $\nu_j^{\pm} = (\nu_i, \bar{\nu}_i) = e^{\pm i \chi_j}$ and associated eigenvectors $\boldsymbol{\mu}_j^{\pm}$. In the following, we can now only work with one representative eigenvector per pair $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_1^+, \boldsymbol{\mu}_2 = \boldsymbol{\mu}_2^+$ and we can again decompose a general ray $\boldsymbol{v}$ in two dimensions as $\boldsymbol{v} = 2\,\mathrm{Re}(a_1 \boldsymbol{\mu}_1 + a_2 \boldsymbol{\mu}_2)$ with complex coefficients a_i. These can be calculated using the symplectic orthogonality as $a_i = \frac{\bar{\boldsymbol{\mu}}_i G \boldsymbol{v}}{2i}$ for a given input ray. These equations will later allow us to construct rays that correspond to one of the eigen-axes of the cavity, such as cartesian indices n, m of a HG mode or radial and angular mode index p, l of a LG mode. The two resulting Gouy phases χ_1, χ_2 now represent the phase accumulated on every round trip for each axis of the mode in a diagonal frame, and it predicts the mode spectrum according to eq. (3.15). When any astigmatism or twist in the

cavity couples the two transverse dimensions, the dynamics of a ray no longer factorize along the original axes but can trace out much more complicated patterns. In the case that the Gouy phases take fractional values $\chi_1 = 2\pi\frac{p}{S}, \chi_2 = 2\pi\frac{q}{S}$ with common denominator S, the ray will come back to itself after S round trips and is then called S-fold degenerate. So far, we have seen how the eigenvalues of classical ray transfer matrices can provide us with the Gouy phase of the fundamental mode, which is a fundamentally a property of optical wave propagation. This treatment allows us to extract many important cavity properties such as the mode waist or spectrum eq. (3.15) that are sufficient for e.g. the design of laser resonators where one is only interested in the fundamental mode. As we want to use the higher order modes of cavity as orbitals of a quantum system, we need to develop an understanding and formalism for treating also these higher order transverse modes of the cavity. To develop an operator formalism that allows us to construct the form of higher order modes algebraically, we can first notice the formal analogy between the paraxial wave equation eq. (3.10) in the form

$$2ik\partial_z u(\boldsymbol{r}, z) = -\nabla_\perp^2 u(\boldsymbol{r}, z) \tag{3.18}$$

for a light field $u(\boldsymbol{r}) = u(x, z)$ and the Schrödinger equation for a single particle

$$i\hbar\partial_t \psi(\boldsymbol{r}, t) = -\frac{\hbar^2}{2m}\nabla_\perp^2 \psi(\boldsymbol{r}, t) \tag{3.19}$$

with wavefunction $\psi(\boldsymbol{r}, t)$ with two spatial dimensions evolving in time t. While the Schrödinger equation usually describes continuous evolution in some external potential $V(\boldsymbol{r}, t)$, the propagation of light in free space according to the paraxial wave equation only corresponds to the kinetic term, whereas potentials in real space are applied *stroboscopically* when the light is reflected or refracted by a mirror or lens. The deeper meaning behind this analogy will be discussed in greater depth later, and might just be a lucky coincidence of nature, however the mathematical analogy is rigorous and we can use much of the machinery developed for

the treatment of quantum mechanics, namely operators and Dirac notation.

The light field $u(\boldsymbol{r}, z)$ at some position z along the cavity can then be promoted to a state vector

$$u(\boldsymbol{r}, z) = \langle \boldsymbol{r} | u(z) \rangle \tag{3.20}$$

and position and momentum are accordingly quantized to be position operator $\hat{\boldsymbol{r}} = (\hat{x}, \hat{y})$ and momentum operator $\hat{\boldsymbol{p}} = (\hat{p_x}, \hat{p_y}) = -i(\partial_x, \partial_y)$.

The quantization we have performed here happened somewhat *ad-hoc*, but it provides the fundamental bridge from the ray into the wave picture we need to connect our transfer matrices M to unitary operators $\hat{U}$ acting on the wavefunction. While special cases of free propagation and lenses can be checked explicitly, we want to motivate the general case here by analogy and later provide a more thorough derivation. In classical mechanics, Ehrenfest's theorem states that the expectation values of quantum operators follow classical trajectories (for linear forces i.e. quadratic Hamiltonians). A classical ray is a vector $\boldsymbol{v} = (x, y, s_x, s_y)^T$ that evolves via a matrix multiplication $\boldsymbol{v}_{out} = M\boldsymbol{v}_{in}$. in the wave picture, a field $|\psi\rangle$ evolves via unitary operators $U_{rt} = \exp(-i\hat{H}_{eff})$ for some time-independent effective Hamiltonian $\hat{H}_{eff}$. To connect both regimes, we require that the center of the wave packet (the expectation value $\langle \hat{\boldsymbol{v}} \rangle$) follows the classical ray path

$$\langle \psi_{out} | \hat{\boldsymbol{v}} | \psi_{out} \rangle = \langle \psi_{in} | U_{rt}^{\dagger} \hat{\boldsymbol{v}} U_{rt} | \psi_{in} \rangle = M \langle \psi_{in} | \hat{\boldsymbol{v}} | \psi_{in} \rangle \tag{3.21}$$

which implies the central operator identity bridging the ray to the wave picture

$$U_{rt}^{\dagger} \hat{\nu} U_{rt} = M \hat{\nu}. \tag{3.22}$$

This resembles exactly the evolution of operators in the Heisenberg picture for position $U^{\dagger}\hat{x}U$ and momentum $U^{\dagger}\hat{p}U$.

Next, we need to find a "lowering" (annihilation) operator $\hat{a}$ that is a linear combination

of the phase space operators $\hat{\nu} = \left(\hat{x}, \hat{y}, \hat{p}_x, \hat{p}_y\right)^T$. We can make an ansatz $\hat{a} = \mathbf{w}^T\hat{\nu} = w_j\hat{\nu}_j$ where $\mathbf{w}$ is a 4-element column vector of complex coefficients. For $\hat{a}$ to be a true eigen-operator of the cavity, it must satisfy the Heisenberg equation of motion upon a round trip. If the classical ray acquires a Gouy phase χ per round trip, the quantum operator must acquire the same phase factor

$$\hat{U}_{rt}\hat{a}\hat{U}_{rt}^\dagger = e^{-i\chi}\hat{a} \tag{3.23}$$

where $\hat{U}_{rt}$ is the unitary round-trip evolution operator. We can then insert this ansatz into eq. (3.22) and get $\hat{U}_{rt}(\mathbf{w}^T\hat{\nu})\hat{U}_{rt}^\dagger = \mathbf{w}^T(\hat{U}_{rt}\hat{\nu}\hat{U}_{rt}^\dagger) = \mathbf{w}^T M\hat{\nu} = e^{-i\phi}(\mathbf{w}^T\hat{\nu})$. This shows that $\mathbf{w}$ needs to be a left eigenvector $\mathbf{w}^T M_{rt} = e^{-i\phi}\mathbf{w}^T$ of the ABCD round trip matrix M_{rt}. We can relate this to the more common right eigenvectors $M\mathbf{u} = e^{-i\phi}\mathbf{u}$ we take the Hermitian conjugate of the right eigenvector equation (noting that M is real, so $M^\dagger = M^T$): $\mathbf{u}^\dagger M^T = e^{i\phi}\mathbf{u}^\dagger$. By using the defining property of the Symplectic Group $Sp(4,\mathbb{R})$ $M^T G M = G \rightarrow M^T G = GM^{-1}$ we can rearrange this to $(\mathbf{u}^\dagger G)M = e^{-i\phi}(\mathbf{u}^\dagger G)$ and by comparing this with the ansatz result $\mathbf{w}^T M = e^{-i\phi}\mathbf{w}^T$, we see that the coefficient vector $\mathbf{w}^T$ is proportional to the symplectic product of the right eigenvector:

$$\mathbf{w}^T \propto \mathbf{u}^\dagger G \tag{3.24}$$

We can find the proper normalization by calculating the commutator

$$\left[\hat{a}, \hat{a}^\dagger\right] = |N|^2 \sum_{j,k,l,m} (u_j^* G_{jk})(-u_m G_{lm}) \left[\hat{\nu}_k, \hat{\nu}_l\right] \tag{3.25}$$

Substitute $[\hat{\nu}_k, \hat{\nu}_l] = i\lambda G_{kl}$:

$$1 = -i\lambda|N|^2(\mathbf{u}^\dagger G)GG^T\mathbf{u} \tag{3.26}$$

Using $GG^T = I$ and $G^2 = -I$:

$$1 = -i\lambda|N|^2(\mathbf{u}^\dagger G\mathbf{u}) \tag{3.27}$$

Solving for N (assuming the normalization convention $\mathbf{u}^\dagger G\mathbf{u} = i \cdot \text{const}$ to ensure N is real):

$$N = \frac{1}{\sqrt{-i\lambda(\mathbf{u}^\dagger G\mathbf{u})}} \tag{3.28}$$

$$\hat{a} = \frac{1}{\sqrt{-i\lambda(\mathbf{u}^\dagger G\mathbf{u})}}\mathbf{u}^\dagger G\hat{\nu} \tag{3.29}$$

eigenmodes of free space maintain their shape eigenstates of HO want ladder operators

3.2.4 Hamiltonian Optics

Evolution is described by a transverse Hamiltonian

$$h(\boldsymbol{p}, \boldsymbol{q}) := -\sqrt{n(\boldsymbol{q})^2 - \boldsymbol{p}^2} = -p_z \tag{3.30}$$

where $\boldsymbol{p}$ and $\boldsymbol{q}$ (taken to be either 1- or 2-dimensional) follow the Hamilton equations

$$\frac{dq_i}{dz} = \frac{\partial h}{\partial p_i}, \qquad \frac{dp_i}{dz} = \frac{\partial h}{\partial q_i}. \tag{3.31}$$

The variables q_i and p_j are canonically conjugate, such that $\{p_i, q_j\} = \delta_{ij}$, and $\{p_i, p_j\} = \{q_i, q_j\} = 0$. The momentum $\boldsymbol{p}$ is related to the ray slope $\boldsymbol{s}$ and the angle θ between the ray and the optical axis by $\boldsymbol{p} = n \sin\theta$. The ray slope $\boldsymbol{s}$ then just expresses θ and its direction in a local coordinate system.

A transverse ray $\boldsymbol{v} = \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix}$ can be viewed as a point in phase space for transverse position $\boldsymbol{q}$ and canonically conjugate transverse momentum $\boldsymbol{p}$. The momentum is typically related to ray slope by $\boldsymbol{p} = n\boldsymbol{s}$, where n is the local index of refraction.

An operator G_f transforms such a ray by

$$G_f : \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix} \to \begin{pmatrix} \boldsymbol{p'} \\ \boldsymbol{q'} \end{pmatrix} \tag{3.32}$$

where $f = f(\boldsymbol{p}, \boldsymbol{q})$ is a polynomial, and the operator G_f is constructed by

$$G_f := \exp\Big(\hat{f}(\boldsymbol{p}, \boldsymbol{q})\Big). \tag{3.33}$$

This exponential structure highlights the fundamental relation of the Symplectic group $\mathrm{Sp}(4, \mathbb{R})$ to its underlying Lie algebra $\mathfrak{sp}(4, \mathbb{R})$ where the polynomials of operators $\hat{f}$ are the generators of group elements G_f.

The hat signifies the operator $\hat{f}$ that is associated with a polynomial f, and it acts by applying the Poisson bracket operator, such that

$$\hat{f}(\boldsymbol{p}, \boldsymbol{q}) := \{f(\boldsymbol{p}, \boldsymbol{q}), \cdot\} = \sum_i \frac{\partial f}{\partial q_i}\frac{\partial}{\partial p_i} - \frac{\partial f}{\partial p_i}\frac{\partial}{\partial q_i} \tag{3.34}$$

The exponential power of a hatted operator is symbolic and practically evaluated as

$$\exp(\hat{f}) := 1 + \{f, \cdot\} + \frac{1}{2!}\{f, \{f, \cdot\}\} + \ldots, \tag{3.35}$$

i.e., a formal expansion of the usual power series of exponentiation. Powers of $\hat{f}$ represent nested Poisson brackets $\{f, \circ\}^m = \{f, \{f, \ldots, \{f, \circ\}\}\}$ (where there are m nested Poisson brackets).

So far this structure looks exactly like the Hamiltonian formalism in classical mechanics, only the Hamiltonian for light $h(\boldsymbol{p}, \boldsymbol{q}) := -\sqrt{n(\boldsymbol{q})^2 - \boldsymbol{p}^2} = -p_z$ has a slightly unusual form. After introducing all this complicated machinery, we want to at least see that we can recover the earlier result of the ABCD matrix formalism. As we have previously said, ABCD matrices represent the linear transformations of phase space caused by quadratic optical elements.

For example, taking the polynomial $f_{\text{prop}} = -\frac{z}{2n}p^2$ for free space propagation this generates the transformation

$$G_{g_{\text{prop}}} : \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix} = e^{\hat{g}_{\text{prop}}} : \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} + \frac{z}{n}\boldsymbol{p} \end{pmatrix},$$

where we used and exactly recover the result of the ABCD matrix.

Similarly, $g_{\text{lens}} = -\frac{1}{2f}q^2$ generates a transformation

$$G_{g_{\text{lens}}} : \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix} = e^{\hat{g}_{\text{lens}}} : \begin{pmatrix} \boldsymbol{p} \\ \boldsymbol{q} \end{pmatrix} = \begin{pmatrix} \boldsymbol{p} - \frac{1}{f}\boldsymbol{q} \\ \boldsymbol{q} \end{pmatrix}, \tag{3.36}$$

which represents the action of an ideal thin lens with focal length f.

Sp(4) -¿ Mpl(4)

3.2.5 Perturbation theory

Fractional fourier transform Aberration "rotation" To benchmark our perturbation theory and the validity of the phase space polynomials, we model the previous twisted, mirror based cavity in fig. 3.2 and fig. 3.3.

Using the operator algebra we are also able to compute mode images by solving $a\,|\psi\rangle = 0$ for the fundamental mode and then applying $\hat{a}^\dagger$ to compute higher order modes algebraically. We compare our perturbation theory to paraxial and experimentally mode pictures in fig. 3.3.

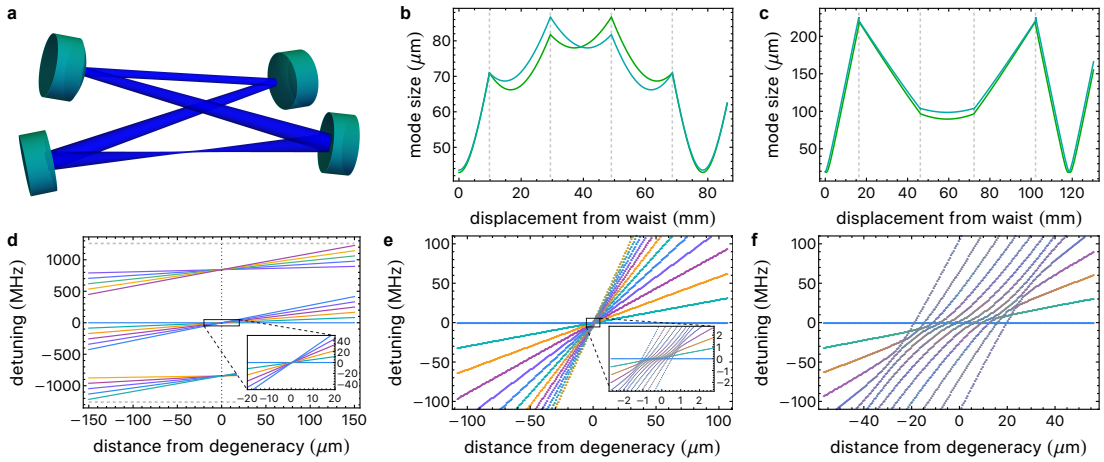


Figure 3.2: Understanding aberrations in the previous generation mirror-based twisted resonator. In order to understand the aberrations caused by non-normal incidents on highly curved mirrors in the previous twisted cavity apparatus, we model the effect of aberrations on the mode spectrum in perturbation theory. In **(b)** and **(c)**, the mode waist along the optical axis is shown for two separate cavities that were used to demonstrate the Landau level for photons [SRG+16]. Within the paraxial approximation **(d)**, we would expect all LG_l modes $l + ns$ to be perfectly degenerate at a single mirror distance for an s-fold degeneracy. Instead, higher-order optical aberrations, such as astigmatism or cubic aberrations, lead to mode coupling and result in a complex structure of avoided crossings **(e, f)** that cause an energy spread and mixing of modes close to degeneracy. The effect of aberrations scales strongly $(\lambda/\omega_0)^3$ with mode waist and exacerbates the spread of the avoided crossings due to the smaller waist of $\omega_0 = 43$ um in **(e)** to $\omega_0 = 19$ um in **(f)**.

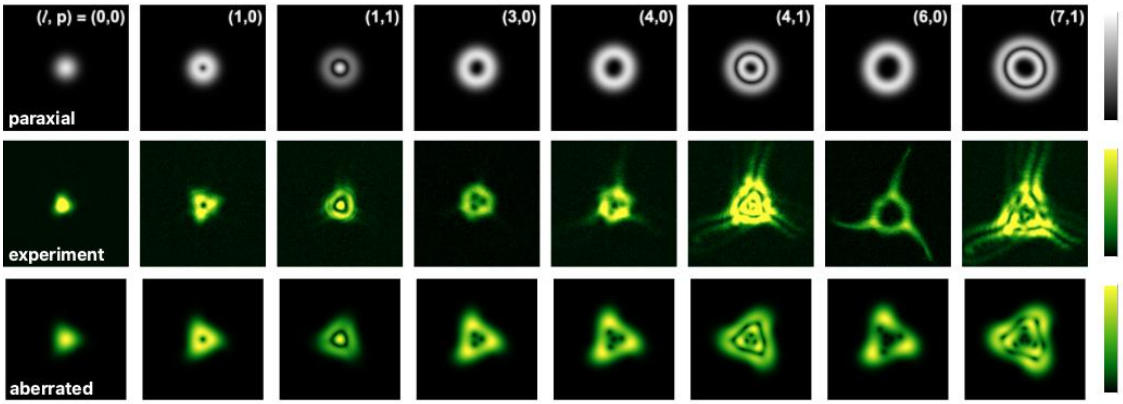


Figure 3.3: Aberrated mirror cavity mode images compared to theory. Mode profiles of the original mirror-based twisted cavity near the $s = 3$ degeneracy. Mode profiles for different LG modes are shown in the paraxial approximation (top row), experimentally measured mode intensities (middle row), and predicted from perturbation theory (bottom row).

In the paraxial regime (top row), one expects perfect LG eigenmodes of the cavity; however, in the experiment (middle row), aberrations from non-normal incidence on the curved mirrors cause strong mode mixing. At the three-fold degeneracy $s = 3$ where astigmatism is spectroscopically suppressed, the primary terms are cubic aberrations that lead to the triangular symmetry of the resulting eigenmodes. As the long tails of the mode images suggest, the mode coupling introduces photon loss towards higher and higher l and makes higher order modes increasingly unstable. The bottom row shows predicted mode profiles from our perturbation theory calculation up to fourth order, capturing the predominant cubic aberration and mode shapes. As it can only be expanded to fourth order, however, it fails to capture the tails involving higher order couplings.

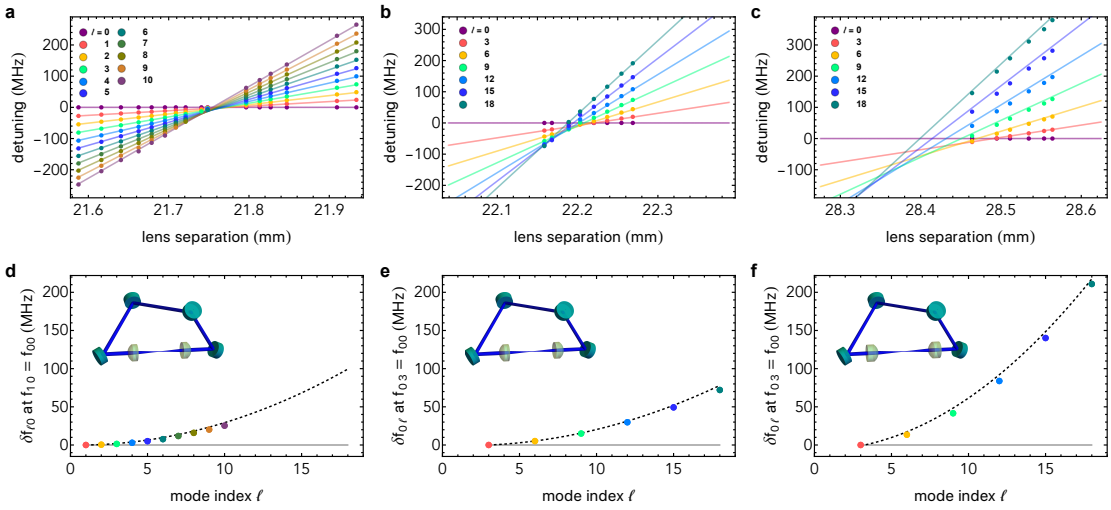


Figure 3.4: Mode dispersion of the lens cavity. Measured mode spectra and energy splittings of a spherical lens-based cavity at different degeneracy points and lens configurations. We tune the transverse mode splitting between different LG modes by fine-tuning the lens separation and thereby Gouy phase. In **(a, b, c)**, we measure the energy of different LG modes at various separations and extract their energy shift through a fit to a straight line. The three columns correspond to a flat $s = 1$ degeneracy, threefold $s = 3$ cone degeneracy, and $s = 3$ degeneracy with the orientation of the lenses reversed. In **(d-f)** we show the dispersion, i.e., the energy splitting between neighboring modes for the lens separation where the fundamental and first excited mode are degenerate. Paraxially, we would expect no mode shift (black line), but we observe a shift in the relative splitting quadratic with mode index $\propto l^2$, perfectly predicted by our parameter-free perturbation theory (dashed line).

Limitations

3.3 Ray tracing

As we have previously explored, there are deep connections between the ray and wave picture of optics. Previously, we have employed the toolkit of operator algebra to derive analytical mode couplings and shifts in a perturbation theory where we ultimately run into an ordering problem of higher order commutators. In this section, we will explore a different avenue that treats the phase space of Hamiltonian optics classically through ray tracing. Classical ray tracing is performed by acting on a ray, a phase space vector by applying the laws of reflection and refraction at every surface or interface. We represent a ray as

$$\mathbf{r} = [\mathbf{x}, \mathbf{k}, \mathrm{OPL}, \lambda_{\mathrm{idx}}, n, \mathbf{p}, E_s, E_p] \tag{3.37}$$

where $\mathbf{x} \in \mathbb{R}^3$ is position, $\mathbf{k} \in \mathbb{R}^3$ is direction (normalized), OPL is optical path length, n is current refractive index, $\mathbf{p}$ is the polarization reference vector, and E_s, E_p are complex Jones vector components. A set of input rays is then propagated through the whole cavity system by concatenating all interfaces

$$\mathbf{r}_{\mathrm{out}} = T_N \circ T_{N-1} \circ \cdots \circ T_1(\mathbf{r}_{\mathrm{in}}; \boldsymbol{\theta}) \tag{3.38}$$

where each T_i represents an optical interface transformation (refraction, reflection, or free-space propagation). The exact update rules

$$\begin{aligned}
\mathbf{x}' &= \mathbf{x}_{\text{intersection}}(\mathbf{x}, \mathbf{k}) \\
\mathbf{k}' &= \mathbf{k}_{\text{refr/refl}}(\mathbf{k}, \mathbf{n}_{\text{surf}}, n_{\text{in}}, n_{\text{out}}) \\
\text{OPL}' &= \text{OPL} + n\|\mathbf{x}' - \mathbf{x}\| \\
\mathbf{p}' &= \mathbf{p}_{\text{basis}}(\mathbf{k}', \mathbf{n}_{\text{surf}}) \\
[E_s', E_p']^T &= \mathbf{M}_{\text{Jones}}(\mathbf{k}, \mathbf{n}_{\text{surf}}, n_{\text{in}}, n_{\text{out}})[E_s, E_p]^T
\end{aligned} \tag{3.39}$$

are implemented analytically for flat or spherical surfaces and intersections with free form aspheres of high degree are solved numerically. While the phase space action of a cavity round trip eq. (3.38) looks like a product of ABCD matrices at first glance, it actually is a series of non-linear functions that describe the full solution to the Eikonal equation without any paraxial or small angle approximation. Instead of computing eigenmodes of the cavity, we can now ask what its *eigen-rays* are which are rays that return back to themselves after s round trips for a s-fold degenerate cavity. Finding the eigenray is solving for the fixed-point condition

$$\mathbf{r}_{\text{eig}} = T(\mathbf{r}_{\text{eig}}; \boldsymbol{\theta}) \tag{3.40}$$

that can be numerically solved as a non-linear root finding problem. This search is greatly aided by a gradient that we can obtain through the Implicit Function Theorem:

Differentiating $\mathbf{F}(\mathbf{r}_{\text{eig}}, \boldsymbol{\theta}) = \mathbf{r}_{\text{eig}} - T(\mathbf{r}_{\text{eig}}; \boldsymbol{\theta}) = \mathbf{0}$ we get

$$\frac{\partial \mathbf{r}_{\text{eig}}}{\partial \boldsymbol{\theta}} = -\left(\mathbf{I} - \left.\frac{\partial T}{\partial \mathbf{r}}\right|_{\mathbf{r}_{\text{eig}}}\right)^{-1} \left.\frac{\partial T}{\partial \boldsymbol{\theta}}\right|_{\mathbf{r}_{\text{eig}}} \tag{3.41}$$

that we can obtain through the AD framework. The ray picture can be connected back to the paraxial ABCD formalism in two ways. From the paraxial eigenvectors $\boldsymbol{\mu}_j$ that we have

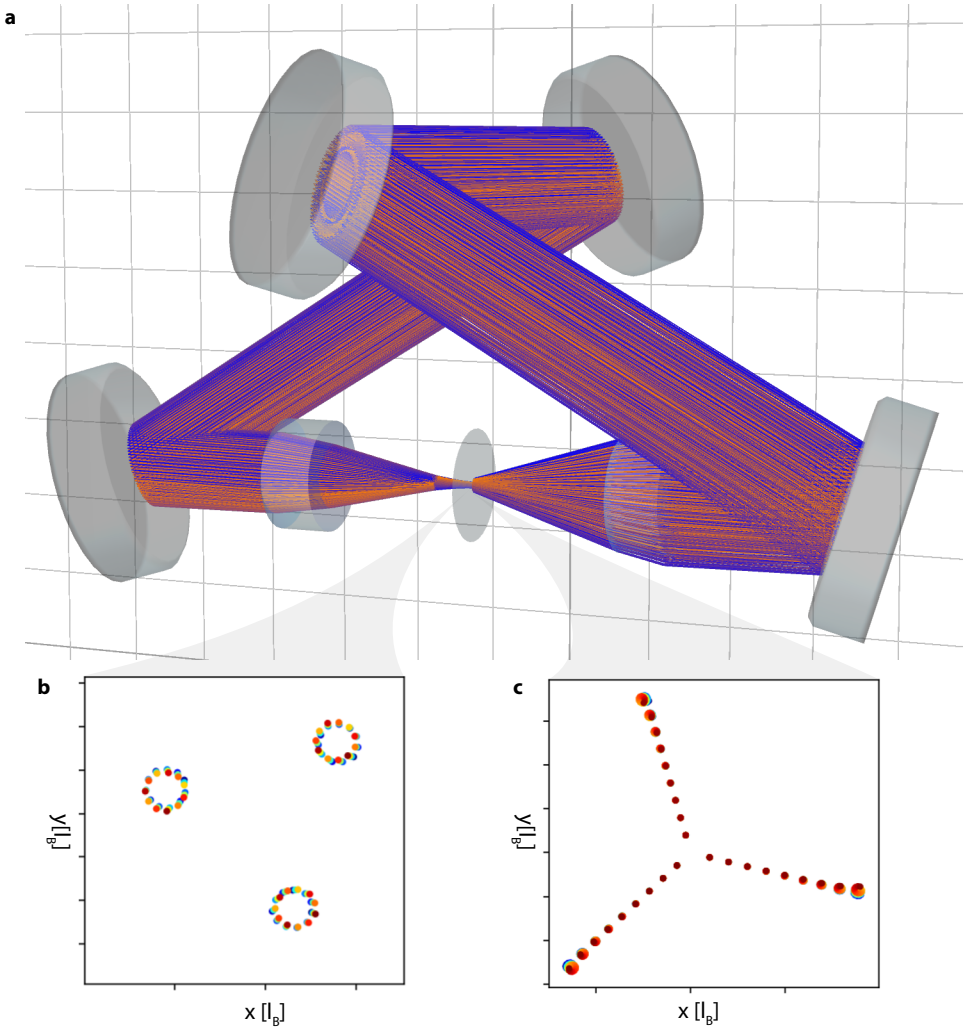


Figure 3.5: Ray tracing semiclassical orbits. Using our custom ray tracer, we can propagate an eigenray representative of a LG_l mode around the resonator thousands of times **(a)** according to the full phase space action of Hamiltonian optics. The resulting phase portrait shown in **(b,c)** in the (x, y) plane represents the classical ray orbits. Depending on which eigenvector $\boldsymbol{\mu}_j$ is chosen to initialize the ray, we can see the effect of both raising operators in the Landau level structure: In **(b)** cyclotron motion in the artificial magnetic field at the $s = 3$ degeneracy can be observed, corresponding to the superposition of different Landau levels related to the $\hat{b}^\dagger$ operator. Conversely, **(c)** shows slow magnetron precession as a function of ray radius $r = |(x, y)|$ resulting from non-degeneracy of the Landau level.

previously treated, we can initialize a ray

$$\boldsymbol{\mu}(n) = a \cdot w_0 \cdot \mathrm{Re}\left[\boldsymbol{\mu}_1 \cdot e^{in\phi_1}\right] \tag{3.42}$$

by mixing the eigenvectors with complex coefficients setting their amplitude and phase. This way, we study the nature of the eigenvectors by plotting their phase portrait in the reference plane over many round trips to visualize the classical evolution, and this way we can also compute robust starting points for the eigenray iteration.

For the reverse connection of the ray picture back the ABCD formalism, we can measure the effect of a small perturbation around an eigenray or any other guide ray to extract the *local* linear transformation matrix. To numerically extract the ABCD matrix relative to a base ray $\boldsymbol{\mu}_0$, we create perturbed "guide rays". Each guide ray is shifted by a small ϵ along one of the available degrees of freedom (spatial coordinates or slopes). We illustrate this here for a simple 1D system, a ray has 2 components, position x and slope s. To extract A and C, we send through a vector perturbed in position $(x+\epsilon, 0)^T$, and subtract the base evaluation at $(x, 0)^T$:

$$\Delta_{\text{out}}^{(x)} = \begin{pmatrix} A(x+\epsilon) \\ C(x+\epsilon) \end{pmatrix} - \begin{pmatrix} Ax \\ Cx \end{pmatrix} = \begin{pmatrix} A \\ C \end{pmatrix} \epsilon \tag{3.43}$$

Dividing by ϵ isolates the first column elements, A and C. Similarly, to extract B and D, we send through a vector perturbed in slope: $(0, s+\epsilon)^T$, and subtract the base evaluation at $(0, s)^T$:

$$\Delta_{\text{out}}^{(s)} = \begin{pmatrix} B(s+\epsilon) \\ D(s+\epsilon) \end{pmatrix} - \begin{pmatrix} Bs \\ Ds \end{pmatrix} = \begin{pmatrix} B \\ D \end{pmatrix} \epsilon \tag{3.44}$$

Dividing by ϵ isolates the second column elements, B and D.

This logic applies symmetrically to the 2D (4×4) case. Let $\hat{e}_j$ be the canonical basis

vector for the j-th phase-space coordinate. The change in the output ray after N_{rt} round trips is computed by perturbing each coordinate $j \in \{x, y, s_x, s_y\}$ independently:

$$\Delta_{\text{out}}^{(j)} = \boldsymbol{\mu}_{\text{perturbed}}(N_{\text{rt}}) - \boldsymbol{\mu}_0(N_{\text{rt}}) = \boldsymbol{\mu}(N_{\text{rt}})\Big|_{\boldsymbol{\mu}_0+\epsilon\cdot\hat{e}_j} - \boldsymbol{\mu}(N_{\text{rt}})\Big|_{\boldsymbol{\mu}_0} \tag{3.45}$$

The j-th column of the extracted ABCD matrix is then approximated via finite differences $\mathbf{M}_{:,j} \approx \frac{\Delta_{\text{out}}^{(j)}}{\epsilon}$.

In practice, finite differences suffer from floating point precision errors and depend on careful tuning of ϵ to reach convergence. Instead, taking the formal limes $\epsilon \to 0$ corresponds to evaluating the mathematical derivative where the output ray $\boldsymbol{\mu}_{\text{out}}$ becomes a strict, continuous function of the input ray $\boldsymbol{\mu}_{\text{in}}$:

$$\mathbf{M} = \lim_{\epsilon\to 0} \frac{\Delta_{\text{out}}}{\epsilon} = \frac{\partial \boldsymbol{\mu}_{\text{out}}}{\partial \boldsymbol{\mu}_{\text{in}}} \tag{3.46}$$

This local ABCD matrix represents the paraxial transformation around an arbitrary optical axis that can be very far from the cavity center and it can therefore be used to measure local Gouy phase, astigmatism and mode spacings which will become very useful in the analysis of cavity arrays.

In the twisted cavity we can use the dynamics of classical rays shown in fig. 3.5 and fig. 3.6 to study the effects of aberrations from semi-classical orbits. While the ray tracing is exact as it treats all optics and propagations without approximation, and it also able to capture the effect of Gouy phase through the dynamics of ray orbits, we need to be careful in relating classical rays to wave-optical modes. There is a deep connection between free space modes and families of rays as formalized in [DA19, Alo14, AD17] where we again make use of the Lie algebra connecting both pictures. As long as we are close to degeneracy and no mixing between modes, but only energy shifts of the different l modes due to even aberrations occurs, we can pick a representative ray out of the bundle representing LG_l and

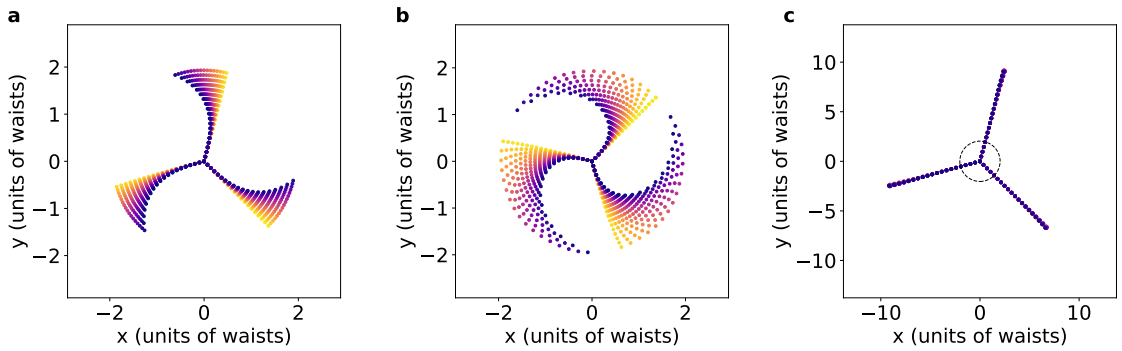


Figure 3.6: Semi-classical ray orbits in lens cavities. In order to study the effects of even order lens aberrations on the lowest Landau level degeneracy, we perform ray tracing of semi-classical orbits to determine the correct asphere coefficients to high order. In **(a)** and **(b)** we show the phase portrait at the s = 3 degeneracy for spherical lenses **(a)** and reversed spherical lenses **(b)** to exacerbate the effect of aberrations. In **(c)**, the magnetron motion is fully compensated by choosing asphere coefficients up to c_8, allowing for a five times greater degenerate radius.

study its dynamics. The radial maximum of the field of an LG_l mode occurs at $r \propto \sqrt{l}$ according to the Landau level wavefunctions eq. (3.13) and we therefore find the eigenray at a radius $r_l = \omega_0\sqrt{l}$ to represent every LG_l mode in fig. 3.6.

3.4 Cavity QED

3.5 Atom coupled to a field

When considering the coupling of an atom to the electromagnetic field of light, we are usually only concerned with its electron whose charge interacts with the field. We can introduce the coupling to a field through a vector potential $\boldsymbol{A}$ by the minimal coupling $\boldsymbol{p} \to \boldsymbol{p} - e\boldsymbol{A}$ which redefines the kinetic term $-i\hbar\nabla \to -i\hbar\left[\nabla - i\frac{e}{\hbar}\boldsymbol{A}\right]$. If we consider a two-level atom in a

single quantized field mode, the Hamiltonian including the dipole interaction is

$$H = H_{\text{atom}} + H_{\text{field}} + H_{\text{int}} \tag{3.47}$$

$$= \hbar\omega_e \left|e\right\rangle\!\left\langle e\right| + \hbar\nu(a^\dagger a) + \frac{1}{2} - e\boldsymbol{E}\boldsymbol{r} + h.c. \tag{3.48}$$

$$= \hbar\omega_e \left|e\right\rangle\!\left\langle e\right| + \hbar\nu(a^\dagger a) + \frac{1}{2} + \hbar\left(g\left|e\right\rangle\!\left\langle g\right| + g^*\left|g\right\rangle\!\left\langle e\right|\right)\left(a + a^\dagger\right) \tag{3.49}$$

in terms of coupling strength $g = -\frac{\boldsymbol{\epsilon}\boldsymbol{d}E}{\hbar}$ with the dipole moment $\boldsymbol{d}$, bosonic raising and lowering operators of the cavity field $a, a^\dagger$ and electric field strength E. Two of the terms in eq. (3.47) $\left|e\right\rangle\!\left\langle g\right| a^\dagger + h.c.$ violate energy conservation as they create (destroy) both a cavity photon and an atomic excitation, and they are neglected in the rotating wave approximation (RWA).

3.6 Light-matter coupling for a quantized field

When considering the coupling of an atom to the electromagnetic field of light, we are usually only concerned with its electron whose charge interacts with the field. We can introduce the coupling to a field through a vector potential $\boldsymbol{A}$ by the minimal coupling $\boldsymbol{p} \to \boldsymbol{p} - e\boldsymbol{A}$ which redefines the kinetic term $-i\hbar\nabla \to -i\hbar\left[\nabla - i\frac{e}{\hbar}\boldsymbol{A}\right]$. If we consider a two-level atom in a single quantized field mode, the Hamiltonian including the dipole interaction is

3.7 Photon Extraction Efficiency and Cooperativity

Beyond power buildup for classical light, a central figure of merit for cavity QED experiments is the photon extraction efficiency with which photons emitted by an atom inside the cavity can be detected. An atom located at the cavity waist can emit photons either into the cavity mode or into free space. For a single-ended cavity, we want photons emitted into the cavity mode to exit through mirror M_1 toward a detector. The total extraction efficiency factorizes

into two contributions:

$$\eta_{\text{total}} = \eta_{\text{cav}} \times \eta_{\text{out}} \tag{3.50}$$

where η_{cav} is the probability that the atom emits into the cavity mode rather than free space, and η_{out} is the probability that a cavity photon exits through the output coupler rather than being absorbed.

The atom-cavity interaction is characterized by the single-photon Rabi frequency g, the cavity field decay rate κ, and the atomic spontaneous emission rate Γ into free space. The cooperativity, defined as

$$C = \frac{4g^2}{\kappa\Gamma}, \tag{3.51}$$

measures the ratio of coherent atom-cavity coupling to dissipation. In the bad-cavity (Purcell) regime where $\kappa \gg g$, the cavity enhances the atomic emission rate into the cavity mode by the Purcell factor $F_P = C$, so the fraction of photons emitted into the cavity mode versus free space is

$$\eta_{\text{cav}} = \frac{F_P\Gamma}{F_P\Gamma + \Gamma} = \frac{C}{C+1}. \tag{3.52}$$

The output coupling efficiency follows from the branching ratio between useful and lossy decay channels. Once the photon is in the cavity mode, it decays through the output coupler at rate $\kappa_1 = T_1\nu_{\text{FSR}}$ or through loss channels at rate $\kappa_{\text{loss}} = L_{\text{tot}}\nu_{\text{FSR}}$, where $\nu_{\text{FSR}} = c/(2\ell)$ is the free spectral range, giving

$$\eta_{\text{out}} = \frac{\kappa_1}{\kappa_1 + \kappa_{\text{loss}}} = \frac{T_1}{T_1 + L_{\text{tot}}}. \tag{3.53}$$

The cooperativity depends on the cavity parameters through the total decay rate $\kappa = (T_1 + L_{\text{tot}})\nu_{\text{FSR}}$, yielding

$$C = \frac{4g^2}{\Gamma(T_1 + L_{\text{tot}})\nu_{\text{FSR}}}. \tag{3.54}$$

Increasing T_1 therefore decreases C through the increased κ, creating a direct tension: η_{out}

increases with T_1 while η_{cav} decreases with T_1 (via κ).

To optimize the total extraction, we define the intrinsic cooperativity $C_0 = 4g^2/(\Gamma L_{\text{tot}} \nu_{\text{FSR}})$ as the value that would obtain if $T_1 = 0$. The operating cooperativity at finite T_1 is then $C(T_1) = C_0/(1 + T_1/L_{\text{tot}})$. Introducing the dimensionless ratio $x = T_1/L_{\text{tot}}$, the total extraction efficiency becomes

$$\eta_{\text{total}}(x) = \frac{C_0}{C_0 + 1 + x} \cdot \frac{x}{1 + x}. \tag{3.55}$$

Differentiating with respect to x and setting the result to zero yields the optimal coupler ratio:

$$T_1^* = L_{\text{tot}} \sqrt{1 + C_0}. \tag{3.56}$$

In the weak-coupling limit where $C_0 \ll 1$, we have $T_1^* \approx L_{\text{tot}}$, recovering the impedance matching condition for power buildup, though $\eta_{\text{total}} \to 0$ because $\eta_{\text{cav}} \to 0$. In the strong-coupling limit where $C_0 \gg 1$, the optimal transmission becomes $T_1^* \approx L_{\text{tot}} \sqrt{C_0}$, placing the cavity deeply in the overcoupled regime with $\eta_{\text{cav}} \to 1$ and $\eta_{\text{out}} \approx 1 - 1/\sqrt{C_0}$. The maximum extraction efficiency at the optimal coupling is

$$\eta_{\text{total}}^{\max} = \frac{C_0}{(\sqrt{1 + C_0} + 1)^2} \xrightarrow{C_0 \gg 1} 1 - \frac{2}{\sqrt{C_0}} + \mathcal{O}(1/C_0). \tag{3.57}$$

The system becomes loss-dominated when $C_0 \lesssim 1$, corresponding to $L_{\text{tot}} \gtrsim 4g^2/(\Gamma \nu_{\text{FSR}})$. In this regime $\eta_{\text{total}} \sim C_0 \propto 1/L_{\text{tot}}$, and every additional increment of loss causes a proportional drop in photon extraction efficiency. Conversely, high-cooperativity systems with $C_0 \gg 1$ automatically operate in the overcoupled regime at optimal extraction and are naturally robust to loss degradation, with fractional efficiency loss suppressed as $1/\sqrt{C_0}$.

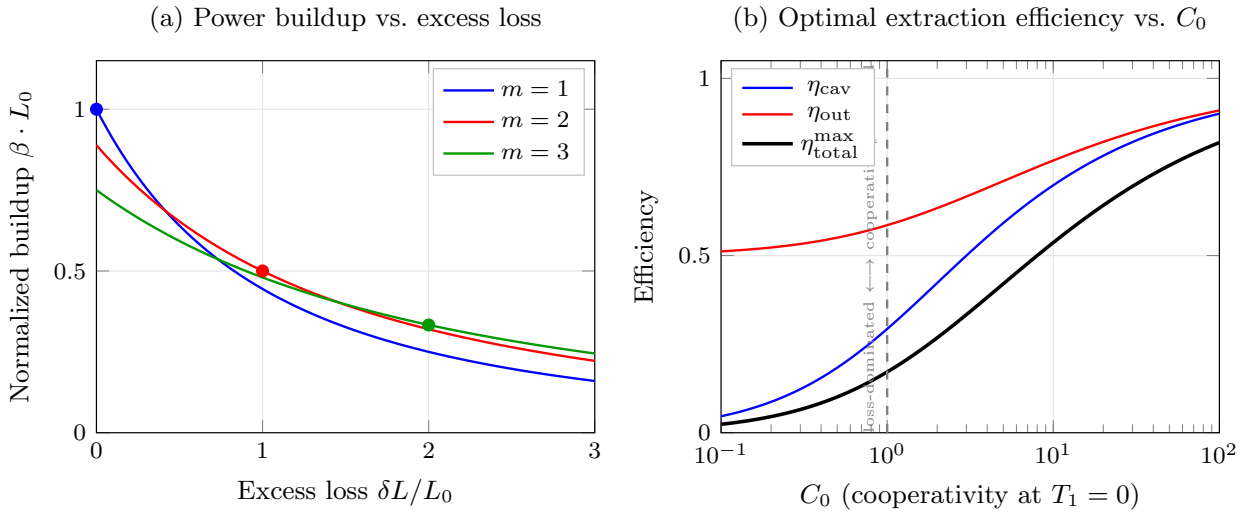


Figure 3.7: Cavity optimization for power buildup and photon extraction. **(a)** Normalized buildup βL_0 versus excess loss $\delta L/L_0$ for different design margins $m = T_1/L_0$. Dots mark where each design reaches critical coupling. The $m = 1$ curve starts highest but degrades fastest. **(b)** Optimal photon extraction efficiency versus intrinsic cooperativity C_0. Below $C_0 \sim 1$ the system is loss-dominated and η_{total} drops rapidly.

3.7.1 Summary of Cavity Design Principles

For power buildup applications, the cavity should be designed overcoupled with $T_1 = L_{\text{tot}}^{\text{worst}}$ to hedge against unforeseen loss increases. This sacrifices some peak buildup compared to critical coupling but ensures robustness against degradation. For photon extraction in cavity QED, the primary goal is to maximize the intrinsic cooperativity C_0 through small mode volume, high finesse, and good atom-cavity spatial overlap. The optimal input coupler transmission is then $T_1 = L_{\text{tot}}\sqrt{1 + C_0}$, which is automatically overcoupled when $C_0 > 1$. The system enters the loss-dominated regime when $L_{\text{tot}} > 4g^2/(\Gamma\nu_{\text{FSR}})$, corresponding to $C_0 < 1$. Above this loss threshold, every additional ppm of internal loss costs linearly in extraction efficiency. High-cooperativity systems with $C_0 \gg 1$ automatically operate in the overcoupled regime at optimal extraction and are naturally robust to loss degradation, with fractional efficiency loss suppressed as $1/\sqrt{C_0}$.

3.8 Jaynes-Cummings model

$$E_{n,\pm}/\hbar = \omega_c(n+1) + \frac{1}{2}\left(\delta_{ec} \pm \sqrt{\delta_{ec}^2 + 4|g|^2(n+1)}\right) \tag{3.58}$$

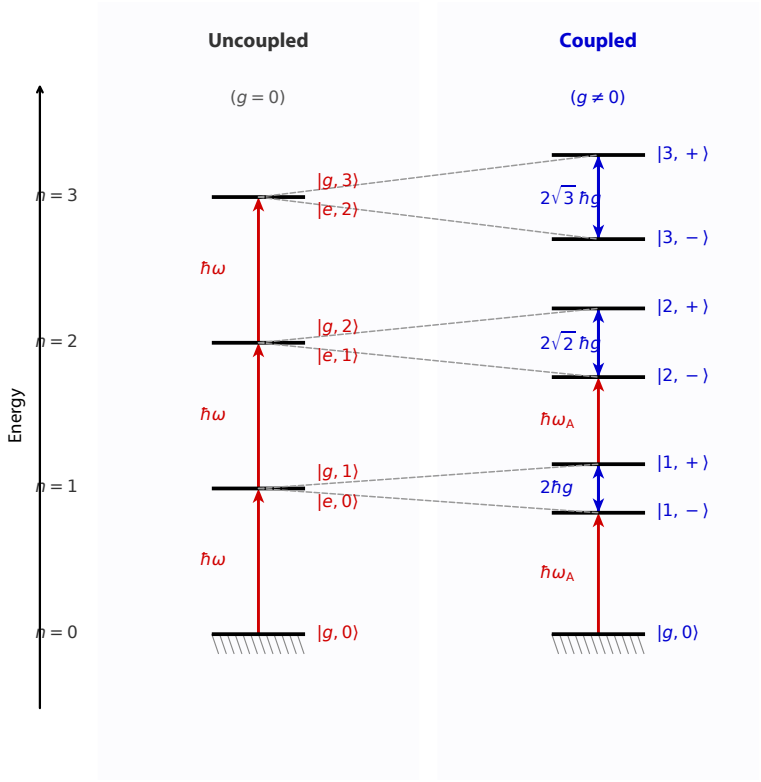


Figure 3.8: Jaynes-Cummings ladder. Eigenstates of the Jaynes-Cummings Hamiltonian. In the uncoupled ($g = 0$) basis states with N total excitations are degenerate between atomic and cavity excitation. In the presence of light-matter coupling g these degenerate pairs split up with a non-linear energy separation $\pm\sqrt{N}g$ providing a strong non-linearity on the single photon level.

For $\omega_e < \omega_c$ (negative δ_{ec})

$$|n,+\rangle = -\cos\frac{\theta}{2}|n+1,g\rangle + \sin\frac{\theta}{2}|n,e\rangle \tag{3.59}$$

$$|n,-\rangle = \sin\frac{\theta}{2}|n+1,g\rangle + \cos\frac{\theta}{2}|n,e\rangle \tag{3.60}$$

$$\tan\theta_i = \frac{2g\sqrt{n+1}}{\delta_{ec}} \tag{3.61}$$

For $\omega_e > \omega_c$ (positive δ_{ec})

$$|n,+\rangle = \sin\frac{\theta}{2}|n+1,g\rangle + \cos\frac{\theta}{2}|n,e\rangle \tag{3.62}$$

$$|n,-\rangle = \cos\frac{\theta}{2}|n+1,g\rangle - \sin\frac{\theta}{2}|n,e\rangle \tag{3.63}$$

3.9 Non-Hermitian perturbation theory

Computing the eigenstates or time evolution of quantum Hamiltonians is hard, especially in a closed-form analytical solution. Only special cases like the Jaynes-Cummings model we just saw can be solved in closed form. Even numerically, quantum systems quickly become intractable as the Hilbert space grows exponentially with the number of constituent particles or modes. In quantum optics systems, such as the cavities that we are treating here, this is even exacerbated: We need to take into account the open nature of these systems. This means we have to treat the full density matrix of a system and not just a single pure state, increasing the Hilbert space dimension from 2^N to 4^N. To rigorously treat open quantum systems, we can employ the master equation formalism that treats decay through Liouvillian jump operators $\hat{L}_k$ for every decay channel γ_k.

$$\dot{\rho} = \mathcal{L}[\rho] \tag{3.64}$$

$$\mathcal{L}[\rho] = -\frac{i}{\hbar}[\hat{H},\rho] + \sum_k \gamma_k \left(\hat{L}_k \rho \hat{L}_k^\dagger - \frac{1}{2}\left\{\hat{L}_k^\dagger \hat{L}_k, \rho\right\}\right) \tag{3.65}$$

The master equation is exact even in the regime of strong driving where atoms can decay back to the ground state through one of the decay channels and be re-excited from there, modeling effects such as saturation. It is a differential equation eq. (3.64) that allows us to compute the time evolution of the system, and spectral properties like transmission spectra can be obtained by Fourier transforming the time response. Finding analytical solutions

to this equation in closed form quickly gets prohibitively complicated, even for a simple three-level atom inside of a cavity.

Instead, we use an alternative formalism called the resolvent method, developed by Cohen-Tannoudji and also referred to as Non-Hermitian perturbation theory in the Simon lab, that allows treatment of open quantum systems in frequency space and allows for a well controlled approximation of a plethora of Hamiltonians.

3.9.1 The Formalism

The name Resolvent Method is inspired by the resolvent from complex analysis, it suggests that the system evolution $U(t)$ is governed by only a few poles of a resolvent operator $G(z)$ in analogy to the principal value theorem.

$$U(t) = \frac{1}{2\pi i} \oint_C e^{-izt} G(z) dz \tag{3.66}$$

In order to model decay, we consider our system coupled to a reservoir with projectors P and Q satisfying $P + Q = 1$. The dynamics of our system (with projector P) that are governed by the resolvent operator G can then be described

$$PG(z)P = \frac{P}{z - H_{\text{eff}}(z)} \tag{3.67}$$

in terms of an *effective*, non-Hermitian Hamiltonian $H_{\text{eff}}(z) = PH_0P + R(z)$ where H_0 is the original, Hermitian Hamiltonian projected onto the system subspace and a so called *level-shift* operator $R(z)$ that describes the coupling to the reservoir. While this coupling can be complicated and for example frequency dependent, in the systems we consider we can capture the decay in terms of a simple linewidth $R(z) = \frac{i}{2}\Gamma$. So far, we haven't gained any computational advantage, as we still need to solve the full system $(z - H)G(z) = 1$ which can be impossible analytically. This is where the unique strength of the Resolvent

method comes in: We consider a weak drive that is applied to our system, as for example, shining a weak probe beam onto a cavity atom system. By treating the driving term V as a perturbation that couples distinct excitation number manifolds, we can expand the full steady state solution formally into a series.

$$|\psi_{\text{ss}}\rangle = \sum_{n=0}^{\infty} \Omega^n |\psi_n\rangle \tag{3.68}$$

$$= |\psi_0\rangle + \Omega |\psi_1\rangle + \Omega^2 |\psi_2\rangle + \Omega^3 |\psi_3\rangle + \ldots \tag{3.69}$$

$$= |\text{⬭ ⬭ ⬭ ⬭}\rangle + \Omega \,|\text{⬭ ● ⬭ ⬭}\rangle + \Omega^2 |\text{● ⬭ ⬭ ●}\rangle + \Omega^3 \,|\text{⬭ ● ● ●}\rangle + \ldots \tag{3.70}$$

The perturbation theory here is really an expansion in the number of excitations in the system, such as photons in a cavity or atoms in an excited state. This makes the Resolvent formalism perfect for treating strongly correlated polariton systems, as there can only ever be at most two excitations in a single mode in the presence of strong blockade. By inserting this into the time-independent Schrödinger equation $(H_{\text{eff}} + \Omega\tilde{V})|\psi_{ss}\rangle = E|\psi_{ss}\rangle$ where E is the eigenenergy of the state, and matching powers in the expansion we obtain the recurrence relation

$$H_{\text{eff}}|\psi_n\rangle + \tilde{V}|\psi_{n-1}\rangle = E|\psi_n\rangle \tag{3.71}$$

where we have factored out the small parameter Ω of the drive term as $V \equiv \Omega\tilde{V}$. The frequency of the driving laser ω_L is included implicitly by working in the frame rotating at ω_L. In this frame, the energy parameter z (or E) in the resolvent $(z - H_{\text{eff}})^{-1}$ becomes the detuning $\Delta = \omega_{sys} - \omega_L$, the "energy" of the propagator is actually the drive detuning. This allows us to calculate the correction to the wavefunction and energy in first order

$$|\psi_1\rangle = \frac{1}{E - H_{\text{eff}}} V|\psi_0\rangle \tag{3.72}$$

and in second order

$$|\psi_2\rangle = \frac{1}{E - H_{\text{eff}}} V \frac{1}{E - H_{\text{eff}}} V |\psi_0\rangle \tag{3.73}$$

which is sufficient to calculate correlation between photons or other parts of the system. There is one extremely important limitation when using this perturbation theory: The higher order corrections can only be obtained if the Hamiltonian is also expanded to the matching number of total excitations. When representing H_{eff} as a matrix, one has to enumerate all possible states of the system in an excitation number restricted basis. We will see an example of that when using the NHPT to predict blockade strength from cavity Rydberg polaritons. There are two important considerations to the applicability of this method: For linear systems such as cavities or harmonic oscillators that both represent Bosonic ladders, the first-order expansion can actually yield the exact closed-form solution, as the system is independent of the number of excitations already in the system. However for non-linear systems in the strong driving limit, the non-Hermitian Hamiltonian neglects repopulation of the ground state as a consequence of the decay. Therefore, it fails to capture effects like saturation that have to be treated using a full Master equation.

In quantum optics experiment such as atoms in cavities, we are usually mostly interested in computing the transmission spectrum of the system, as this is what we can observe experimentally, the microscopics of the atomic excitations are often invisible or uninteresting. Instead of diagonalizing the full Hamiltonian matrix, in NHPT [CDG24] we only have to calculate a matrix inverse, which can be a lot faster and yield simpler analytical expressions especially when the matrix H is sparse which is often the case for hierarchical Hamiltonians.

3.9.2 VRS

For a two-level atom in a cavity, the effective non-Hermitian Hamiltonian in the basis of $|0\rangle$ (ground), $|c\rangle$ (cavity photon), and $|e\rangle$ (excited atom) is represented as:

$$H_{\text{eff}} + V = \begin{pmatrix} 0 & v & 0 \\ v & \Delta_c - i\frac{\kappa}{2} & g\sqrt{N} \\ 0 & g\sqrt{N} & \Delta_a - i\frac{\Gamma}{2} \end{pmatrix} \tag{3.74}$$

where Δ_c and Δ_a are the cavity and atomic detunings, g is the light-matter coupling, and the factor of $1/2$ ensures that the probability amplitudes decay correctly as $e^{-\Gamma t/2}$ so that probabilities decay as $e^{-\Gamma t}$. In practice, the vacuum state $|0\rangle$ doesn't have to be included explicitly as the response is linear in v

$$T(\omega_L) \propto |\langle C|\psi_{ss}\rangle|^2 = \left|\langle C|\frac{1}{\tilde{H}(\omega_L)}|0\rangle\right|^2 \tag{3.75}$$

3.9.3 Collective States

3.10 Rydberg EIT

As we have previously seen, coupling the cavity to an ensemble of many atoms enhances the collective cooperativity like $\sqrt{N}$ and thereby enables strong coupling beyond the intrinsic cavity finesse and waist. However, it also destroys the Jaynes-Cummings non-linearity that a single atom in a cavity forms and to regain strong interactions between cavity photons we have to employ a different kind of nonlinearity. Rydberg blockade provides the strongest available interaction between neutral cold atoms at optical (i.e. μm) separations $d > \lambda$ where we can neglect molecular potentials and light assisted collisions. Strong and long range Van-der-Waals interaction between highly excited Rydberg electrons prevent all further Rydberg excitations in the so called blockade radius R_b around an existing one. By choosing a highly

enough excited principal quantum number, we can match the blockade radius $R_b \approx \omega_0$ to the cavity waist and thereby realize a nonlinearity for polaritons that can allow only one polariton per cavity waist.

The exact mechanism we use to bridge the energy separation between cavity photons and atomic Rydberg state is known as electromagnetically induced transparency (EIT) which will be explained in the following.

3.10.1 Rydberg Polaritons

By coupling our cavity to an atom by making its frequency resonant with the transition from the ground $|g\rangle = |5S_{1/2}\rangle$ to excited $|e\rangle = |5P_{3/2}\rangle$ atomic state we were able to observe VRS. Now, we introduce the Rydberg state $|r\rangle$ as a third level in the EIT ladder allowing excitation of the Rydberg state through a two-photon process involving a cavity photon at 780 nm and an additional laser field at 480 nm that we call the control field. The exact Rydberg state can be chosen through wavelength and polarization according to the desired interaction strength but we usually operate with $|nD_{5/2}\rangle$ to maximize available Rabi frequency. Using the same collective operators we defined in the previous section the Hamiltonian is now in a frame rotating with the cavity and laser frequency. To gain some analytical insights and intuition about the eigenstates of this system, we first assume uniform couplings to the cavity g_i and control field Ω_i and also ignore the coarse graining onto discrete atoms. We can then write down another non-Hermitian Hamiltonian that effectively describes the collective EIT

$$\mathcal{H}_0 = \begin{pmatrix} \delta_c - \frac{i}{2}\kappa & g & 0 \\ g & \delta_e - \frac{i}{2}\Gamma & \Omega \\ 0 & \Omega & \delta_r - \frac{i}{2}\gamma_r \end{pmatrix} \tag{3.76}$$

with δ_x the detunings and associated linewidths of cavity κ, atomic excited state Γ and Rydberg state γ_r.

Before considering the full spectrum including linewidths, we first ignore the imaginary decay coefficients which allows us to find the eigenstates and energies in closed form by diagonalizing $H|\psi\rangle = \lambda|\psi\rangle$. On resonance ($\delta_c = \delta_r$) we find that the 3×3 Hamiltonian has three eigenvectors, the so called Polaritons. At zero energy we find the dark Polariton

$$\psi_D = \frac{1}{\sqrt{g^2 + \Omega^2}} \begin{pmatrix} \Omega \\ 0 \\ -g \end{pmatrix} = \begin{pmatrix} \cos\theta_D \\ 0 \\ -\sin\theta_D \end{pmatrix} \tag{3.77}$$

that only consists of cavity photon and Rydberg state but is *dark* to intermediate p state due to quantum mechanical interference between the different possible excitation paths. This is defining feature of EIT and enables the whole hybrid Polariton experiment! It is longer lived than the bare cavity photons but still inherits the effect of the gauge field, and is imbued by strong interactions through its Rydberg component. The exact composition is given by the ratio between cavity coupling g and control Rabi frequency Ω and can be summarized in a single number, the dark state rotation angle $\theta_D = \arctan(g/\Omega)$. Even though the intermediate state $|e\rangle$ is lossy, it is not populated at all on two-photon resonance. The other two eigenstates are

$$\psi_{B\pm} = \frac{1}{\sqrt{2}} \begin{pmatrix} \sin\theta_D \\ \pm 1 \\ \cos\theta_D \end{pmatrix} \tag{3.78}$$

are so called bright polaritons as they have a large population in the atomic state that can decay into free space. In fig. 3.9 (a) and (b) we can see these eigenstates in a transmission spectrum computed using NHPT by sweeping the frequency of a probe over the full spectrum of the non-Hermitian $\mathcal{H}_0$. The two bright polaritons are split out by $\sqrt{g^2 + \Omega^2} \approx g$ and are wide in frequency as they have a large atomic contribution. At zero frequency the dark polariton ψ_D appears much narrower and its linewidth can be tuned (see fig. 3.9 (b))

through the dark state rotation angle θ_D approaching the bare cavity linewidth when made more photon-like.

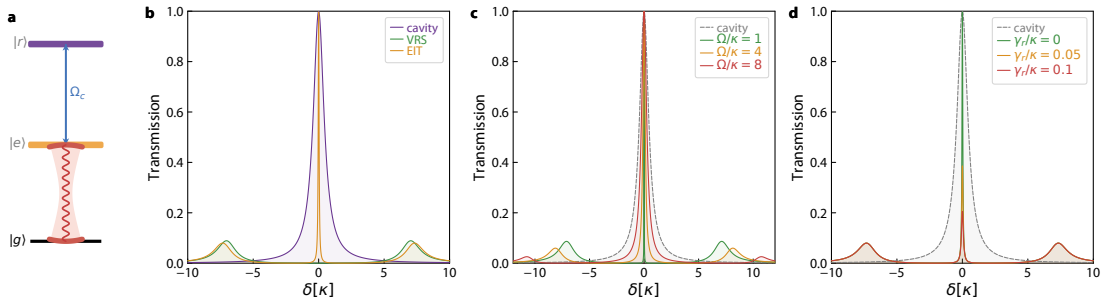


Figure 3.9: Cavity Rydberg EIT Level scheme **(a)** The cavity is resonant with the $|g\rangle = |5S_{1/2}\rangle$ to $|e\rangle = |5P_{3/2}\rangle$ transition and a classical control beam Ω_c excites to the Rydberg state $|r\rangle = |nD_{5/2}\rangle$. **(b)** Transmission spectrum as a function of probe laser detuning δ of the bare cavity (purple), vacuum Rabi splitting (VRS) in the presence of atoms (green) and narrow EIT feature in the presence of the Rydberg laser. All frequencies are normalized in units of the cavity linewidth κ. **(c)** Increase control field strength Ω_c change the character of the polaritons to be more photon like. **(d)** The presence of any Rydberg decay due to limited lifetime or decoherence channels rapidly reduces transmission compared to the bare cavity.

The dark polariton linewidth is a key metric as it sets the lifetime of the quantum particles in our system. In order to quantify it, we can make use of the non-Hermitian Hamiltonian in a different way and expand the characteristic polynomial $\det(\mathcal{H} - \lambda\mathbb{1}) \stackrel{!}{=} 0$ up to second order. We know that the dark polariton is the eigenstate at zero energy so the imaginary part of the characteristic polynomial is the linewidth.

$$\gamma_D = \kappa \cos^2\theta_D + \gamma_R \sin^2\theta_D + a(\delta_c - \delta_r)^2. \tag{3.79}$$

To first order, the dark polariton linewidth $\gamma_D = \kappa \cos^2\theta_D + \gamma_R \sin^2\theta_D$ is a simple mixture of the cavity linewidth κ and Rydberg state decay rate γ_r weighted by their superposition weights from θ_D. The effect of Rydberg linewidth γ_r is shown in fig. 3.9 (c): For infinitely long lived Rydberg states, the dark polariton could be made arbitrarily narrow while always maintaining full transmission T. Any finite γ_r that summarizes Rydberg state lifetime, but also a series of possible collective state decay channels, will lead to a reduction in

transmission and reduced dark polariton lifetime beyond the simple cavity decay. Much of the experimental effort is therefore focused on minimizing γ_r by eliminating stray electric fields and inhomogeneous as well as atomic doppler broadening. A quantitative comparison of all possible decay channels is given in [GSS18].

The second order correction $a = 4\Omega^2 g^2 \Gamma/(\Omega^2 + g^2)^3$ is more subtle. When the single photon detuning $d\delta_1 = \delta_e = \delta_c - \delta_r > 0$ is broken, the atomic $|e\rangle$ state is no longer perfectly dark but has a small occupation introducing additional loss. The effect of this is shown in fig. 3.10 (b) for different transverse cavity modes (a) that are not quite resonant with the $|e\rangle$ state, a situation we often encounter when the science cavity isn't perfectly degenerate. For the central mode $\Delta_1 = 0$ the transmission is unity but with increasing detuning Δ_1 the dark polariton linewidth increases and transmission decreases. This effect is also called the EIT window and it follows a Lorentzian shape

$$T_d \sim \frac{1}{\left|1 + \frac{\delta_c - \delta_r}{\Gamma_w/2}\right|^2} \tag{3.80}$$

with EIT window width Γ_w shown in grey in fig. 3.10 (b).

While the EIT window width naively depends on the blue Rabi frequency Ω, it is fundamentally limited by the collective cooperativity $\mathcal{C}$ and ultimately determines the number of discrete polaritons of different cavity modes that can be resolved. One might assume that increasing the control field Ω would make the polaritons more photon like and thereby space them out more in frequency following the cavity modes, however this will also increase the polariton linewidth to the same extent leading to the same number of modes being resolved.

3.10.2 Structured Hamiltonians

As we have seen the NHPT formalism can be used to easily treat increasingly complicated atomic level schemes. As such, it can easily extend to multiple Rydberg levels for microwave

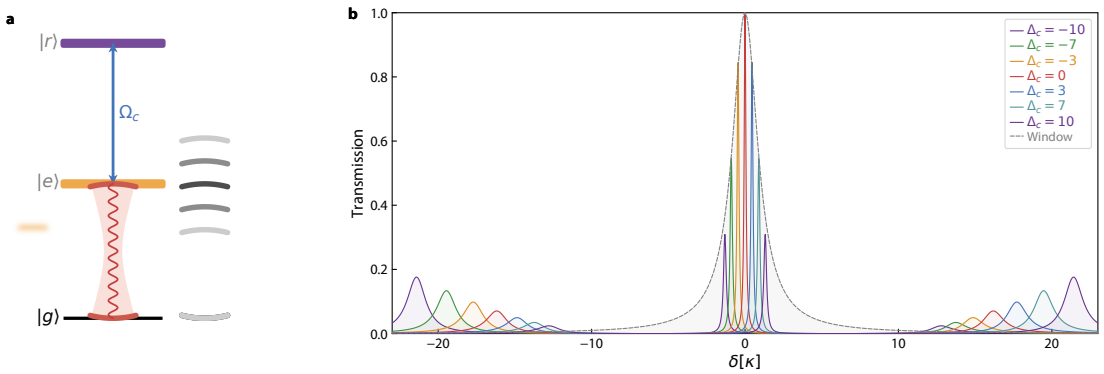


Figure 3.10: Cavity Rydberg EIT Window(a) When modes of the cavity are not on resonance with the $|g\rangle = |5S_{1/2}\rangle$ to $|e\rangle = |5P_{3/2}\rangle$ transition, a single photon detuning Δ_1 breaks the EIT condition (grey cavity modes). **(b)** Apart from a frequency shift corresponding to the photon component of the dark polariton (colors), the intermediate state $|e\rangle$ is not dark anymore and induces additional loss reducing transmission. This shows the so called EIT window (grey) whose width is determined by single photon detuning Δ_1 and which the polariton transmission reduces to half and is limited by the collective cooperativity $\mathcal{C}$.

coupling or multiple magnetic sublevels to study the effect of imperfect polarization. Instead of listing all of these very similar Hamiltonians, we will develop a generic procedure here that applies to any combination of atomic excited and Rydberg states. We consider a basis of a single cavity mode c, n intermediate states P and m Rydberg states R with the Hamiltonian

$$\hat{H} = \left(\begin{array}{c|cc} c & G & 0 \\ \hline G^* & A & M \\ 0 & M^* & R \end{array}\right) \tag{3.81}$$

where G is a $1 \times n$ row vector describing the coupling from the cavity to the n intermediate P states, A is an $n \times n$ square matrix containing the internal energies and couplings of the P state manifold, M is a $n \times p$ matrix describing the control laser coupling from the n intermediate states to the p Rydberg states and R is a $p \times p$ square matrix containing the couplings and decay of the Rydberg manifold. Most imperfections of our EIT scheme except for multiple ground states due to imperfect optical pumping can be described in this form. When computing the matrix inverse in eq. (3.81) we can exploit the block structure

of Hamiltonian $\hat{H}$ and use Schur's complement formula giving

$$X_{11} = c - VD^{-1}V^{\dagger} \tag{3.82}$$

with $X_{11_0} = [c]$, $V = \begin{pmatrix} G & 0_{1\times p} \end{pmatrix}$ and $D = \begin{pmatrix} A & M \\ M^* & R \end{pmatrix}$. Due to the hierarchical coupling in V we only need the top left element $\tilde{A}$ in $D^{-1} = \begin{pmatrix} \tilde{A} & \tilde{M} \\ \tilde{M}^* & \tilde{R} \end{pmatrix}$ and arrive at the exact analytical expression for the effective cavity Hamiltonian

$$X_{11} = c - G(A - MR^{-1}M^*)^{-1}G^*. \tag{3.83}$$

This allows us to write the cavity transmission t in the usual form but now in terms of matrices

$$t(\omega) = \frac{\Omega_p}{(\omega_c - \omega) - G\left((\omega_e - \omega) - \Omega \cdot (\tilde{\omega}_r - \omega)^{-1} \cdot \Omega^{\dagger}\right)^{-1} G^{\dagger}} \tag{3.84}$$

with $A = (\omega_a + \frac{i}{2}\Gamma - \omega)\mathbb{1}$ and ω the probe frequency.

3.11 Rydberg interactions

So far we have only discussed the Rydberg interactions as a VdW force that gives rise to Rydberg blockade due to an energy shift C_6/R^6, the real picture is more nuanced however and will be necessary to later understand the physics of the Förster resonance. The intrinsic interaction strength between two Rydberg atoms is determined by the dipole coefficient C_3 and the Förster defect Δ_F. Two Rydberg states with identical quantum numbers experience no direct dipole-dipole ($1/r^3$) interaction because of the odd parity of the dipole operator. The interaction is therefore of a Van-der-Waals type ($1/r^6$) and can be understood as a second order interaction via a virtual pair state, which is symbolized in fig. 3.11 (b) as a

Feynman type diagram. The initial pair state (here, of identical $|d\rangle$ states) couples to a pair of atoms that conserve total angular momentum, via the dipole operator. The strength of the second order VdW interaction is given by $C_6 = C_3^2/\Delta_F$.

To understand the origin of this form we model the interaction between a target Rydberg pair state, $|rr\rangle$, and a coupled pair state, $|r'r''\rangle$, using a simple effective Hamiltonian in a single channel model. Here we consider the target state to be a D state pair, $|rr\rangle = |dd\rangle$, coupled to a P-F state pair, $|r'r''\rangle = |pf\rangle$ as this will be closest to our experiment but in principle any pair of states, even magnetic sublevels that conserve total angular momentum are allowed. The energy difference between these unperturbed pair states is $\Delta_F = E_{r'} + E_{r''} - 2E_r$ which is known as the Förster defect. This defect depends sensitively on the atomic defects of the atomic species that appear in the Rydberg formula and that are a function of the microscopic structure of the filled electron shells. In fig. 3.11 (d) this defect is shown for ^{87}Rb for two channels possible in our experiment. While both permutations of the unlike pair state must to be accounted for, only the symmetric state $|\tilde{pf}\rangle = (|pf\rangle + |fp\rangle)/\sqrt{2}$ couples to $|dd\rangle$ via the dipolar interaction, leading to the model Hamiltonian:

$$\mathcal{H} = \begin{pmatrix} 0 & V(R) \\ V(R) & \Delta_F \end{pmatrix}, \tag{3.85}$$

with the eigenstates

$$V_{\pm}(R) = E_{\text{avg}} \pm \sqrt{\left(\frac{\Delta_F}{2}\right)^2 + \left(\frac{C_3}{R^3}\right)^2} \tag{3.86}$$

where $V(R)$ is the dipole-dipole interaction. The behavior of the eigenvalues falls into two

distinct regimes:

$$V_{\pm}(R) \approx \begin{cases} -\dfrac{C_6}{R^6} & \text{for } |\Delta_F| \gg |V(R)| \quad \text{(Van der Waals regime)} \\ \\ \pm\dfrac{C_3}{R^3} & \text{for } \Delta_F \to 0 \quad \text{(Resonant regime)} \end{cases} \tag{3.87}$$

In the far-detuned (Van der Waals) regime, the energy shift is described by the C_6 coefficient, calculated by summing over the contributions from all coupled pair states:

$$C_6 = \sum_{r'r''} \frac{\left|\langle r'r''|V(R)|rr\rangle\right|^2 R^6}{\Delta_{r',r''}} \tag{3.88}$$

In the case of resonance with a particular pair state $\Delta_{r',r''} = 0$, the eigenstates are split symmetrically by the resonant dipole-dipole interaction, which scales as $1/R^3$. This crossover between the two regimes is shown in fig. 3.11 (e) where the levels indicate the zero-, one- and two-excitation sectors and the shading indicates the eigenstate composition in terms of the bare Rydberg pair states. The x axis denotes an external tuning field that can be used to match the shift exactly to the defect Δ_F where the two base states become degenerate. We will later utilize this to access a Förster resonance around $|n = 44, D_5/2\rangle$ in ^{87}Rb to enhance photon blockade in the cavity.

The previous treatment has only considered a pair of individual Rydberg atoms and their interaction, as is common in optical tweezer experiments with well defined separations [RLB$^+$14]. Since our Rydberg polaritons however are collective excitations, where a single excitation is shared amongst the whole atomic ensemble as a superposition, the situation gets much more complicated. To first order, all possible pairs of atoms in the cloud that are within the cavity field lead to a distribution of possible distances from almost zero, for two atoms that sit right next to each other, to separations of $2\omega_0$ at opposite ends of the cloud. This makes two Rydberg polaritons that are excited at the same time to not have a well

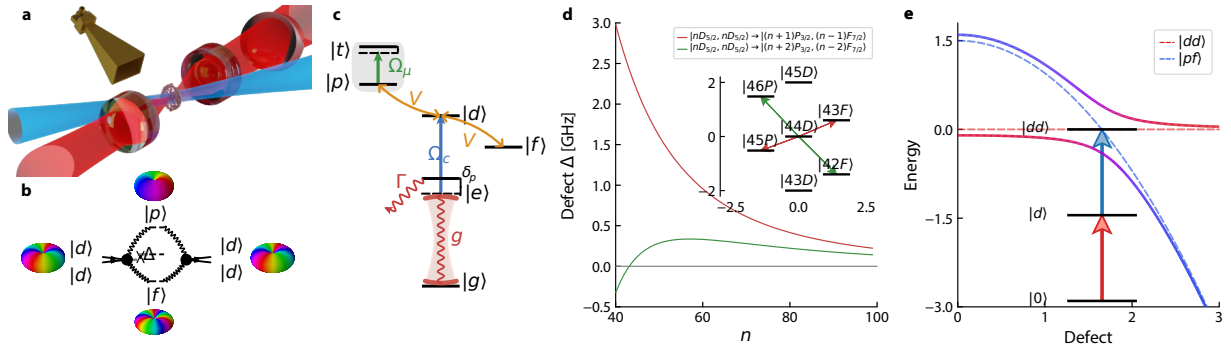


Figure 3.11: Rydberg interactions and Förster defect (a) Schematic of the experimental setup: a cloud of ultracold ^{87}Rb atoms is transported into a tightly focused cavity at 780 nm (red) while a control field at 480 nm (blue) excites to the Rydberg state. A microwave horn (gold) is used to selectively shift different Rydberg states into a Förster resonance. **(b)** The atomic ground $|g\rangle$, excited $|e\rangle$, and Rydberg $|d\rangle$ state are coupled by a cavity mode with collective coupling g and a classical control field Ω_c to realize cavity electromagnetically induced transparency (EIT). The Rydberg state $|d\rangle$ that is dipole-dipole coupled to two other Rydberg states $|p\rangle$ and $|f\rangle$ with strength V, where $|p\rangle$ is AC Stark shifted by driving a tuning transition to a state $|t\rangle$ far off resonance with microwave Rabi frequency Ω_μ. **(c)** Förster defect Δ_F as a function of principal quantum number n for the two dominant channels shown in red and green (see inset for $n = 44$). At $n = 43$ a natural Förster near resonance $\Delta_F \approx 8$ MHz occurs. **(d)** Blockade mechanism for tuned Förster resonance: In the Van-der-Waals regime at zero tuning field, the defect between pair states $|dd\rangle$ (red) and $|pf\rangle$ (blue) is large and leads to only a small interaction shift (solid line) compared to the non-interacting case (dashed line). With increasing tuning strength and resulting $|pf\rangle$ shift, the interaction becomes dipolar and the two pair states show an avoided crossing at $\Delta_F = 0$; the splitting is given by the dipolar interaction strength and the eigenstates are symmetric superpositions, indicated by the color. While the single-excited Rydberg state $|d\rangle$ is unperturbed, the second Rydberg excitation is now strongly blockaded from entering the system.

defined interaction energy but a range of energies over all pairs. Even worse, the interaction between two nearby atoms can be strong enough to cause a large two-photon detuning δ_2 for a second polariton entering the system, which will move it so far out of the EIT window that it becomes very lossy. For even shorter distances, the interaction enters the so called *spaghetti* regime where we can no longer consider a single pair state channel but many different states get coupled. At the same time, atoms that are shifted by the Rydberg blockade from the first polariton in the system will not be excited when the EIT condition and the real physics is an interplay of these two effects which must be treated in a self-consistent manner. This requires a rigorous field theory treatment in a diagrammatic expansion that is performed in [GTW$^+$16] for one-dimensional systems. Here, we only reproduce the results important for our experiment. Also, it is unclear if the results hold for a two-dimensional system as we consider it. The interaction potential is independent of the experiment geometry and should therefore also apply for our 2D system. The dynamics however must be modified for which a heuristic effective theory was attempted in [GSS18].

We can treat our thin ensemble of atoms as a continuum of Rydberg excitation $r^\dagger(\boldsymbol{r})$ in two dimensions

$$H_{int} = \frac{1}{2}\iint d^2\boldsymbol{r}\, d^2\boldsymbol{r'}\, r^\dagger(\boldsymbol{r}) r^\dagger(\boldsymbol{r'}) U(\boldsymbol{r}-\boldsymbol{r'}) r(\boldsymbol{r'}) r(\boldsymbol{r}) \tag{3.89}$$

where $U(x-x') = C_6/|x-x'|^6$ is the bare VdW interaction that we derived previously. By projecting this Hamiltonian onto the dark polariton manifold and assuming two-photon resonance we get [BCF$^+$14, SBS15]

$$H_{int} = \sin^4\theta_D \int d\mathbf{r} d\mathbf{r'} \psi_D^\dagger(\mathbf{r}) \psi_D^\dagger(\mathbf{r'}) U_{eff}(\mathbf{r}-\mathbf{r'}) \psi_D(\mathbf{r}) \psi_D(\mathbf{r'}) \tag{3.90}$$

yielding the effective interaction potential for dark polaritons

$$U_{eff}(x - x') = \frac{C_6}{C_6\chi + |x - x'|^6}, \tag{3.91}$$

with

$$\chi = \frac{\delta_{1p} - \frac{i}{2}\Gamma}{\Omega^2} - \frac{1}{\delta_{1p} - \frac{i}{2}\Gamma}. \tag{3.92}$$

This can be thought of as a screened or soft-core interaction whose character depends strongly on the one photon detuning δ_{1p} from the intermediate P state. In fig. 3.12 (inset) this detuning is shown in the EIT level scheme and (a)-(d) depict the transmission spectrum for on resonance (b) and for negative (a) and positive (c,d) detunings. As the transmission is computed in linear NHPT, the dark polariton linewidth doesn't change but the bright polaritons become asymmetric in splitting and height. The corresponding effective potentials according to eq. (3.91) are shown in fig. 3.12 (lower) vs detuning. From the functional form eq. (3.92) it is apparent that the interaction is complex, and both real dispersion (blue) as well as the imaginary component (red, dashed) are shown. The imaginary part signifies the loss that the interaction causes due to breaking of the EIT condition due to Rydberg blockade at small distances. The loss is most pronounced on p state resonance (f) and for negative detuning (e), in the latter case even some scattering resonance can be seen. With one (g) or two (h) Γ of p state detuning the interaction becomes more real at small distances resembling an extended contact interaction and the imaginary component is significantly reduced.

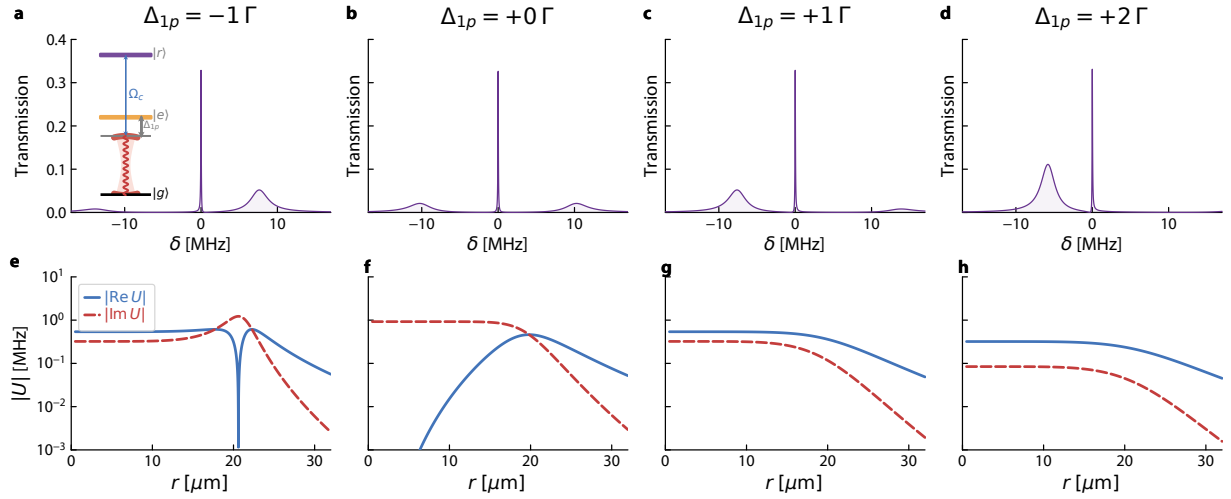


Figure 3.12: Cavity Rydberg Polariton Potential A two-photon detuning Δ_2 from the intermediate atomic $|e\rangle$ state **(inset)** can drastically change the polariton interactions. In the NHPT transmission spectra as a function of probe laser frequency δ **(a - d)**, the bright polaritons (VRS) move closer to the dark polariton and cause a Fano like lineshape, but stay resolved for $\Delta \ll g$. As NHPT assumes uniform interactions for the collective Rydberg excitations and only captures linear effects, the transmission is unchanged. A rigorous calculation [BCF+14] of effective Rydberg interaction in continuous ensembles **(e-h)** reveals a complex interaction potential U leading to a lossy *imaginary* interaction and short distances r of two-photon resonance $\Delta_2 = 0$. Detuning from the intermediate state can reduce this loss and make the interaction more repulsive **(g,h)** or introduce resonances and bound states **(e)** depending on the sign of Δ_1 relative to the bare interaction.

CHAPTER 4

MULTI-MODE CAVITY QED: METHODS

This chapter describes the details of the experiment focusing on the cavity structure in the vacuum chamber, laser systems and light modulation and preparation of the atomic sample. Included are all the steps involved in realizing the cavity design outlined in chapter 3 into a hybrid experimental platform for realizing quantum Hall states motivated in chapter 1.

4.1 Cavity Structure

The cavity structure sits at the heart of the experimental platform: it integrates the twisted cavity, that realizes the lowest Landau level for photons as a degenerate family of cavity modes, with all the other ingredients for atomic control and in-vacuum alignment, including the intra-cavity aspheres. Integrating all of these components in a small space in the vacuum chamber constrained by the MOT while maintaining optical access for all trapping and cooling beams, high mechanical stiffness to prevent vibrational resonances, and maintaining key alignment degrees of freedom in vacuum was the major challenge. In the following, an overview over all the major components and their co-design between underlying physics and hardware is presented:

Twisted Science Cavity

The twisted cavity for 780 nm photons that realizes the artificial gauge field for photons had to be re-designed completely compared to early cavity iterations . All previous mirror-only cavities could not reach a stable degeneracy due to optical aberrations (as detailed in chapter 3) incurred from the off-axis reflections on curved mirrors that can't be avoided in a non-planar cavity. Using carefully designed aspheres inside the cavity solved this problem by decoupling the non-planar geometry using only flat mirrors and providing the Gouy phase

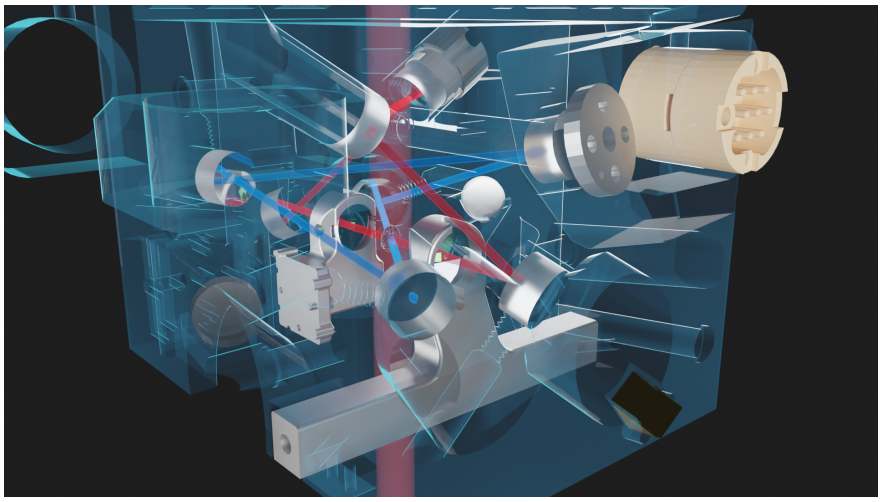


Figure 4.1: X-ray image of the cavity structure. The image shows the two aspheric lenses (A) and the four mirrors (M) forming the twisted science cavity (red). The blue cavity mirrors form the Rydberg buildup cavity (blue) and the atoms are shown schematically as metallic spheres.

and confinement for a stable cavity using lenses. As the lenses even in the smallest available size of 7.75 mm take up considerable space close to the waist, the established symmetric tetrahedral geometry would clip on the lenses for any reasonable opening angle. Instead, a new parametrization was chosen

$$\mathbf{x}_{\text{lens}} = \begin{pmatrix} \pm d_{\text{lens}}/2 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{x}_{\text{in/out}} = \begin{pmatrix} \pm L_1/2 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{x}_{\text{upper}} = \begin{pmatrix} \pm (L_2/2)\cos\theta \\ d \\ \pm (L_2/2)\sin\theta \end{pmatrix} \tag{4.1}$$

where d_{lens} is the separation between the two lens surfaces, L_1 is the length of the horizontal arm between the input and output mirrors, L_2 is the length of the upper arm between the upper mirrors, d is the vertical base distance, and θ is the twist angle of the upper mirrors. Starting from an isosceles trapezoid for the flat geometry when the twist angle $\theta = 0$, the lenses are placed on the lower base and the distance d can be increased until the cavity axis isn't clipping on the lenses with an additional safety factor for higher order modes. The geometrical twist angle corresponding to the image rotation upon one round-trip depends sensitively on all of these parameters and is determined numerically in the cavity

design chapter 3. The design was then optimized for a small total round-trip length, keeping the incident angles of the mirrors as small as possible and providing enough clearance between the mirrors for other optics. Compared to the old designs, the angle on the mirrors is still comparatively large, leading to birefringence of the Bragg stack of the dielectric HR coatings. While the twist of the cavity also favors circular polarization as the rotationally symmetric eigenstates, the birefringence would have made the polarization eigenmodes elliptic and thereby drive unwanted atomic transitions. The mirror coatings were therefore especially designed to compensate birefringence at the design angles, courtesy of the coating expertise of FiveNines optics. On the upper mirrors, the incident angle is especially large but an additional feature minimizes the birefringence: The image rotation between these two mirrors is almost exactly 90 degrees, so any phase incurred for e.g. horizontal polarization on the first mirror will be added to the vertical one on the second mirror and thereby cancel.

Experience from previous cavity iterations has shown that the mounting structure itself can have many mechanical resonances at acoustical frequencies if it is assembled from different pieces with screws or has particularly susceptible shapes like a tuning fork formed by two posts on a plate. Every mechanical resonance adds an additional 180° of phase to the transfer function of the piezo to cavity length feedback loop, thereby limiting the feedback bandwidth to well under the resonance frequency. While attempts to mitigate these resonances using matched filters have been made , the exact resonance frequencies drift with piezo voltage and over time due to changes in the piezoelectric material. The best solution is to design the mounting structure to be as stiff as possible by making it monolithic and push piezo resonances to high frequencies. The mirror attached to a piezo forms a textbook example of a mass on a spring. However, the piezo material itself is actually quite stiff ($\approx$ 50 GPa, almost as stiff as Fused silica at 73 GPa) so the glue used to mount mirror and piezo is usually the most flexible element. While more advanced designs using springs (,) exist, a reasonably high resonance frequency can be achieved by making the glue layer cover the

full area of the joint and making it very thin through a combination of minimal application and high clamping force (first discussed in). Fabrication from a single piece of metal puts more stringent requirements on manufacturability and one has to design the structure with machining on a 5 axis CNC machine in mind. The final cavity structure was still so complex and had small features that it could only be milled from aluminium. While there were some concerns aluminum is actually more vacuum compatible due to lower hydrogen diffusion and being non-magnetic. Only the intra-cavity lens holders are machined separately as they need to be carefully aligned to the cavity axis defined by the rotation axis of the cavity set by the exact mirror positions and tilts. As they do not change the total length of the running wave cavity, small vibrations do not change the resonant frequency. They are also machined out of 316 stainless steel as this has been historically proven as a Rydberg compatible material. The lenses and mirrors are held in their mounts by a very small amount of glue (Masterbond) that was selected for dimensional accuracy while curing and low outgassing. Even before being glued, the mirrors and lenses are constrained by making their mounting holes very precise. The tolerance on the hole diameter is specified exactly to give a slip fit (H7/h6) with the substrates, such that they form a stable cavity without any alignment [1].

Rydberg Buildup Cavity

As we want to realize EIT for a large number of modes and also operate at high principal quantum number, a large Rydberg excitation beam with high optical power is needed. For 25 LG_l modes the Rydberg beam waist needs to be at least 5 times the waist of the science cavity, and to provide a uniform Rabi frequency even larger. While Rydberg excitation using a free-space beam would be possible, it would require more optical power than our 480 nm laser system can provide with 250 mW of available power. Therefore, a second cavity is used to build up the 480 nm light and allow for large blue Rabi frequencies. Since this cavity also

1. when one of the mirrors is swapped for a concave one to stabilize a cavity mode

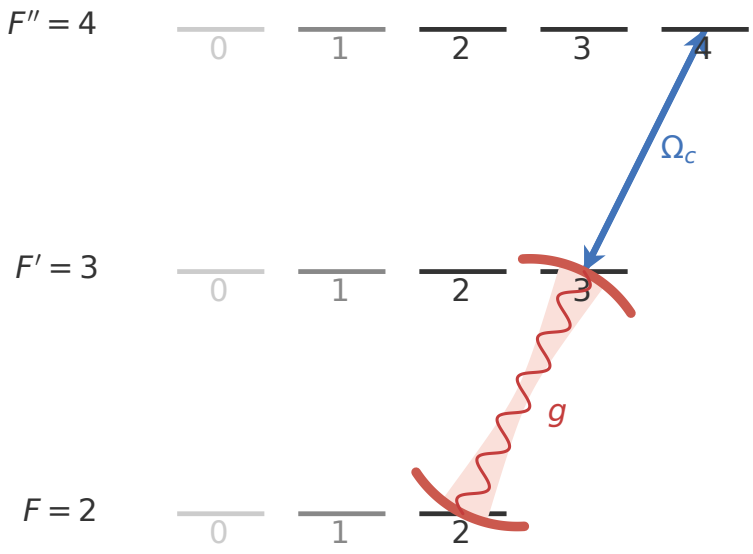


Figure 4.2: EIT level scheme. The twisted cavity is circularly polarized with the $|5S_{\frac{1}{2}}, F = 2, m_F = 2\rangle \rightarrow |5P_{\frac{3}{2}}, F = 3, m_F = 3\rangle$ transition, from where a classical control field takes excitation up to the Rydberg manifold via the $|5P_{\frac{3}{2}}, F = 3, m_F = 3\rangle \rightarrow |nD_{\frac{5}{2}}, J = 5/2, m_J = 5/2\rangle$ transition. Operating on the *stretched* transitions maximizes cooperativity and Rabi frequency as these transitions have maximum Clebsch-Gordan coefficients.

needs to be running-wave and have circularly polarized eigenmodes to drive the stretched p-to Rydberg-state transition $|5P_{\frac{3}{2}}, F = 3, m_F = 3\rangle \rightarrow |nD_{\frac{5}{2}}, J = 5/2, m_J = 5/2\rangle$, this cavity is also designed to be non-planar four mirror cavity. It is however not designed to be degenerate and as the small astigmatism of the fundamental mode due to non-normal incidence on the mirrors can be ignored. The blue cavity geometry follows the more traditional tetrahedron with the opening angle chosen such that the mirrors fit around the science cavity mirrors. The two cavities would ideally be counter-propagating to minimize the influence of Doppler broadening on the EIT resonance, however they cannot share the same lenses or mirrors. Therefore, the blue cavity mode passes by the aspheric lenses as close as possible without clipping and the lens holders are notched out to not add additional margin.

Optical Alignment Features

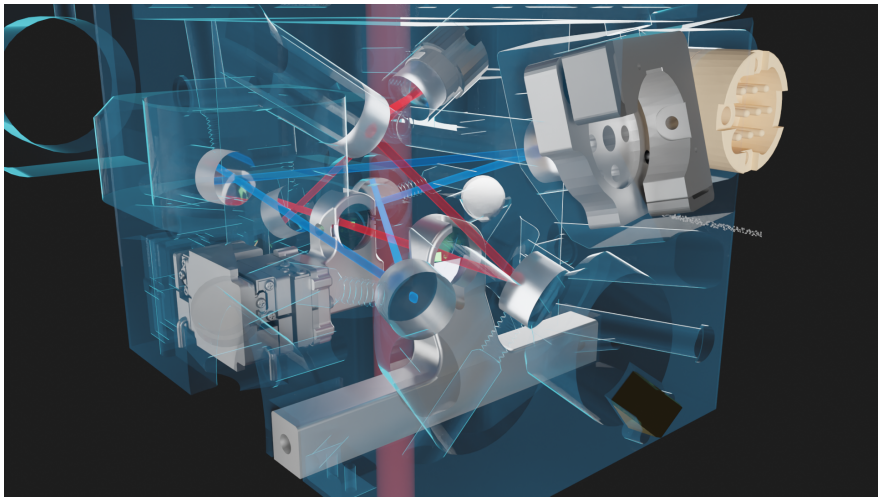


Figure 4.3: X-ray image of the cavity structure with the in-vacuum SmarAct stage. The stage is mounted on the main cavity structure and holds the blue cavity mirror. It allows for precise alignment of the blue cavity with respect to the science cavity.

A complex cavity design as the two twisted cavities have too many degrees of freedom to be aligned by hand. As much of the alignment as possible is therefore achieved by the machining of the cavity structure with tight tolerances, and the lenses are aligned to the

cavity with two 5 axis stages outside vacuum before being fixed with glue. Furthermore, the science and Rydberg cavity need to be overlapped exactly at the location of the atoms that will sit at the waist of the red cavity. In order to perform this alignment in the absence of atoms, a SNOM (scanning near field microscopy) fiber tip was used. A crosshair of thin tungsten wire was used for this purpose historically, however any object that is comparable to the waist size of 18 um will spoil the cavity even when inserted partially. The SNOM tip though is tiny with a diameter of only 50 nm and can therefore be inserted into the cavity, with the additional advantage that the light field at the tip will be evanescent coupled into the fiber and can be used to scan the cavity mode. Even better, once the waist of the science cavity is found the blue light coupled into the fiber can be maximized for optimum overlap. Even though the two cavities can be overlapped very precisely initially, the last iteration of the cavity suffered from a substantial misalignment caused by the vacuum bake out presumably due to differential thermal expansion. This caused the two cavities to drift apart by several 100 um and could only be fixed by using a higher order blue cavity mode at the price of a substantial loss in power. The new cavity structure therefore hosts an in-vacuum translation stage for alignment of the blue cavity. Initially it was considered to have two separate structures for the red and blue cavity and to translate them coaxially using a PicoMotor or similar actuator, but the complexity and space that would take up was prohibitive. Instead, as only small movements of the cavity waist of a few 100 um are necessary, a more elegant solution is to tip and tilt one of the cavity mirrors itself. Intuitively, adding a linear phase gradient across a mirror re-defines the cavity axis to where this phase will be compensated by the quadratic curvature of other mirrors. Tilting a mirror moves the eigen-axis everywhere in the cavity, so in order to find the mirror that provides maximum actuation at the atom location we used both a 5×5 ABCD misalignment formalism and full ray tracing to relate tilt angles θ_i, ϕ_i on mirror i to the sagittal and transversal positions in the

cavity waist. The chosen tip-tilt stage [2] has a compact footprint of only $20\times20\times10$ mm and can be disconnected when not in use to not add mechanical noise. For the two aspheric lenses of the science cavity, the alignment is the most critical and sensitive on the order of a few *nm*: In chapter 3 aberrations were carefully compensated through aspheric coefficients to make as many modes as possible degenerate before higher order aberrations take over. All these calculations assume perfectly aligned lenses, so any small misalignment will re-introduce low order aberrations that we fought so hard to avoid. To lowest order, de-centering of the lenses introduces astigmatism of the form $x^2 - y^2$ which exactly provides the dreaded coupling between modes $LG_{p,l}$ with $\Delta l = \pm 2$ destabilizing the cavity. As this is the dominant effect, we need to be able to correct the relative lens alignment *in-situ* and therefore mount one of the lenses on a three axis translation stage to also tune the lens spacing between different degeneracy points. As we need fine resolution for the transverse degrees of freedom and simultaneously large travel distance of 1 mm to change between the $s = 3$ and $s = 1$ degeneracy point and tune away from degeneracy, we chose a slip-stick piezo translation stage [3] also from SmarAct. Its extremely small size of only $11 \times 11 \times 5.2$ mm allowed us to fit three axes into the cavity structure right next to one of the lenses. Slip-stick actuators have a number of advantages for our application: The Rydberg atoms at high quantum numbers are extremely susceptible to electric fields. While the cavity locking piezos were buried in the metal structure for shielding, this would be impractical for all the stages and would introduce a constantly changing background field when the cavity needs to be aligned. The slip-stick stages on the other hand can be operated in a *linear* mode when driven slowly to scan across the full travel range of the piezo of ≈ 1.2 um. With the optimal position identified, small steps in the *slip-stick* mode are made to that point, after which the piezo can be completely turned off. Due to the stochastic nature of the slip-stick mechanism, the step size is not

2. SmarAct STTSTT-12.7 https://www.smaract.com/en/opto-mechanics/product/stt-12-7-tip-tilt-mirror-mount

3. SL-0610 https://www.smaract.com/en/linear-stages/product/sl-0610-compact-linear-stage

always identical but varies with travel direction, axis, time of last actuation, load and also just randomly. This makes it hard to know the exact position, and some out-of-loop physics signal like the transverse mode splitting has to be used to infer the precise position. For coarse position, all steps are logged by a custom controller GUI to an InfluxDB database as it is otherwise easy to lose track during large moves or between different days. On the moveable lens mount itself a small mirror is affixed that can be used in an interferometer to continuously track the z position with interferometric precision, however it proved tricky to track the position through the very fast slip motion. Therefore, this path was only used for the initial piezo and controller calibration.

Electric field control

Rydberg atoms are among the most sensitive electric field sensors as their electric polarizability scales as n^7 with principal quantum number n. In the experiment, a variety of electric potentials from the cavity locking piezos as well as charges on dielectric surfaces caused by adsorbates or patch charges cause stray electric fields. These shift the Rydberg transition frequency or can also broaden them when they vary across the atomic cloud and cause incoherence of the collective excitation. We therefore use a set of in-vacuum electrodes to compensate the field E_i as well as first order derivatives $\partial E_i/\partial x_j$ with $i, j \in \{x, y, z\}$ at the location of the atoms. It is important that these electrodes are in vacuum as otherwise adjusting these fields will cause mobile charges on windows or glass cells to migrate and thereby change the electric field in an uncontrolled way. Naively, these 12 components for fields and gradients need to be controlled individually, however in free space without charges Maxwells equations

$$\nabla \cdot \mathbf{E} = 0, \nabla \times \mathbf{E} = 0 \tag{4.2}$$

provide 4 constraints reducing the number of electrodes required to 8. As space in the cavity structure is at a premium, the electrodes could not be placed in a particularly symmetric

pattern but were arranged in roughly opposite pairs wherever there was some space. The fields and gradients caused by a voltage V_i on electrode i are then calculated in a finite element simulation of the whole structure in Ansys Maxwell. The resulting matrix

$$M\mathbf{V} = \begin{pmatrix} \mathbf{E}_1 & \dots & \mathbf{E}_N \\ \nabla\mathbf{E}_1 & \dots & \nabla\mathbf{E}_N \end{pmatrix} \begin{pmatrix} V_1 \\ \vdots \\ V_N \end{pmatrix} \tag{4.3}$$

then predicts the fields for a given set of electrode voltages $\mathbf{V}$. Optimizing those electrode voltages individually is the wrong choice of basis as they are coupled and would result in a highly dimensional search space. Instead, we want to scan the fields and gradients directly as they are linear and can be related to the shift of the Rydberg transition through the Stark Hamiltonian. We therefore invert the matrix M to calculate all voltages V_i based on the desired fields. In the real experiment, some electrodes are shielded more than others by e.g. the lens mounts reducing their field strength resulting in an ill-conditioned matrix and potentially unrealistic voltages on those electrodes. To mitigate this problem, the cavity structure itself is treated as a 9th electrode that is isolated from the grounded vacuum chamber through Macor ceramic spacers. This now makes the matrix M 12x9 and overdetermined, and instead of the inverse we apply a pseudo-inverse solving the least squares problem. The new lens based cavity exposes more dielectric surface area due to the lenses close to the atoms, which necessitated much higher compensation voltages than in the previous generation. A custom 9 channel opamp based amplifier with an output voltage of ± 22 V was built that also heavily filters the analog voltages from the DAC in the control room. At high quantum numbers Rydberg atoms are extremely sensitive and electric fields on the order of mV/cm would already shift and for AC voltages broaden the EIT peak by more than its linewidth. Therefore an active, second order lowpass filter with a cutoff frequency of 1 kHz was implemented on each channel to suppress any noise picked up from the lab. As

noise in the RF range of a few MHz is ubiquitous in the lab from AOM drivers, an MFB filter is chosen for its superior performance at high frequencies [4]. In hardware, the in-vacuum

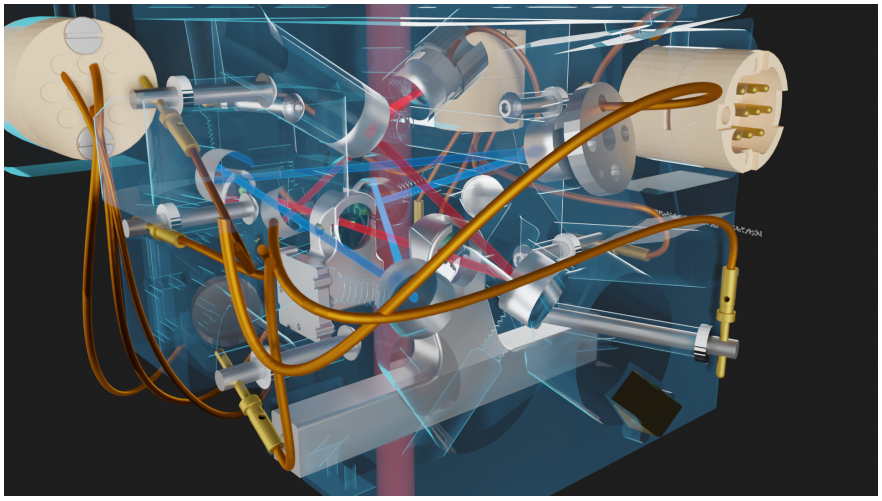


Figure 4.4: X-ray image of the cavity structure with electric field electrodes. The electrodes are formed by stainless steel dome head screws and isolated from the main structure by Macor ceramic washers. The electrodes are connected to Accuglass Sub-C vacuum connectors by press-fitting pre-crimped pins into precision holes in the electrode screws.

electrodes are formed by stainless steel dome head screws and isolated from the main structure by Macor ceramic washers. As dielectrics this close to the atoms were a concern copper washers were stamped out of shim stock to cover the Macor washer from line of sight. The screws are connected to the vacuum feed-throughs by using pre-fabricated Sub-C Kapton wires from Accuglass with pins already crimped on both sides and ordered with custom lengths to minimize excess wire length in the vacuum chamber that could block laser beams. These pins are then press fit into holes that were drilled into the electrode screws on a jig before assembly, providing a strong electrical and mechanical connection without the use of any special vacuum solder or clamping nuts that could become loose during the bake. While it's possible to crimp the pins yourself using the matching tool, the failure rate compared to the factory made ones was much higher during testing as the stripping tool can easily

4. The more common Sallen-Key filter suffers from somewhat surprising non ideal performance https://www.ti.com/lit/an/sloa049d/sloa049d.pdf?ts=1767922303030 that was initially observed on a spectrum analyzer.

damage the thin conductor. This method was therefore chosen for all electrical connections to the cavity structure, and only the piezo wires were soldered at the factory and shipped cleaned.

4.1.1 Laser System

Many different laser sources are used in the experiment (see fig. 4.5), and they can largely be classified into two groups: The majority of lasers is devoted for preparation of the atomic sample from MOT, PGC, trapping and cooling lasers. A special laser is the ultra-narrow 780 nm laser used to probe the hybrid atom cavity system, and these photons are the quantum particles that we make materials out of.

The Rydberg EIT is completed by a control laser at 480 nm driving either the $|5P_{3/2}, F' = 3\rangle \rightarrow |nS_{1/2}\rangle$ or $|5P_{3/2}, F' = 3\rangle \rightarrow |nS_{1/2}\rangle$ transition to Rydberg states using either σ^- or σ^+ polarization. Optical pumping benefits from a smaller number of hyperfine sublevels $|F', m'_F\rangle$ and is therefore performed on the D1 line $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{1/2}, F' = 2\rangle$ using a 795 nm laser with an additional repumper from the lower hyperfine groundstate $|5S_{1/2}, F = 1\rangle \rightarrow |5P_{1/2}, F' = 2\rangle$. The infrared transition $|5P_{3/2}\rangle \rightarrow |4D_{5/2}\rangle$ can be used to shift or modulate the atomic P state using off-resonant AC Stark shifts such as for Floquet modulation, or resonantly in 4-wave mixing schemes for mode conversion and it is coincidentally a convenient telecom band. While there is no true *magic* wavelength in Rubidium that causes equal Stark shifts for S and P state, we can use a near-resonant 1013 nm laser closely red-detuned from the $|6P_{1/2}\rangle \rightarrow |nD_{5/2}\rangle$ transition to provide simultaneous trapping to the group and Rydberg state.

MOT lasers

For the D2 line in Rubidium, convenient diode lasers are available at the wavelength of 780 nanometers. We use three IPS volumetric-Bragg-grating stabilized lasers. One laser is locked

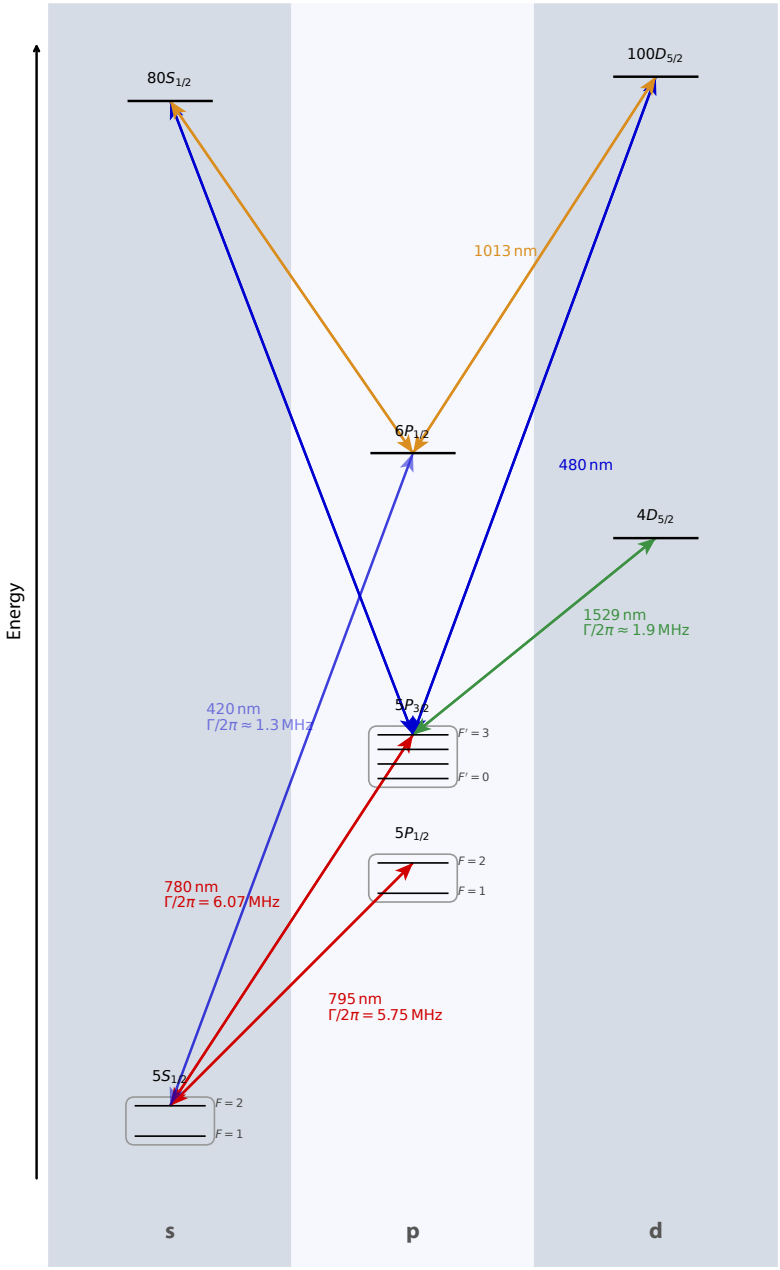


Figure 4.5: Electronic level structure of ^{87}Rb. All of the relevant states and transitions used in the experiment are shown schematically each with their respective laser wavelength and transition linewidth. For Rydberg state transitions the linewidth, as the inverse of the state lifetime that can be long-lived, depends strongly on the exact state and environment that can cause decay through black body radiation. The most prominent transition is the D2 line $|5S_{1/2}, F = 2\rangle \rightarrow |5P_{3/2}, F' = 3\rangle$ with the strongest dipole moment that is used to couple atoms to the cavity as well as provide cooling during MOT and PGC.

to a rubidium transition using Doppler free FM spectroscopy to provide absolute frequency stability called the master laser. Then, MOT and re-pumper laser are both locked to this master laser using frequency offset locks where the offset frequency can be controlled using DDS synthesizer. The initial generation of the experiment used Photodigm DFB laser diodes driven by Thorlabs lasers and TC controllers. These have now mostly been replaced by IPS diodes that have a much narrower line width of 100 kHz compared to the 1-2 MHz of the Photodigm ones. We have also switched to small integrated OEM laser controller modules from Koheron that can be integrated in the laser mount with much smaller footprint and less current noise. Light for the three different MOT axes is amplified by a Thorlabs tapered amplifier (TA) to a power of 1W before distribution. MOT and repumper light are combined into the same optical fibers to simplify alignment.

Lattice lasers

The main lattice laser used for the whole experiment is a Toptical ECDL diode laser at 784 nm amplified by a home-built TA using Eagleyard chips. The wavelength is chosen as a compromise between detuning to reduce resonant scattering and keeping optical powers manageable and ease the requirements on laser safety and beam dumping needed for 100 W dipole trap lasers. The total output of the laser is only around 1.5 W and after AOM and fiber losses the VLAT uses 0.5 W maximum power. To optimize lattice depth for the available power, we use a retro-reflected lattice configuration opposed to two separate paths. As the path length difference introduced by the complicated retro-configuration for transport introduces considerable path length difference, the laser needs to be narrow linewidth to not introduce excessive heating. Therefore, the ECDL is also locked to the ultra-stable cavity that has a linewidth of 50 kHz and allows narrowing of the laser to a sub-kHz linewidth using both piezo and current feedback using a custom analog high-bandwidth lock.

To form the 3D lattice required for dRSC, an additional lattice in the horizontal plane

(HLAT) is added. Since this lattice is only needed during cooling, it is realized using a free-running IPS laser diode and detuned from the D2 line by only 30 GHz as the resulting heating from off-resonant scattering is counteracted by the cooling itself.

Cooling and pumping

Optical pumping is provided by a 795 nm laser operating on the D1 line to the $5P_{\frac{1}{2}}$, which has fewer excited hyperfine states (and crucially as many mF sublevels as the $5S_{\frac{1}{2}}, F = 2$ ground state). A Photodigm DFB diode laser is locked using Doppler-free absorptions spectroscopy to an additional Rubidium reference cell on the $|5S_{\frac{1}{2}}, F = 2\rangle \rightarrow |5P_{\frac{1}{2}}, F' = 1/2\rangle$ crossover feature. Two sidebands are then derived using a fiber EOM to provide tones for the pump and repump from the lower hyperfine $F = 1$ state that can be controlled independently.

Rydberg excitation laser

Atoms are excited to the Rydberg state in a ladder EIT scheme where the first leg is the science cavity at 780 nm and the second leg is a *classical* control field at 480 nm. This light is generated by a Toptica SHG system consisting of an ECDL and TA at 960 nm which is then doubled in a cavity with doubling crystal to provide up to 300 mW of blue light. The oscillator at 960 nm is first locked to the ultra-stable cavity using both a slow (analog lockbox, see appendix A.7) and fast (Toptica FALC 110) feedback to provide a narrow linewidth for coherent Rydberg excitation. This seed light is then amplified by an internal TA Pro to a power of about 1 W and subsequently frequency doubled, where the SHG cavity is locked to the now stable oscillator. The wavelength of the ECDL can be tuned over a wide range of several nm to access a wide range of Rydberg states from as low as $n = 37$ up to $n = 111$. As the level spacing scales as n^{-3} especially at low quantum numbers the alignment of the laser itself and into the fiber needs to be re-optimized after changing the wavelength.

Magic laser

The lattice at 784 nm can introduce significant inhomogeneous broadening to the dark polaritons, as only the atomic ground- but not Rydberg-state is trapped leading to a light shift that depends on the atom position in the trap. While no true magic wavelength exists for Rubidium, the differential shift can be canceled by using a 1010 nm laser that is close to the $6P$ to Rydberg transition and still traps the $5S_{\frac{1}{2}}$ ground state far off-resonantly. As no narrow fiber-based lasers are available at this wavelength, we built a custom tapered laser building on a previous attempt but significantly re-designed to make it stable and useable. The laser is designed around a special tapered amplifier chip [5] that is both amplifier and seed laser in one, and compared to a pure TA chip have a partially reflective coating on the output facet. This partial reflector makes the whole active region part of the feedback cavity that is completed by a frequency selective filter and cateye retro-reflector. The initial seedlight is the amplified spontaneous emission from the waveguide region of the TA chip itself, the rest works like a regular external cavity laser where the narrow interference filter selects a single cavity FSRs by reducing gain for all other modes. This feedback light needs to be coupled exactly back into the waveguide for stable and single mode operation and the previous design was extremely sensitive to any misalignment of the end mirror or thermal drifts of the TA chip position itself due to the several W of heat it generates. In the new iteration of the design, this was solved by making the feedback cavity a 4f system that has several advantages: Most importantly, it re-images the light exactly back onto the waveguide independent of small mirror misalignment, and the interference filter can be placed in a collimated part of the beam improving selectivity. Combined with careful thermal management and machining the laser body out of Invar, the laser is now running single-mode without mode hops and can also be locked to the ultra-stable cavity that is still partially reflective at 1010 nm.

5. DILAS TAL-1010-2000, now sold by Coherent

Ultra-narrow probe laser

The multimode polariton experiment aims to assemble quantum materials out of photons. Here, the cavity photons itself are the quantum particles in contrast to e.g. quantum gas microscopes where atoms are the quantum resource. This requires the photons to be exactly identical and not distinguishable by variations in wavelength due to laser noise, meaning the probe laser at 780 nm therefore needs to be narrow As very narrow fiber lasers are available in the telecom band at 1560 nm a fiber seed and fiber amplifier is used to create up to 2 W of 1560 nm light that is then doubled in a PPLN SHG crystal to 780 nm. For the seed a Orbits Lightwave Ethernal was used [6] that was amplified by a Nufern fiber amplifier. A small portion of the 1560 nm light is picked off and locked to the ultra-stable cavity and is also used for locking of the science cavity. Unfortunately both lasers died in rapid succession due to a power bump and have since been replaced with a NKT Koheras E15 HP erbium fiber laser that can output up to 10 W. It comes as a separate laser and amplifier module and was unfortunately not as narrow free-running as the old Orbits laser was. Therefore, the seed is located right next to the ultra-stable cavity to minimize delay through long optical fibers and cables and enable high-bandwidth feedback through a custom analog lockbox appendix A.7 modified to provide differential output. Still, this laser is not as narrow as the Lightwave Ethernal and has two sharp peaks in its noise spectrum, one caused by the internal switching power supply at around ≈ 125 kHz that could be eliminated by powering it from an external power supply via the extension connector, and a second one at 40 kHz from unknown origin. Luckily, the total frequency excursion caused by this frequency noise component is smaller than the cavity linewidth and a *noise eater* using a fiber EOM and fast feedback from the error signal could be implemented.

6. The *virtual ring laser* employed is quite interesting and can be found in a patent by Orbits, and served as inspiration for a two-mirror running wave cavity

4.2 Cavity locking

We lock many of our lasers to a common ultra-stable cavity to transfer coherence between all the lasers and stabilize them at an absolute frequency. This is especially important for the 780 nm probe and 480 nm Rydberg laser that need to be coherent over the duration of the experimental sequence to not decohere the Rydberg polaritons. Conversely, we also lock the science and blue buildup cavity to the respective lasers at the fundamental wavelength before doubling to stabilize their length. For all of these locks we generate an error signal using the established Pound-Drever-Hall method by using an EOM to add phase modulation at a frequency Ω_1

$$E(t) = E_0 e^{i(\omega_c t + \beta_1 \sin(\Omega_1 t))} = E_0 \sum_{n=-\infty}^{\infty} J_n(\beta_1) e^{i(\omega_c + n\Omega_1)t} \tag{4.4}$$

creating sidebands $\omega_c \pm n\Omega_1$ around the carrier ω_c according to the Jacobi-Anger expansion. As PDH locking is a standard technique, we won't repeat all the details here, however we have developed a few variations to this technique that are worth mentioning. For many lasers the fixed length ultrastable cavity doesn't have a resonance at the exact frequency we want to operate the laser at, and we therefore have to create a feature offset from the laser carrier by a frequency Ω_2. This can be achieved using a two tone modulation, creating what we call *sidebands on sidebands* or three separate PDH features with a tunable spacing. By applying two RF tones to the EOM that are simply summed using a diplexer or splitter, not mixed

$$E(t) = E_0 e^{i(\omega_c t + \beta_1 \sin(\Omega_1 t) + \beta_2 \sin(\Omega_2 t))} \tag{4.5}$$

we can now apply the Jacobi-Anger expansion twice resulting in a double sum

$$E(t) = E_0 \sum_{m=-\infty}^{\infty} J_m(\beta_2) e^{i(\omega_c + m\Omega_2)t} \sum_{n=-\infty}^{\infty} J_n(\beta_1) e^{in\Omega_1 t} \tag{4.6}$$

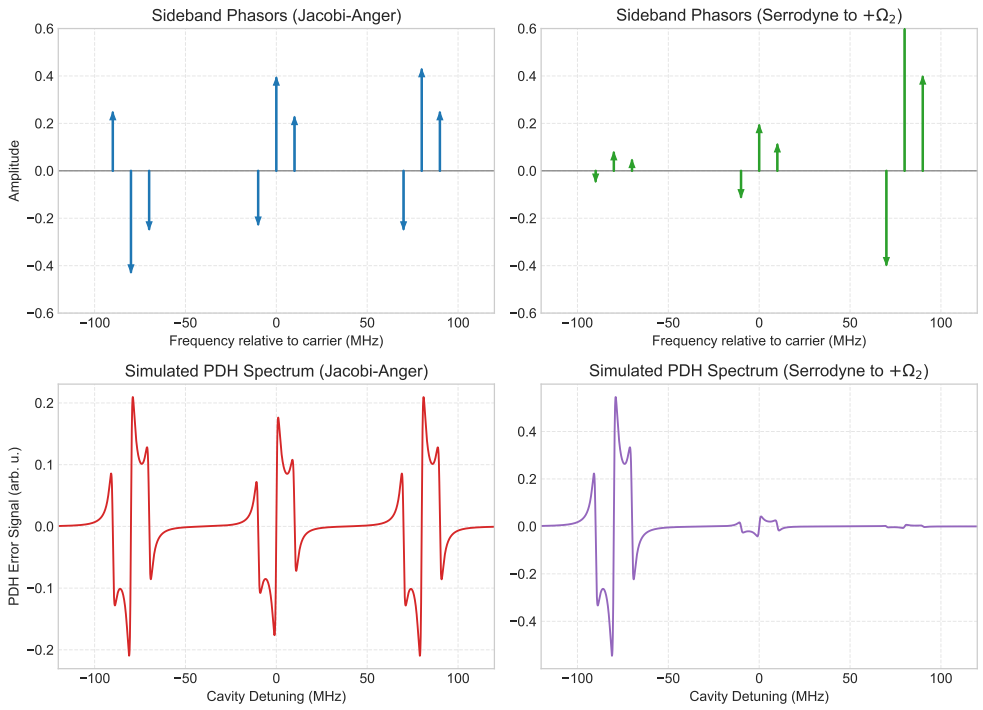


Figure 4.6: PDH locking using sideband on sideband technique. Using harmonic RF tones **(a)** generates three families of PDH locking features with identical slope and similar amplitude. Using serrodyne modulation **(b)** all power can be concentrated in a single sideband allowing for unambiguous and automatic locks.

where the total field is a sum of components with frequencies $\omega_c + \omega_{n,m}$, where $\omega_{n,m} = n\Omega_1 + m\Omega_2$, and amplitudes $C_{n,m} = J_n(\beta_1)J_m(\beta_2)$. When reflecting off the cavity, all the different sidebands experience a complex reflection coefficient $F(\omega_{n,m})$ depending on their position relative to the cavity resonance. The Pound-Drever-Hall error signal is generated by measuring the power $P_{ref} = |E_{ref}|^2$ and demodulating it at the primary modulation frequency Ω_1. This isolates beat notes between components separated by exactly Ω_1 (e.g., between n and $n-1$). Factoring the complex amplitude of the beat note X we find

$$X_m \propto J_1(\beta_1)J_0(\beta_1)\mathbf{J_m}(\beta_\mathbf{2})^\mathbf{2}\left[F_{1,m}F^*_{0,m} - F_{0,m}F^*_{-1,m}\right] \tag{4.7}$$

which is the usual form of the PDH error signal of which we usually measure the imaginary part $Im[X]$, but multiplied with the intensity $\mathbf{J_m}(\beta_\mathbf{2})^\mathbf{2}$ of the offset frequency sideband. This means that in addition to the central one we always get two additional, symmetric copies of the PDH signal at $\pm\Omega_2$ as shown in fig. 4.6. As amplitude of the copies are also Bessel functions we can never eliminate the carrier without creating higher order sidebands nor eliminate one of the copies, often leading to ambiguity when locking a cavity or laser. This can be solved by replacing the harmonic tone Ω_2 with a serrodyne (sawtooth) modulation. Equivalent to a constantly rising voltage when the amplitude is matched to $2 \times V_p i$ this creates a phase that rises constantly in time and thereby acts like a displacement in frequency space. Serrodyne modulation is designed to place all the power into the $m = +1$ sideband, the coefficients for the second modulation become complex numbers $E(t) = E_0 \sum_{m=-\infty}^{\infty} c_m e^{i(\omega_c + m\Omega_2)t} \sum_{n=-\infty}^{\infty} J_n(\beta_1)e^{in\Omega_1 t}$ and for a perfect sawtooth $c_1 = 1$ and $c_i = 0 \forall i \neq 1$. Many analog methods for generating sawtooth shaped RF tones have been developed such as analog oscillators or non-linear microwave transmission lines, our RFSoc based DDS is uniquely capable of directly synthesizing a sharp sawtooth up to ~ 1 GHz. This places almost all optical power into a single sideband as shown in fig. 4.6 (b) allowing for unambiguous and automatic locks.

4.3 Light shaping with DMD

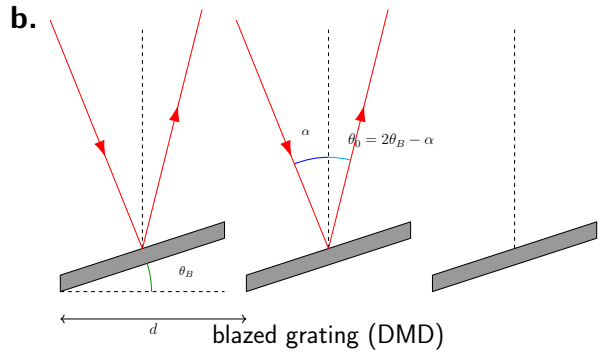
To realize many-body physics and inject light into the different cavity modes that are the orbitals of the LLL, the light sent into the cavity must be shaped into LG_l^p modes or a superposition of them. For this a spatial light modulator (SLM) like an LCos SLM or a Digital Micromirror Device (DMD) can be used. We choose a DMD as for our application we only need to generate very weak beams of very few photons so low diffraction efficiency is not a concern.

DMDs control the amplitude and the phase of the reflected light by switching an array of tiny mirrors between two angles. While in the configuration that is used for white light projectors, it seems it is an amplitude-only device, the fact that the mirrors form a blazed grating makes it possible to control both amplitude and phase in one of the higher diffraction orders. DMDs have the advantage that the pattern can be switched very fast at multiple kHz, and additionally, the mirror configuration can be completely static and don't need a refresh cycle like liquid crystal modulators do, eliminating amplitude noise. Compared to LCos SLMs that are phase-only devices, the hologram for DMDs can be generated in closed form and don't need iterative solvers like Gerchberg-Saxton substantially speeding up hologram generation.

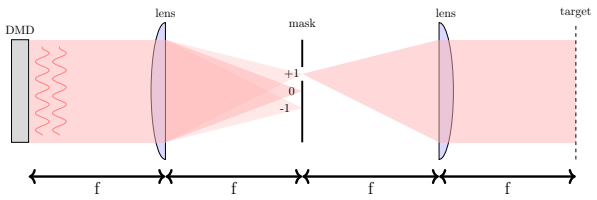
When used with monochromatic light for beam shaping, the grating structure of the DMD chip imposes an additional constraint that needs to be fulfilled in order to achieve reasonable diffraction efficiency. Light is reflected from a single mirror by the law of reflection with the incoming angle $\pm$ the tilt angle. The whole chip, however, is a blazed grating, and the direction of the prompt reflection needs to overlap with a diffraction order of this grating which sensitively depends on pitch and wavelength.

A DMD consists of a pixelated array where each pixel is a mirror tilted by an angle θ_B. When used with coherent light, the device acts as a blazed diffraction grating. The efficiency of the device depends on the alignment of the diffraction orders (determined by the pixel

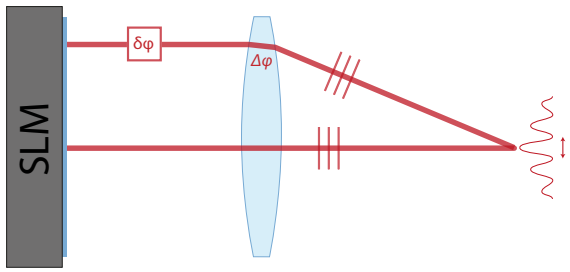
periodicity d) with the envelope of the single-pixel reflection (determined by the tilt angle θ_B).



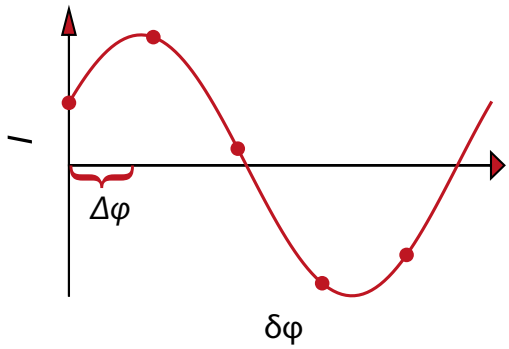
(a) Blazed grating geometry of a DMD showing the blaze angle θ_B and diffraction angles.



(b) 4f optical system for Lee holography using a DMD.



(c) Aberration measurement scheme using a reference beam on the optical axis and a scanning patch probing local phase shifts $\Delta\phi$.



(d) Robust phase reconstruction with phase shift interferometry.

Figure 4.7: Main caption for the whole figure

When this so-called blazing condition is fulfilled, light reflected from the different mirrors interferes constructively along the reflection direction, and the diffraction efficiency is maximized.

The condition for maximum efficiency (the "blazed condition") in this 2D configuration is given by:

$$\sin(2\theta_B - \alpha) + \sin(\alpha) = 2\sin(\theta_B)\cos(\theta_B - \alpha) = p\frac{\sqrt{2}\lambda}{d} \tag{4.8}$$

Where:

θ_B is the micromirror tilt angle. α is the angle of incidence (relative to the normal of the DMD surface). λ is the optical wavelength. d is the pixel pitch. p is the diffraction order (an integer for peak efficiency).

from this a *blazing-number* $\mu = \|2p \bmod 2 - 1\|$ can be defined that is shown in . When p is an integer (blazed), $2p$ is an even integer, so $2p \pmod 2 = 0$, and $\mu = |-1| = 1$. When p is a half-integer (off-blaze), $2p$ is an odd integer, $2p \pmod 2 = 1$, and $\mu = 0$.

The DMD is rotated by 45° in the lab frame to align the mirror tilt axis (diagonal of the diamond pixel geometry) vertically to keep the beam path horizontal. This rotation, combined with viewing the array at an oblique angle β from normal, produces a shear transformation between pixel coordinates (c, r) and lab coordinates (x, y):

$$\begin{pmatrix} x \\ y \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 + \cos\beta \\ \cos\beta - 1 \end{pmatrix} (c - c_0) + \frac{1}{2} \begin{pmatrix} \cos\beta - 1 \\ 1 + \cos\beta \end{pmatrix} (r - r_0) \tag{4.9}$$

where (c_0, r_0) is the array center and $\beta = 25°$ is the output beam angle. The off-diagonal terms arise because the projection compresses distances along the tilt axis (the lab-frame anti-diagonal) by a factor of $\cos\beta$, while distances perpendicular to this axis remain unchanged.

Since we want to generate LG_l^p modes and control their amplitude and phase precisely at the location of the cavity waist, we operate the DMD in the so-called Fourier plane. This means that the DMD is placed in a plane where the generated field is related to that in the cavity waist plane by a Fourier transform that is realized by one of the cavity aspheres itself. Simple beams like Gaussian or Laguerre-Gaussian are self-similar under propagation, so it is sufficient to display the hologram for an LG mode on the DMD. For more complex patterns like two translated Gaussians, we numerically perform an inverse Fourier transform of the desired field before generating the hologram. When trying to populate the different modes of the LLL, especially at degeneracy, mode matching of the input light to the cavity

directly determines the fidelity with which single modes can be addressed. Initial coarse mode matching is performed by walking a pair of mirrors to maximize transmission of the fundamental mode, like for usual cavities. However, to simultaneously optimize coupling for the large number of higher order transverse modes this is not precise enough, and we use the DMD itself to perform the final mode matching. Because of the Fourier plane operation, translating the center of the pattern on the DMD corresponds to a tilt into the cavity, and displaying a phase gradient on the DMD changes the beam position in the cavity plane. Furthermore, the size of the pattern and a quadratic phase mask for de-focus are used to optimize the beam's q-parameter exactly to that of the cavity mode. Especially when performing cavity alignment, but also on a regular basis because of thermal drifts, the mode matching needs to be optimized. The tip-tilt of the x and y axis as well as the waist and defocus are coupled, which means that these are two-dimensional scans that need to be performed with high accuracy. This means that hundreds of holograms need to be computed on-the-fly and sent to the DMD in real-time, as this has been proven to be a severe bottleneck using a previous DMD chip. A completely new DMD system using the Texas Instruments DLC-900 controller and the DLP-6500 DMD chip was therefore designed that supports pattern display over the built-in display port interface. The hologram generation and display was implemented from scratch with full GPU support using the CUDA programming framework to speed up the hologram calculation. For a full-HD pattern that previously took up to a minute to compute, now 100 patterns can be computed in just under a second. To achieve this high-speed, copying of holograms between GPU and CPU for display is fully avoided by using memory mapped GL texture buffers to transfer the images straight from GPU memory.

As shown in the fig. 4.7, even at the correct blazing angle the DMD has multiple diffraction orders. We use a 4f telescope to filter out only the desired +1 diffraction order using an iris before sending it into the cavity. This optical system, including the DMD chip itself

due to not being perfectly flat, adds aberrations that distort the LG modes and prevent good mode-matching. Luckily, we can use the DMD itself for aberration compensation by measuring the aberration phase and displaying a correction. For this, two small patches are displayed on the DMD: One on the optical axis that serves as a phase reference, and a second one that is rastered over the full surface of the DMD. These two patches generate beamlets that get focused in the Fourier plane and generate an interference pattern whose position is then proportional to the local aberration phase. To improve the robustness of the phase reconstruction to amplitude variations from the illumination envelope or variations in fringe spacing, we use phase shift interferometry where a number of offset phases is applied to the reference path and the phase ϕ for a single pixel of the camera can be calculated as

$$\phi = \arctan\left(\frac{\sum_{n=0}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right)}{\sum_{n=0}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right)}\right) \tag{4.10}$$

where I_n is the intensity at the n-th phase step, N is the total number of phase steps. To further increase the signal-to-noise ratio, the phase is reconstructed over a small ROI of the whole camera image, and a linear phase is fit to the reconstructed 2D phase. The piston term then gives the resulting averaged phase offset in the center. From the measurement grid, the phase is interpolated onto the full HD resolution pixel grid using either rectilinear splines or Zernike polynomials.

4.4 Detection path

The materials made of light leak out of the cavity through the photonic part of the polaritons. In contrast to tweezer experiments or quantum gas microscopes, where many photons are scattered from a single atom in imaging, we only get a single photon per polariton. That is in the ideal case without any cavity or detection path loss, which necessitates the use of single photon sensitive detectors at the output of the cavity more akin to quantum optics

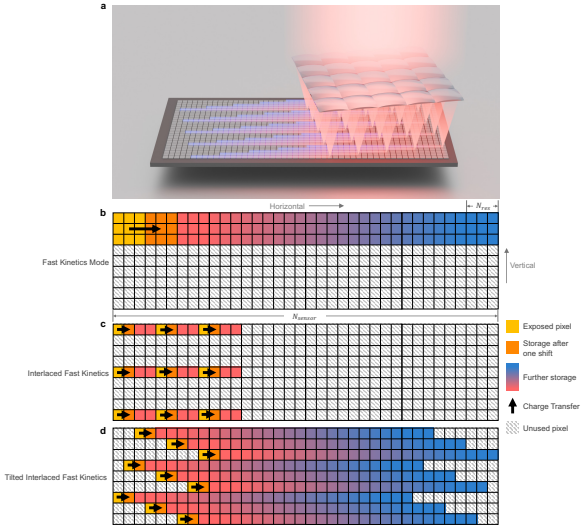
experiments. A series of detectors can be used for different experiments or diagnostics: For coarse cavity alignment or cavity in-coupling, an APD and analog camera measure the transmitted cavity light for strong, classical beams and can be viewed in real time. For any experiments involving atoms, only *single photons* with a rate from a few kHz to MHz can be used and are detected at the cavity output using single photon counter modules (SPCM) from Excelitas [7]. To provide spatial filtering of background light or mode selectivity, a multimode- or single-mode fiber is coupled to the cavity only transmitting the LG_0^0 mode (SM) or a number of higher order modes (MM) for a total field. In the previous iteration of the experiment, cavity modes were not actually degenerate and could therefore be sorted in energy with a series of mode-sorter cavities. Now for a truly degenerate cavity this is no longer possible, and the different cavity modes need to be distinguished through their fields carrying OAM. In principle mode sorter using a hologram/SLM converting OAM into different angles and an array of SPCMs would be possible, but this scales poorly in size, price and efficiency with mode number. Instead, the new apparatus uses a single-photon sensitive EMCCD camera to spatially resolve the position of the photons which is equivalent to the momentum space LG_l due to the special properties of the LLL.

4.5 Ultra-fast EMCCD

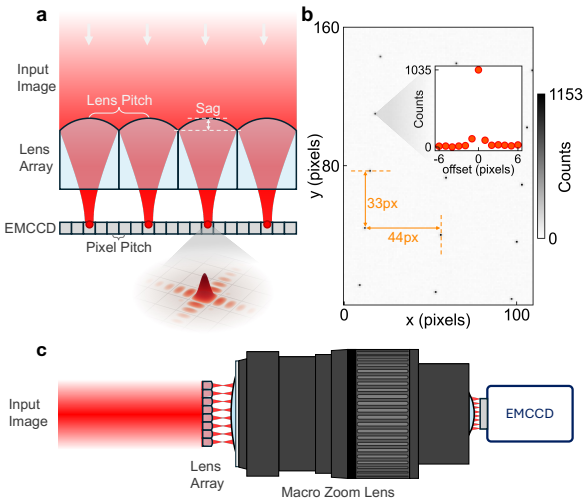
Much like SPCMs or photon multipliers, EMCCD cameras rely on an electron avalanche driven by a high voltage gradient to amplify a single photo-electron into a measurable current. In an EMCCD sensor this happens in special EM-gain registers before any digitization is performed. They excel at discriminating single photons from no photons with low background counts making them ideal for the sparse, *photon-starved* regime of $0.01-0.1\ \mathrm{el}^-/pix/s$ our experiment operates in. While more modern camera technologies like qCMOS are available, they are better at discriminating photon numbers, but still have a worse SNR at such

7. SPCM-AQRH-13-FC-ND, SPCM-AQRH-14-FC-ND

low photon flux. Another alternative would be an SPAD camera with relatively small pixel count (usually 32×32) that can provide integrated time-tagging for some modules, however they still suffer from low quantum efficiencies even with integrated micro-lenses.



(a) Optical architecture using MLA **(a)** and frame transfer dynamics of the CCD in **(b)** fast kinetics **(c)** interlaced fast kinetics and **(d)** with rotated MLA.



(b) Optical MLA requirements for sub-pixel focus size **(a)** and resulting spillover **(b)** with custom re-imaging objective **(c)**.

Figure 4.8: High level interleaved fast kinetics EMCCD idea and optical design. Interleaved fast kinetics exploits the frame transfer dynamics of the CCD sensor to store data on the sensor itself trading spatial for temporal resolution. The technique requires focusing of each lenset to a sub-pixel PSF to minimize spatial and temporal crosstalk.

4.5.1 Optical Design

We expand on existing EMCCD technology through a special optical setup to turn it into an ultra-fast camera exploiting the intrinsic row-transfer or *shifting* dynamics to achieve high time-resolution similar to streaking. The central idea is illustrated in fig. 4.8a: In Fast Kinetics mode, the camera can expose a small ROI (yellow) of $N_{res} \times N_{res}$ pixel multiple times in rapid succession by shifting it N_{res} between every exposure. This still limits the maximum image rate to f_s/N_{res} where the transfer clock $f_s \approx 2$ MHz. The new idea is to use a micro lens array (MLA) in front of the sensor to map the incident field onto a sparser

grid of pixels trading in spatial for temporal resolution. Now, a new image can be acquired after shifting over the sensor pixels by only one column enabling acquisition at the full f_s. In the simple, axis aligned grid many rows go unused, so rotating the MLA precisely for every pixel to occupy its own row increases the storage capacity to the full N_{sensor} images. The

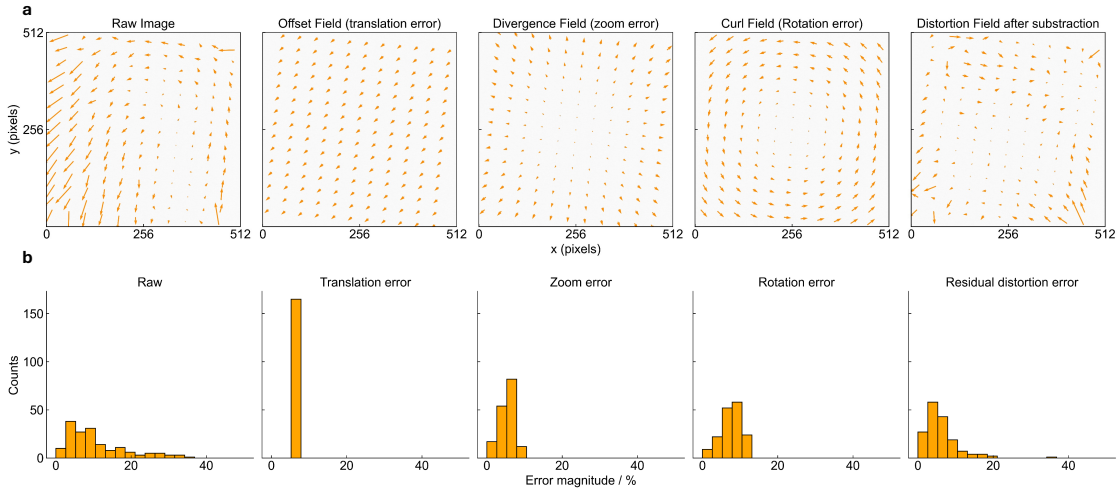


Figure 4.9: Fast EMCCD Shack-Hartman alignment. In order for the interlaced fast kinetics on the EMC-CD camera to work, the foci of the MLA need to lay exactly on a single pixel of the sensor grid. As this has to be true for every lens lit corresponding to a macro pixel, the zoom, rotation, and translation of the MLA have to be aligned exactly to the underlying grid. Any misalignment will cause light to leak out into a neighboring pixel. We use finite differences of the nearest neighbor pixels to compute a misalignment vector, similar to a Shack-Hartman wavefront sensor. The raw disorder vector field in brackets, left column, can be decomposed into: translation (offset), zoom (divergence), rotation (curl) and a residual random error. This is performed in real time and allows for precise alignment of all these degrees of freedom.

optical setup to achieve this is shown in fig. 4.8b: Light from a single lenslet needs to be focused to below the pixel pitch to not introduce any spatial (rows) or temporal (columns) bleed out and reduction of contrast. Also, the MLA pitch must be exactly matched to an integer number of pixels taking into account the rotation of one row per column spacing. Custom MLAs with these properties that also have a large sag are too expensive for a proof-of-principle experiment, which is why we have to use a re-imaging system and an off-the-shelf MLA. As the MLA creates an array of diffraction limited foci, the re-imaging system must admit a high NA to not lose any light and simultaneously have a large field of view to image

the whole array uniformly onto the sensor. Furthermore, the flange distance of the EMCCD camera is 17.75 mm, longer than the focal length of the MLA chosen (Edmund Optics 72238; 17.6×18 mm) also requiring a relay system. A 4f re-imaging system with a large field of view that allows maintaining a diffraction limited focus across the whole array is challenging and cannot be achieved with only two spherical lenses. Instead, we use a commercial macro zoom lens with adjustable magnification (Canon MP-E 65 mm f/2.8 1–5x Macro) to re-image the focal points onto the sensor. Because the array is much larger than the sensor, we use the objective in reverse configuration to *de-magnify* the image by $\approx 0.3\times$.

4.5.2 Noise characteristics

4.6 Experimental Sequence

In our hybrid experiment platform, atoms are necessary to mediate any kind of interaction between the cavity photons, without them the experiment would be linear and could only probe *single particle* physics. Hypothetically, if a non-linear crystal exhibiting a significant non-linearity at the single photon level existed, we could simply place it inside the cavity, and we wouldn't need to operate in a vacuum chamber with all the attached difficulties. However, even for the best currently available non-linear crystal materials such as LN, BBO or KTP the non-linear $\chi^{(2)}$ coefficients are only on the order of requiring millions of photons to show non-linear behaviour. While this enables mean-field interactions in some applications such as solitons , it is far from the strong interactions required to reach true many-body physics. The strongest non-linearities achievable at optical photon energies are provided by single atoms. A critical part of the experiment is therefore to prepare a cloud of ultra-cold atoms located at the cavity waist to achieve strong collective atom-photon coupling and thereby strong photon-photon interactions. The preparation of this cold sample is a sequence involving many steps where the atoms are first slowed, trapped and cooled from a thermal background

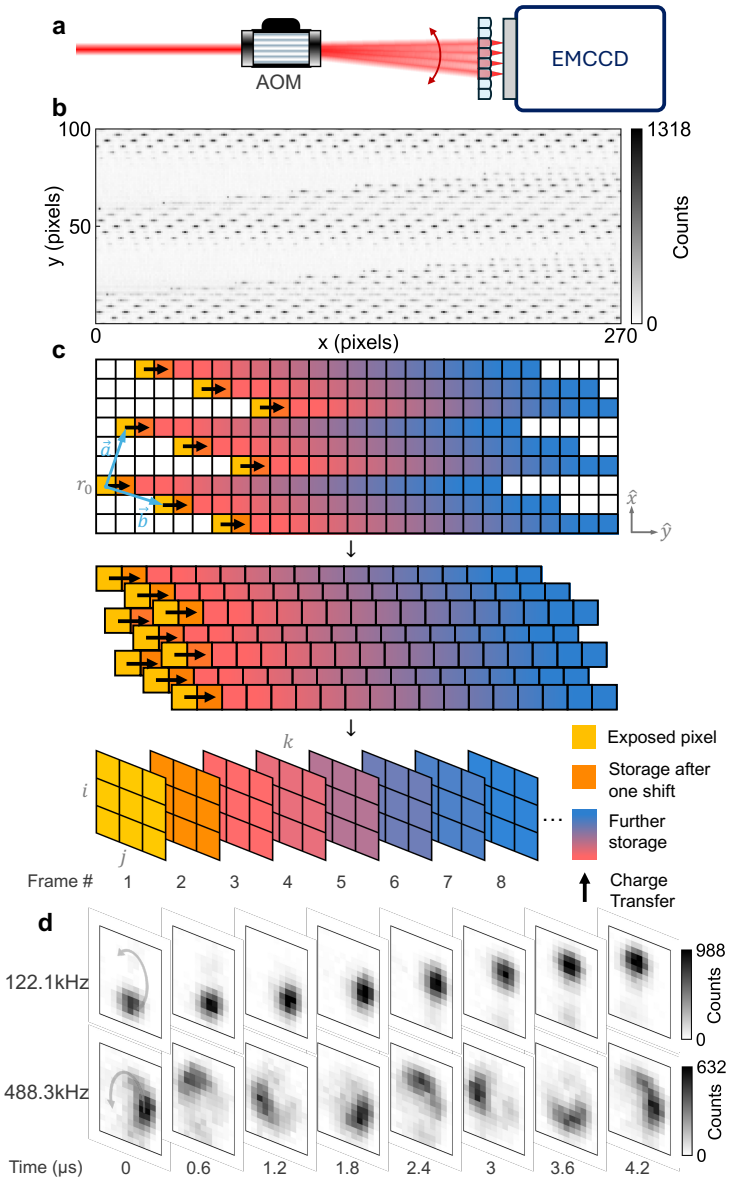


Figure 4.10: All optical demonstration of the FAST EMCCD camera. In order to benchmark the spatial and temporal contrast of the interlaced fast kinetic scheme, we create rotating beam in two dimensions using crossed AODs. A resulting single exposure of the whole sensor **(b)** shows the raw streaked image. By reconstructing the movie through pixel reordering **(c)** we can obtain the final movie **(d)** at different rotation speeds.

gas of Rubidium vapour, transported into the cavity waist, and after further cooling and optically pumping ready to be interrogated through the hybrid atom-cavity system. During one run (*shot*) of the experiment, the same sample can be probed many times as the lifetime of the atoms on the order of several 100 ms is over an order of magnitude longer than the polariton lifetime on the order of 10 us. This way, the loading time of the atoms can be offset by a long probe time to reach a reasonable duty cycle of data collection. Compared to more conventional cold atom platforms such as BEC or Quantum Gas Microscope experiments, we don't need to reach quantum degeneracy of our gas. This eliminates the need for the most time-intensive part of such experiments, the evaporation, that can easily last for up to 30 s. Instead, we can reach the desired temperature of ≈ 500 nK and required phase space densities using sideband cooling in a few ms, leading to a total run time of only 300 ms and a high repetition rate of 3 Hz.

4.6.1 Steps

MOT

A typical experimental sequence begins by loading an initial sample of ^{87}Rb atoms from a thermal background gas in the vacuum chamber using a 3D magneto-optical trap (MOT). We use a standard 6 beam MOT using 3 retro-reflected beams for increased optical power and compactness. The magnetic field gradient responsible for the spatially trapping character of the MOT is generated by a pair of in-vacuum coils in an anti-Helmholtz configuration at a current of 4.5 A. The use of coils inside the vacuum chamber is non-traditional but allows for much smaller coils not obstructing optical viewports, alleviates the need of water cooling at lower currents, and has less inductance and reduced eddy currents in the vacuum chamber enabling faster field switching. We load around atoms at a temperature of uK although the exact atom number depends on the time spent loading in between two shots of the experiment and the loading rate determined by the getter current. In the steady

state, our MOT is around 1 mm in size and density-limited by re-scattering of photons inside the cloud. Both the MOT temperature and loading rate are measured in fluorescence imaging and reveal a vacuum-limited atom lifetime of ≈ 2 s which is sufficient for our kind of experiments.

dMOT

The radial waist of our transport lattice of around 100 um is significantly smaller than the size of the MOT. Therefore, the maximum number of atoms that can be loaded into the lattice isn't so much dependent on the total size of the MOT but much more on its density in the center. The previously described re-scattering of photons limits this density and can only be reduced by lowering the MOT laser intensity, which would come at the price of a reduced capture range. Historically, several techniques have been employed to circumvent this problem all utilizing the concept of a *dark* MOT (dMOT). This is most often achieved by spatially occluding the repumper laser in the very center of the MOT, where after many scattering events these now coldest atoms decay and accumulate in the lower hyperfine ground state that is dark to the MOT light. Usually this requires imaging a mask onto the MOT which is difficult in our 6 beam configuration where the MOT and repumper light is combined in all the fibers. Instead, we utilize a temporal dark MOT that gradually lowers the laser intensities to keep most of the already captured atoms but increase their density in the center of the MOT. After carefully tuning those parameters, the number of atoms loaded into the lattice can be increased by 50%.

PGC

In order to load a macroscopic fraction of the MOT into an optical transport lattice, the temperature of the atoms has to be reduced further to sub-Doppler temperatures. We use polarization-gradient-cooling (PGC) using the same set of lasers although at larger detuning

and lower intensities and turn off the magnetic quadrupole field. The counter-propagating σ^+ and σ^- polarized beams create a complex, position dependent polarization gradient that the atoms move in. Once they have 'climbed' to the peak of one potential they are more likely to scatter a photon and decay into the opposite Zeeman sublevel losing kinetic energy and in a 'Sisyphus' like manner repeat this process over and over leading to a net cooling down to the recoil limit. The exact process is complicated and depends sensitively on the laser detuning and power and especially the absence of any magnetic fields requiring careful nulling of any background fields. With careful optimization, we can reach temperatures down to 5 uK leading to an efficient loading of $\approx 50\%$ of all MOT atoms into the optical transport lattice.

Transport

As the MOT can not be prepared inside of the optical cavity structure due to limited optical access, the atoms need to be transported into the cavity waist after the initial cooling over a macroscopic distance of 45 mm. The transport lattice is formed by a retro-reflected 784 nm laser with a power of 400 mW compromising a relatively close detuning for a relatively low optical power. The lattice is turned into a *conveyor-belt* by dynamically detuning the retro-reflected beam through two double-passed AOMs shifting to frequency of the second beam by up to 10 MHz. The atoms trapped at the anti-nodes as points of constant phase $\phi = 2kx - 2\pi\delta ft$ for a detuning δf therefore move with a velocity $v = \frac{\delta f \lambda}{2}$. The detuning is ramped in a trapezoidal profile to limit the peak acceleration and velocity beyond which the atoms can be lost from the traps. Experimentally no atoms are lost for accelerations of 200 g and more allow for rapid transport in only 15ms. The exact transport distance can be fine-tuned by optimizing the collective coupling strength g which overlaps the cavity mode with the densest part of the atomic sample.

dRSC

Once in the cavity structure, the atoms are further cooled using 3D degenerate Raman sideband cooling (dRSC). The transport lattice in vertical direction (VLAT) is combined with an additional 2D lattice in horizontal direction (HLAT) to form a 3D lattice with matched trap frequencies along all three axes. For the *degenerate* part of dRSC the vibrational level spacing given by the trap frequency is matched to the Zeeman splitting between magnetic hyper-fine sublevels. The HLAT is created by a 780 nm laser detuned from the D2 line by only 30 GHz and with one of the retro-beams polarizations rotated, it enables Raman transitions between levels $|Zeeman\rangle \otimes |vibrational\rangle = |m, n+1\rangle \rightarrow |m-1, n\rangle$. The actual cooling is then provided by an optical pumping beam on the D1 line at 795 nm with well-defined σ^+ polarization. $\eta = k\sqrt{\frac{\hbar}{2m\omega}} = \frac{2\pi}{\lambda}x_0$ where ω is the trap frequency and $x_0 = \sqrt{\frac{\hbar}{2m\omega}}$ is the ground state wavefunction spread (harmonic oscillator length). The Lamb-Dicke regime condition $\eta\sqrt{\langle(a+a^\dagger)^2\rangle} \ll 1$ the dipole operator can be expanded as $e^{ik\hat{x}} \approx 1 + i\eta(a+a^\dagger)$.

Lattices & Trap Ramps

Several trap geometries for the atoms can be used depending on whether a 0D, 1D, 2D, or 3D system should be realized. The *pseudo*[8]-dimensionality is determined by the overlap of the atomic cloud with the cavity mode. The cavity modes quantize the two transverse degrees of freedom of the polaritons, while the size of the atomic cloud along the cavity axis limits their longitudinal degrees of freedom. Utilizing only the fundamental LG_0^0 mode of the cavity, a 1D system of polaritons can be realized for a cloud thicker $\sigma_z \gg \omega_0$ than the cavity waist or a 0D 'quantum dot' for $\sigma_z < \omega_0$. Higher order LG_l^p modes then add two additional transverse degrees of freedom although the geometry is constrained through the lowest Landau level. As the QHE effect only occurs in 2D we need to confine the atoms to a

8. Statements about reduced dimensionality can be somewhat controversial as our physical world is always three-dimensional. However, if the size of the sample along some direction is smaller than all other characteristic length scales of the system, an effective reduced dimensionality can be realized.

thin 'pancake' $\sigma_z < \omega_0 < R_b$ along the cavity axis where the real limit is set by the Rydberg blockade radius R_b to allow only a single longitudinal excitation. The atomic sample after transport and dRSC is roughly isotropic and of the scale of the VLAT lattice waist. In order to compress the cloud while avoiding the complications of occupying only a single anti-node of a real lattice along the cavity, we instead use an *elliptical* lattice (ELAT) with highly anisotropic waist sizes. The ELAT is generated from the same 784 nm laser as the VLAT, and is derived from the undiffracted order of the VLAT AOM so the two can share power and smoothly hand off. The ELAT is detuned by 160 MHz to avoid any interference between the two that could introduce phase noise and heating due to the different paths. In order to achieve a thin atomic sample, the aspect ratio of the ELAT is expanded by a cylindrical telescope on one axis to fill the subsequent 1 inch optics. This astigmatic beam gets focused by the lens on the top of the chamber down to a waist of ≈ 7.5 μm significantly smaller than the cavity waist. As both lattices have to share the same path, they are multiplexed onto the atomic cloud in polarization using PBSs on the top and bottom of the chamber to also allow two different retro-reflection paths. While the VLAT is retro-reflected by two-AOM double pass cat eye, the ELAT is retro-reflected by an interference filter as a static mirror. This reduces lattice heating from phase noise and doesn't clip the large numerical aperture beam. The ELAT beam is re-imaged onto the atomic cloud through a 4f imaging system to make sure forward and backwards pass are always aligned. In practice, the two lattices are intentionally misaligned by a small angle to make sure the other lattice leakage from polarization imperfections is not retro-reflected in the wrong path. Because the high NA of the ELAT beam the depth of focus is very short with a Rayleigh range of $z_R \approx 140$ μm, and to be able to adjust the focus exactly to the atoms an intermediate 4f re-imaging system is used where one lens can be electronically actuated to detune the telescope and tune the focus position.

In order to transfer the atoms from the VLAT into the ELAT, the lattices are ramped

simultaneously using an adiabatic ramp. As the axis along the cavity is compressed considerably, some atoms are lost where there is no overlap and the transferred atoms increase in temperature. A second dRSC step is therefore performed in the ELAT to dissipate this kinetic energy but the cooling is less efficient due to the much more inhomogeneous trap frequencies. Finally, the lattice power is slowly ramped to the trap depth used in the probe cycle, which needs to be carefully balanced to be deep enough to not lose atoms due to scattering but also not introduce excessive inhomogeneous broadening. As no magic wavelength is available in $^{8}7$Rb for the $5S_{\frac{1}{2}}$ to $5P_{\frac{3}{2}}$ transition, this ultimately limits what collective coupling g due to atom number is achievable when not increasing dark polariton linewidth beyond other decoherence channels.

Optical Pumping

All atoms in the cloud need to be initialized in the same state in order to realize the clean three-level EIT required for long-lived polaritons. If some atoms were to occupy different Zeeman sublevels, they could readily absorb cavity photons and thereby broaden the EIT peak substantially. We therefore pump all atoms into the $|5S_{\frac{1}{2}}, F = 2, m_F = 2\rangle$ hyperfine groundstate using the same 795 nm laser as for cooling. Ideally, we would want both the cavity and Rydberg excitation light to be perfectly circular to not drive transitions to other intermediate or Rydberg levels. Because of the intra-cavity lenses however, the two beams can not be perfectly counter-propagating but subtend a small angle. For reasons that subtly depend on the different coupling schemes for low-lying electronic and Rydberg states [9], we pick the quantization axis to point along the blue cavity axis. The blue cavity mirrors are designed to be transparent for the pumping wavelength allowing the two beams to be co-aligned. The purity of the pumping beams σ^+ polarization determines the final state occupation fidelity and is therefore carefully optimized using motorized waveplates where

9. Low-lying states couple the electrons total angular momentum J to the nuclear spin I result in eigenstates of the hyperfine structure F, whereas Rydberg states remain J coupled

both a $\lambda/4$ and $\lambda/2$ waveplate compensate any birefringence in the dichroic or cavity mirrors. In addition to defining the quantization axis, a bias magnetic field of 5 G is applied to lift the degenerate Zeeman sublevels and provide some spectroscopic protection from exciting the wrong Rydberg sublevels. The bias coils are carefully ramped over 10 ms followed by an additional 15 ms settling time for any eddy currents to decay.

Probe

After the atom state preparation, the probe is the part of the sequence where the actual science happens. The details are highly specific to the type of experiment but two general regimes can be identified. Our hybrid atom-cavity system can either be probed in the steady-state or in a time-dependent fashion: A steady-state probe is the most straight forward mode of operation and the cavity is continuously interrogated by the probe laser. This continuously creates new polaritons in the system if the laser is resonant with either the dark or bright polaritons, and the photonic part of the polaritons leaking out of the cavity can be detected as a *transmission* signal. By sweeping the probe laser frequency over the duration of the probe, a spectroscopy of the system is performed (under the assumption that the sweep is much slower than the lifetime of any quasi particle in the system). If the spectrum and thereby the frequency of the VRS peaks i.e. bright polaritons is known, the probe laser can be set to a slope of these features and thereby provide a dispersive readout of e.g. the atom number over time. Alternatively, time-resolved experiments can be performed by pulsing the probe laser rapidly using a EOM and a fast RF switch. The photons leaking out of the cavity are still detected by SPCMs but the single photon events are time-tagged using the photon timer with ns time resolution. This way, dynamics on the timescale of the polariton lifetime can be resolved such as Rabi oscillations or ring-down measurements. As the dark polariton lifetime is 10 us or less, a large number of repetitions can be performed in a single probe cycle to dramatically increase data rate. In both cases, atoms inside the cavity scatter photons

which can slowly depolarize the sample. We therefore perform periodical pumping and repumping to pump atoms back to the initial state and recover atoms that have decayed to the $F = 1$ lower hyperfine state. The probe sequence is therefore divided into several *repetitions* with short interleaved OP/REP periods. Those *repetitions* can be further divided into a number of *sub-repetitions* that are usually chosen to be on the order of 50 us long to allow enough time for Rydberg excitations to decay.

4.7 Characterization and Optimization of a Hybrid Cavity Rydberg Polariton experiment

The hybrid nature of our experiment combining an exotic twisted cavity with highly excited Rydberg atoms allows us to combine their properties in a novel way and enables reaching the challenging FQHE regime. While the synthesis of cavity QED and strong interactions enables us to access a unique regime of parameters, it also comes with a unique set of challenges in that it requires optimal operation of both the optical resonator as well as a cold atomic sample for Rydberg blockade. The ideal model of a LLL for interacting photons as it is described in chapter 3 is only valid if many possible imperfections originating from misalignment of the cavity, atomic cloud or sources of incoherence are carefully eliminated. This section provides an overview and guide on how to characterize and optimize every aspect of the experiment that is required for long lived dark polaritons on many cavity modes. To serve as a guide for future generations of students taking on this experiment we describe the key steps and techniques that require physical insights, while some purely technical tricks are deferred to the appendix A.

As this experiment has grown in complexity over the years the biggest overall challenge is the curse of dimensionality: Even when every subsystem is working for 80% of the time or with 80% of fidelity, the large number of cooling, trapping, pumping, stray field optimization, cavity alignment and detection efficiency optimization steps can quickly lead to a small

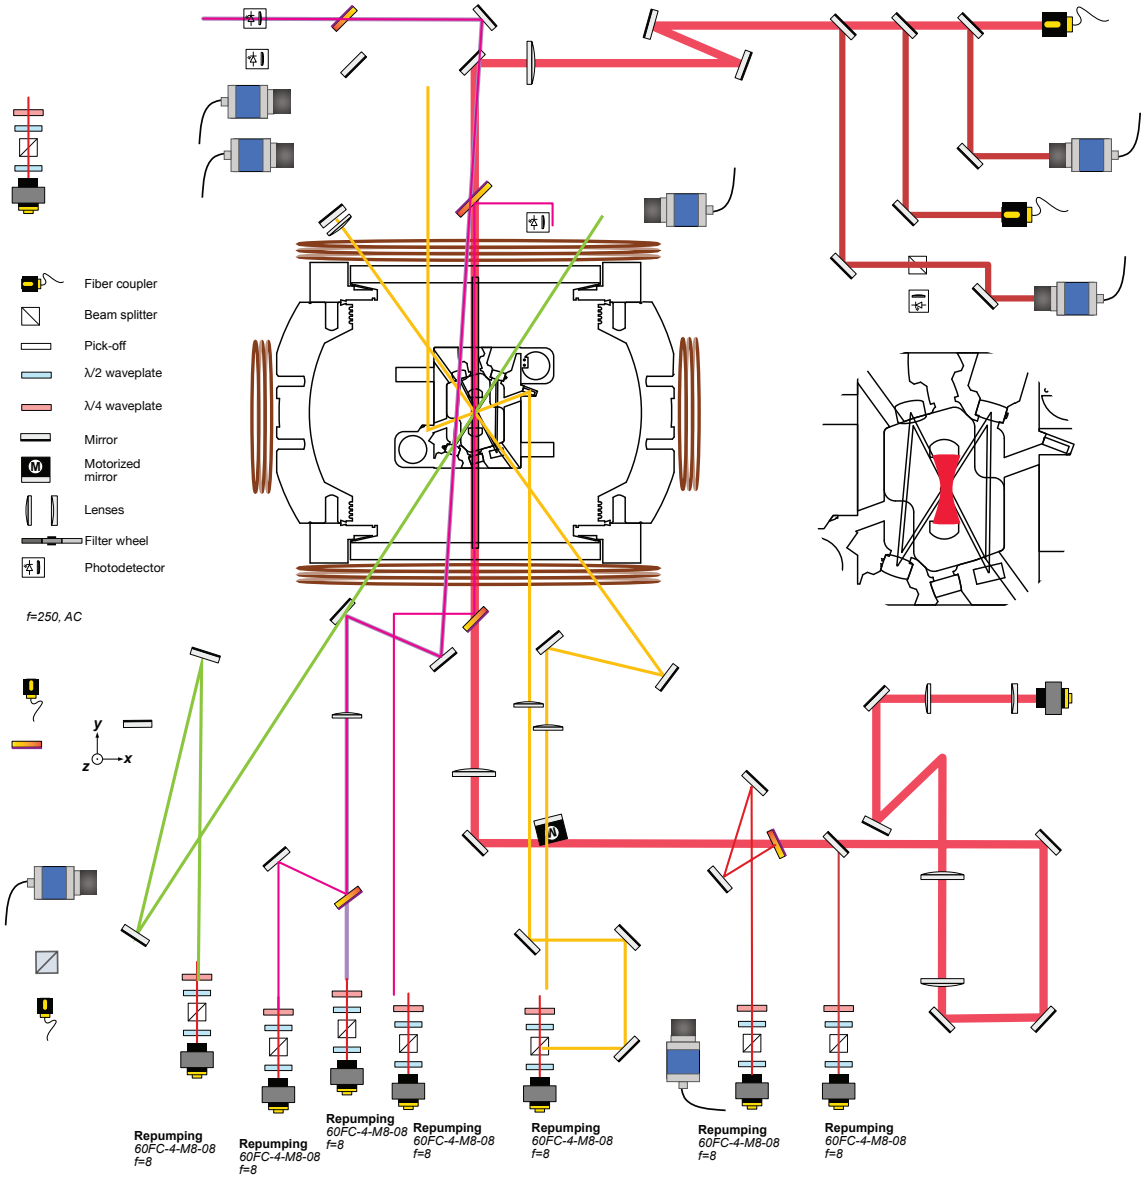


Figure 4.11: Main optical setup. Central cavity structure with asphere based science and blue Rydberg buildup cavity in the vacuum chamber. The main breadboards host all optics for cavity probing and locking as well as atom sample preparation. To excite and detect many higher order *LG* modes of the cavity the main cavity beam (red) from the DMD to the different detection paths (single- and multi-mode fiber, EMCCD, CMOS/PD) requires the highest NA. Both cavities at wavelengths (780 nm and 480 nm) are stabilized in length using locking light at the second harmonic (1560 nm and 960 nm) available from the respective laser source before doubling.

window to perform actual science since $0.8^{10} \sim 0.1$. This management of complexity is the key challenge an experimenter faces and it is therefore crucial to make every system as robust as possible and come up with clear, quantitative metrics to be optimized. At the same time, the machine has to be flexible enough to quickly iterate on new ideas and make modifications without too much technical overhead.

In the following sections we will show experimental data for every major step of the atom preparation sequence and outline how to optimize its performance.

4.7.1 MOT



Achieving a high collective cooperativity in the cavity starts by preparing a large enough and dense sample of cold atoms as they can only be lost during the subsequent steps. In an image of the MOT is shown obtained using Fluorescence imaging showing a cloud with size of um. To achieve a large capture range and load enough atoms from the background pressure of Rb provided by the getters, we expand the MOT beams to 1" and use 150mW of optical power in the hMOT and 200mW of power in the shared dMOT path. Since MOT is a fairly standard technique, we will not describe all the details to optimize it, but only the improvements we have made to it. As we have to transport the atomic sample into the cavity using the VLAT transport lattice, the key to achieving a large collective g is not to maximize the *size* but the *central density* of the MOT as the lattice only cuts out the very central, cigar shaped section of it. There is an inherent tradeoff coming with a large scattering range between large capture range of the MOT to trap many atoms that however at the same time limits the central density due to re-scattering of MOT photons inside the cloud once it gets optically dense. This tradeoff in the steady state can be overcome by implementing a *dark* MOT that masks the repump beam in the cloud center or a *compressed* MOT that dynamically ramps the magnetic field gradient of the anti-Helmholtz coils to compress the

sample after initial loading. Both techniques are difficult to implement in our experiment as the in-vacuum MOT coils can not handle the heat load of a drastically increased current and the repump beam is added on the laser table to the all the fibers carrying MOT light. Instead, we implement a temporal dark MOT where intensity and detuning of the cooling laser are adjusted dynamically to reduce overall scattering and thereby re-scattering increasing the central density of the MOT and improving the number of loaded atoms by over 50%.

4.7.2 PGC



Atoms in the MOT are still too hot to load them into the optical transport lattice with a depth of only $\sim$ 220 uK. Also the lattice is a conservative potential and in order to populate the lower vibrational levels we need to add dissipation in the form of friction. This is both achieved using PGC with the same set of lasers as the MOT, but with any magnetic fields carefully nulled using external bias coils. We can measure the temperature of the atoms using time-of-flight (TOF) and achieve a temperature of down to $\sim$ 5 uK after careful optimization. Empirically we have found that the number of atoms loaded into the lattice can be further improved by turning off the repump laser from the $F = 1$ manifold early by a few ms. This is theorized to act similar to a dark MOT where atoms that have undergone a lot of scattering events can fall down to and build up in the $F = 1$ hyperfine manifold where they become dark to the PGC beam and not contribute to rescattering. As these atoms will also be dark to the imaging light, repumping using the imaging beam or optical pumping in the cavity needs to be performed in order to notice this improvement.

4.7.3 Transport



With the help of cooling and optical molasses during PGC, a large fraction of the MOT can be successfully loaded into the transport lattice. As shown in (c), by optimizing the temporal dark MOT and early stopping of repumping during PGC, we can transfer up to 50% of the total number of atoms from MOT into the transport lattice leaving behind a cloud of similar size. Since the transport lattice is very deep and we are transporting against gravity along the direction of strong confinement in the lattice, we can transport with an acceleration of up to $200g$, completing transport over 40 mm in only 15 ms.

4.7.4 dRSC



Once transported to inside the cavity structure, the cloud of atoms undergoes another cooling stage to reduce its temperature before probing and minimize Doppler induced broadening of the dark polariton caused by the phase vector mismatch of cavity and control field. As the primary role of the atoms is mediating interactions we do not need to reach degeneracy as BEC experiments do, but a low temperature still helps in achieving a high phase space and thereby optical density. The result of round of dRSC cooling is shown in fig. 4.12 as both an absorption image post cooling (a) and measuring the temperature via TOF to be $T_x = 277$ nK, $T_y = 374$ nK. This temperature seems to be close to the limit of what is achievable in ^{87}Rb [UVR+19].

4.7.5 Optical Pumping



Our two twisted cavities have been carefully engineered through compensating birefringence at large angles via custom coatings. The eigenmodes of the science and blue build-up cavity therefore are circularly polarized with high fidelity, allowing clean driving of the

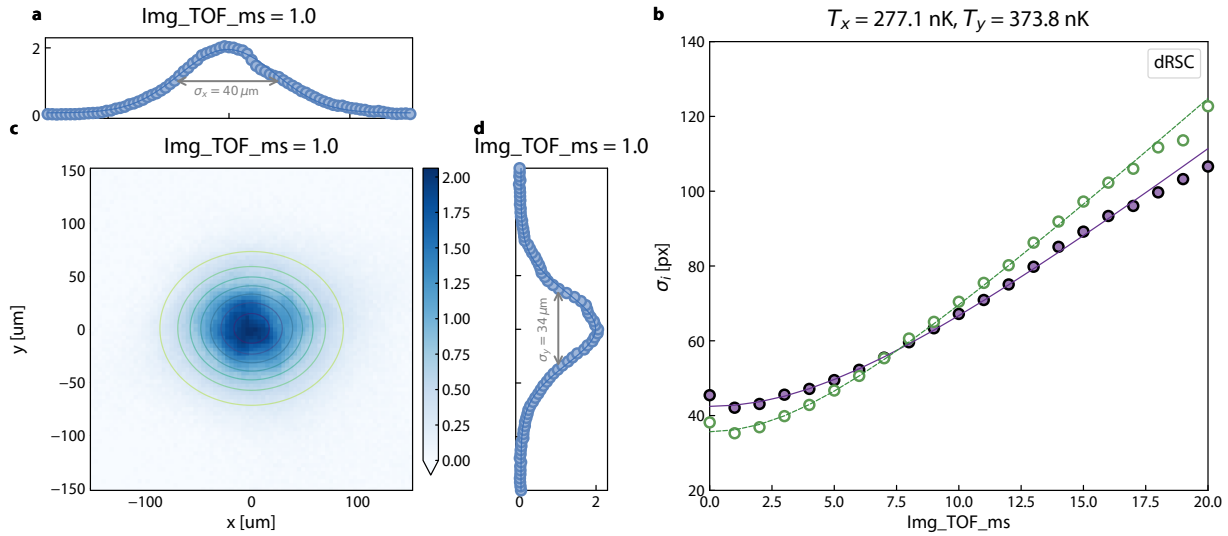


Figure 4.12: Degenerate Raman sideband cooling. The atomic cloud inside the cavity structure is cooled in a 3D lattice formed by the 1D VLAT and the 2D HLAT in a Zeeman field matched to the vibrational level spacing. Panel **(a)** shows an absorption image of the atomic cloud from the top of the chamber taken along the VLAT direction. The radial extent after cooling measures $\omega_x = 34$ um, $\omega_y = 34$ um from a Gaussian fit (line cuts shown in **(b)** and **(c)**). In **(d)** we measure the temperature of the atomic sample in time of flight to be $T_x = 277$ nK, $T_y = 374$ nK with a small anisotropy attributed to the non-orthogonal HLAT.

stretched EIT transition fig. 5.6. In order to realize narrow and long-lived dark polaritons, the atomic sample also needs to be extremely well optically pumped to prepare all atoms in the same hyperfine sublevel, as they could otherwise absorb photons from the cavity as the EIT resonance is broken intentionally for other Zeeman sublevels using a magnetic field. While optical pumping can be performed on the D2 or D1 line of ^{87}Rb , we choose the $|5S_{1/2}\rangle \rightarrow |5P_{1/2}\rangle$ transition due to the smaller number of magnetic sub-level hyperfine levels in the excited state, providing a cleaner pumping. For perfect σ^+ polarization, the state $|5S_{1/2}, F = 2, m_F = 2\rangle$ is dark as there is no state $|5P_{1/2}, F' = 2, m'_F = 3\rangle$ to scatter off (fig. 4.13). Any imperfect polarization of the optical pumping beam caused by elliptical polarization or magnetic field pointing errors relative to the pump beam (light blue), will break the dark state and lead to imperfect atomic polarization. Experimentally, we can optimize the pump beam polarization and magnetic field by applying strong pumping and turning off the re-pump (green), where any imperfect π or σ^- component will scatter photons and lead to the atoms accumulating in the lower hyperfine $F = 1$ level, where they are dark to our VRS.

4.7.6 Trap Ramp



The quantum Hall physics we want to ultimately realize only occurs in two dimensions. We therefore need to confine the geometry of the polaritons to the 2D plane by making the cloud of atoms thin compared to the cavity waist and blockade radius. To this extent, we transfer the atomic sample from the transport VLAT into a highly elliptical ELAT lattice with the major axis perpendicular to the cavity direction, as described in section 4.6.1. In order to transfer the maximum number of atoms and not heat the sample excessively, we use an adiabatic ramp to hand off the trapping between the two lattices. Furthermore, they are detuned from each other by 160 MHz to not mutually interfere and lead to excess heating

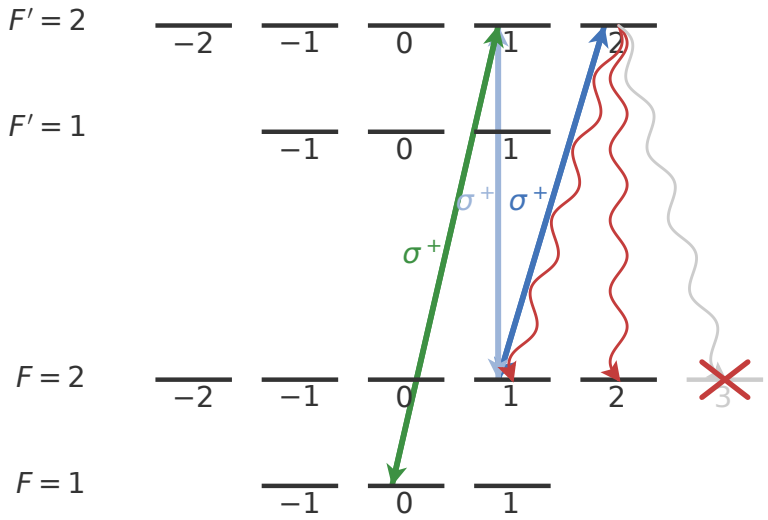


Figure 4.13: Optical pumping level scheme and optimization. We perform optical pumping on the $|5S_{1/2}\rangle \rightarrow |5P_{1/2}\rangle$ transition to polarize the entire atomic sample in state $|5S_{1/2}, F = 2, m_F = 2\rangle$. For perfect σ^+ polarized light (blue) the target state $|2, 2\rangle$ is dark and a repumper (green) prevents atoms from accumulating in the lower hyperfine state $F = 1$.

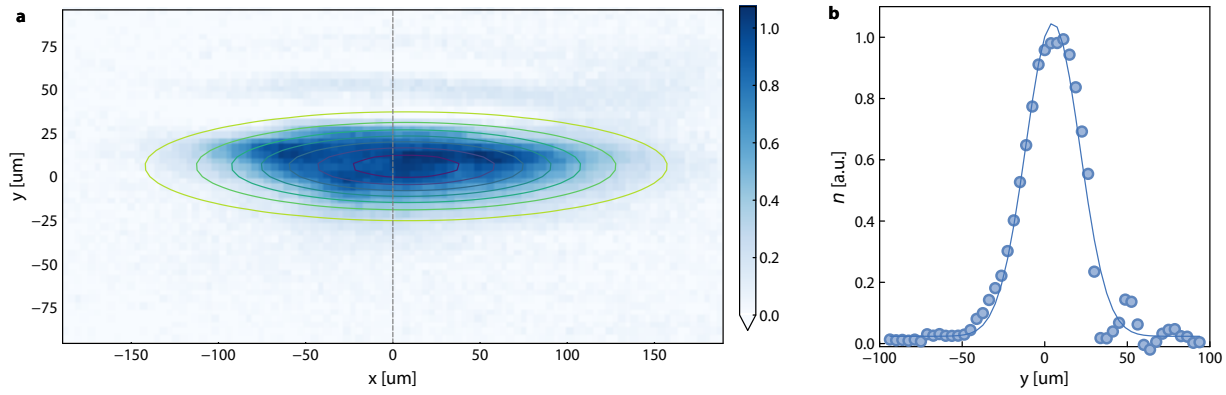


Figure 4.14: ELAT absorption image. Realizing 2D physics like the FQHE requires an effective 2D system with only a single excitation in longitudinal direction. We can make the polariton gas 2D by squeezing the atomic cloud to be thinner than the magnetic length l_B given by the cavity waist ω_0 in a special elliptical lattice (ELAT). An absorption image from the top **(a)** shows the atomic sample is thin along (y axis) and large perpendicular (x axis) to the cavity axis.

from phase noise between the two different paths. As the spatial overlap between the round VLAT and thin ELAT is not perfect, we lose some fraction of atoms and retain only $50-60\%$ and also heat them back up to a few uK from the phase space compression. As both lattices are conservative potentials, this transfer can also induce a sloshing oscillation of the cloud. We dampen this kinetic energy by applying another round of dRSC in the ELAT, however, the achievable temperatures are not as low as in fig. 4.12 due to the highly inhomogeneous trap frequencies in the elliptical beam. The resulting atomic sample is shown in fig. 4.14 where the cavity mode is propagating along the y axis.

CHAPTER 5

MULTI-MODE CAVITY QED: RESULTS

5.1 Cavity

The twisted, asphere based cavity that generates the artificial magnetic field for photons is the heart of the multimode cQED platform. We have already explored the calculation of the optimal asphere shape in chapter 3 and the mounting structure and alignment methods in chapter 4, but the full performance of the cavity degeneracy could only be measured after installing it in the vacuum chamber, as vibrations and air currents are eliminated here. Most degrees of freedom of the cavity, importantly the tilt of the lenses and centering relative to the optical axis defined by the image rotation of the twist were performed outside of vacuum before the bake. While this already results in a stable cavity, the alignment in order to minimize aberrations such as astigmatism is sensitive on the nanometer level and therefore needs to be performed in situ using the slip-stick three-axis translation stages built into the cavity structure. After alignment to degeneracy, the cavity is stable for multiple hours, mostly limited by thermal drifts outside of the vacuum chamber, which allows us to perform a detailed characterization of the cavity spectrum in the following.

5.1.1 Cavity spectroscopy

To probe the cavity transmission with resolution both in frequency and spatial mode space, we use an electro-optic modulator to sweep our narrow 780 nm probe laser over the cavity resonances, and use the DMD to generate a range of different input fields, mostly LG modes. Most of the optical aberrations in the excitation path have already been mapped out and compensated using a reference camera right in front of the vacuum chamber as described in section 4.3. Some of the Zernike polynomial terms such as tip and tilt, as well as exactly matching the Q parameter of the cavity in terms of waist size and wavefront curvature,

can only be fixed by using the cavity itself. Also, the phase measurement path that is not common mode with the full cavity path, such as the vacuum window, can re-introduce small aberrations that we need to compensate for by adjusting the higher order Zernike terms based on cavity transmission. The main challenge is to not conflate imperfect mode matching, that leads to excitations of several LG modes, with the intrinsic alignment of the cavity, which can also provide mode coupling that can be hard to tell apart. These two effects can be disentangled by considering that changing the input field from the DMD can only change the relative amplitude of different cavity modes, while alignment of the lenses actually shifts those resonances in frequency as the phase inside of the cavity is modified. As mode matching using the DMD and cavity alignment using the longitudinal and transversal lens displacement are high-dimensional parameter spaces, we have developed a series of techniques to perform mode matching of all 20 target modes and alignment to degeneracy in a quantitative and reliable way. To optimize the excitation fidelities we first need a clean system to excite, so we move the cavity spacing to a transverse mode splitting of at least 30 MHz which suppressed aberrations but keeps all target modes within the observable frequency range. Subsequently, we optimize the transverse lens positions using the Fine Scan mode of the piezo stages and make use of the fact that there is a global minimum of the Gouy phase as shown in fig. 5.3 (a). This can be understood by computing the mode shift in the presence of Astigmatism using a perturbative treatment

$$\Delta E_{l,p} \propto \frac{1}{\nu_{\mathrm{TMS}}}(l+1)(p+1) \tag{5.1}$$

which shows that the shift of mode LG_m scales as $m+1$ and is also present for the fundamental mode. This enables coarse alignment over several steps of the piezo stages by observing the shift of the 00 mode and the turn around point on an oscilloscope, while the fine alignment is performed using the experimental control system. In fig. 5.3 (a) this effect is shown for several LG_l modes and their minima coincide when both the x and y axis are

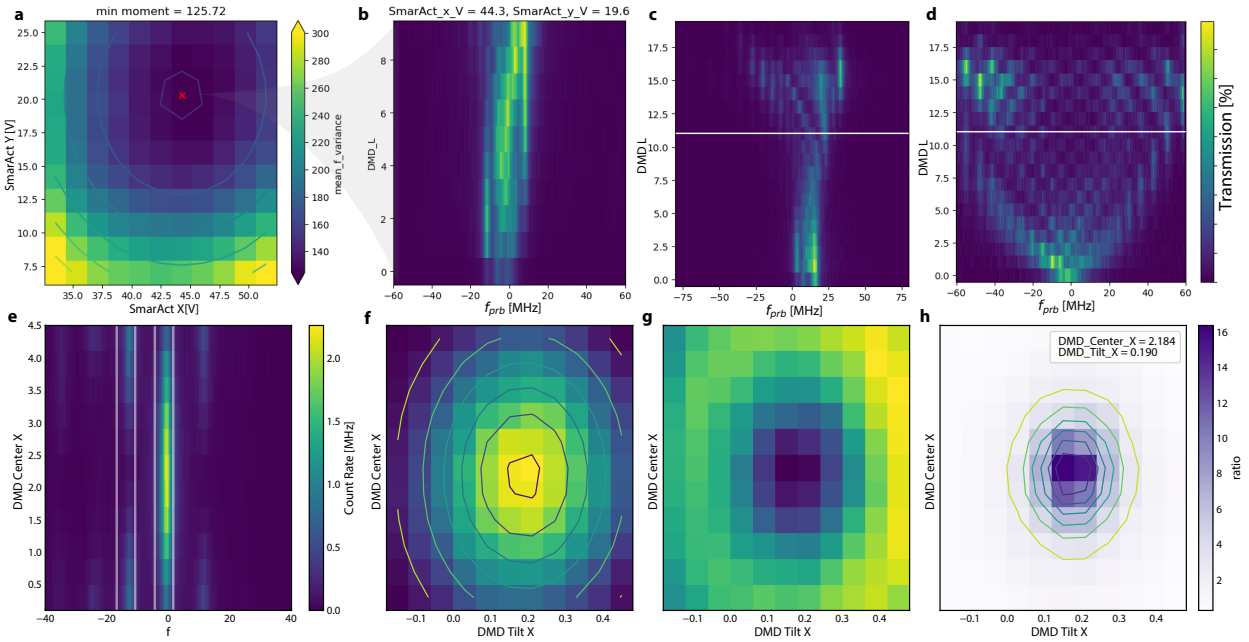


Figure 5.1: Cavity alignment and mode matching. Cavity alignment (top row) and mode matching using the DMD (bottom row). At the degeneracy, the only metric to align the lenses transversely is the full mode spectrum itself. In order to find the best degeneracy point, we perform a fine scan of the two piezo voltages of the SmarAct stage $V_{x/y}$ **(a)**, and compute the second moment of the resulting spectrum versus excited mode l. The best degeneracy point **(b)** including more modes **(c)** minimizes the spread from astigmatism **(d)**. Final mode matching is performed using the DMD itself, (bottom row) shows the transmission spectrum **(e)** and the intensity of the target **(f)** as well as undesired mode of opposite parity **(g)**. Computing their ratio **(h)** gives a sharp condition, and the optimal tip tilt parameters are extracted from a 2D fit.

aligned well. The total scan voltage range of 100 V corresponds to a displacement of $\sim$ 1 um and is therefore sensitive on the 10 nm level. For final alignment, a 2d scan is performed as there could be some coupling between the axis and the final operating point is derived from a 2d Gaussian fit. With the cavity perfectly aligned, mode matching using the DMD can be optimized by scanning individual Zernike terms to optimize transmission in the target mode. As Zernike polynomials form an orthogonal basis a series of N 1D scans can be performed saving an exponential amount of time compared to an ND scan. However for some parameters, such as waist size and defocus, the transmission of a single mode can be deceiving: Changing the defocus also changes the number of active pixels on the DMD and therefore the overall intensity. This can be circumvented by optimizing the ratio of two subsequent

modes l and l+1, and also tip-tilt alignment is much more sensitive when observing the ratio of even to odd modes.

Such a spectrum fig. 5.1 (b) at optimal degeneracy shows the first 10 LG_l modes degenerate within a 10 MHz window with two additional side lobes.

After mode matching, we can probe the full cavity spectrum by exciting different input l modes on the DMD and sweeping across all resonances. The resulting 2D spectrum in fig. 5.3 (c) allows us to compute the full cross talk matrix $T(l_{in}, l_{out})$ as the modes are well resolved in frequency. Subsequently the DMD parameters should not be changed anymore, as cavity modes can no longer be distinguished in frequency when approaching degeneracy. The degeneracy point of the cavity has to be very carefully approached by taking small steps of a few micrometers along z at a time, because the effect of aberration is no longer suppressed eq. (5.1) by the mode splitting and becomes extremely sensitive. After every longitudinal step in z, the transverse alignment has to be optimized again using the minimum in Gouy phase to not get completely lost in the mode landscape. Once at degeneracy or a few cavity line widths away, the only way to optimize the lens alignment is by directly observing the spectrum as shown in fig. 5.3 (b, insets) and minimizing the mode spread and dispersion. As this can be somewhat subjective, we have found that computing the second moment or inverse participation ratio gives results very close to the optimal point determined by eye. Sometimes the spectrum shows side-lobes disconnected from the main family of modes that we attribute to local potential minima caused by imperfections of the lens surface.

The approach to degeneracy is shown in fig. 5.2 (a) in frequency space as an average over all excited l modes, and the spatial modes (c -f) are still well defined Laguerre Gaussian. Any variation in the amplitude of the different LG_l modes is caused by the increased number of active DMD pixels for larger modes and inhomogeneous coupling into the multi-mode detection fiber. Once at degeneracy, the residual trap that was previously given by the transverse mode splitting is now completely flat, and since kinetic energy is frozen out in

the LLL, the only energy scale left in the system is that of disorder. The disorder is caused by small imperfections and form errors in the surface of the lenses that are only on the order of nm. However, the thousand round trips around the cavity can amplify these small phase deviations to mode shifts on the order of MHz. In the presence of this disorder, the eigenmodes are now delocalized over the 2D plane and have a random shapes as shown in fig. 5.2 (g-j). We therefore change bases from the LG basis into one of localized, Gaussian excitations $g(x, y)$ in the 2d plane. Since our possible wave functions are restricted to the LLL, they can only have a minimum size of the magnetic length l_B and are also further constrained. The

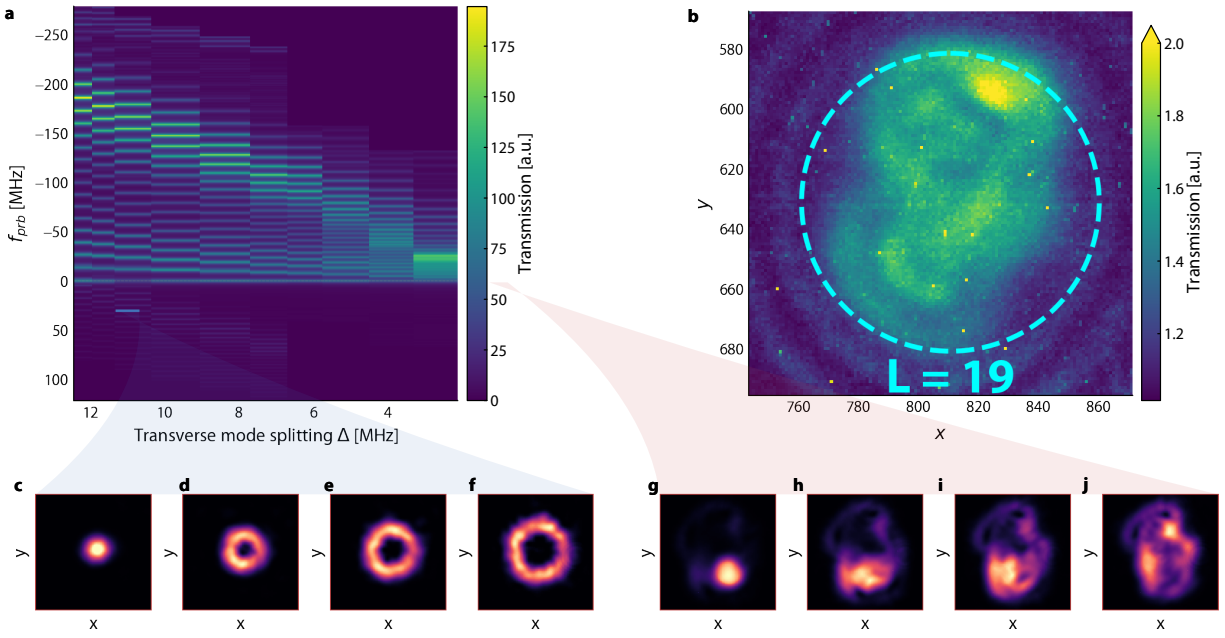


Figure 5.2: Cavity away from and at degeneracy. Cavity transmission spectrum **(a)** as a function of probe laser frequency f_{prb}. Adjusting the exact lens spacing L changes the transverse mode splitting ν_{TMS} and allows to bring all modes into degeneracy. Away from degeneracy, the modes are well defined LG modes **(c -f)** imaged as cavity transmission on a camera. At degeneracy, kinetic energy is frozen out and disorder is the only energy scale, leading to mixing between the eigenmodes **(g -j)**. Hence, LG modes are no longer a good basis, but instead we can probe the cavity locally using magnetically translated states revealing the size of the LLL puddle in an image **b** averaged over all excitation positions.

5.1.2 Landau level characterization

After having previously outlined the cavity alignment and mode-matching procedure in section 5.1.1 and fig. 5.1, we now characterize the cavity spectrum close to the degeneracy to measure the impact of residual aberrations on the mode spectrum and determine the usable size of the 2D cavity system. By exciting different L modes using the DMD and sweeping the probe laser over the full span of modes, we obtain a transmission spectrum fig. 5.3 (a) and (d). In (b), we fit every dominant mode using a Lorentzian line shape to measure the exact mode frequency (green) and linewidth (purple). Apart from an accidental degeneracy close to $l = 1$, the cavity linewidth γ is constant and independent of mode index l, indicating our cavity is highly stable even close to the degeneracy. The spectrum is highly linear, and a fit of the first ten modes to a straight line shows a mode splitting of 12 MHz. In (c) we plot the residual, showing a perfectly flat dispersion within a window of ± 2 MHz up to $L = 16$, after which the effect of higher order aberrations can't be compensated anymore which forms the edge of the system at $l = 19$. For realizing a number of orthogonal dark polaritons, the cavity modes must be pure and addressable with minimal crosstalk. In (f), we show the crosstalk matrix between the input mode excited on the DMD and the output mode that, at this splitting, is still resolvable in frequency. The matrix is obtained through a simultaneous fit for all 20 modes and normalized per row. The exact quantitative crosstalk depends on which mode was used for the optimization, here we can see mode fidelities of 90% and higher around $L = 15$. The non-uniform intensity of the different LG modes is a combination of increasing area of active DMD pixels and a reduction in coupling to the multimode fiber at high L which can be calibrated away in the experiment control. This highly linear spectrum confirms the validity of our detailed theory calculations using perturbation theory as well as ray tracing in chapter 3. The cavity supports a large system size of twenty modes that can be filled with ten particles to realize the $\nu = \frac{1}{2}$ Laughlin state.

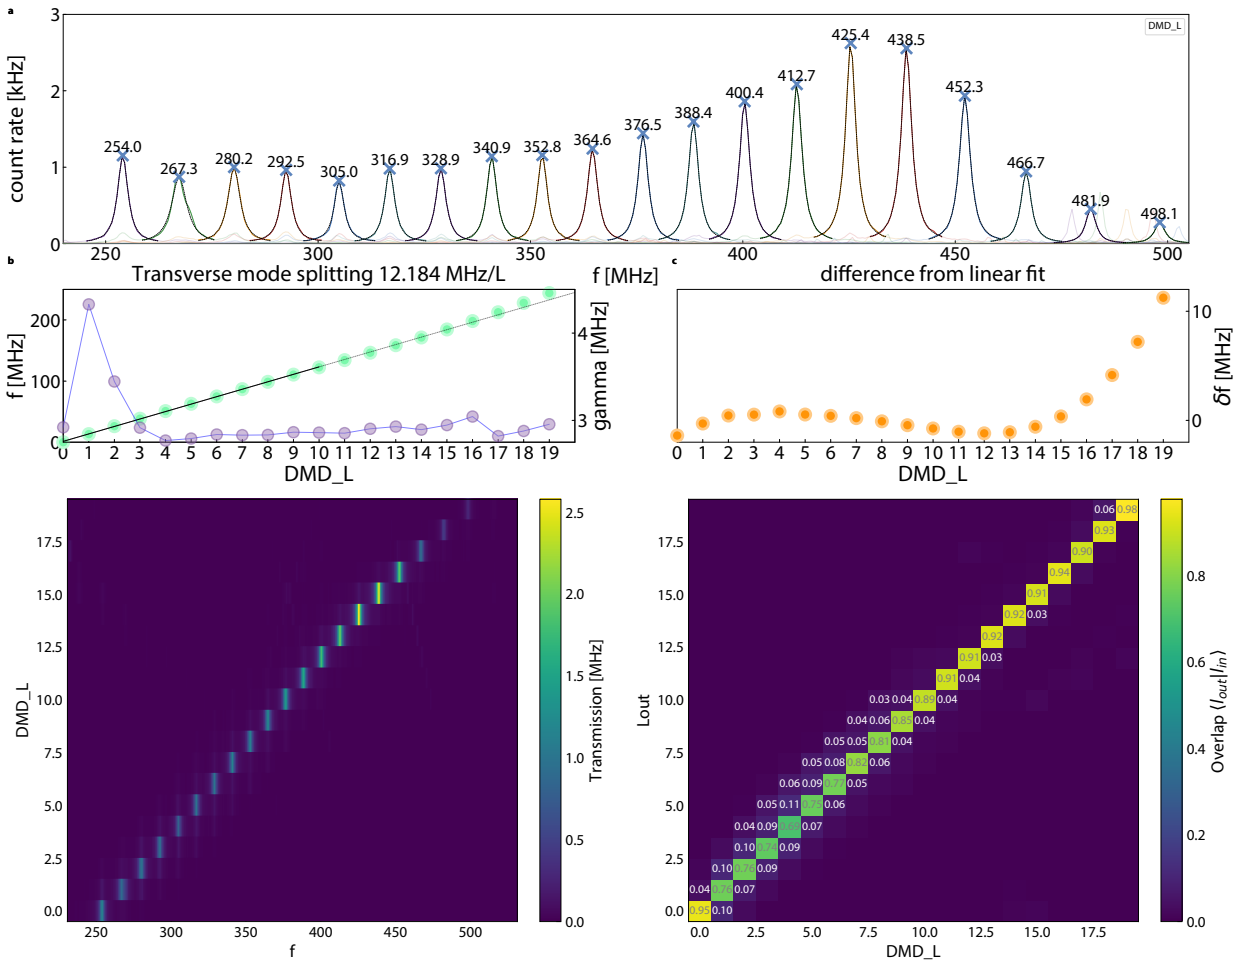


Figure 5.3: Landau Level characterization Mode spectrum of the asphere based cavity close to degeneracy for the first 20 LG $|l\rangle$ modes **(a)** after cavity and DMD incoupling optimization. We fit every mode of the transmission spectrum **(a)** with a Lorentzian lineshape to in **(b)** extract mode frequency (left axis) and linewidth (right axis) for every l. Owing to the careful aberration correction, the spectrum is extremely linear and **(c)** shows the residual from a linear fit over the first 10 modes. The edge of the degenerate region starting at $l = 17$ is due to a competition of higher order aberrations with the compensated coefficients c_4, c_6. We measure the mode mixing and crosstalk from combined cavity misalignment and imperfect mode matching **(e)** by performing a simultaneous fit of the spectrum **(d)**.

5.1.3 Disorder measurement

The highly linear spectrum described in section 5.1.2 inspires confidence that the cavity can be made degenerate with a small frequency spread. After aligning the cavity exactly to degeneracy and minimizing aberrations through transverse lens alignment, as shown in fig. 5.1, we can now probe the quality of the degeneracy. As all kinetic energy is frozen out for a perfectly flat band, the only energy scale is now given by imperfections and form errors in the lens surface that cause a disordered potential. While the exact deviations in surface height act like a local phase mask, they are each applied in a different Fourier plane with respect to the cavity waist, and since the modes are fields, some annular averaging is also performed over the disorder. The cavity field, when projected to the lowest Landau level, travels in a self-similar manner which makes all the disorder act like a local potential to lowest order. We can then map out the shape of this disorder potential by performing a full cavity tomography. By sweeping the probe laser over the full support of modes and recording the transmission at each frequency on a camera, we obtain images as in fig. 5.4 (c). The previously well defined LG modes fig. 5.4 (b) get fully mixed by the disorder and form a new basis. We interpret these new eigenmodes as the eigen rays mapping out the equipotential lines of the disorder. By plotting the 50% quantile of transmission at each frequency, we reconstruct the potential shown in fig. 5.4 (a). To mitigate the effect of vortices we average the transmission shown in fig. 5.4 (a) over a range of different l modes excited using the DMD. Apart from a distinct valley (bottom right) and peak (top left), the modes at degeneracy span a disorder width of $\approx 10 - 20$ MHz. While this corresponds to several times the cavity linewidth, the relevant energy scale to compare to, however, is the interaction energy U from Rydberg blockade. Furthermore, the bare cavity spectrum will be compressed due to the dark state rotation to span only $\approx 100 - 200$ kHz. As long as the interactions are stronger than the dark state compressed disorder, the fractional quantum Hall state should be robust.

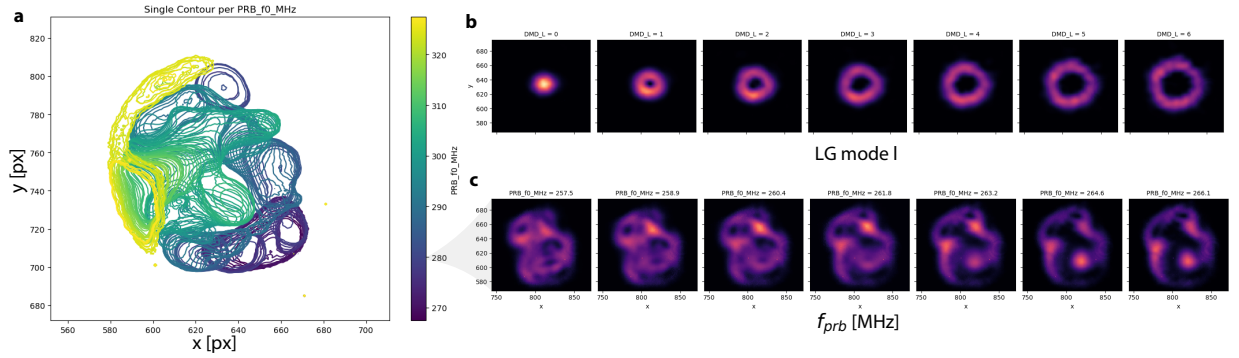


Figure 5.4: Cavity disorder tomography at degeneracy. At degeneracy, the only energy scale remaining is the disorder potential caused by form error of the aspheric lens surface. In **(a)** we reconstruct this potential in real space and frequency by showing equipotential transmission of the 50% quantile at every probe frequency f_{prb}. Away from degeneracy, we can spectroscopically resolve LG modes $|l\rangle$ as shown in **(b)** as a function of mode excited using the DMD. At degeneracy, the eigenmodes are completely mixed **(c)** and can no longer be resolved according to their spatial wavefunction but only frequency.

5.2 Single-mode VRS and EIT

In the previous sections, we have described how the atomic sample and the cavity are prepared and aligned. With the atoms transported into the cavity, we are now able to probe and characterize the light-matter coupling by probing the combined cavity-atom system in a way resembling quantum optics experiments through transmission spectra. First, we'll measure the cooperativity and lifetime of the achievable dark polariton on a single, fundamental spatial mode of the cavity away from degeneracy. Subsequently, we will bring higher order LG modes of the cavity close to resonance such that we can study VRS and EIT on many higher order $|l\rangle$ modes. After cooling, shaping, and optical pumping of the atomic sample, we interrogated the sample during the probe sequence by sweeping the narrow probe laser over a span of 100 MHz around the atomic resonance and collecting the transmitted light on single-photon-sensitive SPCM detectors. As the atomic ensemble absorbs and re-emits photons from the cavity field, the atomic population can be de-pumped or fall into the lower hyperfine state. We therefore repeat the optical pumping step section 4.6.1 periodically and divide the probe sequence into a number of 10ms probe periods interleaved with 100 μs

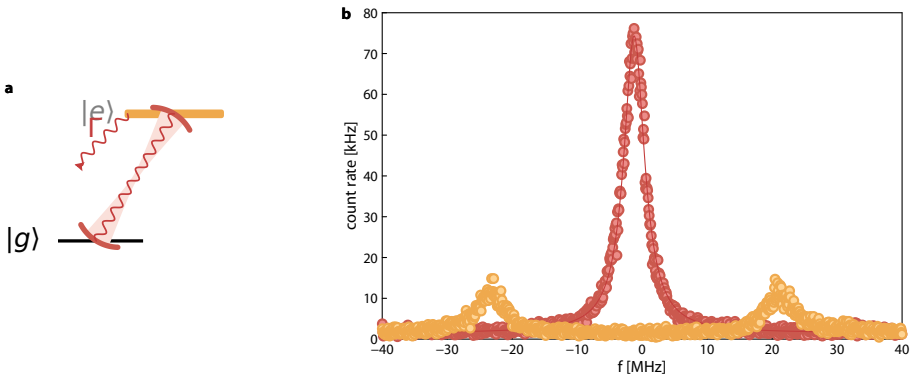


Figure 5.5: Vacuum Rabi splitting level scheme and transmission spectrum. Level diagram **(a)** and measured transmission spectrum **(b)** of the vacuum Rabi splitting. Placing an ensemble of atoms in the cavity leads to a splitting of the cavity resonance (red) into two polaritons (yellow). Solid lines are a fit to NHPT theory calculations.

pump and re-pump periods. This highlights the quantum optical nature of our experiment, where the same sample of atoms can be probed repeatedly over 200 ms before any significant atom loss occurs. This is in contrast to cold atom experiments, where an atomic sample can only be interrogated once due to destructive imaging. During the probe sequence, we can either probe the system spectroscopically in the steady state or in a time-resolved regime by pulsing an excitation into the cavity and recording the leaking photons. In the steady state, the sweep rate of the laser is much slower than the lifetime of the dark polaritons, and therefore the system finds an equilibrium polariton number between drive and decay at every frequency. Conversely, in the time domain, we can inject π or other pulses by switching the RF driving the probe EOM rapidly. This can reveal Rabi oscillations of a blockaded polariton mode and allows us to measure decay rates or time resolved movies using the fast EMCCD.

We begin by probing the atom-cavity system where the cavity is made resonant with the $|5S_{\frac{1}{2}}, F = 2, m_F = 2\rangle \rightarrow |5P_{\frac{3}{2}}, F = 3, m_F = 3\rangle$ transition. This results in a vacuum Rabi splitting shown in fig. 5.5 as predicted by the Tavis-Cummings Hamiltonian. The two bright polaritons are equal superpositions of cavity photon and atomic P state excitation and are separated by the collective coupling g. As $g = g_0\sqrt{N}$ is proportional to the number of

atoms in the cavity mode, it is a great proxy to optimize overlap between the sample and cavity mode. Even though our relatively large mode waist and modest finesse attributable to the special constraints of the lens cavity lead to a small single atom coupling strength g_0, we can compensate this by optimizing our atom loading and cooling and we regularly achieve collective couplings up to $g \sim 50$ MHz. Due to the large linewidth of the atomic P state contribution, bright polaritons are not useful for building up correlated quantum states. In order to obtain the long-lived and interacting dark polaritons, we now add an additional laser where a classical control field takes excitation up to the Rydberg manifold via the $|5P_{\frac{3}{2}}, F = 3, m_F = 3\rangle \rightarrow |nD_{\frac{5}{2}}, J = 5/2, m_J = 5/2\rangle$ transition as shown in the level diagram fig. 5.6 (a). Now the spectrum shown in fig. 5.6 shows an additional narrow feature

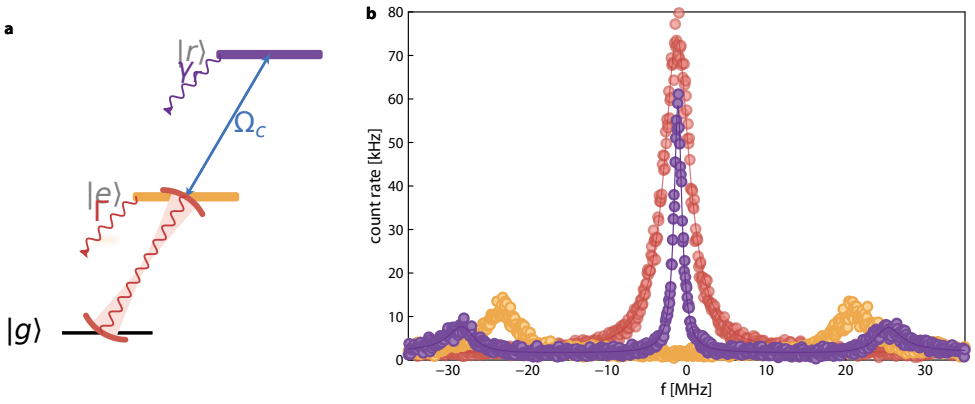


Figure 5.6: EIT level scheme and transmission spectrum. Level diagram **(a)** and measured transmission spectrum **(b)** of cavity Rydberg EIT. Adding a classical control field Ω (blue) to the hybrid cavity atom system leads to Rydberg EIT, evident by the narrow peak in the middle of the spectrum showing the dark polariton.

at two-photon resonance which is the EIT peak corresponding to the dark polariton. The reduction in transmission compared to the bare cavity is caused by the finite lifetime of the Rydberg state as well as other decoherence mechanisms as discussed in section 3.10 combined in an effective linewidth γ_r. One of the main limitations in the old apparatus was the limited available blue laser power that was caused by a degradation of the buildup cavity mirrors. In the new apparatus, the Rydberg buildup cavity was designed with a high enough finesse of

approximately $\mathcal{F} \approx 3000$ to enable a 1000-fold power buildup of the available 480 nm laser light. We need this much optical power since the Rydberg excitation beam needs to have a large waist to cover all higher order modes of the science cavity and provide a uniform Rabi frequency Ω_c. In fig. 5.7 we show the effect of varying the control field strength over a large

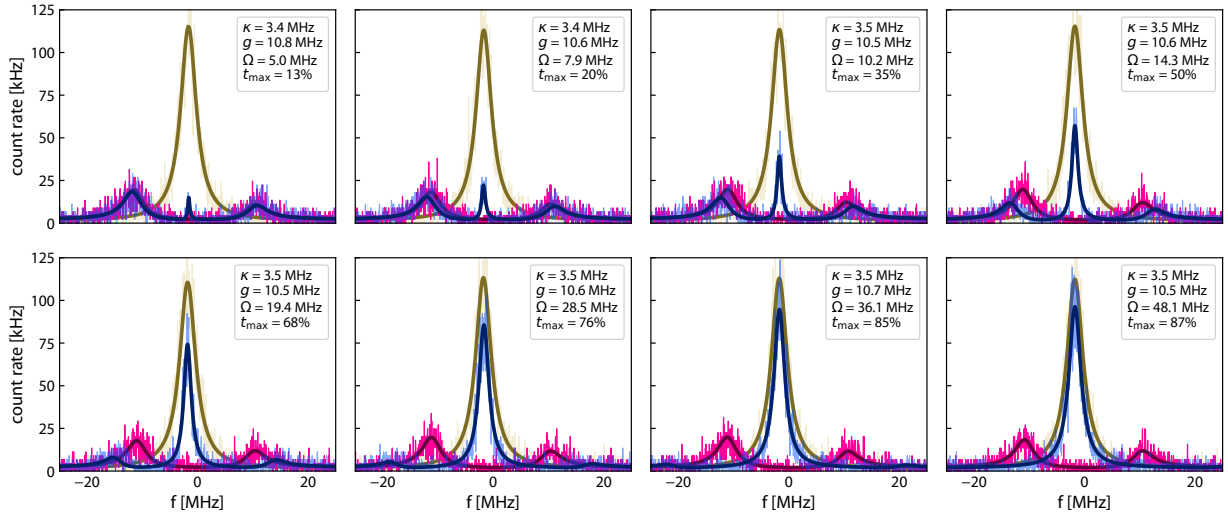


Figure 5.7: Cavity Rydberg polaritons vs dark state rotation angle. The nature of the dark polariton can be tuned via the dark state rotation angle θ_D between mostly Rydberg like **(a)** to mostly cavity like **(f)** by changing the strength of the Rydberg control field. The bare cavity transmission (yellow), VRS (pink), and EIT transmission (blue) spectra are shown as a function of the control field Rabi frequency Ω. With increasing control field strength **(a-h)**, the dark polariton resonance approaches the bare cavity transmission in both transmission overall transmission and line width. Legends show fit coefficients to an NHPT theory.

range, tuning the character of the polaritons from almost only Rydberg-like to almost fully cavity-like. As evident in fig. 5.7 (h) the new buildup cavity allows us to achieve high blue Rabi frequencies of $\Omega = 48.1$ MHz, on par with the collective coupling strength.

5.2.1 Single-mode Photon Blockade

5.3 Multi-mode VRS and EIT

The key metric of our platform is how well we can realize the ideal model of a LLL for interacting polaritons from chapter 3, i.e. how many cavity modes we can achieve good dark

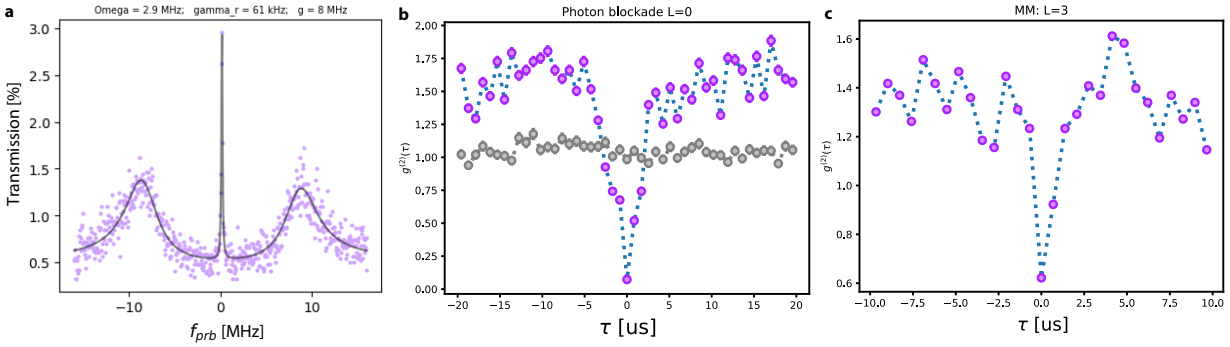


Figure 5.8: Single mode photon blockade. At a principal quantum number of $n = 98$ Rydberg interactions of the $|98D_{5/2}, 5.2\rangle$ state are strong enough to blockade the fundamental cavity waist $R_b \approx \omega_0$. A representative EIT spectrum **(a)** shows coupling $g = 8$ MHz, control field Rabi frequency $\Omega = 2.9$ MHz. Using a single-mode fiber splitter we can measure the $g^{(2)}$ correlation function **(b)** showing strong blockade (purple). The bare cavity light without atoms (grey) is uncorrelated $g^{(2)} \approx 1$. On a higher order mode with $l = 3$ **(c)** blockade is less strong due to the increased separation between atom pairs.

polaritons on before non-uniformity of the atomic cloud or Rydberg control field lead to a drop of cooperativity or decoherence. To characterize this, we first probe VRS and EIT on many l modes with the cavity aligned away from degeneracy using a transverse mode splitting ν_{TMS} of many linewidths.

5.3.1 Atomic lensing in VRS

With a well aligned cavity, all its $|l\rangle$ modes and therefore the resulting polaritons should be orthogonal to each other and not couple to each other in a transmission spectrum. In fig. 5.9 (b) the transmission spectrum when exciting mode $|l = 2\rangle$ on the DMD shows the characteristic VRS, however many additional transmission peaks can be seen that resemble a detuned VRS of neighbouring cavity modes. As the collective coupling $g \approx 45$ MHz is comparable or even larger than the mode splitting ν_{TMS}, coupling between the different polaritons is no longer spectroscopically suppressed but any inhomogeneity of the atomic sample can lead to mode crosstalk. We can study the effect of any imperfections in the atomic cloud or misalignment of the control field quantitatively using an atomistic non-

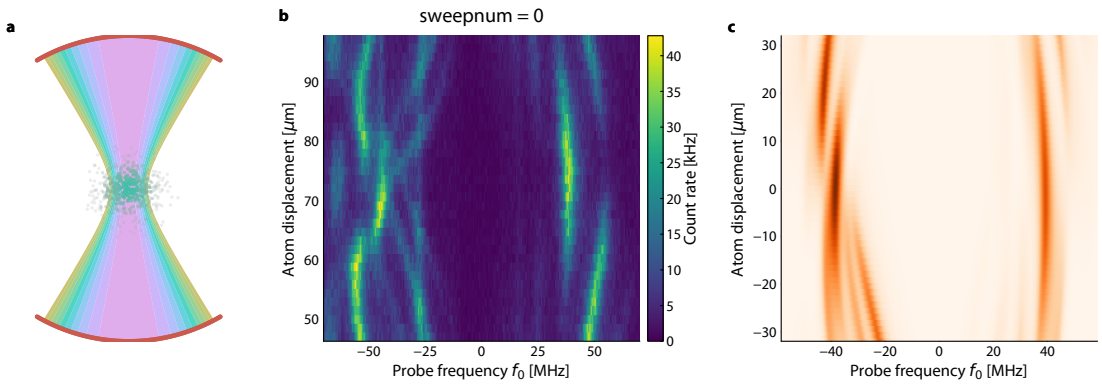


Figure 5.9: Intra-cavity lensing from atom misalignment. Ideally, the atomic sample is centered on all the higher-order cavity modes **(a)** and is much larger than the highest cavity mode L to provide a uniform coupling constant g_l. Any misalignment **(b)** leads to coupling between different L modes due to lensing resulting in mixing of bright polaritons. **(c)** theoretical model using linear and NHPT, which predicts the key features.

Hermitian perturbation theory (NHPT) [CJS+19] that treats N_{at} individual atoms in N_{cav} modes. The cavity field C_n of mode n can be solved to in first order

$$C_n^* = \delta_c^n A_c^n + \sum_m^{N_{at}} \frac{|g_{mn}|^2}{\frac{|\Omega_m^b|^2}{\delta_2} - \delta_e} A_c^n + \sum_{k\neq n}^{N_{cav}} \sum_m^{N_{at}} \frac{g_{mn}^* g_{mk}}{\frac{|\Omega_m^b|^2}{\delta_2} - \delta_e} A_c^k \tag{5.2}$$

where A_c^n is the input field on mode n, g_{mn} is the coupling of atom m to mode n and the explicit amplitudes of atom n have been eliminated in closed form. We evaluate the NHPT for a set of N_cav modes while replacing the sum over atom locations by a coarse grained continuum versions with integrals over a Gaussian atomic density

$$C_n^* = (\delta_c + n\delta_t)\, A_c^n + \frac{g_0^2}{-\delta_e} \iint_{-\infty}^{\infty} \rho(\mathbf{r})\, |\psi_n(\mathbf{r})|^2\, d^2r\, A_c^n + \sum_{k\neq n}^{N_{cav}} \frac{g_0^2}{-\delta_e} \iint_{-\infty}^{\infty} \rho(\mathbf{r}) \psi_n^*(\mathbf{r}) \psi_k(\mathbf{r})\, d^2r\, A_c^k \tag{5.3}$$

where $\delta_c = -\omega_p - i\frac{\kappa}{2}$ is the bare cavity detuning including linewidth, δ_t the transverse mode spacing, and $\delta_e = -\omega_p - i\frac{\gamma}{2}$ the atomic detuning with decay. The last term only drops out for a homogeneous sample and in the thermodynamic limit. We can then compute the cavity transmission in linear response by solving for $MA_c = C_{in}^*$ for a photon drive C_{in}.

When the atoms are not centered transversally within the cavity, they can act as a lens and couple light into different modes. In fig. 5.9 (b) the effect of this lensing is shown in a transmission spectrum as a function of atom position (y axis). For a centered cloud at 72 um, only the two desired bright polaritons are visible at ±45 MHz, but when the cloud is displaced additional detuned VRS spectra start to overlap. These couplings would destroy the ideal picture of independent bosonic modes in the absence of the Rydberg control field and therefore need to be suppressed. Along the transport direction, this is simple as the atom position in the transport lattice can be changed easily. Transverse to the cavity and transport axis, the optical lattice location needs to be translated while keeping overlap with the MOT and retro-reflected beam. This is achieved using motorized control of all lattice and retro mirrors and a first-order model of the optical system to change all eight parameters simultaneously. The simple first order model in fig. 5.9 (c) reproduces this effect qualitatively but fails to capture some of the higher order processes that lead to an even more complicated transmission in fig. 5.9 (b). While the alignment transverse to the cavity mode is most critical to suppress lensing, the same motorized system can also be used to translate the atomic cloud longitudinally to maximize overlap with the Rydberg control field that intersects the cavity mode at an angle of $\sim 30°$.

5.3.2 Spatially resolved cooperativity

5.3.3 EIT for higher order modes

A second key metric determining the number of useful polariton modes is the size of the EIT window

$$\Delta = \frac{\Omega^2/\Gamma}{\sqrt{\eta}\cos(\theta_D)} \tag{5.4}$$

where Ω is the control field Rabi frequency, Γ is the decay rate of the intermediate state, η is the collective cooperativity, and θ_D is the dark state rotation angle. In this form, it seems

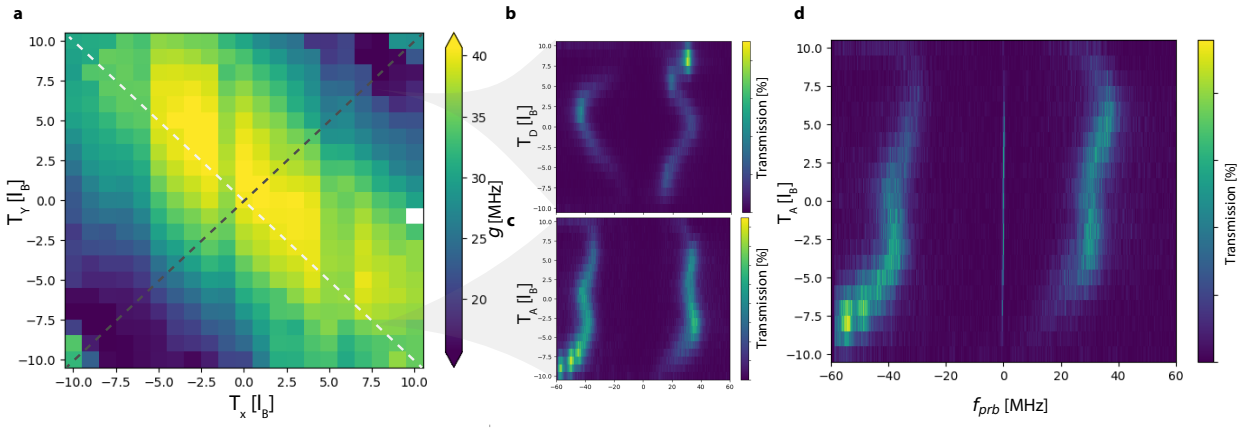


Figure 5.10: Spatially resolved VRS at degeneracy Light-matter coupling constant $g(x, y)$ measured over the full degenerate cavity system as a function of magnetic translation in T_x and T_y direction **(a)**. Every point corresponds to a VRS spectrum and cuts along the diagonal **(c)** and anti-diagonal **(b)** reveal local asymmetry in the VRS due to cavity detuning from disorder. The elliptical shape in **(a)** exactly reflects the almond shape of the atomic cloud in the VLAT and the rotation by 45 ∘ is caused by the DMD rotation in the lab frame (see fig. 4.7). Magnetically translated EIT spectrum **(d)** shows dark state compression and usable system size of $\pm 7 l_B$.

that the EIT window can be made larger with increased blue Rabi frequency Ω, however this also changes the dark state rotation angle. This makes the polaritons more photon like and increases their linewidth, leading to the same ratio of EIT window size to dark polariton linewidth, leaving the number of resolved modes unchanged. The only way to increase to number of resolved polariton modes is to increase the collective coupling strength g and conversely cooperativity η by loading more atoms.

We have derived the

While for realizing FQHE states the width of the EIT window needs to be larger than the dark state rotated disorder potential which is easily satisfied for , operating with a harmonic trap however makes this more difficult as the modes are spaced out linearly.

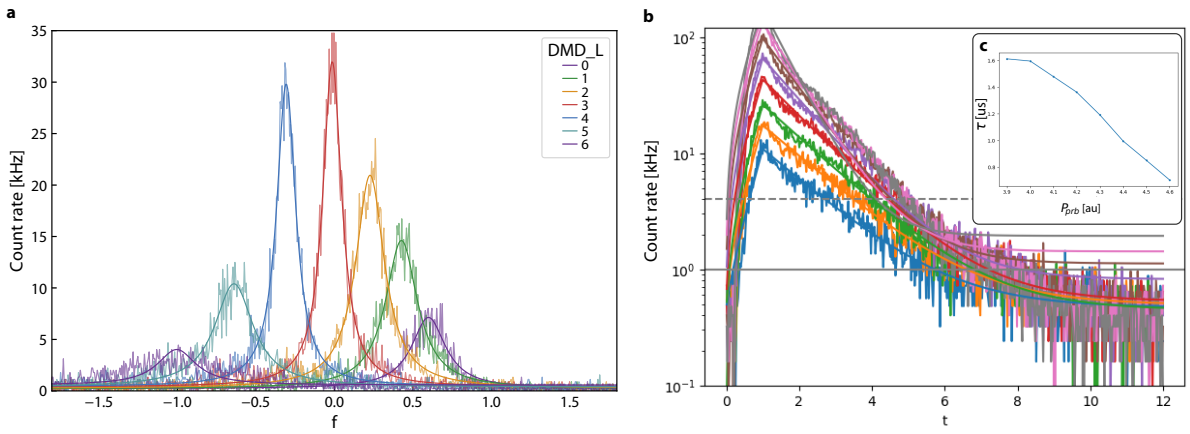


Figure 5.11: Magnetically translated dark polaritons

5.3.4 Spatially resolved Polaritons

To benchmark the degeneracy of our photonic LLL, spectroscopy is no longer suitable in the presence of mode mixing. Instead, we directly probe the quality of the degeneracy by exciting magnetically translated polaritons away from the cavity center, which are a superposition of all $LG_l^{p=0}$ modes. A well-defined translated polariton can only be formed if all modes interfere with the correct amplitudes and phase and this interference would be suppressed by disorder exceeding the polariton linewidth.

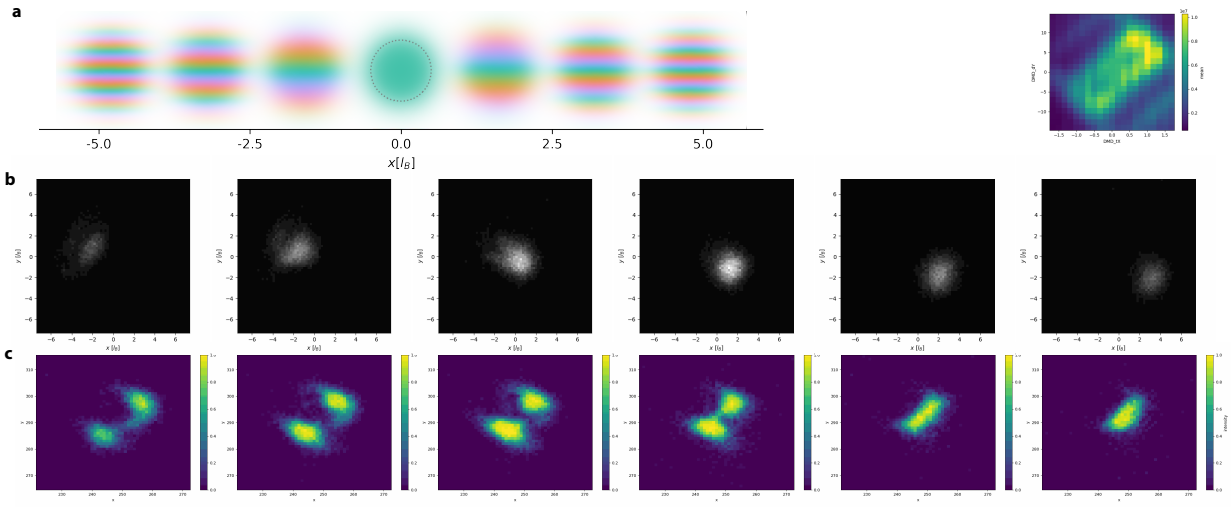


Figure 5.12: Magnetically translated dark polaritons

Due to the special *quantum geometry* [Hal] of the LLL, the smallest feature that can be made is still given by the magnetic length l_B and magnetically translated polaritons still look like LG_0^0 mode, albeit at a different location and with a phase gradient to remain in the LLL.

The magnetic translation operation is operated on the DMD that is operated in a Fourier plane with respect to the cavity waist. For a given displacement $d_i, i \in \{x, y\}$ in units of the cavity waist, we scan the tilt angle $t_j, j \in \{y, x\}$ along the orthogonal direction reflecting the non-commutative geometry. We fit the recorded transmission $T(t_j)$ with a Gaussian for every d_i and thereby extract the magnetic factor in experimental units.

With the DMD in the Fourier plane, we can then compute the inverse Fourier transform of the desired target field before computing the hologram. In (c) we then show an example of two displaced polaritons excited in the system.

5.4 Time-resolved dynamics in the LLL

In contrast to the average images shown in section 5.3.4 recorded in the steady state, detecting single photon correlations requires time resolved measurements of single polaritons. To benchmark the high speed of our specially developed ultra-fast EMCCD camera [LPJ+25], we demonstrate a time resolved measurement of dynamics in the LLL. By slightly changing the lens spacing away from degeneracy, they provide weak harmonic trap for photons.

Analogous to electrons in a magnetic field, photons in the twisted cavity undergo cyclotron and magnetron orbits. Since we only populate a single Landau level, the faster cyclotron motion is frozen out but a magnetron motion akin to the $\mathbf{E} \times \mathbf{B}$ drift is caused by the radial force of the harmonic trap.

In the experiment, a displaced wavepacket is excited by pulsing on an EOM for ~ 1 us exciting a coherent polariton state. For its lifetime of $\tau \approx 2-5$ us, this state rotates with ω_M in the trap while its photonic part leaks out of the cavity and is recorded on the EMCCD. The

large Rydberg component of the polariton does not provide significant interactions as the principal quantum number in this experiment $n = 37$ is low, however it slows the dynamics sufficiently to be recorded at 2 Mfps.

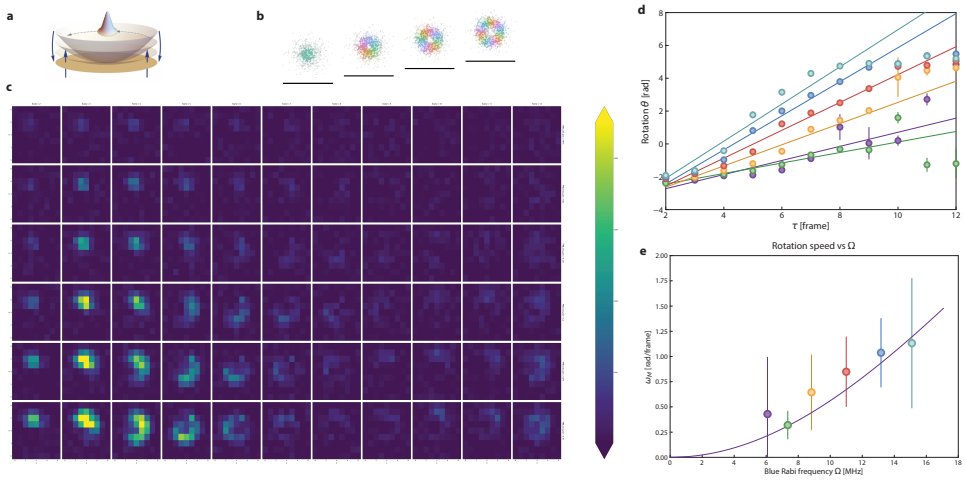


Figure 5.13: Magnetron motion of a magnetically translated polariton in the LLL. **(a)** Schematic of wavepacket in the residual trap. **(b)** Magnetic translation to stay in the LLL **(c)** Frames of the movie for one control field strength. **(d)** Speed of the rotation vs Ω. **(e)** EIT window tbd

A key strength of our hybrid platform is that the character of the quasi-particles can be changed dynamically. As the strength of the Rydberg control field determines the composition of the polaritons, the speed of the magnetron motion can be controlled through the blue Rabi frequency Ω. In fig. 5.13 (c), the speed of rotation is extracted by fitting a two-dimensional Gaussian to the wave packet in each frame and plotting its angle over time. The resulting magnetron frequency ω_M is then shown in fig. 5.13 (d) as a function of Ω and matches the theoretical prediction.

Even at low principal quantum number, there is still some Rydberg interaction present that leads to mean field effects: While we would like to excite a coherent state $|\alpha(z)\rangle$ with large average occupation α to boost the SNR of the camera, our interactions are significantly imaginary leading to a faster decay of the wave packet at higher probe powers. Because in our EIT scheme, the intermediate $|5P_{3/2}\rangle$ state is resonant, this intermediate state is only dark for the first polariton in a mode. Due to interactions, the Rydberg state is shifted for all

subsequent photons trying to enter the system, which leads to a small population of the lossy P state and therefore makes this interaction lossy. The exact ratio and spatial dependence of the real and imaginary part of the interaction potential is detailed in appendix A.6.1. As can be seen in fig. 5.13 (e), the lifetime at low probe powers (polaritons) is us and is reduced to us at high probe power (polaritons).

This demonstrates: high cooperativity, large EIT window, fast camera limitations: CIC induced noise, signal only $0.01 e^-/pix/image$ as the flux is only 20.0 kHz/mode.

5.5 Flux lattice

5.5.1 Background: mean field

To treat the mean-field physics in the LLL we start from the well known Gross-Pitaevskii Energy functional for a 2D boson gas in a rotating frame (or equivalently, with a synthetic magnetic field) including complex interactions in the continuum.

$$\mathcal{H} = \psi^* \left[\frac{(\mathbf{p} - \mathbf{A})^2}{2m} + V_{\text{trap}}(\mathbf{r}) - \Omega L_z \right] \psi + \frac{g}{2}|\psi|^4 \tag{5.5}$$

with the gauge vector potential $\mathbf{A} = \frac{B}{2}(-y, x)$ in symmetric form, Ω the rotation frequency and g the interaction parameter. As Rydberg interactions between Polaritons can lead to both shifts of the resonance as well as two-body loss depending on intermediate state detuning (see appendix A.6.1), we model g as a complex number $g \equiv g_R + i g_I$ with g_R representing elastic interactions and $g_I < 0$ a two-body loss rate (nonlinear dissipation).

In the twisted cavity close to degeneracy ($\Omega \to \omega_c/2$) the kinetic term is frozen out and all dynamics are constrained to the LLL. We can therefore expand the macroscopic wavefunction ψ in terms of basis functions

$$\phi_l(z) = \frac{1}{\sqrt{\pi l!} w^{l+1}} z^l e^{-|z|^2/(2w^2)} \tag{5.6}$$

where the cavity waist w takes the role of the magnetic length scale in appendix A.4. This transforms the GPE of a continuum wavefunction to

$$\psi(z,t) = \sum_{l=0}^{\infty} c_l(t)\phi_l(z) = \frac{e^{-|z|^2/2w^2}}{\sqrt{\pi}w} \sum_{l=0}^{\infty} \frac{c_l(t)}{\sqrt{l!}} \left(\frac{z}{w}\right)^l \tag{5.7}$$

to a differential equation in the complex coefficients $c_l(t)$ representing the field of every cavity mode l. The dynamics in the flat band are driven by contact interactions, which can be projected into the LLL in the usual form

$$V_{mnpq} = \frac{g}{2\pi w^2} \frac{(m+n)!}{\sqrt{m!n!p!q!}} \frac{1}{2^{m+n+1}} \delta_{m+n,p+q} \tag{5.8}$$

conserving angular momentum $m+n = p+q$. The equation of motion in the absence of any harmonic trap, disorder or drive is then

$$i\hbar \dot{c}_l = \sum_{m,n,p} V_{lmnp} c_m^* c_n c_p \tag{5.9}$$

and can be efficiently solved numerically.

Depending on the ansatz or the physics one is interested in, the wavefunction in the LLL can be written in two equivalent ways: Either the coefficients c_l in the Fock basis can be specified (e.g. when driving specific angular momentum modes), or given the position of vortices z_i in the complex plane the wavefunction can be written in first-quantized form

$$\psi(z) \propto e^{-|z|^2/2w^2} \prod_{k=1}^{N_v} (z - z_k) \tag{5.10}$$

that is convenient for describing regular lattices or crystals. To simulate the GPE based on eq. (5.9) we need a map $\{z_i\} \to \{c_l\}$ for the vortex positions. This can be achieved by

expanding the polynomial

$$P(z) = \prod_{k=1}^{N}(z - z_k) \stackrel{!}{=} \sum_{l=0}^{N} a_l z^l \tag{5.11}$$

and computing the coefficients a_l of the polynomial having roots $Z = \{z_1, z_2, \ldots, z_{N_\nu}\}$. This relation can be done with Vieta's formulas

$$a_{N-k} = (-1)^k \sum_{1 \le j_1 < \cdots < j_k \le N} z_{j_1} \ldots z_{j_k} \tag{5.12}$$

or simply numerically with `numpy.polynomial.polynomial.polyfromroots` and scaling by the factorial normalization $c_l = a_l\sqrt{l!}$.

5.6 Förster resonance

To achieve state-selective control, we employ an AC microwave field. The dynamic polarizability $\alpha(\omega)$ accounts for the photon energy $\hbar\omega$ in the detuning denominator:

$$\alpha(\omega) = \sum_{k \neq r} \frac{|\langle k|\hat{\mu}|r\rangle|^2}{E_k - E_r - \hbar\omega}. \tag{5.13}$$

We aim to tune the Förster defect Δ_F between the initial and coupled pair states to zero by precisely shifting one of the participating $|r'\rangle$,$|r''\rangle$ states while minimally impacting the $|r\rangle$ state. By selecting a microwave frequency ω near a specific transition—for instance, close to a transition of the $|p\rangle$ state but far detuned from the $|d\rangle$ and $|f\rangle$ states—we can induce a large AC Stark shift on the $|pf\rangle$ manifold while leaving the $|dd\rangle$ target state relatively unperturbed.

Our states of interest are therefore (a) two atoms in our "target' state $|dd\rangle \equiv |44D_{5/2}, m_j = 5/2;\ 44D_{5/2}, m_j = 5/2\rangle$ and (b) the closest coupled pair state $|pf\rangle \equiv |46P_{3/2}, m_j = 3/2;\ 42F_{7/2}, m_j = 7/2\rangle$ that conserves both total $\hat{L}$ and $\hat{L_z}$. The second channel $|pf\rangle \equiv |46P_{3/2}, m_j = 3/2;\ 42F_{7/2}, m_j = 7/2\rangle$ has a large detuning and can therefore be ignored.

This differential shift allows us to sweep Δ through zero, reaching the Förster resonance dynamically (see fig. 5.15). At the same time, small background field fluctuations are suppressed $(E + \delta E)^2 = 0 + \mathcal{O}(\delta E^2)$ quadratically when the DC fields are carefully nulled.

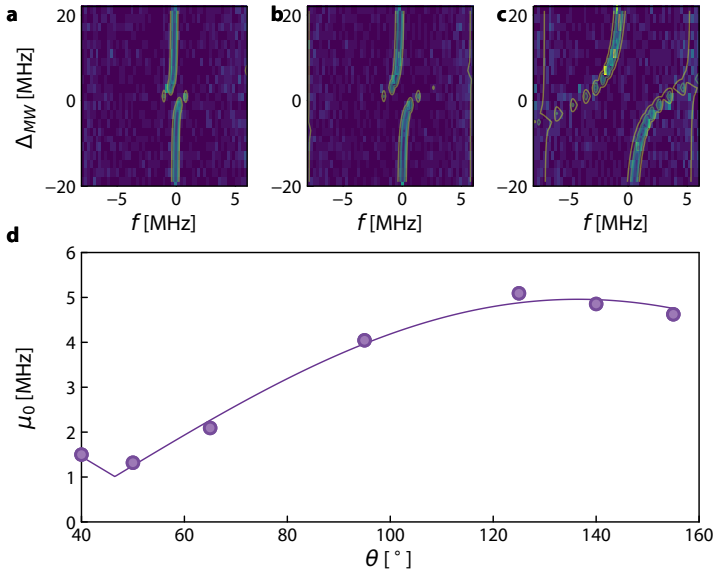


Figure 5.14: Microwave polarization optimization. (a) - (c) Autler-Townes splitting of the EIT peak due to applied microwaves for three different horn orientations. Scanning the detuning Δ_{MW} of the applied microwave frequency reveals splitting of the dark polariton peak with increasing microwave drive power, analytical fit overlaid as contours. **(d)** Extracted microwave Rabi frequency μ_0 as a function of Horn rotation θ.

To determine the optimal frequency and polarization for the RF driving field, we compute the AC stark shift of all participating states including their fine structure in fig. 5.15 using the *ARC* [ŠPAW17] library, see section 3.11 for details of the numerical calculations. The Förster defect for $|44\ D_{5/2,5/2}\rangle$ is $\Delta = 65$ MHz so we need to shift either the P- or F-state to lower energies using a microwave drive. As the Rydberg manifold has a plethora of transitions, we need to carefully choose our microwave frequency to be as far off-resonant as possible while not driving other unwanted transitions. We choose the P-state, as it has fewer dipole allowed transitions and is further separated from a D-state resonance than the F-state as shown in fig. 5.15 (c). It is important to only generate scalar AC stark shifts as vector shifts would act as a virtual transverse magnetic field coupling between different magnetic

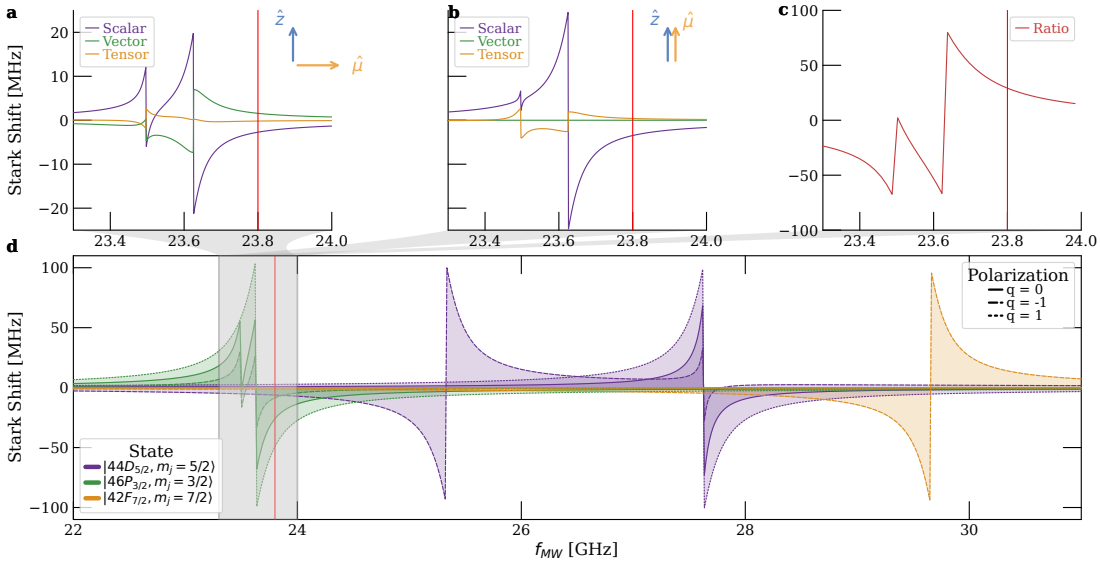


Figure 5.15: AC stark shift calculation We explore the optimal way of tuning the pair state $|r'r''\rangle = |(n+2)P_{3/2}, (n-2)F_{7/2}\rangle$ into resonance with our target state $|rr\rangle = |nD_{5/2}, nD_{5/2}\rangle$ while minimizing the shift of the latter one. We compute the AC stark shift for the three Rydberg states (color) over a wide range of microwave driving frequencies for pi- **(a)** and sigma **(b)** polarization at a field strength of $E = 0.1$ V/m, see A.6.1. Driving with π polarization is optimal as it maximizes the frequency separation of the $|(n-2)P_{3/2}\rangle$ from the target state and eliminates any vector shifts that would cause a transverse coupling between magnetic sublevels. The microwave driving frequency used in the experiment of $f_{MW} = 23.8$ GHz is indicated as a vertical red line.

sublevels and breaking the simple single-channel model. In fig. 5.15 we compute the scalar, vector and tensor shifts for σ^- (a), π (b), σ^+ (c) polarization of the RF field, where only (b) shows no vectorial shifts. Linear polarization also avoids an unwanted resonance of D-state that occurs closer to the P-state resonance for σ^+ polarization.

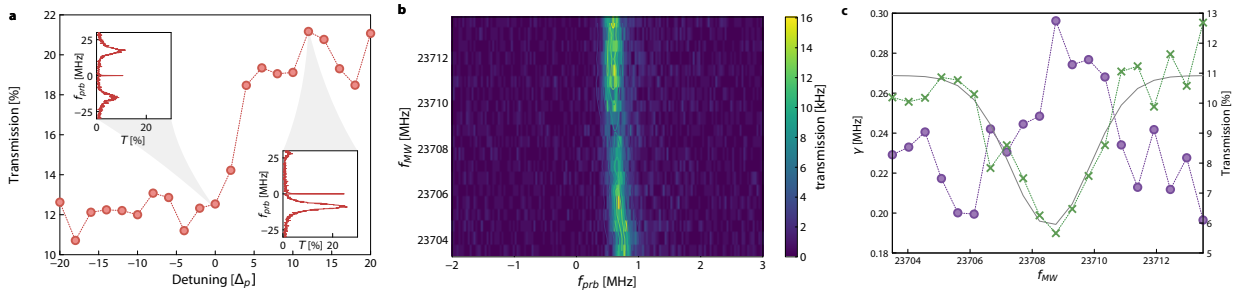


Figure 5.16: Locating the Förster Resonance: Mean field non-linearity. The condition for the Förster resonance $\Delta_F = 0$ is narrow and needs careful tuning of the microwave frequency and power. To achieve this, we measure the mean-field interaction between Rydberg polaritons. **(a)** Depending on the intermediate state detuning Δ_e in our EIT scheme, the character of the interaction can be changed. On resonance ($\Delta_e = 0$) the interaction potential is mostly imaginary leading to loss of additional polaritons, while at positive detuning the interactions are predominantly real and lead to a dispersive shift of the transmission peak and increased overall transmission (see inset). For maximum contrast we probe at an intermediate state detuning of $\Delta_e = 12$ MHz. We measure the exact resonance condition by observing a reduction in transmission **(b)** and an increase in dark polariton linewidth **(c)** as we vary the microwave drive frequency.

To experimentally tune into a Förster resonance, precise calibration of microwave power, polarization, and frequency is necessary (see fig. 5.14). While 2-photon correlations provide an unambiguous signature of photon blockade, they are time-consuming to measure and hence a very inefficient method to find the Förster resonance. Instead, we locate the resonance by using the mean-field nonlinearity of the transmitted probe light as a proxy for interactions. The general scheme of finding the Förster resonance is as follows: We probe the system at a 780 nm cavity photon flux that ensures that in steady-state multiple polaritons (around 3) are excited. When fine-tuning into resonance, the Rydberg blockade radius becomes larger (due to increasing interaction strength), and hence fewer polaritons can be simultaneously excited, resulting in a reduced cavity transmission.

In order to increase the transmission contrast, we furthermore operate detuned from the

intermediate $5P_{3/2}$-state. For our ladder EIT scheme, the exact interaction potential of two dark polaritons depends upon the intermediate state detuning Δ_e as detailed in section 3.11. With a detuning off the intermediate state $\Delta_e \gg \Gamma$ the photon is mostly reflected off the cavity as in a real valued interaction.

We operate with a small detuning $\Delta_e \approx 2\Gamma$ where the interaction between polaritons leads both to a mean field shift of the two-photon resonance as well as an increased loss broadening the polariton linewidth. In fig. 5.16 (a) the increase in transmission for a positive p-state detuning Δ_p is seen to double the cavity transmission while maintaining large detuning from the bright polariton (see lower inset).

Due to technical details specific to our setup (microwave power can only be tuned in coarse steps on the source), we operate at a fixed microwave power that gives us the calculated d-state shift and then vary the microwave frequency f_{AC} around the point indicated in fig. 5.15 to tune exactly onto the resonance. The resulting dark polariton signal is shown in fig. 5.16 (b). By fitting Lorentzians to this data, we extract the transmission and linewidth as a function of the microwave frequency. As shown in fig. 5.16 (c) a reduction to half the transmission (green crosses) and simultaneous increase in linewidth (purple circles) for a narrow range in microwave drive frequency can be observed. From this we extract an experimental microwave frequency of 23708.57(2) MHz, showing a slight deviation from the calculated parameters.

The true hallmark of strong single-photon-level non-linearity is the observation of photon blockade. We quantify the blockade by measuring correlations in the transmitted photons using single-photon-counting-modules (SPCM) and time tag the photon arrival times to compute coincidences as in fig. 5.8.

The measured $g^{(2)}$ correlation function is shown in fig. 5.17 without and with Förster resonance in (a) and (b), respectively. To rule out any microwave dressing or admixture effects, the microwave drive was turned on in both cases but shifted from resonance by

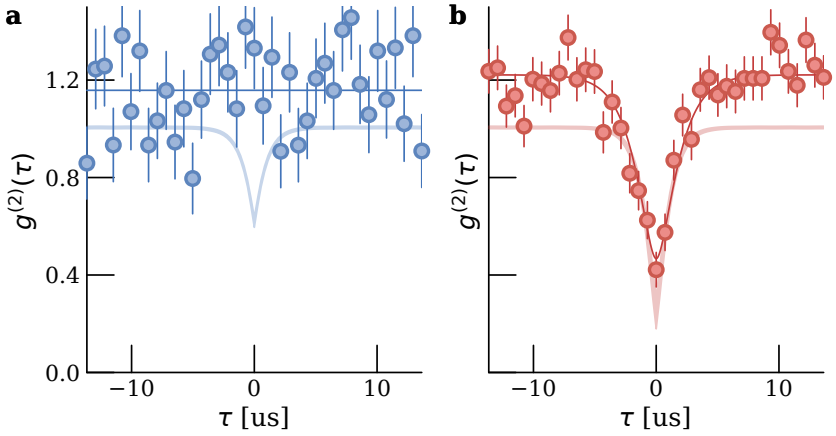


Figure 5.17: Observation of photon blockade via Förster resonance. We measure the temporal correlation function $g^{(2)}(\tau)$ of photons transmitted through the cavity QED system. **(a)** Away from the resonance, interactions in the $|44D\rangle$ state are insufficient to blockade the $LG_{0,0}$ mode, yielding $g^{(2)}(0) = 1.00(16)$. **(b)** At the Förster resonance, interactions are enhanced, resulting in a suppression of multi-photon events with $g^{(2)}(0) = 0.38(11)$. The atomic cloud is confined to a quasi-2D geometry ($\sigma < R_b \lessapprox \omega_0$), ensuring that the blockade radius covers the transverse optical mode. Shaded bands indicate *ab-initio* atomistic calculations (see appendix A.6.1) using experimentally measured cloud thickness σ_z. In **(a)** the microwave drive is still present but intentionally detuned in frequency to exclude effects of state admixture.

18 MHz in (a). This small change in frequency breaks the exact resonance condition but would change off resonant effects from other states only slightly, it also keeps the $|dd\rangle$ state and thereby EIT resonance unchanged. While no antibunching dip is visible for the detuned microwave drive (a), a clear dip is visible at Förster resonance (b). Using a simple analytical model [CSB+20] $G^{\text{mod}}_{mn}(\tau) = G^{\text{bkg}}_{mn}\left(1 + (1 - G^{\text{min}}_{mn})(e^{-\gamma|\tau|} - 2e^{-\gamma|\tau|/2})\right)$, we extract a $g^{(2)}(0) = 1.00(16)$ in (a), and a $g^{(2)}(0) = 0.38(11)$ in (b), corresponding to an interaction strength of $U = 1.80(36)$ MHz. The width of the correlation dip is determined by the dark polariton linewidth $\gamma_{DP} = \cos^2(\vartheta)\kappa + \sin^2(\vartheta)\gamma_r$ as derived in section 3.10.

We compare the measured correlations with an atomistic, non-Hermitian perturbation theory that estimates $g^{(2)}$ based on the theoretically calculated [ŠPAW17] C_6 and C_3 values. The results are shown in fig. 5.17 as shaded bands reflecting uncertainty in the exact cloud thickness, yielding $g^{(2)}(0) = 1.00(16)$ and $g^{(2)}(0) = 0.38(11)$ for the two cases, respectively. Both measured $g^{(2)}$ are compatible with the ab-initio calculations within the experimental

errors, however the experimentally measured values of $g^{(2)}(\tau)$ at longer τ deviate from the expected value of 1 for a coherent state. These extraneous correlations can be caused by atom number fluctuations between- and within- experimental runs. Note, however, these classical fluctuations can only increase the $g^{(2)}$ value.

From the theoretical model, we can furthermore deduce that the Förster enhanced interactions at $n = 44$ lead to a similar blockade strength to what we anticipate for Van-der-Waals interaction at $n = 78$, while requiring less blue power (scaling $P \propto n^3$, indicates 5.6 times less) for the same Rabi frequency, and reducing the DC polarizability (and thus field sensitivity) by a factor of 55.

CHAPTER 6

SMALL WAIST CAVITIES AND ARRAY

Having demonstrated a finesse of $\mathcal{F} = 33$ k in a lens based cavity using two spherical lenses in the setup for fig. 3.4 breaks the conventional wisdom that no optics other than mirrors can be used for high-finesse cavities. Including optical elements such as lenses, waveplates and crystals into cavities opens up a whole new field of cavity engineering using all the tools of optics.

First of all, decoupling the geometrical twist due to non-planar reflections from the focusing action in chapter 5 led to drastically lower optical aberrations and enabled a truly degenerate lowest Landau level cavity. In the previous iteration of this apparatus, only curved mirrors had to both provide a narrow focus and non-planar geometry making non-normal incidence causing strong odd-order terms like cubic aberrations. Using aspheres in the new cavity the same small mode waist of 18 um could be achieved and as the mirrors can consequently be made flat they do not introduce any aberration at arbitrarily large angles of incidence.

Furthermore, lenses in cavities allow to decouple the roles of propagation and focusing as they can be placed anywhere along the optical axis, whereas mirrors are always constrained to sit at the mode boundary.

6.1 Lens based degenerate cavities

For Fabry-Perot cavities both the confocal and concentric regime are desirable as the fundamental mode waist becomes minimal in these two limits leading to a higher cooperativity. More elaborate schemes [KML$^+$23] seek to reach higher cooperativity through the interference of multiple higher order modes in a *super mode*. The confocal cavity is also interesting as it enables non-local connectivity in the central plane due to the two-fold symmetry

[KML+23]. Many quantum information experiments such as the emerging field of cavity readout of single atoms in tweezers [HSW+24] therefore want to operate in the concentric or confocal regime, however as they both form the edges of the stability region they are only marginally stable and can be disturbed by mirror misalignment on the pm level. The stability diagram in fig. 6.1 (b) shows this mode waist as a function of cavity length i.e. mirror separation L. Building a cavity using lenses now allows for a new type of cavity

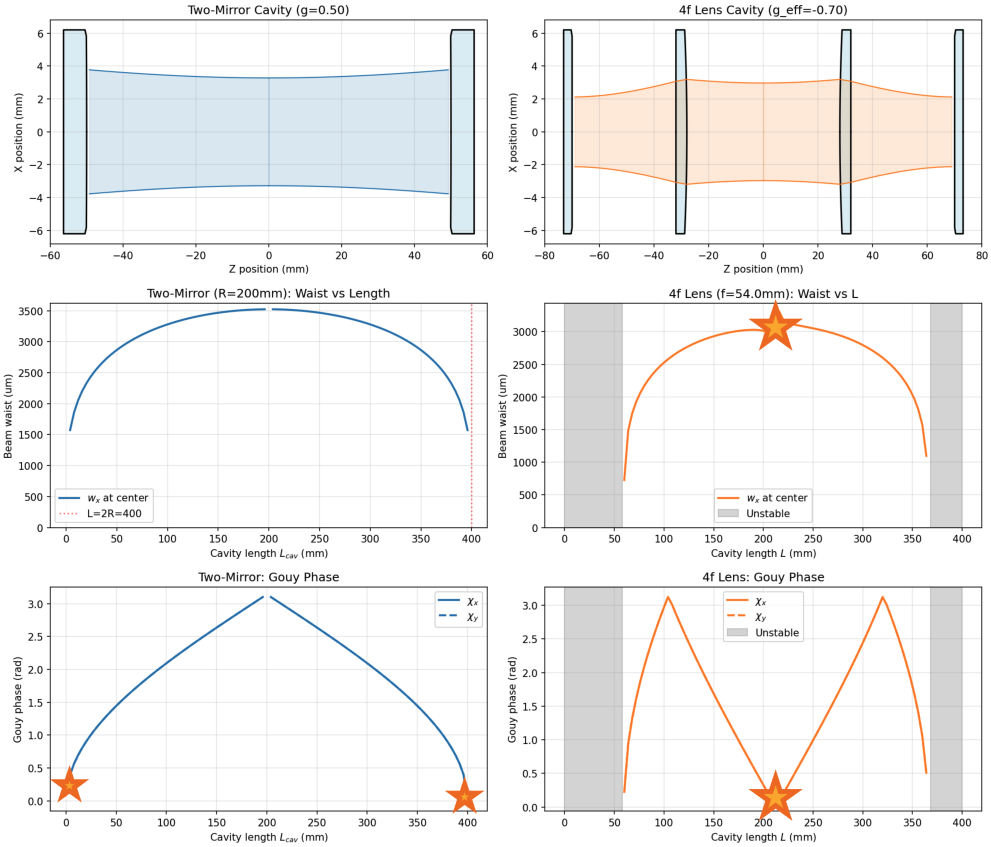


Figure 6.1: Comparison of a two mirror and two lens cavity. A regular Fabry-Perot cavity **(a)** using two curved mirrors can only be made degenerate at the edges of its stability region, in the confocal and concentric limit. Conversely, a $4f$-cavity using two lenses and flat mirrors **(b)** can be degenerate right in the center of its stability region, enabling stable fully degenerate cavities.

that is shown in fig. 6.1 (right column): We call it the $4f$ cavity in analogy to $4f$ imaging systems comprised of two lenses separated by exactly their cumulative focal lengths and two mirrors also one focal length away from each lens. Usually we are interested in the imaging

properties of this often termed *relay* system that can also provide magnification and cancels certain aberrations due to symmetry. In a cavity however, this same property responsible for imaging now leads to a phase condition due to the Gouy phase of the lens pair that makes this cavity degenerate exactly in the center of its stability diagram. As shown in fig. 6.1 (f) the transverse mode splitting given by the Gouy phase is exactly $mod(\pi, \pi) = 0$ at $L = 4f$ while the waist is maximal at the center of the stability diagram. This allows to realize a fully degenerate cavity that is unconditionally stable that can be used to realize arbitrarily small local modes or exotic physics like optical wormholes [Tan].

6.2 GRIN lens cavity for sub-micro waist

During the quest to minimize aberrations in multi-mode cavities for chapter 5 the Landau level cavity, we came across a very different kind of lens candidate in gradient-index (GRIN) lenses. Compared to regular curved or spherical lenses, that achieve their focusing action through a curved surface that causes a phase accumulation which is to lowest order quadratic in the radius of a ray, GRIN optics achieve the same through variations in the index of refraction with radial distance. This has several key advantages: The optical surfaces of GRIN optics are flat, which reduces spherical but also other aberrations and notably chromatic aberrations as well, which is why they are often used in *in-vitro* bio imaging applications using white light. Also, the flat surfaces can be polished to a very low surface roughness using traditional polishing techniques and are not limited in their focal length by the leading MRF machine available. (The $f = 10$ mm lenses used in chapter 5 are the most curved lens any company could super-polish using MRF). GRIN optics are manufactured by exchanging Sodium ions in silicate glass by either Silver or Lithium ions that change the refractive index of the glass. This is usually achieved through diffusion in a molten salt bath which makes the radial concentration gradient that is approximately quadratic. By changing the diffusion constant over time via temperature or concentration, more complex radial index

profiles can be realized. Therefore, GRIN lenses start out much closer to the ideal quadratic phase profile than spherical lenses do. As the profile is written via diffusion, it is perfectly coaxially aligned with the glass cylinder and orthogonal to the end faces which significantly simplifies alignment. The refractive index profile falls off quadratically

$$n(r) = n_0 \left(1 - \frac{g^2 r^2}{2}\right) \tag{6.1}$$

from the central value n_0 with an empirical coefficient g. In this quadratic *duct* [Sie86] light rays undergo periodic orbits and the exact action of the lens is determined by its length or *pitch* $P = \frac{2\pi}{g}$ where for a focusing lens we want $L_{1/4} = \frac{\pi}{2g}$ for an infinite conjugate ratio. The NA of a GRIN lens is limited by the index contrast and can be as high as 0.5 making them well suited for multi-mode cavities. The continuous focusing action of a GRIN lens is

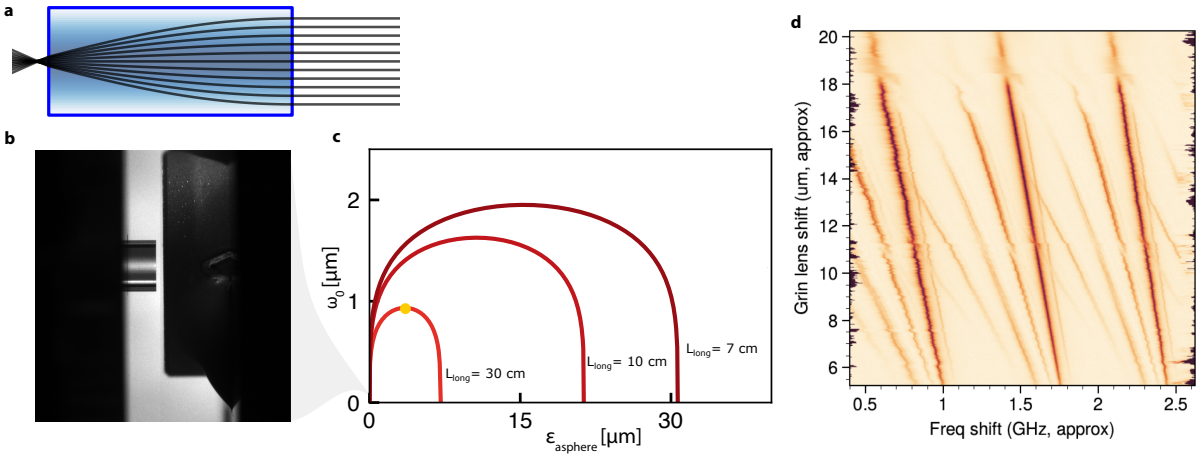


Figure 6.2: GRIN lens cavity. A gradient-index GRIN lens **(a)** focuses light through a quadratic refractive index, not curved surfaces. Due to the short focal length and working distance the lens has to be placed 230 um in front of the endmirror (b). Fundamental mode waist as a function of second mirror distance L_{long} shows waist approaching diffraction limit. Cavity spectrum over full stability region (d) shows higher order modes and small avoided crossings.

shown in fig. 6.2 (a) by integrating the propagation as a differential equation.

6.2.1 Cavity Spectroscopy

To form a stable cavity, a GRIN lens is placed between two mirrors with the lens placed $f + \epsilon$ away from a flat mirror surface to make the cavity stable. The separation ϵ can then be tuned to change the Gouy phase and scan over the whole stability diagram.

Due to their small size desirable for biology application and high NA, most commercially available GRIN lenses have a very short focal length. The lens shown in fig. 6.2 (b) has a diameter of 1.8 mm and focal length of only 1.74 mm. As the focusing action does not happen at a discrete curved surface but is distributed over the whole length of the lens, this makes the working distance very short with only 0.23 mm. The lens mirror spacing is therefore controlled by a flexure piezo stage that provides sufficient throw to scan the whole stability region and carefully approach the mirror without damaging any optical surface. The stability diagram fig. 6.2 (c) shows that the fundamental waist of this cavity can be reduced to 1 um and below by increasing the separation of the second mirror also called the long arm distance L_{long}.

As we are interested in studying the effect of aberrations in this cavity, we scan across the whole stability region while sweeping the laser frequency to obtain spectra as a function of mode spacing as shown in fig. 6.2 (d).

6.2.2 Scanning defect waist microscopy

A number of effects such as mode mixing, non-paraxial propagation, vector effects such as longitudinal polarization could affect the true waist size for a diffraction limited cavity. We therefore use a form of scanning probe microscopy in order to verify the experimental waist of the GRIN and asphere based cavity.

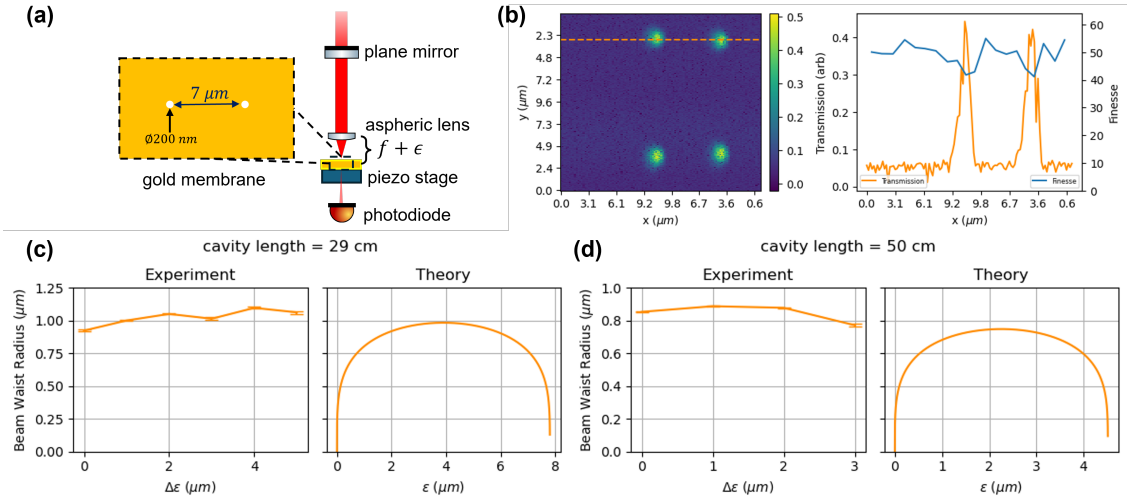


Figure 6.3: Scanning probe microscopy of diffraction limited waist. We experimentally validate the diffraction limited waist of the GRIN and asphere based cavity *in-situ* using a sub-wavelength defect. As any object the size of λ would completely spoil the cavity, we use a nano-fabricated membrane with sub-wavelength size holes to act both as a cavity end mirror and small defect (a). The size of the nanopore is experimentally chosen at $d = 200$ nm to induce sufficient loss to reduce the cavity finesse only slightly (b) but evanescently transmit a small fraction of the cavity field. (c) and (d) comparison of measured and theoretical cavity waist for two different L_{long} shows a close match of measured waist to paraxial theory and validate diffraction limited performance.

6.3 A first cavity array

6.3.1 Mode picture for cavity array

To fully understand the idea of the cavity array geometry, one has to first consider a fully degenerate cavity that can be realized in a stable way using intra-cavity lenses.

The eigenmodes of a Fabry-Perot cavity are fundamentally Gaussian modes, as they are eigenstates of the free space propagation together with the boundary conditions given by the curved mirrors. The higher order modes can be equivalently described in a Hermite-Gaussian (HG) or Laguerre-Gaussian (LG) basis where this symmetry is broken either if birefringence in an optic breaks the equivalence between x- and y-axis favoring HG, or if any image rotation is present favoring the rotationally symmetric LG modes. In cases where astigmatism is present, the modes can also be of a more general family called Ince- or General-

Gauss modes that interpolate between LG and HG. For many applications such as lasers, photonic whispering gallery ring-resonators, fiber cavities for quantum transduction etc., only the fundamental GG_00 mode is used. In photonic ring-resonator arrays, multi-mode systems are constructed by fabricating an array of cavities close to each other to provide a tunnel coupling and realize a tight-binding Hamiltonian. This requires a macroscopic number of cavities that is only feasible in fabricated nano-photonics and not using macroscopic optics.

In some specialized settings, the higher order modes of a single cavity can themselves be used to realize quantum many-body systems. A premier example is the realization of a Landau level for light by creating an artificial magnetic field for photons through a twisted cavity geometry. In this specific case, the LG eigenmodes of the resonator are exactly equivalent to the Landau level wavefunctions, the eigenstates of an electron in a strong magnetic field, in the symmetric gauge.

While this enables boutique quantum simulation experiments of continuum systems, the higher order HG and LG modes are a poor match for the emergent digital quantum simulation platforms utilizing arrays of single atoms in optical tweezers or lattices.

As the HG and LG eigenmodes are equivalent to eigenstates of the 2D quantum harmonic oscillator, they have a field maximum at the classical turning points but are fundamentally delocalized in real space. This is reflective of the fact that they are eigenstates of the momentum operator in the Helmholtz equation. With increasing mode number, their field gets more and more spread out with an increasingly complicated envelope with a growing number of nodes. Placing an array of atoms inside a cavity, any higher order mode will couple to an increasing amount of atoms making readout of a single atom impossible. While a series of more complicated techniques including hiding the atom in internal states, or placing them along the cavity axis and using multiple longitudinal modes together with a longitudinal Zeeman field gradient, attempt to fix this shortcoming, the only scalable solution is to re-engineer the cavity itself: By expanding the cavity toolbox and allowing for more

intra-cavity optics, the natural basis of the cavity itself can be re-shaped to a localized basis ideally suited for local addressing and readout.

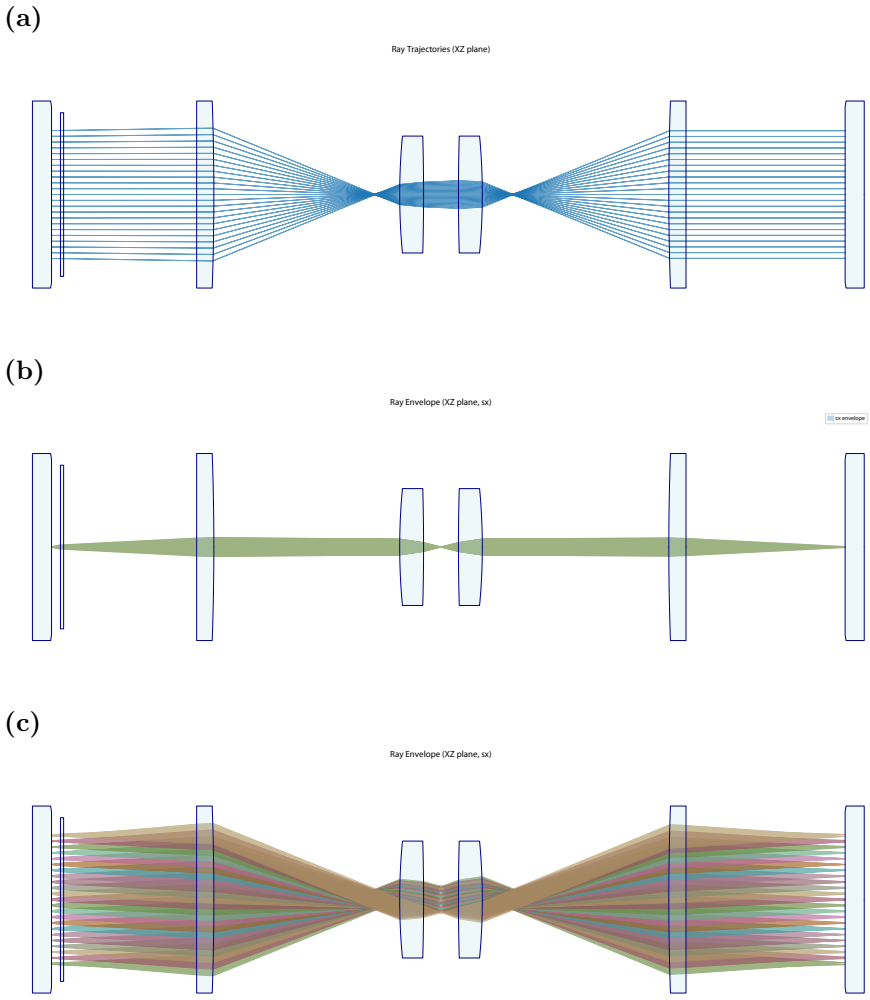


Figure 6.4: CAM $8f$ cavity geometry. The fully degenerate $8f$ cavity geometry at the heart of the cavity array microscope fulfills two simultaneous imaging conditions: (a) The position of the local cavity axis (blue) is imaged from the input (left mirror) to the atom plane with demagnification and then re-imaged by the right $4f$ system. A single high NA beam (green) gets focused to a small waist in the atom plane and is also imaged onto the two end-mirrors. (c) A microlens array (MLA) right next to the input mirror stabilizes the local cavities from disorder and allows an array of tightly focused waists in the atom plane.

Only when all the higher order modes of the cavity are degenerate can they interfere to form localized modes away from the optical axis. In a simple two-mirror Fabry-Perot resonator, full degeneracy is only reached in the concentric configuration $R = L/2$ where

the Gouy phase is $\chi = 2\pi$. As this mirror spacing is right on the edge of the stability region and only marginally stable, in practice even pm deviations and misalignment make this configuration unstable and hard to maintain in practice. The mode waist inside of the cavity is also very small while conversely it is large on the location of the end mirrors.

Instead, the landscape of cavity geometries can be significantly expanded when lenses are introduced inside the cavity, as they allow decoupling of Gouy phase/focusing from propagation distance. In fig. 6.1, the stability diagram, showing waist as a function of cavity length, and the Gouy phase are shown for the case of a two-mirror cavity (a) - (c) and for a two-lens cavity (d) - (f) using flat mirrors.

Opposed to the mirror cavity, the lens cavity actually reaches full degeneracy $\chi = 0$ right in the center of its stability diagram (e) with a finite waist and being unconditionally stable. This point exactly occurs when $L = 4f$, i.e. the lenses form exactly a 4f-telescope that is, a telecentric imaging system. This allows another intuitive explanation of the degeneracy. A cavity is stable when a ray returns onto itself after an integer number of round trips, making it an eigenray. Usually this is only true on the optical axis, and a displaced ray will undergo more complicated dynamics equivalent to the motion of a particle in a harmonic potential. For the 4f degenerate lens cavity however, every input ray gets imaged back onto itself due to the two passes through the 4f imaging system, making any input ray an eigenray.

This means the cavity can be excited by a Gaussian beam at any location in the input plane, which in turn gets imaged exactly onto the second mirror. Also, the size of this excitation is no longer fixed as more modes can interfere constructively to make a smaller feature. A degenerate 4f lens cavity is therefore ideal for interfacing to single atoms, as the light can be localized exactly to a single atom's location.

The perfect imaging condition of the 4f telescope and the corresponding degeneracy of the cavity up to infinitely high order modes is of course only true in the paraxial limit: In practice, spherical aberrations as the dominant aberration in this design introduce a

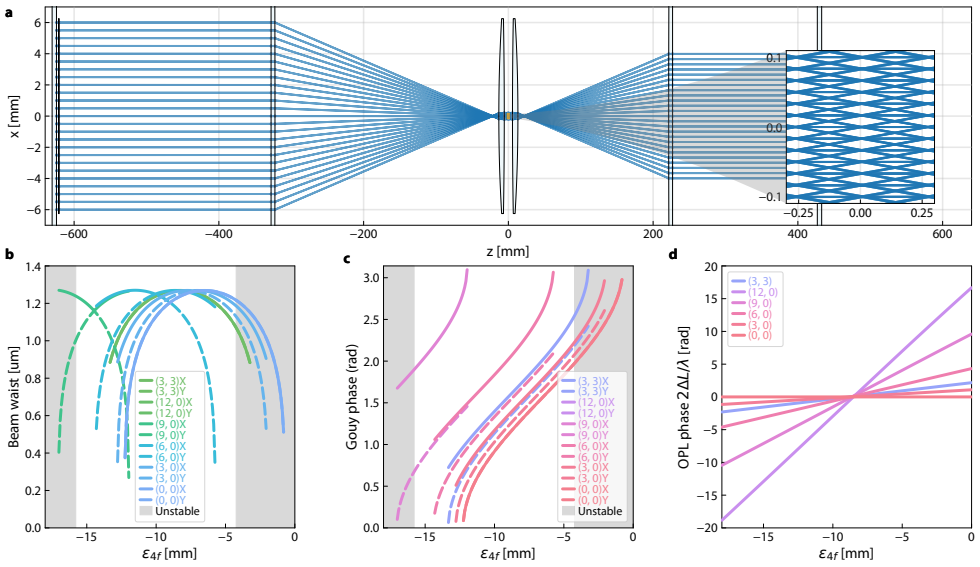


Figure 6.5: Paraxial analysis of CAM local cavities. The degenerate $8f$ cavity together with the MLA can stabilize an array of local cavity modes (a) with small waists at the plane of atoms (z=0) providing high cooperativity for individual atom readout and manipulation. The perfect $4f$ imaging condition and thereby degeneracy of all local cavities is broken by optical aberrations for real optics: We extract the shift in stability region (b), Gouy phase (c) and optical path length (OPL) for each individual cavity by differential ray tracing.

quadratic r^2 phase shift breaking degeneracy for modes at larger radii r and limiting the size of the addressable region. Furthermore, any irregularity in the optics shape such as form error, roughness or strain introduce disorder, leading to mode mixing and smearing out the Gaussian input beams into deformed blobs.

Based on the picture developed in the context of the twisted cavity earlier, small detunings of the lens spacing from degeneracy generate a mode dispersion acting as a potential for photons. Instead of detuning the 4f cavity, which would break the degeneracy and introduce a single, global potential, a micro-lens array can be introduced into the cavity to act as a local potential, stabilizing an array of local fundamental modes protected from disorder.

6.3.2 Aberrations

6.3.3 Cavity resonance condition

The cavity is resonant at frequencies ν where

$$\nu(L,\chi) = \frac{c}{2L}\left[q + \frac{n+m+1}{\pi}\chi\right] \tag{6.2}$$

where $\frac{c}{2L}$ corresponds to the Free Spectral Range (FSR) and total mode index $N = n+m+1$. To determine the sensitivity to small fluctuations in cavity length (δL) and Gouy phase ($\delta\chi$), we calculate the total differential $d\nu$:

$$\delta\nu = \frac{\partial\nu}{\partial L}\delta L + \frac{\partial\nu}{\partial\chi}\delta\chi \tag{6.3}$$

giving

$$\delta\nu = -\frac{\nu}{L}\delta L + \frac{\nu_{FSR}}{\pi}N\,\delta\chi \tag{6.4}$$

we can rewrite this using $\frac{\partial\nu}{\partial L} = -\frac{\nu}{L} = -\frac{c}{\lambda L} = -\left(\frac{c}{2L}\right)\frac{2}{\lambda} = -\nu_{FSR}\frac{2}{\lambda}$ to

$$\delta\nu = \nu_{FSR}\left(\frac{N}{\pi}\delta\chi - \frac{2\delta L}{\lambda}\right) \tag{6.5}$$

OPL, local Gouy phase...

In the ABCD matrix formalism, rays are always considered close to the optical axis and only the quadratic action of surfaces or propagation is considered. The paraxial properties such as stability, waist and Gouy phase can therefore be only calculated analytically for the central cavity, whereas any off-center modes are significantly displaced from the optical axis and hit lenses at large angles. Furthermore, the additional telecentric condition of the 4f telescopes imposes an additional constraint not relevant for the central cavity, which narrows the stability region with increasing distance from the center. To model the non-paraxial

effects from spherical lenses and aberrations of the aspheres, we therefore perform numerical ray tracing of the full cavity without any approximations and find the local eigenmodes in the array by finding the local eigen-rays that are closed upon a round trip. We then make use of our custom differentiable ray tracer to extract local ABCD matrices around each eigen-ray to capture the non-paraxial effects on Gouy phase and waist.

This allows us to bridge the ray picture for off-axis behaviour with the wave picture for diffraction effects and extract the resonance frequency of each local cavity mode

$$\nu(L, \chi) = \frac{c}{2L}\left[q + \frac{n+m+1}{\pi}\chi\right] \tag{6.6}$$

that depends on optical path length L, mode indices q, n, m and Gouy phase χ. Beyond the purely classical optical path length L and associated phase, an additional phase is acquired due to diffraction of the Gaussian mode and even for the fundamental mode $n = m = 0$ the Gouy phase χ shifts the fundamental mode energy. The total shift for small changes in OPL δL and χ can be written as

$$\delta\nu = \nu_{FSR}\left(\frac{N}{\pi}\delta\chi - \frac{2\delta L}{\lambda}\right) \tag{6.7}$$

highlighting that both contributions have an opposite effect on the modes resonant frequency. In order to achieve degeneracy over the whole array, these two terms must therefore be carefully balanced.

CHAPTER 7

A CAVITY AFM

7.1 Introduction

Atomic force microscopy [BQG86] is one of the techniques that opened up a window into the atomic scale of our world. AFM is a ubiquitous tool across the sciences: It enables imaging of cells, biomolecules and proteins in biology and medicine [NA21], allows for characterization of surface properties and study of nanoscale features in material science, facilitates study of electronic properties, magnetic structures and vortices in superconductors in condensed matter physics [AL20, Gie03]. It achieves sub-optical, atomic resolution by sensing the surface morphology through interaction of the sample with a sharp tip deflecting a cantilever. The exact properties of the tip determine the minimum lateral resolution. In addition to simply probing the surface profile it also allows for measuring of magnetic or biological forces like single molecule forces through functionalization of the tip [JCL23]. It is also one of the few techniques that can be used for active manipulation, assembly or probing of cells, single molecules or single atoms. Since its inception, AFM has matured as a technique and has seen numerous technical optimizations. Every part of an AFM has received intense engineering over the years. Nowadays, cantilever parameters and tip sizes have been pushed to a few atoms, the scanning systems have evolved from tube piezos to optimized high speed scanning piezos and complicated scanning waveforms and control loops are employed to increase accuracy and scanning speed. Also the detection mode has evolved from the original contact mode to various tapping and non-contact frequency- or amplitude modulation techniques to measure not only surface profile but also surface properties or torsional forces.

All of the techniques using modulation make use of the fact that the cantilever is a mechanical harmonic oscillator, providing signal gain close to its resonance frequency and the presence of surface forces on the cantilever cause a shift of the oscillator's eigenfrequency

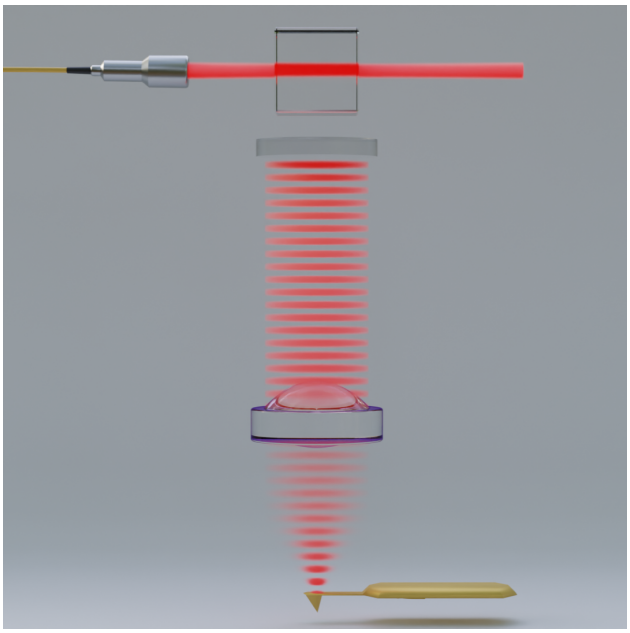


Figure 7.1: Resonator Design. The heart of the apparatus, shown in **(a)**, is a cavity formed by the reflective cantilever back surface, a flat mirror and a high-NA aspheric lens (focal length f) that focuses light down to a micron size spot and makes the cavity configuration stable. A cavity transmission spectrum in **(b)** shows a (fitted) finesse $F = 80(2)$; for a 1.1 GHz free spectral range (FSR), the linewidth is thus $\kappa = 2\pi \times 12.4(7)$ MHz. The distance of the cantilever from the lens is $f + \epsilon$ while the long arm length L is chosen to provide a waist size of 2.5 um limiting clipping and diffraction loss from the limited beam width even for small, high frequency cantilevers.

which can be measured very precisely. This provides the necessary amplification to measure surfaces on the sub nm scale or forces on the nN scale using a relatively simple optical readout. Even with all this innovation, the premier method of measuring cantilever deflection is still via optical displacement of a laser beam that is reflected off the back of the cantilever and measured on a split photodetector. A number of different readout schemes like piezoelectric detection, capacitive readout or laser Doppler vibrometry have been developed but have not found broad use and remain mostly used in specialized applications. The most sensitive tools developed to measure distances with light use interferometry (for example the LIGO detector can measure fractional length changes on the order of $10^{-23}/\sqrt{Hz}$ due to gravitational waves [Mar16]). It therefore seems intuitive to use a Michelson or similar interferometer to read out the cantilever deflection, and there have been many proposals and realizations [RF16, And18]

thereof. However, these methods are very sensitive to laser amplitude noise and wavelength drift. More surprisingly even, the optical deflection method using a split photo detector has the same sensitivity as a simple Michelson interferometer [PdvG92] while avoiding a lot of the technical problems through an inherently matched path length. In our realized readout method, an optical cavity of finesse F enhances the sensitivity of the z measurement by the finesse of F that can be in excess of 100,000 for optical cavities [JMM$^+$22]. Simple, Fabry-Pérot type cavities have been used previously [KÇO20], usually involving a curved fiber end face as one of the mirrors, but realizations have been limited to relatively low finesse or simple readout methods as reflected power. We report on a new cavity geometry using an intra-cavity asphere, enabling the use of off the shelf cantilevers and flat, high quality mirrors. The achievable size of the cavity waist is usually limited by the overall length of the cavity, and a small waist requires a highly curved mirror and operation close to the edge of the stability region. Our lens based cavity enables waist sizes of few um or even sub 1 um making it compatible with small, high frequency cantilevers and de-coupling other parameters such as cavity FSR and linewidth.

The combined cantilever and optical cavity system constitutes two coupled harmonic oscillators. This is the prototypical optomechanical system [KV08] where the cantilever acts as a movable endmirror that is subject to photon pressure from the cavity field, which in turn provides back-action by changing the cavity length. The optomechanical system itself is well studied [KV08], but we can harness its properties to realize a new AFM readout method: Probing the coupled cavity cantilever system reveals sidebands around the resonances of the optical cavity spaced by the mechanical oscillator frequency. Driving the cavity with a blue detuned sideband $\omega_c + \omega_m$ pumps energy into the system causing excitations of the mechanical degree of freedom. With increasing power of the sideband drive, the system undergoes a bifurcation phase transition and starts mechanical lasing at the eigenfrequency of the cantilever.

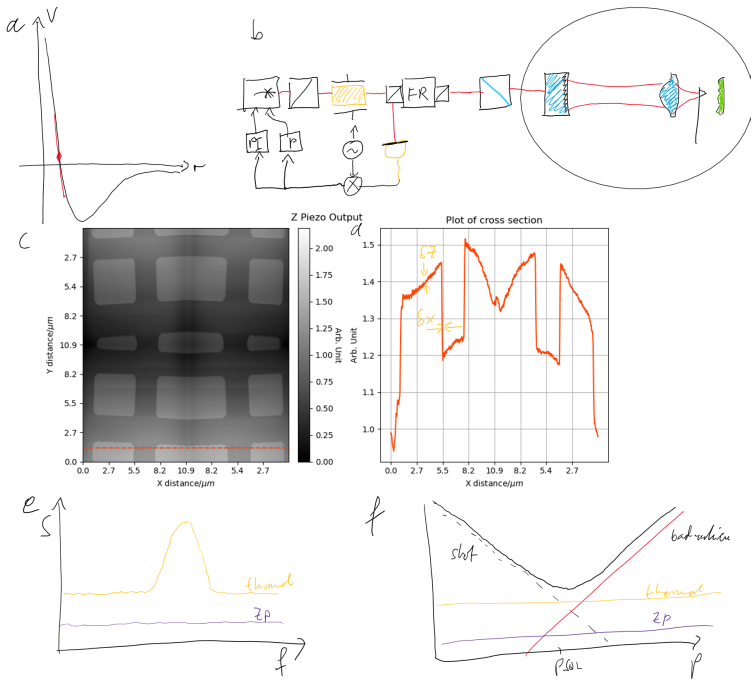


Figure 7.2: Contact mode AFM (a) Van-der-Waals potential between the tip and the sample surface as a function of separation z **(b)** Full schematic of the optical setup used in the CM-AFM comprised of a diode laser, electro-optical modulator (EOM) and reflection measurement using a Faraday Isolator. The laser frequency is stabilized to the cavity resonance using a Pound-Drever-Hall error signal and a PI controller. **(c)** AFM image of a Budget Sensor HS-20MG calibration target obtained using this setup at a resolution of 512×512 points and line scan frequency of 1.2 Hz. **(d)** Cross-section of **(c)** (indicated by red dashed line) reveals axial and lateral resolution of 123 nm and 123 nm respectively. **(e)**

Forces from the sample surface acting on the tip change the eigenfrequency of the mechanical oscillator which can be read out through a change in the cavity frequency. The hybrid system therefore combines the sensitive deflection measurement of the cavity with the ability for completely passive FM readout. FM operation in conventional AFMs is achieved through a mechanical drive of the cantilever via piezos or thermal excitation using pulsed lasers, and feedback of the readout signal to the drive via an external PLL. Two challenges here are the limited bandwidth of the feedback due to the loop bandwidth and delays in readout and piezo actuation, and the excitation of unwanted higher order modes of the cantilever. Our purely optomechanical excitation scheme solves both of these simultaneously, as the excitation of the cantilever at its resonance frequency is autonomous and happens

within the physical system, and it only couples to relevant mechanical modes of the beam that change the cavity length. Thereby it provides extremely high bandwidth and spectrally clean excitation. All of the sophisticated microscopy modes in AFM (AM-AFM, FM-AFM, ...) still rely on the original optical deflection method on a split photo detector, that does not optimally use the available photons.

In this work we present, for the first time, an all optical AFM readout and excitation scheme. Using optical cavities to measure cantilever displacement in AFM has been previously proposed [HLZ20]. This enables pushing the speed while maintaining high sensitivity in the sub nm range. In Section 7.2 we introduce the cavity geometry and optical setup. In Section 7.3 we describe how to operate the cavity AFM in the contact mode regime and show images and noise spectra. In Section 7.4 we demonstrate optomechanical excitation and readout first without and with a sample and use it to image samples in FM mode.

7.2 Cantilever Lens Cavity

The main advantage of using an optical cavity is the possibility to measure the exact position of the end mirror by accruing a position dependent phase shift on the intracavity field for the same photons over many round trips. For that, we desire a high cavity finesse F that will linearly improve the sensitivity of the readout over the deflection method by F for a fixed photon number and interrogation time. A key challenge in using an AFM cantilever as a cavity end mirror is their small size of $w = 35$ um for low frequency, contact mode cantilevers down to $w = 5$ um for high frequency cantilevers. To not suffer from clipping or diffraction losses when the reflective surface is that small, the size of the focus must be several times smaller than the cantilever width or length. Using a simple clipping model, a beam waist of $\omega = 1$ um limits the achievable finesse to $F_{max} = erfc(\frac{\omega}{d})$. In two mirror, Fabry-Pérot type cavities this is achieved using short cavities with a highly curved end mirror in addition to the cantilever. This produces large free spectral ranges (FSR) and cavity linewidths, that

can make it technically hard to lock a laser as their modehop free tuning range is usually limited to a few GHz. In contrast, we achieve a small cavity waist using a novel geometry using an intra-cavity high-NA aspheric lens (focal length f) limiting clipping and diffraction loss from the limited cantilever width even for small, high frequency cantilevers. See A.6.1 for more details on the cavity stability diagram.

We measure cavity properties such as FSR and finesse by sweeping a laser (IPS I0780S50100B) over several FSRs and recording the reflection spectrum. The fitted linewidth $\kappa = 2\pi \times$ 12.4(7) MHz and FSR $FSR = 1.14(1)$ GHz resulting in a finesse of $F = 80(2)$. Currently the finesse is limited by the finite reflectivity of the commercially available gold coating on the back of the cantilever. As the reflectivity is not paramount for deflection detectors, the coating thickness is only 70 nm leading to a reduced reflectivity of %, falling short to the bulk value of 99.5 % because of evanescent effects. The finesse could therefore be readily improved with thicker metallic or dielectric coatings in the future.

7.3 Contact mode

In contact mode (CM-AFM) the tip is in constant contact with the sample surface so if we keep track of the cantilever mirror displacement we obtain the surface morphology. To this end, we either need to maintain a constant cavity length through locking the cavity to a stable reference laser by feeding back on the cavity end mirror, or equivalently, we can follow the cavity resonance by changing the laser frequency. One of the main challenges in cavity AFM readout is the narrow linear region where the error signal is proportional to cantilever displacement, which is narrower by a factor of the finesse $\approx \frac{1}{\lambda \mathcal{F}}$ compared to $\approx \frac{1}{\lambda}$ for a simple interferometer. Differences in sample height of more than a few nm therefore require high feedback bandwidth to stay on the locking fringe. The first approach of locking the cavity to the laser requires an already narrow and stable reference laser and can directly benefit from the absence of frequency fluctuations that through the cavity get transferred

into amplitude or error signal fluctuations, however the achievable feedback bandwidth is limited by mechanical resonances of the piezo actuated mirror that usually occur at 10s of kHz for mount or piezo resonances. We therefore choose the second approach and lock the laser frequency to the cavity which provides a much higher feedback bandwidth of up to 1 MHz through the diode current. This can add noise to the readout signal but since our laser has an intrinsic linewidth of $\approx$ 100 kHz compared to the cavity linewidth of 10 MHz this contribution is much smaller than other noise sources in air and at room temperature. For FM readout with high frequency cantilevers, the laser noise is also far separated in frequency and can therefore be effectively suppressed by the lock without influencing the cantilever motion. In 7.2 (b) the optical setup for contact mode is shown: To achieve a high lock bandwidth, we choose to lock the laser to the cavity using a Pound-Drever-Hall (PDH) scheme and a PI controller on the laser current versus locking the cavity to the laser and being limited by piezos and mechanical resonances. The lock output is then proportional to the cavity length and thereby the cantilever displacement. Because of the high cavity finesse and the resulting large change in cavity frequency δf per cantilever displacement δz of $\delta f = \frac{\delta z}{\lambda/2} F\kappa$ a large capture range is required even for modest z step heights on the sample of 20 nm. We choose a PDH modulation frequency of up to 300 MHz to achieve a wide error signal (see A.6.1) and keep the laser in lock during a scan. The PDH scheme makes the readout insensitive to amplitude fluctuations of the laser and we are mainly limited by vibrations and acoustic noise. Therefore we isolate the whole setup on a MinusK vibration isolation box and enclose it in a soundproof box. The output of the PI controller is then proportional to a change in cavity length directly corresponding to the sample height. Optionally, we can also feed back on the sample z position to operate in constant force mode and in this case we record the output signal of the z feedback PI loop. We scan the sample (Budget Sensor HS-20MG calibration target) using a 2d piezo flexure stage driven by our FPGA based acquisition system using Fourier limited triangle waveforms with a line scan

rate of 1.2 Hz to not excite mechanical modes of our optical setup. In 7.2 (c) the resulting image with a resolution of 2048×512 pixel and in (d) a 1d linecut is shown.

In order to determine the sensitivity of the cavity AFM, we perform a simple contact mode readout of a ContGD-G cantilever ($f_0 = 13$ kHz, $L = 450$ μm, $k = 0.2N \cdot m^{-1}$).

Several methods of measuring the absolute cavity to mechanical coupling can be used [PLN+21]. We use the method described in [HPS+06] using the fluctuation dissipation theorem to provide an absolute calibration of the readout sensitivity. In order to avoid interference with the laser lock, we chose to perform the sensitivity measurement with a Olympus AC55TS high frequency cantilever ($f = 1.6$ MHz) at room temperature. The thermal excitation spectrum shown in 7.2 (e) is fit with the functional form $S(f) = P_{white} + \frac{P_{dc} f_0^4}{(f^2 - f_0^2)^2 + \frac{f^2 f_0^2}{Q^2}}$ to reveal a Q factor of and thermal noise floor . From this, the inverse optical lever sensitivity InvOLS can be calculated using $InvOLS = \sqrt{\frac{2k_B T}{\pi k f_0 P_{DC} Q}}$. .

We follow the derivation in [PdvG92] to compute the theoretical SNR for a cavity AFM readout. The change in the number of photons at the quadrature point is $\Delta N = \frac{2\pi}{\lambda} N_{tot} t \delta z$ and $SNR = \frac{\Delta N}{\sqrt{tN_{tot}/2}} = \sqrt{8\pi^2 t N_{tot}} \frac{2\delta z}{\lambda}$. [Bla01]

In the hybrid opto-mechanical system, the signal to noise ratio does not monotonically improve with increased readout power, instead the optical back action of the fluctuating intra-cavity field onto the cantilever mirror produces a fluctuating force. As shown in 7.2 (f), there is an optimal power of the readout signal that approaches the thermal noise floor in air, or in vacuum and with cooling can reach the standard quantum limit (SQL) for quantum limited readout. Our cavity AFM method is fully compatible with vacuum operation as only optical access is necessary.

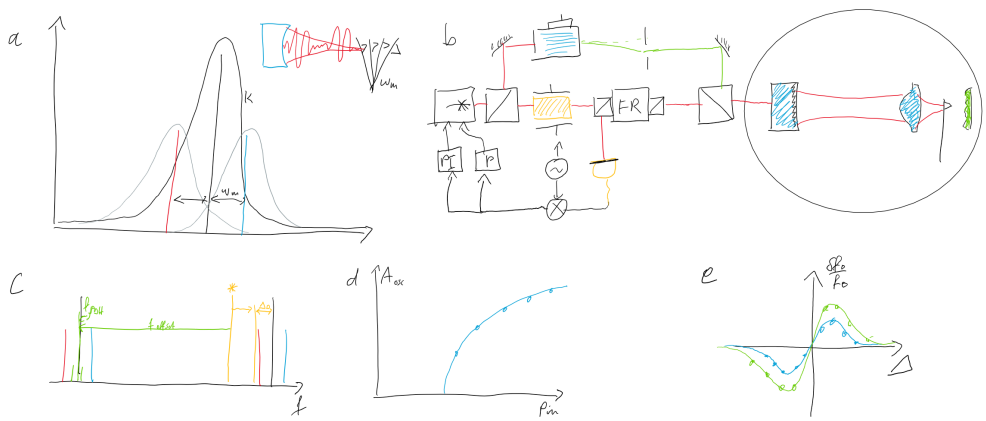


Figure 7.3: FM-AFM. **(a)** Optomechanical coupling between cavity and cantilever with mechanical eigenfrequency ω_m leads to sidebands $\omega_0 \pm \omega_m$ on the cavity resonance ω_0 that can either add (blue sideband) or subtract (red sideband) motional quanta from the mechanical oscillator. **(b)** Schematic optical setup for FM readout: In addition to 7.2 **(b)** an acousto optical modulator (AOM) provides a tunable strong pump tone with polarization orthogonal to the probe field to optically drive mechanical sidebands. **(c)** In order to realize a pump with tunable power and detuning, the AOM is operated at a fixed frequency but with variable RF power. The probe/locking light is detuned by one FSR to avoid interference with the pump light while the offset frequency $f_{offset} = FSR - \Delta$ allows positive and negative detunings Δ of the pump light. **(d)** Ramping up the pump power at a fixed detuning reveals the onset of stable oscillations at an equivalent input power of $1 - 10$ mW whose amplitude grows initially but is then limited by the PDH readout. **(e)** Changing the detuning of the pump leads to the optical spring effect [SGM+04] that can modify the mechanical oscillators properties.

7.4 FM mode

To fully utilize the power of the hybrid cavity cantilever system for AFM operation, we demonstrate how non-equilibrium cavity optomechanics [KV08] can be harnessed to realize an autonomous FM cavity excitation that fully replaces external phase locked loops (PLL) or similar feedback systems. In FM mode, the surface forces are read out in the frequency domain as the change in the mechanical oscillator's eigenfrequency due to the external force. In conventional AFM, its resonance frequency must be tracked externally through a PLL that feeds the measured PSD back to the excitation detector with a phase shift of 90 degrees to stay on resonance. The bandwidth of this feedback loop is limited by the limited PSD bandwidth, delays and mechanical resonances in actuating the piezo and finite loop filter

bandwidth of the PLL limiting the readout speed achievable in FM mode. Different photo acoustic driving methods have been proposed [AL20] that rely on photo thermal [MGR$^+$20] or laser-induced photoacoustic (PA) waves [NKKN09]. In the cavity AFM, this excitation of the cantilever at its eigenfrequency is achieved by continuously driving a mechanical sideband on the optical spectrum of the cavity making it completely intrinsic to the physical system and its bandwidth is only limited by the optical-mechanical coupling strength g_0.

The mechanical oscillator as well as the readout and pump cavity field can be described as three bosonic modes with a Hamiltonian.

$$H = \hbar\omega_c a^\dagger a + \hbar\omega_m b^\dagger b + g(b^\dagger a + a^\dagger b) \tag{7.1}$$

with coupling strength $g = \sqrt{\frac{\omega_c}{L}} x_{ZPF}$. The quantum mechanical treatment provides a simple picture for how a steady state excitation of the cantilever can be created from a continuous drive, but the same results also follow from a classical treatment [KV08] using Doppler shifts. The optomechanical coupling g creates two sidebands $\omega_c \pm \omega_m$ on a cavity resonance ω_c shown in 7.3, where the r̈ed(Stokes) sideband $\omega_c - \omega_m$ can be used to cool the mechanical degree of freedom while the b̈lue(Anti-Stokes) sideband provides excitation at the mechanical eigenfrequency. While in the bad cavity limit $\kappa > \omega_m$ cooling won't reach low temperatures, excitation is not affected by the unresolved sidebands. The full optical setup for FM operation is depicted in 7.3. We derive the excitation tone using an acousto-optic modulator (AOM) for control of the excitation power and multiplex the excitation and readout light onto the cavity using opposite polarizations to prevent interference. We also offset the readout light by one FSR using a fiber EOM to further reduce added noise due to leakage of excitation light from imperfect polarization (see 7.3 (c).

First we demonstrate strong optomechanical coupling, mechanical lasing and the optical spring effect of the cantilever in air and then probe new non-linear dynamics in the presence of a sample surface potential.

The readout light is locked exactly to the cavity frequency ($\Delta_r = 0$) and therefore provides a continuous measurement of the cantilever position. When increasing the pump power at a blue detuning of $\Delta_p = 1\text{MHz}$, the system undergoes a Hopf bifurcation into a stable limit cycle as predicted by the theoretical treatment. In fig. 7.3 (d) the amplitude of the mechanical oscillation derived from the PDH error signal is shown as a function of the driving field optical power P_{in}.

Apart from the possibility to excite or cool collective modes of the system, the presence of the intra-cavity field also modifies the properties of the mechanical oscillator. The oscillator's eigenfrequency and damping change as a function of detuning Δ, which we also demonstrate experimentally in fig. 7.3 (e).

A key advantage of the optomechanical excitation is the spectral purity of the cantilever excitation because the cavity frequency only couples to mechanical modes of the cantilever beam that change the length of the cavity. Opposed to excitation using a shaker piezo, that can couple two mount-, torsional- and higher-order mechanical resonances (see e.g. [RMT08]), the spectrum of the optically excited cantilever shows only the fundamental mechanical frequency and for large oscillation amplitudes, harmonics because of the limited PDH bandwidth.

In conventional AFM the scanning speed is fundamentally limited by the Q factor of the cantilever [YDZJ23] and it is constrained by two competing requirements: a high Q factor provides a high sensitivity to fractional frequency changes of the cantilever in FM mode or equivalently amplitude changes in AM mode which are both attributed to a sharp phase response of the harmonic oscillator around its resonant frequency. Conversely, the small damping responsible for the high quality factor means that the oscillator is slow to respond to external changes in force. This limits the scan speed because at every position a minimum time inverse to the mechanical linewidth γ must be spent to converge to the new external force. Therefore conventional AFM must always trade in sensitivity for speed by choosing a

lower Q factor cantilever. Our cavity enhanced FM-AFM has a sensitivity enhanced by the cavity finesse $\mathcal{F}$ and could therefore operate with cantilevers of much lower Q factor enabling higher scanning speeds. Furthermore, the optical spring effect could be used to dampen the cantilever even when operated in vacuum as the cavity can provide a frictional force.

CHAPTER 8

OUTLOOK

8.1 Summary

Quantum simulation of the FQHE has seen a resurgence in recent experiments [CSB$^+$20, LKK$^+$23, LHR$^+$24, WLC$^+$24, IHK$^+$25] all demonstrating two-particle Laughlin states. In order to resolve quasihole excitations in the form of Anyons [KL09] these experiments need to be scaled to $N \approx 10 - 15$ particles in order to resolve the statistical angle. We have demonstrated a new twisted cavity apparatus chapter 5 hosting up to 20 highly degenerate cavity modes that will be able to host $N = 10$ dark polaritons for realizing the $\nu = 1/2$ Laughlin state.

While the hybrid cavity platform using dark polaritons offers unique advantages in system size and large gap between Landau levels, it suffers from the open nature of the combined atom-cavity system limiting the coherence time. In fig. 8.1 and table 8.1 we compare the key metrics for different cold atom and photonic platforms.

8.2 Laughlin state preparation

8.2.1 Steady state driving

Having demonstrated the capabilities of the new experimental platform with a mesoscopic number of degenerate modes and high collective cooperativity, one of the key questions that remains is what scheme offers the most promising path to prepare Laughlin states of more than two particles. In the previous proof-of-principle demonstration [CSB$^+$20] a two-particle state equivalent to

$$|\psi^{\text{target}}_{Laughlin}\rangle = \frac{1}{\sqrt{2}}\left(|1,0,1\rangle - |0,2,0\rangle\right) \tag{8.1}$$

Platform	$\min(J,U)/h$	τ_{coh}	$\sqrt{\mathcal{C}}$	Δ_{MB}/k_BT	Fidelity	N_{part}	Modes
Rotating Trap [LHR+24]	9.9 kHz	191.0 ms	1891	7.00	96%	2	3
Quantum Gas Microscope [LKK+23]	37 Hz	1.25 s	46	0.38	43%	2	16
Circuit QED [WLC+24]	5.0 MHz	6.2 μs	31	12.00	60%	2	16
Cavity Polaritons [CSB+20]	4.0 MHz	2.7 μs	11	3.70	76%	2	3
Next-Gen Cavity (Proj.)	4.0 MHz	5.3 μs	21	6.00	–	2	16
Floquet Lattices [IHK+25]	71 Hz	26.0 ms	1.8	0.50	–	24	48

Table 8.1: Quantitative comparison of experimental platforms for realizing Fractional Quantum Hall (FQH) states. The coherent collision number is given by $\mathcal{C} = \min(J/h, U/h) \times \tau_{\text{coh}}$, representing how many times particles can interact within the coherence time of the system. The gap over temperature ratio $\Delta_{\text{MB}}/k_BT_{\text{eff}}$ indicates protection against thermal and drive-induced excitations. Values for the next-generation cavity are projected.

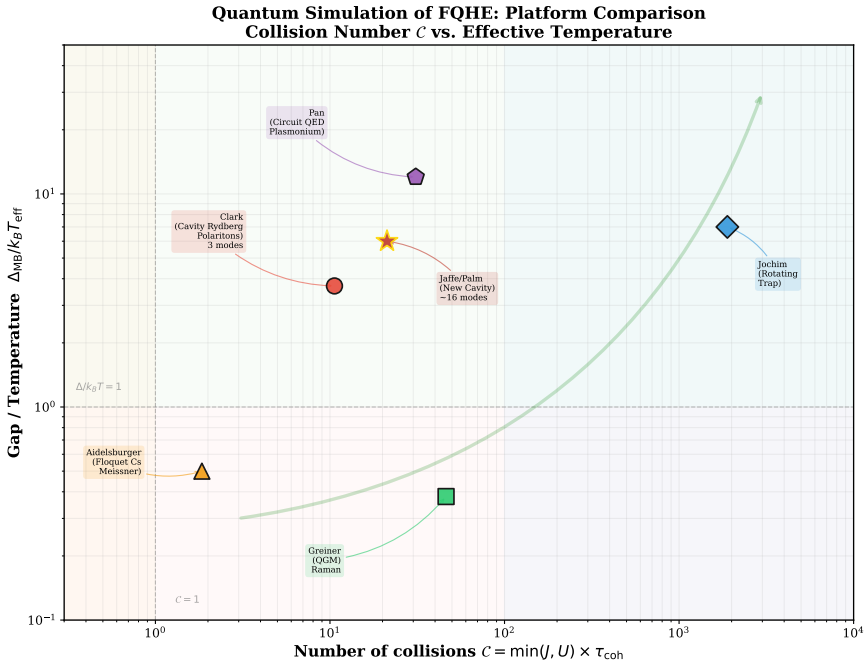


Figure 8.1: Quantum Hall Platform comparison. Comparison of recent experiments realizing FQHE states in terms of key many-body metrics: The size of correlated states such as Laughlin states is limited by the number of collision $\mathcal{C}$ between quantum particles limited by interaction strength U or tunneling rate J depending on the platform and coherence time τ_{coh}. Fidelity of QH ground states such as Laughlin state is limited by the small many-body gap Δ_{MB} and finite temperature limited by heating in Raman or Floquet driven systems.

(if we translate the state on the cone to the flat disk) was effectively prepared written in the excitation number basis $|n_0, n_1, n_2\rangle$. This was facilitated by the fact that $|\psi^{\text{target}}_{Laughlin}\rangle$ exactly minimizes the interaction from the Haldane pseudopotential V_1. Somewhat subtly but almost more importantly, the steady state driving scheme generated a mixed state of different number of excitations N inside the cavity and only by using two SPCM modules and detecting coincidences (or lack thereof, but conditioned on a photon) effectively post-selected the states to the two-particle sector. As the Laughlin state was well separated in energy, the resulting fidelity

$$F_L = \frac{\langle \psi^{\text{target}}_{Laughlin} | \rho_{ss} | \psi^{\text{target}}_{Laughlin} \rangle}{\text{Tr}(P_2 \rho_{ss})} \tag{8.2}$$

was still high with 76(18)%. Going to larger N Laughlin states in the same way will suffer from two limitations: To reach a N-particle Laughlin state $|\psi^{(N)}_{Laughlin}\rangle$, an N-photon tran-

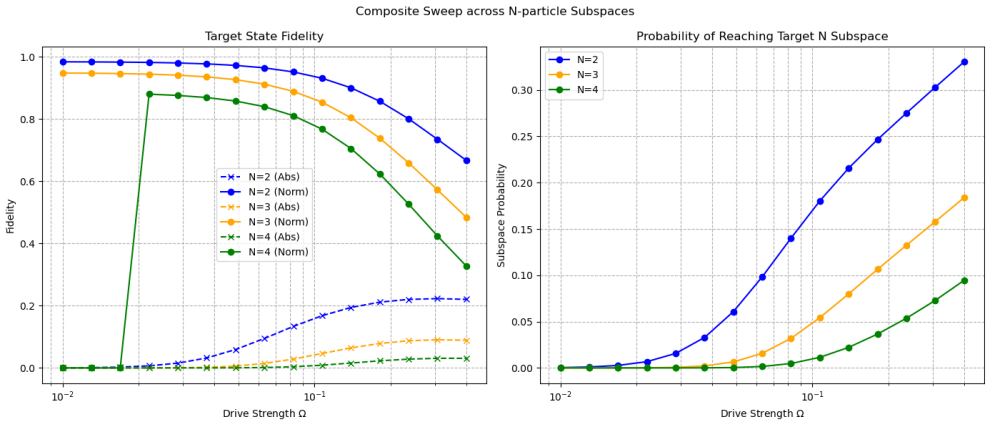


Figure 8.2: Steady state Laughlin state preparation. The fidelity with the N particle Laughlin state $|\psi^{(N)}_{Laughlin}\rangle$ as a function of drive strength Ω with a single driving mode $l_{drive} = N-1$ diminishes **(a)**, while the rates for the N-photon process increase exponentially with Ω **(b)**. In order to observe a target state $|\psi^{(N)}\rangle$ an appropriate driving strength interpolating between fidelity and N-particle subspace population must be chosen **(c)**.

sition is required creating a mixed state involving all intermediate number of particles with decaying N-subspace populations. Furthermore, the many-body gap Δ_{mb} gets exponentially smaller with system size, causing more unwanted excitations of e.g. Quasihole states in the higher N manifolds. In fig. 8.2, these two competing requirements are shown by calculating both the fidelity according to $eq : LL_fidelity$ and the population for the $\mathrm{Tr}(P_N\rho_{ss})$ subspace. These numerical results are derived using a full Master equation simulation in an excitation number restricted basis considering effective bosonic Rydberg polariton modes interacting via contact interactions and including polariton decay through a Liouvillian jump operator $\mathcal{L}[\sqrt{\ }](\gamma_r)$. Here, we consider the simplest driving protocol on a single mode l_{drive} for comparison with [CSB+20] where we have to consider which l_{pump} is appropriate for $|\psi^{(N)}_{Laughlin}\rangle$. The total angular momentum is $L_{\text{target}} = p\frac{N(N-1)}{2}$ for the Laughlin state of inverse filling fraction p (for the $\nu = 1/2$ state we consider $p = 2$). For N pump photons in the driving mode l_{drive} the total angular momentum is $L_{\text{pump}} = N \times l_{\text{drive}}$ and by as the

interactions causing scattering between l modes have to conserve total L they have to match

$$N \times l_{\text{drive}} = p\frac{N(N-1)}{2} \implies l_{\text{drive}} = \frac{p(N-1)}{2}. \tag{8.3}$$

As the highest occupied orbital is $l_{max} = p(N-1)$ we always have to drive on half this maximal angular momentum $l_{\text{drive}} = \frac{l_{max}}{2}$. In our case ($p = 2$):

- $N = 2$: $l_{max} = 2$. Modes are $[0, 1, 2]$. The exact central mode is $l_{drive} = 1$.
- $N = 3$: $l_{max} = 4$. Modes are $[0, 1, 2, 3, 4]$. The exact central mode is $l_{drive} = 2$.
- $N = 4$: $l_{max} = 6$. Modes are $[0, 1, 2, 3, 4, 5, 6]$. The exact central mode is $l_{drive} = 3$.

The key tradeoff is apparent in fig. 8.2 (c): In order to realize a $N = 4$ Laughlin state, we generate a much larger background of $N = 1, 2, 3$ states that are also increasingly impure Laughlin states in their respective N subspace. This will become a big challenge in detection as post selecting on N-fold coincidences will lead to exponentially diminishing data rates especially in the presence of additional cavity loss and detector efficiencies. While not as definitive a signature as correlations, an alternative, much simpler observable could be imaging the steady state density of the cavity output on our EMCCD camera where Laughlin states form a characteristic disk shape [Gem07]. However also imaging the density suffers from the same mixed state problem and the clear disk shape will be obscured by the exponentially stronger signal from lower N sectors and will require careful subtraction and image decomposition. A way to potentially improve the fidelity with the Laughlin state is to come up with smarter driving terms that also utilizes the flexibility in input field afforded by our DMD. An alternative way to arrive at the Laughlin wavefunction apart from the usual first quantization form is given in the Mathematical framework of Jack polynomials [Jac22]. While a full derivation is mostly of mathematical interest, it allows for a systematic construction of all terms $\{n_l\}$ in the second quantized waveform by applying a series of

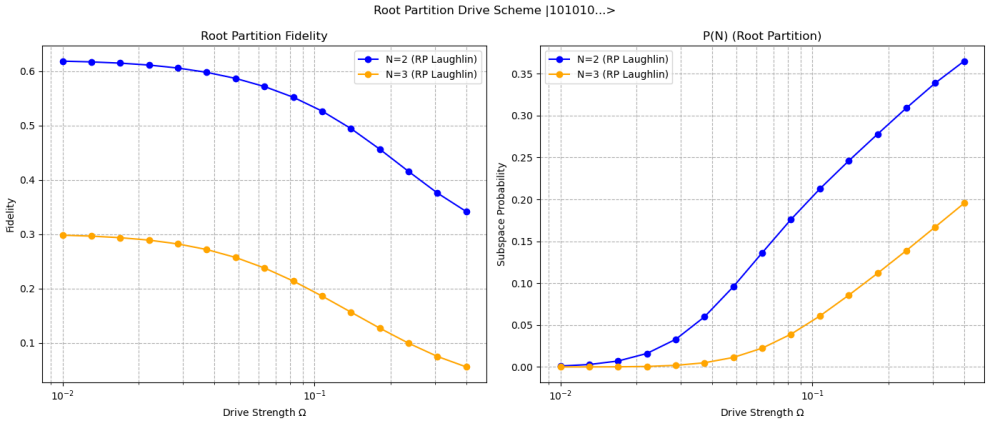


Figure 8.3: Root partition Laughlin state preparation. The fidelity with the N particle Laughlin state $|\psi^{(N)}_{Laughlin}\rangle$ as a function of drive strength Ω with a single driving mode $l_{drive} = N-1$ diminishes **(a)**, while the rates for the N-photon process increase exponentially with Ω **(b)**. In order to observe a target state $|\psi^{(N)}\rangle$ an appropriate driving strength interpolating between fidelity and N-particle subspace population must be chosen **(c)**.

squeezing rules to a so-called root partition. For the $p = 2$ bosonic Laughlin state this is $|\psi_{root}\rangle = |1, 0, 1, 0, 1, 0, \dots\rangle$ i.e. a state with every other site occupied. We can immediately see that this state has filling fraction $\nu = 1/2$ and all other terms of the wavefunction can be generated from it by the squeezing rules that correspond exactly to the physical scattering between modes V_{ijkm} from the lowest Haldane pseudo potential i.e. contact interactions. By driving using this root partition, we can potentially excite to the N particle Laughlin state $|\psi^{(N)}_{Laughlin}\rangle$ much faster and with higher fidelity. In fig. 8.3 the same protocol as fig. 8.2 is considered but now displaying $|\psi_{drive}\rangle = |1, 0, 1, 0, 1, 0, \dots\rangle$ every other l mode on the DMD. Furthermore, the root partition can be written down also for the Quasi-hole states

$$\Psi^{(QH)}(z_1, \dots, z_N; \eta) \propto \prod_{1 \leq i < j \leq N} (z_i - \eta)(z_i - z_j)^2 \exp\left(-\sum_i \frac{|z_i|^2}{4\ell_B^2}\right) \tag{8.4}$$

that are generated by the root partition $|\psi^{(QH)}_{root}\rangle = |0, 1, 0, 1, 0, 1, 0, \dots\rangle$ where the first mode $l = 0$ is unoccupied.

8.3 Dissipative state preparation

As previously discussed, preparing large entangled states like the Laughlin state in an open system using Rydberg polaritons faces unique challenges compared to e.g. quantum gas microscopes, however the open nature also opens up new opportunities. Particles in the system are lost when the photonic component of the Rydberg polariton leaks out of the cavity, but new polaritons can be injected readily using an external coherent drive. This is in stark contrast to optical lattice experiments where new atoms can only be prepared by undergoing a full cooling and evaporation cycle once they are lost from the system. In order to stabilize true quantum many-body states and not just random mixed states however, our drive has to be highly selective both in frequency as well as particle number. The theoretical feasibility to stabilize Laughlin states in driven-dissipative systems has been demonstrated in [UC17, USC21] using sophisticated Master equation calculations. Quantum Hall states such as the Laughlin state are incompressible, and therefore if we can realize a hard wall potential in real space can create an energy gap in particle number ΔE_N. To fill up the system using polaritons we need a non-reversible drive as a simple coherent drive would simply drive Rabi oscillations. In this section, I will lay out a path to actually realize such a non-reversible or *parametric* drive in our hybrid system by leveraging a second species ^{8}5Rb coupled to the same science cavity to form a superradiant drive.

The collective emission of light by an ensemble of emitters, known as superradiance, represents a fundamental concept in quantum optics, explored both theoretically [BL78, KRKDT19] and experimentally in various systems, particularly cavity QED setups demonstrating collective phenomena like self-organization and dynamical phase transitions [MDL$^+$19, ZDL$^+$19, LDZ$^+$21, MBL$^+$20, DBL$^+$22]. Steady-state superradiance, often achieved using atoms with narrow optical transitions coupled to a cavity mode, can lead to laser-like emission with extremely narrow linewidths, offering potential for metrology and quantum technologies [HKF$^+$93, MH10, BCW$^+$12, NLC$^+$18, TXM$^+$17, KK18, SFR$^+$23, GXF$^+$16]. Mod-

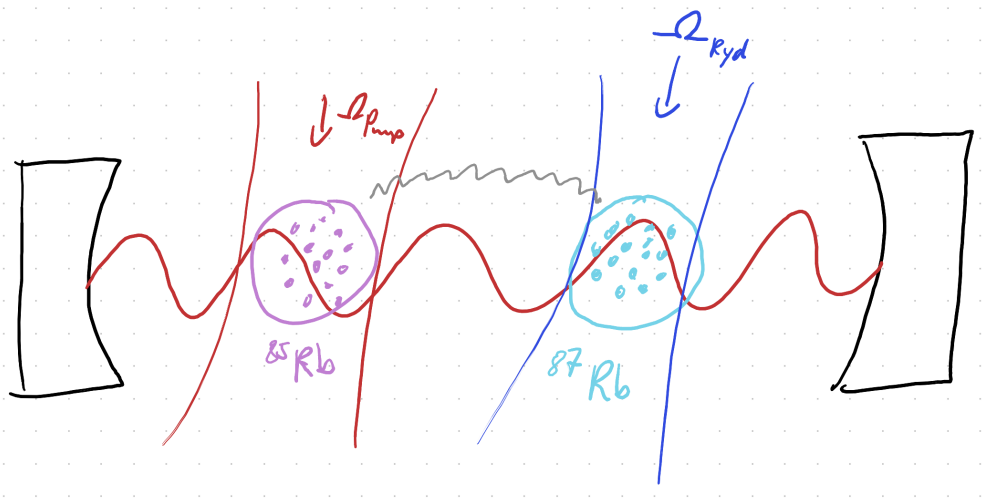
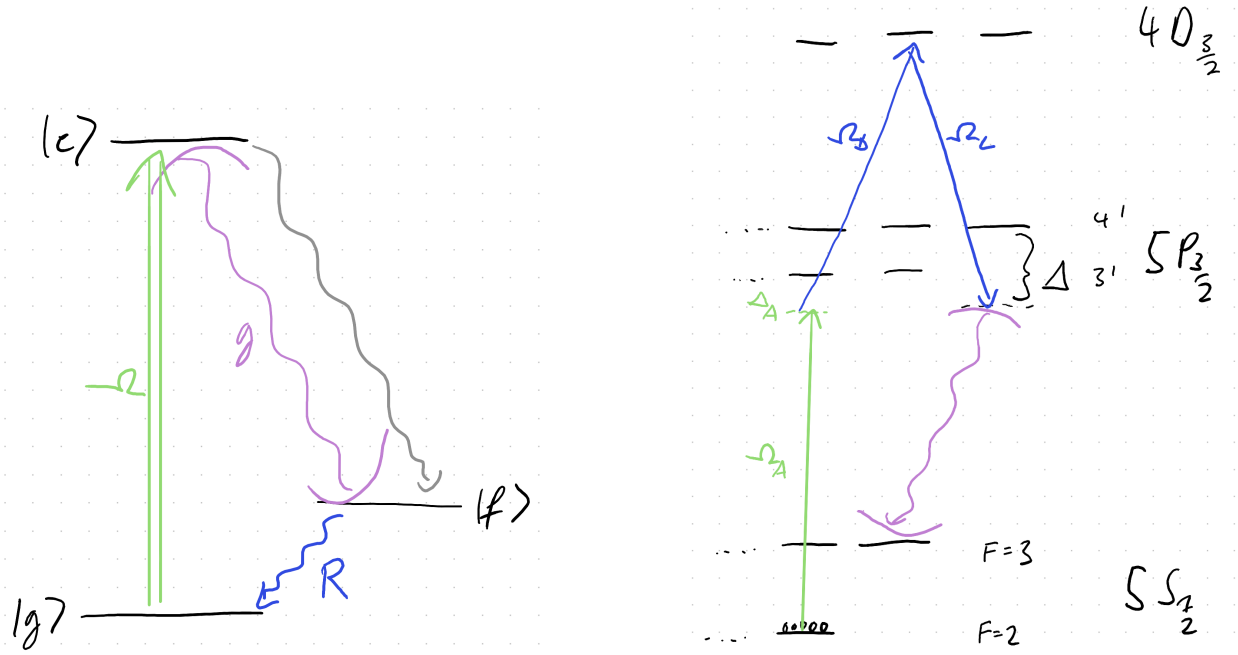


Figure 8.4: Design of proposed scheme. In order to enable dissipative state engineering for cavity QED systems, a frequency selective gain/non-reversible drive has to be included inside of the cavity. Our proposed scheme for the dissipative preparation of Cavity Rydberg Polaritons consists of one ensemble of atoms (here ^{87}Rb) in an EIT configuration where the ground to excited state transition is resonant with an optical cavity and the second transition up to a high-lying Rydberg state is driven by a classical control field Ω_{ryd}. This system hosts long-lived *dark polaritons*, hybrid quantum particles that inherit both strong interactions as well as trapping or artificial gauge fields from the photonic degree of freedom of the cavity. We propose to use a second atomic ensemble (here ^{85}Rb) to realize a frequency selective/non-reversible drive of photons into the cavity that in turn creates new dark polaritons. The two ensembles share the same cavity field acting as a photon "bus".

eling the complex dynamics of driven-dissipative many-body systems, especially those exhibiting permutation symmetry, often requires sophisticated theoretical tools beyond simple mean-field approximations [Dog19]. Techniques like quantum cumulant expansions provide a systematic way to truncate the hierarchy of operator expectation values, allowing for the inclusion of higher-order correlations crucial for describing quantum fluctuation dynamics [SSF20, PHR22, KK18, SFR$^+$23]. Other approaches, such as diagrammatic methods, can be tailored for specific interaction regimes like Rydberg blockade [GGB18]. Combining the strong, long-range interactions characteristic of Rydberg atoms with the enhanced light-matter coupling provided by optical cavities gives rise to hybrid quasi-particles known as Rydberg polaritons [GXF$^+$16, GGB18]. These polaritons inherit strong optical nonlin-

earities from their Rydberg component, enabling the exploration and creation of strongly correlated states of light, such as photonic Laughlin states [CSB+20]. Engineered dissipation offers a powerful paradigm for quantum state engineering, providing a route to autonomously stabilize complex quantum many-body states against inherent system losses and decoherence [GFP+20, USC21]. This approach, pioneered in circuit QED systems [MSO+19], has been successfully applied to realize phases like photonic Mott insulators and explore phenomena such as limit cycles and dynamical phase transitions in various driven-dissipative platforms [MSO+19, CSB+20, SKK+24, MBL+20, DBL+22, Dog19, NLC+18, KRKDT19].



(a) Single atom laser level scheme. **(b)** Realistic dual species driving scheme.

Figure 8.5: Non-reversible driving schemes. **(a)** Single atom laser level scheme. **(b)** Realistic dual species driving scheme.

For hybrid, cavity QED based quantum simulators, that are uniquely suited for realization of FQHE states, traditional adiabatic preparation routines might not be optimal due to the limited lifetimes of the underlying quasi-particles and the possibility to dynamically grow increasingly correlated states due to the open nature of the system. In [USC21] the authors show that the Laughlin state can be prepared in a dissipative way due to its incompressible

nature that, together with an edge in the system, causes a gap both to higher excited and higher particle number states. The key ingredient for this is a non-reversible drive that can inject new particles into the system to initially fill the system or replace decaying ones. A classical drive with an external coherent state would simply drive Rabi oscillations of the dark polariton and every mode would be at most populated with 0.5 particles on average. Instead, single photons, that then form dark polaritons, must be generated inside of the cavity while the time-reversed process must be forbidden. This is incompatible with unitary evolution and we therefore need to make use of atomic decay/dissipation. In [MSO+19] using circuit QED, the authors used a two-photon transition in a special site with a large difference in decay rate between the $|2\rangle \to |1\rangle$ and $|1\rangle \to |0\rangle$ transition. The natural transition lifetime in neutral atoms can't be modified in this way and usually are ordered in an opposite hierarchy with higher excited states longer lived, however we can make use of the rich internal level structure to engineer a similar mechanism.

8.3.1 Flux threading and dynamical preparation

A different appealing route of assembling a Laughlin state of increasing size in the open cavity system would be to use flux threading. A single polariton can readily be inserted on e.g. the $l = 0$ mode using a π-pulse in the presence of photon blockade. This polariton can then be *stirred up* by transferring it to a higher orbital angular momentum mode using a Raman process (as demonstrated in [BJP+23] for our setup), equivalent to threading a single quantum of flux through the plane. Subsequently, a second polariton can be inserted at the origin again using a second π-pulse that will interact with the one already present in the system. When the polariton insertion is slow enough such that the drive linewidth stays below the many-body gap, it can be used to address the Laughlin state spectroscopically as it is always the state minimizing interaction energy. This scheme is described in [ILSF18] in great detail, however the threading operation suffers from a major challenge in practice. As

all cavity modes are degenerate, the STIRAP step would immediately transfer population from the initial $l = 0$ mode to the highest mode l_{max} as all the couplings would be identical. This can be prevented by inserting a different *internal* intermediate state into the flux threading step, however this intermediate state must also support the same Laughlin state itself which rules out using e.g. the opposite helicity state of the cavity as it is too far separated in frequency space. One possibility would be to use EIT using the lower hyperfine level $F = 1$ if the cavity FSR can be made commensurate with the HFS splitting but this could lead to other unforeseen effects for the regular EIT.

A second route would be an adiabatic ramp protocol akin to cold atom experiments [LKK$^+$23, PKB$^+$24, PGP20, LHR$^+$24]. Here, a product state of N particles is initialized and the Hamiltonian is subsequently transformed by ramping up tunneling or trap rotation and ellipticity until a number of small many-body gaps have been traversed to reach the target ground state. In our twisted cavity experiment, the amount of angular momentum is fixed per cavity mode due to static geometry, and the equivalent way of preparation would be to prepare N polaritons in a staggered configuration like $|\psi_{root}^{(QH)}\rangle = |0, 1, 0, 1, 0, 1, 0, \dots\rangle$ in the absence of interactions, and then adiabatically ramping up the interactions strength while staying in the many-body ground state. Our platform is not particularly well suited for adiabatic preparation routines due to limited coherence time of the dark polaritons and the challenge to control the system Hamiltonian on the very fast interaction timescales of under 1 μs. The residual photonic trap frequency can only be adjusted on slow, mechanically limited timescales and controlling the Rydberg interactions would require the use of a Förster zero to reduce them sufficiently. Also while our DMD has a high frame rate of up to 10 kHz, this is still too slow for the μs timescales set by the Rydberg interactions and multiple DMDs gated with AOMs would be necessary to dynamically pulse in different l modes.

8.3.2 Non reversible driving scheme

The simplest non-reversible drive can be devised with three internal levels of an atom as shown in fig. 8.5a which resembles the simplest level scheme needed for inversion in a laser. A classical pump beam with Rabi frequency Ω excites the atom from its ground state $|g\rangle$ to an excited state $|e\rangle$ from where it can spontaneously decay. The transition $|e\rangle \rightarrow |f\rangle$ to the metastable state $|f\rangle$ is exactly resonant with the cavity and due to strong atom-cavity coupling g the emission is Purcell enhanced into the cavity by a factor $\frac{\mathcal{C}}{\mathcal{C}+1}$ while emission into free space only happens with a small probability $\frac{1}{\mathcal{C}+1}$. From $|f\rangle$, the atom is rapidly repumped to the ground state $|g\rangle$ via an intermediate excited state (not shown) to complete the cycle. The key feature here is that the repumping happens incoherently, as the atom emits another photon into free space. This separation of timescales between the coherent process $|g\rangle \xrightarrow{\Omega} |e\rangle \xrightarrow{g} |f\rangle$ and much faster repump $|f\rangle \rightsquigarrow |g\rangle$ returns the atom to its initial state before it can re-absorb the photon from the cavity and thereby breaks the time-reversal symmetry. One key challenge to implement this in the experiment is a large enough separation in frequency space to implement this drive while not disturbing the sensitive dark polaritons forming the many-body state. Furthermore, the polaritons are collective excitations of all atoms in the ^{8}7Rb ensemble and the superposition between cavity and Rydberg state would be destroyed if one tried to implement the drive using the same ensemble of atoms. In theory, it would be possible to place a second cloud of atoms into the cavity, for example in the upper waist of the cavity, but due to the large mode waist there it will be hard to achieve a high cooperativity and trying to address a second ensemble in real space requires extreme care to not have any crosstalk from the excitation beams. Instead, we can make use of a previously hidden resource of our apparatus and leveraging a different nuclear isotope ^{8}5Rb of Rubidium that is naturally more abundant and also present in our atomic dispensers. These two *species* have very desirable properties, as their transition wavelengths are close enough ~ 3 GHz to both be trapped by the same off-resonant 784 nm

lattices and traps, but separated far enough that they can be separately addressed using resonant light with high fidelity due to minimal off-resonant scattering. Now both ensembles can occupy the same space and we can keep our simple transport lattice but the transverse pump beam Ω can be applied only to the ^{8}5Rb which emits light into the shared cavity as shown in fig. 8.4. While the isotope shift allows us to address both species separately, we then need to bridge the difference in transition frequency by making the cavity resonant with the transition, which will make it too far detuned from the excited state and cause excessive off-resonant scattering. Instead, we envision a four-wave mixing scheme fig. 8.5b involving the $|^{85}\mathrm{Rb}, 4D_{3/2}\rangle$ state that can be reached using a convenient 1529 nm transition. This will use an existing laser previously used in [BJP$^+$23] and allow a frequency conversion of single photons.

8.4 Many-mode cavity QED

The cavity array microscope could have much broader uses beyond just providing strong coupling to single atoms for readout and addressing. Every single atom coupled to a local cavity forms a cavity QED system in itself with a potentially strong Jaynes-Cummings non-linearity. If the different cavity sites can be coupled to each other, either through an optically designed coupler using AOMs or custom mirrors in the cavity, or using the atoms themselves as a Rydberg many-body simulator, a multi-mode cavity QED platform could be realized. While the photonic excitations in this system are bosons in nature, rich phase diagrams have been theorized [MCT$^+$08, GTCH06, IMV$^+$20, PdT$^+$22].

The periodic cavity field in the atom plane of a tightly focused gaussian beam per site will lead to the Talbot effect inside the cavity. In fig. 8.6, the CAM intra-cavity field shows a revival of the periodic, localized field of the original atom plane every Z_T, and an intermediate, shifted revival every $(n + 1/2)Z_T$. Atoms placed on the different planes can either provide: on-site interactions (left column), nearest neighbour tunneling (center column) or

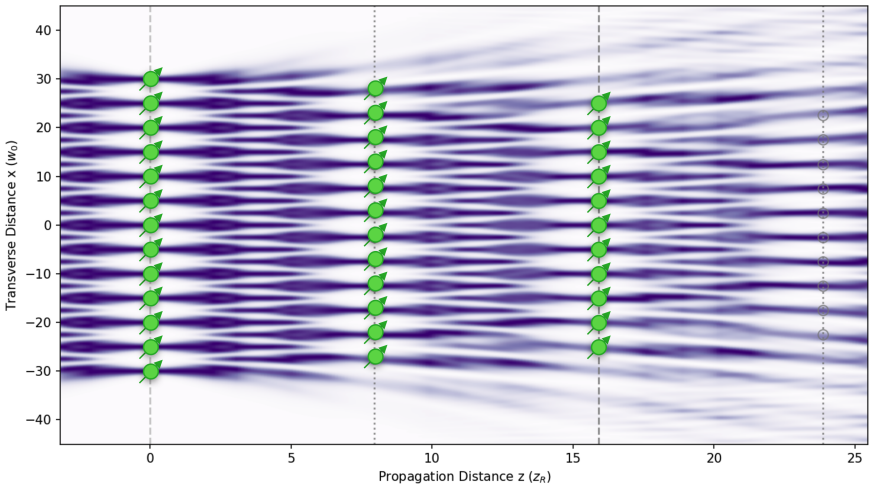


Figure 8.6: Talbot effect of the CAM intra-cavity field shows a revival of the periodic, localized field of the original atom plane every Z_T, and an intermediate, shifted revival every $(n + 1/2)Z_T$. Atoms placed on the different planes can either provide: on-site interactions (left column), nearest neighbour tunneling (center column) or next-nearest neighbour tunneling (right column).

next-nearest neighbour tunneling (right column). As controlling of single atoms is an established technique nowadays, the atoms in the different planes could be selectively populated using optical tweezers and using local stark shifting beams their coupling to the cavity could be tuned between resonant (for interactions) and dispersive (to mediate tunnel couplings).

The characteristic length scale for this effect is the Talbot Length (Z_T)

$$Z_T = \frac{2d^2}{\lambda} = \frac{2(25\,\mu\text{m}^2)}{0.78\,\mu\text{m}} \approx \mathbf{64.1}\,\mu\text{m} \tag{8.5}$$

that depends on array pitch in the atom plane d (assumed to be 5μm here) and wavelength λ. To estimate the number of revivals this has to be compared to the Rayleigh range z_R:

$$z_R = \frac{\pi w_0^2}{\lambda} = \frac{\pi(2.25\,\mu\text{m}^2)}{0.78\,\mu\text{m}} \approx \mathbf{9.1}\,\mu\text{m} \tag{8.6}$$

where a cavity waist of $w_0 = 1.5\mu$m has been assumed. Comparing the Talbot length to the Rayleigh range reveals the diffractive nature of the revivals. Note that λ cancels out in this

ratio:

$$\frac{Z_T}{z_R} = \frac{2d^2/\lambda}{\pi w_0^2/\lambda} = \frac{2}{\pi}\left(\frac{d}{w_0}\right)^2 \tag{8.7}$$

Substituting the geometric fill factor ($d/w_0 \approx 3.33$) yields $\frac{Z_T}{z_R} \approx 0.6366 \times (3.33)^2 \approx \mathbf{7.1}$.

Since $Z_T \approx 7z_R$, the individual spots diffract significantly between revivals, confirming the "breathing mode" behavior where spots blur into a continuum before re-phasing. Due to diffraction the cavity field will still drop off with distance from the central atoms plane limiting the number of planes with useful cooperativity. The critical order M_{crit} where intensity drops significantly is independent of wavelength:

$$M_{crit} \approx \frac{\pi}{4} N \left(\frac{w_0}{d}\right) \approx \mathbf{47} \text{ orders} \tag{8.8}$$

The total useful propagation distance is $Z_{max} \approx M_{crit} \cdot Z_T \approx 47 \times 64.1\,\mu\text{m} \approx \mathbf{3.0}\,\text{mm}$ which is a lot of space to place atoms into talbot planes. This could also provide an avenue to trap 3D arrays of atoms in a single CAM when different layers can be loaded by addressing them using the sideband technique from [SKS+25].

8.5 Dynamical potentials using SLM

Even though the fractional quantum Hall effect should be robust to the amount of disorder from lens imperfections in our system, it would still be desirable to further reduce this disorder or instead realize arbitrary potentials for the photonic fluid. This could include localized defects to study the robustness of edge modes in the system. One straightforward way of realizing such local potentials would be to incorporate a spatial light modulator into the cavity; however, off-the-shelf SLMs are limited in both the fill factor, mirror reflectivity, and loss from diffraction from the pixel grid, that are on the order of $95 - 99\%$ which is prohibitive for high finesse cavities.

At the same time, SLMs usually provide a phase shift of 2π, which would be equivalent to

shifting the frequency of the cavity by a whole FSR, which is more than the 20 to 50 MHz shift we require. Therefore, we propose to employ coupled cavities [LWS+10, XZZ19, JZG+23] to trade in frequency, lower frequency actuation for less loss.

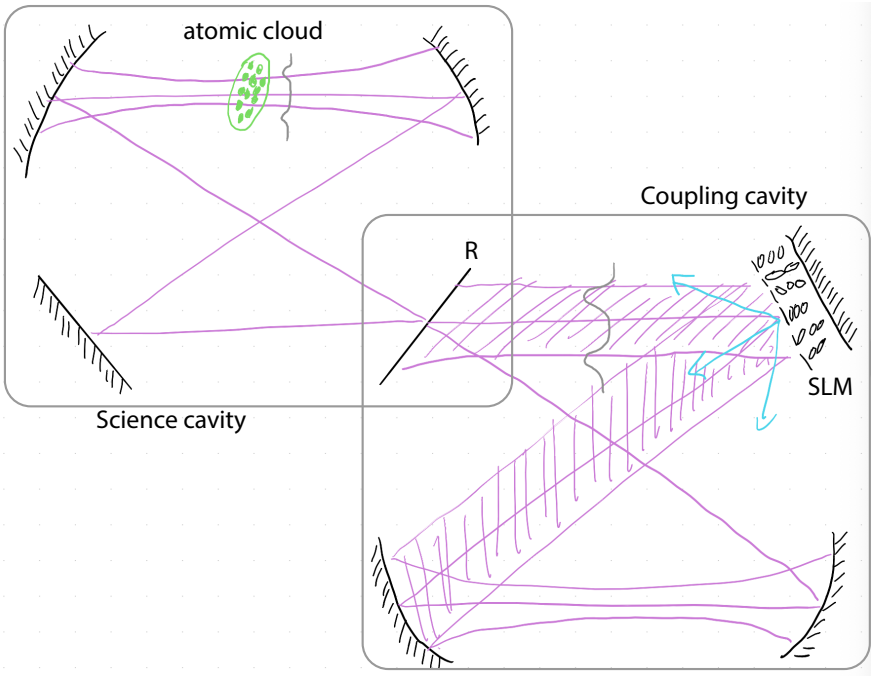


Figure 8.7: Coupled cavities for creating local potentials and compensating disorder using an SLM. Two identical twisted cavities are coupled through a partially transmissive mirror R where the lower coupling cavity contains an SLM in place of a flat mirror. The coupled cavity scheme allows to trade loss due to the limited SLM reflectivity for a reduced frequency shift.

In fig. 8.7 the proposed Coupled cavities for creating local potentials and compensating disorder using an SLM are shown. Two identical twisted cavities are coupled through a partially transmissive mirror R where the lower coupling cavity contains an SLM in place of a flat mirror. The coupled cavity scheme allows to trade loss due to the limited SLM reflectivity for a reduced frequency shift. Assuming otherwise ideal mirrors, the whole effect of the lower coupling cavity can be condensed into an effective reflectivity

$$r_{eff}(\omega) = \frac{r - r_2 e^{i\phi_2}}{1 - r r_2 e^{i\phi_2}} \tag{8.9}$$

where r_2 is the amplitude reflectivity of the end mirror of C_2, and $\phi_2 = 2\omega L_2/c$ is the

round-trip phase in C_2. The global resonant modes of the coupled system must satisfy the round-trip condition for C_1: $r_1 r_{eff}(\omega) e^{i\phi_1} = 1$. In the lossless limit ($r_2 = 1$), the resonance condition algebraically reduces to a transcendental equation governing the coupled eigenfrequencies

$$\cos\left(\frac{\phi_1 + \phi_2}{2}\right) = r\cos\left(\frac{\phi_1 - \phi_2}{2}\right). \tag{8.10}$$

To observe the crossover from strong hybridization to a dispersive shift, we expand this condition around the bare cavity resonances ω_1 and ω_2. Assuming highly reflective mirrors ($r \approx 1 - |t^*|^2/2$), the exact condition maps to the eigenvalue equation of a coupled two-level system $(\omega - \omega_1)(\omega - \omega_2) = J^2$. where $J = c|t^*|/2L$ is the inter-cavity coupling rate (assuming $L_1 \approx L_2 = L$). Solving for the global eigenfrequencies yields:

$$\omega_\pm = \frac{\omega_1 + \omega_2}{2} \pm \sqrt{\left(\frac{\Delta}{2}\right)^2 + J^2} \tag{8.11}$$

where $\Delta = \omega_1 - \omega_2$ is the detuning. In the dispersive limit, C_2 acts not as a coupled mode, but as a boundary potential. The effective phase shift $\theta_{eff} = \arg(r_{eff})$ modifies the boundary condition of C_1, shifting the resonant frequency by $\delta\omega = \frac{c|t^*|^2}{4L_1}\cot\left(\frac{\omega_0 L_2}{c}\right)$. To realize programmable, real-space potentials $V(x, y)$, we replace the end mirror of C_2 with an SLM (e.g., LCOS). However, SLMs suffer from intrinsic losses and a limited filling fraction ($\sim 95\%$).

Direct intra-cavity placement of an SLM typically destroys cavity finesse. By embedding the SLM inside the secondary, detuned cavity C_2, we leverage the middle mirror to demagnify both the imparted phase and the associated loss.

The effective phase θ_{eff} scales with the middle mirror transmission $|t^*|^2$. By choosing $|t^*|^2 \approx 0.1$, a full 2π phase shift at the SLM is scaled down to a small fraction of the C_1 Free Spectral Range (FSR). Consequently, the substantial insertion loss of the SLM is proportionately reduced, allowing C_1 to maintain a high finesse while still being subjected

to the potential $V_{eff}(x, y)$.

A critical challenge of SLM integration is the non-active pixel grid (filling fraction), which acts as a static, lossy diffraction grating. We mitigate this not via spatial filtering, but via spectroscopic suppression utilizing the coupled-cavity dynamics.

The SLM is backed by a high-reflectivity dielectric mirror. The total cavity length L_2 is strictly tuned such that the unmodulated light reflecting off the "dead zones" is perfectly anti-resonant with C_2 $\Phi_{unmod} = \frac{2\omega_0 L_2}{c} + \phi_0 = \pi \pmod{2\pi}$. At anti-resonance, $r_{eff}(\pi)$ is purely real and positive. The intra-cavity field in C_1 undergoes destructive interference upon attempting to enter C_2 at these spatial coordinates. Consequently, the field does not penetrate the dead zones, extinguishing scattering and absorption losses associated with the pixel grid. Also, the dead zones impart zero dispersive shift ($V_{eff} = 0$). Conversely, a global digital bias phase ϕ_{bias} is applied to all active LCOS pixels, pulling them away from anti-resonance and into the desired dispersive regime:

$$\Phi_{mod}(x, y) = \pi + \phi_{bias} + \delta\phi(x, y) \tag{8.12}$$

As a result, the cavity field only reaches C_2 at the precise spatial coordinates of the active pixels. The photon in C_1 is entirely blind to the microscopic pixel grid, experiencing only the smoothed, continuous, and highly efficient potential landscape $V_{eff}(x, y)$ dictated by the active SLM area. Even with this spectroscopic protection, the finesse will still ultimately be limited by the interinsic loss of the LCOS material. As this is not specified by any manufacturer (and presumably not even known) it remains to be seen when finesse can be achieved with this scheme experimantally.

8.6 Experiment Upgrades

As detailed in appendix A.1, presence of the strong 480 nm light in the blue buildup cavity required for Rydberg excitations causes two major problems: Through the photoelectric effect on Rubidium thin films coating the chamber and cavity structure, it causes patch charges exacerbating the challenge of stray electric fields significantly. Even more severely, it causes laser induced contamination appendix A.1 on the 480 nm Rydberg buildup cavity itself, significantly limiting the available Rabi frequency and duty cycle of the experiment, significantly impeding data rate and unsupervised operation of the machine.

As the high energy of the blue photons of approximately 2.58 eV is at the heart of both of these problems, one potential route for future improvement would be to switch to a three photon excitation scheme [BBC$^+$24, BYS$^+$24, TMA$^+$19, RBT$^+$11, BRV$^+$] for the Rydberg EIT. While we want to keep the D2 transition at 780 nm resonant with the twisted cavity as the lower leg in order to maximize cooperativity, the transition to the Rydberg state would be facilitated by two infrared photons and an intermediate atomic state. A true dark-state, such as the dark polariton which has no population in the lossy, intermediate atomic state, can only be achieved in three-level EIT. In order to minimize additional loss from intermediate states and maintain the long lifetime of the dark polariton in three-photon EIT, the last two legs need to be very far detuned from the additional intermediate state. This large detuning would ultimately limit the achievable Rabi frequency Ω to the Rydberg state when compared to the collective cavity coupling g. Therefore, the exact path and intermediate states need to be chosen carefully to trade off between Rabi frequency and decoherence. Three different pathways can be identified from the literature: The path $5P_{3/2} \rightarrow 5D_{5/2} \rightarrow nF_{7/2}$ uses convenient laser wavelengths at 776 nm and 1260 nm that are available as diode lasers and TAs. Similarly, the $5P_{3/2} \rightarrow 4D_{5/2} \rightarrow nF_{7/2}$ path using 1529 nm and 699 nm are available with high power using EDFA fiber and diode lasers. However, the F Rydberg states have a strong angular dependence of the interaction strength that could be disadvantageous in

the specific 2D geometry in the experiment. Furthermore, the F states have no quantum defect, making them more sensitive to external electric field. Transitions to the P Rydberg states are possible but their matrix elements are smaller due to the quantum defect. The alternative $5P_{3/2} \rightarrow 6S_{1/2} \rightarrow nP_{3/2}$ pathway also uses convenient wavelengths at 1367 nm and 743 nm and offers several interesting Förster resonance candidates.

To suppress decoherence from intermediate states, the intermediate optical fields are intentionally detuned from exact one-photon resonance. Let Δ_1 be the detuning from state $|1\rangle$, and Δ_2 be the total detuning from state $|2\rangle$. When these intermediate detunings are chosen to be much larger than both the individual single-photon Rabi frequencies ($\Delta_1, \Delta_2 \gg g, \Omega_2, \Omega_3$) and the natural linewidths of the excited states ($\Delta \gg \Gamma$), the system enters a dispersive regime.

$$\Omega_{eff} = \frac{g\Omega_2\Omega_3}{4\Delta_1\Delta_2} \tag{8.13}$$

As the Rabi frequencies Ω_i need to be strictly smaller than the detunings Δ_i to justify the adiabatic elimination, the resulting effective Rabi frequency Ω_{eff} could become prohibitively small (on the order of ~ 40 kHz for realistic parameters). This could be improved by reducing the first leg detuning Δ_1 to close to resonance and relying on the same dark state as in the two-photon case as long as the second intermediate state stays far detuned.

A distinct and potentially promising regime occurs when the second transition is driven extremely strongly. This is a scenario easily achieved utilizing the 1367 nm pathway, where the immense dipole moment allows for $\Omega_2 \gg g, \Omega_3, \Delta_1, \Delta_2$. In this limit, the intense intermediate coupling drastically dresses the intermediate atomic states, creating a pronounced Autler-Townes doublet separated by a frequency splitting equivalent to Ω_2. When the first and third lasers are tuned to exact resonance with the bare, undressed states, the massive dressing field created by the intermediate laser actively pushes the intermediate probability amplitudes to near zero. Under these specific conditions of strong intermediate coupling (Ω_2) and moderate outer coupling (g, Ω_3), the standard adiabatic elimination formula is modified,

and the effective three-photon Rabi frequency converges [BBC+24] to

$$\Omega_{eff} = \frac{\Omega_1 \Omega_3}{\Omega_2}. \tag{8.14}$$

This could improve the effective Rabi frequency up to $\sim$ 400 kHz and has the additional advantage of eliminating spatially dependent Stark shift from finite beam sizes.

An additional advantage using the three photon excitation scheme would be the ability to completely eliminate Doppler broadening of the EIT transition. With three beams (satisfying the triangle inequality) the angles can be chosen such that the total sum of all wave-vectors vanishes. This will make the transition insensitive to atom temperature and has been demonstrated even in hot vapor cells. For our hybrid system this comes with a big caveat, however. The polaritons, which are collective atomic states, could still be de-phased by the atomic motion and decay outside of the cavity, representing a decoherence channel. However, this dephasing is depressed by the constant exchange between atomic excitation and cavity mode phase as described in [GSS18], and the quantitative improvement from a three-photon excitation has to be calculated using atomistic NHPT.

8.6.1 Magic lattice

Currently one of the main potential sources of decoherence of the dark polaritons is the effect of inhomogeneous broadening on the EIT resonance. Only atoms in the ground state are trapped by the 784 nm lattice, creating a position dependent detuning through the AC Stark shift that depends on the atoms position in the trap. Depending on the temperature of the atomic cloud, atoms explore different parts of the trap breaking the exact single photon resonance condition away from the absolute trap minimum. This inhomogeneous shift is especially pronounced in the ELAT due to the tight confinement along the cavity direction even for temperatures of around $1 - 5$ μK. The lattice depth needs to be shallow enough

to not introduce excess broadening beyond the Rydberg lifetime which needs to be carefully traded off against atom loss over the course of the probe sequence as atoms heat up from scattering cavity photons.

This decoherence channel could be eliminated by use of a magic wavelength trap that has the same light shift for both the ground and Rydberg state. While in Rubidium no true magic wavelength exists, trapping of the Rydberg state can be achieved with an intermediate red detuning from the $6P_{3/2} \rightarrow R$ transition at 1012 nm [**?**, ANS$^+$25]. The same light simultaneously acts like a far-off-resonant trap for the ground state and by carefully choosing the detuning both trap depths can be matched.

As this is not a true magic trap where the atomic polarizabilities cross but the fine-structure of the $6P_{3/2}$ state is resolved, the magic detuning differs for running or standing waves. Since a real optical lattice has both a periodic interference and static background term for both forward and backwards beam, the two conditions for a purely static or periodic field can not be satisfied at the same time [ANS$^+$25]. While an optimal detuning can still be found, it comes at a finite differential light shift on the 10 kHz scale that is already comparable to the Rydberg state lifetime. In our attempted realization, suboptimal coatings of the vertical vacuum viewports and $\sigma^+ + \sigma^-$ polarization further increased the minimum differential shift and thereby inhomogeneous broadening to a point where it did not justify the additional technical complexity.

In a future iteration of the apparatus, the 1012 nm laser that we successfully built should instead be used to realize an optical dipole trap replacing the ELAT. A running wave dipole trap would require more power from an additional TA but provide two main advantages: Currently, the atom density that can be estimated from the collective coupling to be 1.20×10^{11} cm^{-3} is locally increased by the lattice to a peak density of $n_{peak} \approx 6.7 \times 10^{11}$ cm^{-3}. At this density, effects such as light assisted two-body collisions or three-body recombination start to limit the maximum atomic density and even with increased loading times we were

not able to achieve more than $g \sim 55$ MHz. An optical dipole trap would allow for higher collective coupling strength as it reduces the peak density and is further off-resonant limiting loss. Simultaneously, a true magic condition exists for a pure running wave that could reduce broadening of the collective Rydberg polaritons and enable longer coherence times. To support against gravity, the dipole trap should ideally be sent in horizontal direction which is currently not possible at high NA due to limited optical access even with the turning mirrors for side imaging.

8.7 Spatio-temporal correlations

In this work we have demonstrate photon blockade in individual cavity modes using SPCM coupled to single mode (for $l = 0$) and multi-mode (higher l modes) fibers, and spatially resolved density images and movies using the novel, ultra-fast EMCCD camera. The ultimate hallmark of simulating fractional quantum Hall states, however, will be the measurement of spatial correlations of the Laughlin state showing a characteristic correlation hole in the 2D plane

$$g^{(2)}(\mathbf{z}_1, \mathbf{r}_2) \propto |\mathbf{r}_1 - \mathbf{r}_2|^4 \quad \text{as} \quad \mathbf{r}_1 \to \mathbf{r}_2. \tag{8.15}$$

As our ultrafast EMCCD camera is sensitive on the single-photon level and using a shift rate of 2 MHz has a time bin width of 500 ns it is uniquely suited to measure spatially- and time-resolved correlation functions. Due to the unique open nature of our hybrid quantum-optical experiment, and assuming that we can stabilize a many-body state in equilibrium using the driven-dissipative preparation scheme, photons that were part of the many-body state will continously leak out of the cavity. If we post-select for photons arriving within a time window much shorter than the characteristic dynamics of the system (the dark polariton lifetime),

we have access to the equal-time spatial correlation function

$$g^{(2)}(\mathbf{z}_1, \mathbf{z}_2; \tau = 0) = \frac{\langle \hat{\psi}^\dagger(\mathbf{z}_1)\hat{\psi}^\dagger(\mathbf{z}_2)\hat{\psi}(\mathbf{z}_2)\hat{\psi}(\mathbf{z}_1)\rangle}{\langle \hat{\psi}^\dagger(\mathbf{z}_1)\hat{\psi}(\mathbf{z}_1)\rangle\langle \hat{\psi}^\dagger(\mathbf{z}_2)\hat{\psi}(\mathbf{z}_2)\rangle} \tag{8.16}$$

as available in e.g. quantum gas microscopes. Beyond this usual observable, correlating photons arriving at different times, $t1$ and $t2$, with a delay $\tau = t2t1$, we can probe the excitation spectrum and the dissipative dynamics of the state. The unequal-time density-density correlation is given by

$$G^{(2)}(\mathbf{z}_1, \tau; \mathbf{z}_2, 0) = \langle \hat{n}(\mathbf{z}_1, \tau)\hat{n}(\mathbf{z}_2, 0)\rangle. \tag{8.17}$$

In an equilibrium-like steady state, taking the spatial and temporal Fourier transform of this correlation function yields the dynamical structure factor $S(\mathbf{q}, \omega)$. This unusual observable in AMO systems could give access to the collective modes of the Laughlin state: We would expect $S(\mathbf{q}, \omega)$ to peak at an energy corresponding to the magneto-roton [TSG82, GMP86, HK10] gap at a finite momentum $q \sim 1/\ell_B$. In the time domain, this manifests as damped oscillations in $g^{(2)}(\mathbf{z}_1, \tau; \mathbf{z}_2, 0)$ at the roton frequency. While these observables would be most interesting to probe the incompressible nature of Laughlin state and the fractional excitations through quasiparticle-quasihole binding, even for mixed or thermal states of polaritons these would already be novel measurements.

Measurement of $g^{(2)}$ correlation function using the camera is limited by two main factors experimentally: While the background electrons from dark currents are exceptionally low, clock induced charges (CIC) are incurred for every row shifting operation and therefore affect every pixel. As we expose every pixel for only the row transfer time itself as an extremely short exposure time, the relatively low CIC per image contributes significantly to the background noise at an equivalent rate of $\sim$ 1 kHz about twice as bad as two SPCM modules. The detection efficiency is limited due to the re-imaging optic for the MLA and the

intrinsic cavity loss from the second port and excess loss of $\sim 2000ppm$ from contamination after the bake. Even though the imaging path is a lot more complex, the increased QE of the EMCCD sensor at 780 nm leads to a comparable overall detection efficiency as the SPCM path. This could be improved by collecting light from both high transmisison cavity ports (lower mirrors) underaa transverse dissipative drive and a custom re-imaging objective.

For a singel pair of pixels we can estimate the signal to noise ratio as

$$\mathrm{SNR} = \sqrt{N_{shots}}\frac{p_s^2}{p_s + p_b}(1 - g^{(2)}_{ideal}) = \sqrt{N_{shots}\tau}\frac{(\eta R)^2}{\eta R + B}(1 - g^{(2)}_{ideal}) \tag{8.18}$$

with B the background and R the signal rate where $p_s = \eta R\tau$ and $p_b = B\tau$ for an integration time τ and detection efficiency η. For several transverse modes up to l_{max} at degeneracy, the number of photon pairs grows quadratically with mode number and will provide significantly more signal photons. The total number of experimental shots for a given SNR can then be estimated as

$$N_{shots} = \mathrm{SNR}^2\frac{1}{N_{pairs}}\frac{(\eta R + B)^2}{(\eta R)^4\tau^2(1 - g^{(2)}_{ideal})^2}. \tag{8.19}$$

Even with high loss of photons in the detection system and a count rate that is limited by having only one dark polariton per mode, detection of correlations should be possible as long as $\eta R \gg B$ the signal dominates the background from CICs. This will require careful setting of the right EM gain for photon counting (as the flux is so weak we can neglect two-photon events) and imaging the cavity field on as few pixels as possible while maintaining sufficient spatial resolution for the LLL. When the detection efficiency is 10%, we will have to integrate 10^4 times longer than with a perfect detector, but the scaling is manageable and such datasets could be obtained on the timescale of 1 day - 1 week.

APPENDIX A

A.1 Laser Induced Contamination

Already in the previous, mirror based cavity apparatus, a slow degradation of the blue buildup cavity finesse over the lifetime of the experiment was observed. As the originally planned HR mirror coating suffered from excess loss, the input mirror was replaced at the last minute with a lower reflectivity mirror from eBay. Therefore, the decline in blue cavity finesse was always attributed to the soft coating on the mirror that, according to conventional wisdom, is more susceptible to oxygen leaching and degradation. After measurement of the Laughlin state, the finesse had degraded so much that even at maximum available 480 nm laser power sufficient EIT transmission could only be obtained at low cooperativity $g \sim$ 15 MHz.

In the new apparatus, we therefore designed the Rydberg enhancement cavity with a large waist to cover the anticipated large number of ~ 20 degenerate modes and sufficiently high power buildup to reach several 10s of MHz of Rabi frequency. While the blue buildup cavity using high quality coatings from FiveNine optics worked excellent after commissioning of the new apparatus, reaching Rabi frequencies up to 60 MHz, the finesse dropped rapidly from its design value of $\mathcal{F} \approx 4000$ over the course of a few weeks. This reduction was not caused by the initial bakeout, which did not change the mirror cavity finesse significantly, but worsened continuously and more rapidly at higher powers of the blue control field.

As oxygen leaching from Ta_2O_5 coating layers had already been identified as a possible mechanism for coating degradation in the literature, we had specified our mirrors with a final SiO_2 layer which should have prevented this process. We therefore decided to build a separate test bench setup to investigate the finesse degradation and possible ramifications without the constraints of the main experiment. A simple two-mirror cavity using one flat and one curved mirror was built using surplus substrates from the same coating run and

designed to replicate the experimental cavity as closely as possible, also in terms of mode waist on the mirrors. Since we initially deemed it unlikely that substantial contamination could be present in the ultra-high vacuum system at a pressure of < 1 nTorr in the main chamber, we chose to only place the test setup in a small vacuum chamber pumped down to turbo pump pressures. The same 480 nm laser as for the main experiment was used to expose the cavity to different power levels and pulse durations to study the onset of finesse degradation.

Most surprisingly, no power level that we exposed the cavity to at ambient conditions, even at 100 mW of input light of CW corresponding to $\sim$ 100 W of intra-cavity power, lead to any kind of finesse degradation. Conversely, a reduction of finesse caused by a rapid increase in excess loss could be observed in vacuum. This rules out damage of the mirror coatings as the root cause of the loss in finesse, but points to the specific conditions of the vacuum in interaction with the high-intensity blue light in the cavity. From the test setup, we could easily extract the cavity mirrors after degradation and study them under an optical microscope and LEXT optical surface profilometer. In fig. A.1, the mirror surface is shown under a white light microscope, revealing a large and amorphous ring shaped contamination with a size multiple that of the waist.

A search of the literature for high power beam-line and UV lasers points to laser induced contamination as the likeliest explanation: Ultra-high vacuum (UHV) chambers and heated vapor cells inevitably contain trace amounts of outgassing organic materials, hydrocarbons, and residual solvents. When subjected to high-intensity blue light (which carries a photon energy of approximately 2.58 eV), these volatile organic molecules undergo multi-photon ionization or photochemical polymerization directly on the surfaces of the optical components. Over time, this localized photochemical reaction deposits an opaque, carbonaceous thin film on the mirrors. This film substantially increases the optical loss of the 480 nm beam, gradually extinguishing the available power buildup and Rabi frequency. As the in-

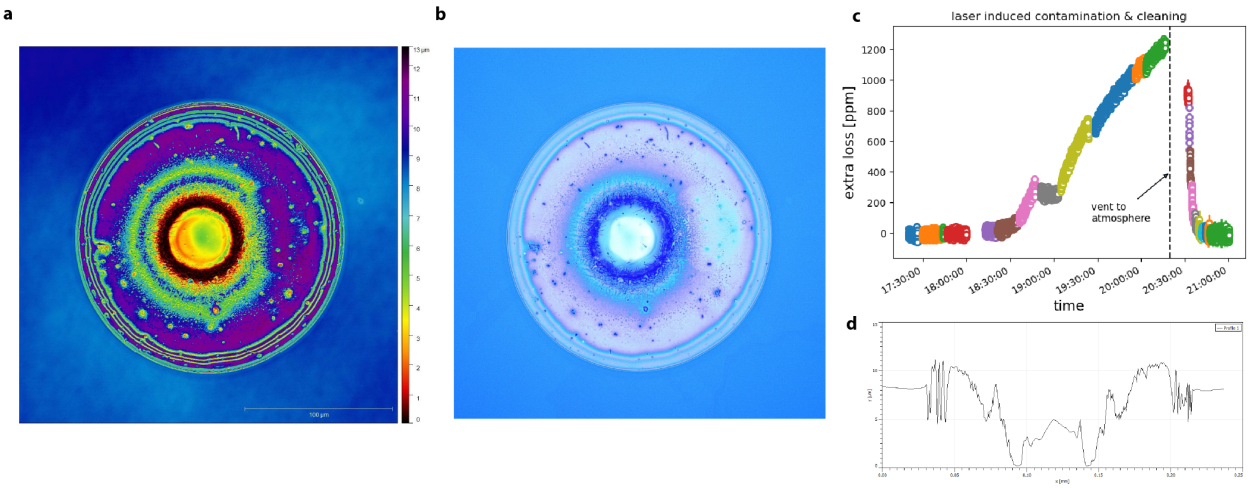


Figure A.1: Laser Induced Contamination (LIC) on the 480 nm Rydberg buildup cavity mirrors causes a buildup of hydrocarbons from a combination of optical trapping and photochemical polymerization. Microscope images with 50x magnification using a LEXT surface profilometer show (a) height map in false color and (b) real space image of the ring shaped buildup and (d) shows a 1D cut (y axis inverted). The contamination causes excess loss from absorption and scattering (c) spoiling the cavity finesse rapidly. This process can be reversed by exposing the cavity to ambient conditions as the oxygen in air with the blue light are able to oxidize the buildup material.

creased absorption also changes the impedance matching of the cavity, even less light can be coupled into it exacerbating this effect. The donut shaped height profile shown in fig. A.1 shows a plateau of several 100 um diameter with a height of buildup larger than in the mode waist itself. It is theorized that this shape is the result of an equilibrium between re-emission of hydrocarbon chains in the beam center with highest intensity and a maximum trapping force of the dipole trap formed by the cavity light at the FWHM of the Gaussian mode.

Two observations are notable: Only the flat cavity mirror of the test cavity was contaminated as explained by the smallest waist in the cavity occurring on the mirror surface, the curved mirror showed no signs of contamination. This effect did not occur in ambient conditions, as either the photochemical reactions do not occur in the same way or the presence of oxygen leads to the oxidation of the long-chain hydrocarbons. In high-power UV lasers and kW scale beam lines this effect is well known as LIC and is usually combated by regular mechanical cleaning of the mirrors, or by the use of highly reactive Ozone generated from Oxygen by UV-C lamps ex- or in-situ. Since neither of these are compatible with our

delicate cavity apparatus using super-polished lenses in a ultra-high vacuum chamber with Rubidium metal, we develop a novel *in-situ* laser cleaning method in a low, partial pressure of oxygen.

In the testbench setup, the cavity finesse was fully restored by exposing it to an increasingly high blue power at ambient conditions. The oxygen in the air, in combination with the 480 nm light itself, seem to reverse the photochemical reaction and oxidize the hydrocarbon polymers in a way that they reduce only into CO_2 and H_2O without leaving any measurable residue on the mirror surface. In the real experimental apparatus, however, venting to atmosphere would require a full bake of the vacuum system and exchange of the Rubidium getters taking up to a month of downtime.

Instead, our novel cleaning method uses a small partial pressure of $\sim 100 - 500$ mTorr of ultra high purity oxygen that is sufficient for cleaning but does not introduce water vapor or Hydrogen into the system. Therefore, the vacuum system does not need to be re-baked and the ion pumps restore ultra-high vacuum pressures within a day. Furthermore, only a slight initial degradation of flux from the Rubidium dispensers was observed, possible from a thin oxide layer on the getter surface that was subsequently reduced by the getter material.

In the partial oxygen atmosphere, the same 480 nm light inducing the contamination is able to reverse the photo-chemical process and oxidize the polymerized hydrocarbons to H_2O and CO_2 that can be both pumped efficiently from the chamber. This method is similar to plasma cleaning where a highly reactive oxygen plasma is used to clean any organic contamination, but it is much more specific as only the illuminated cavity mirrors are cleaned and the finesse of the science cavity did not change.

The progress during and result of the *in-situ* cleaning is shown in fig. A.2 over time. Using the real-time analysis system appendix A.12 we continuously measured the blue cavity finesse and the derived excess loss. Initially, the 480 nm laser power is carefully increased until the onset of cleaning can be observed and is subsequently only increased when the cleaning

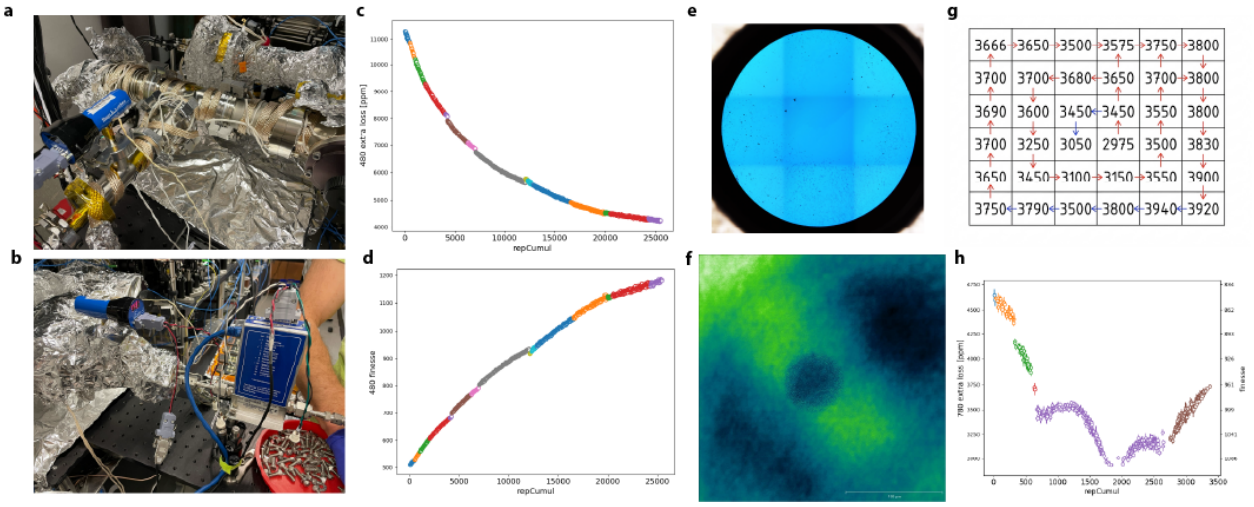


Figure A.2: LIC in-situ cleaning procedure is able to reverse the buildup of hydrocarbons (c, d) using a partial pressure of $100 - 500$ mTorr of ultra-high purity oxygen delivered via a closed loop feed system (a, b) comprised of an MKS mass-flow controller and Lesker Pirani vacuum gauge. After cleaning, almost no visible residue is visible (e, f) on the mirror surface, only in the enhanced false-color height map (f) on the nm level. Finesse of the blue cavity was restored over the whole mirror surface (g) by rastering the cavity waist using the tip-tilt stage. On the 780 nm science cavity, exposure to blue light had little effect but the oxygen atmosphere also lead to a temporary improvement in finesse.

plateaus.

In an ideal case, the oxygen would be flowed in and pump from opposite ends of the vacuum system, however we only had a single CF gate valve attached to the system. The oxygen was flowed in right at the gate valve using an MKS mass flow controller in a closed loop with a pressure transducer to establish a constant oxygen pressure in the presence of constant pumping from a turbo pump (Pfeiffer vacuum).

While cleaning the mirrors in this way, we have found that if the optical power used exceeds a certain threshold, the contamination becomes non-reversible and gets "baked in" (either from full polymerization or carbonization) and can only be removed at full ambient pressure. It is therefore paramount to choose the optical power to achieve a constant cleaning rate in ppm/min and only increase the optical power once the cleaning plateaus.

This cleaning method was successfully used to clean the full mirror surface and remove any LIC areas previously caused by operating the cavity at different locations and restore the

finesse uniformly to $\mathcal{F} \approx 4000$. Even though effective, the cleaning procedure takes several days of manual labor and is very prone to user error and could potentially vent the whole vacuum system to atmosphere if the gas feed system is not operated carefully or fails.

For future cavity designs, it would therefore be highly desirable to circumvent the chance of LIC by eliminating any source of hydrocarbons through new cleaning procedures and glue-free mounting structures. While no single material could be identified as the source of the hydrocarbons (even in the glue-less test cavity LIC occurred), any epoxy or other glues should be eliminated as they can be a source of solvents like Toluene that have been identified on residual gas analyzer traces of similar vacuum systems. Even though we have carefully chosen certified low-outgassing epoxies from Masterbond and Torrseal, a silver epoxy from Epotek that was necessary to electrically connect the lens mount to the rest of the cavity structure could be the source. More likely, however, every material that is not metal, glass or ceramic, such as Kapton wire insulation and PEEK connectors or glue could be contributing hydrocarbons and even residual solvents from the vacuum cleaning procedure could be the cause.

To eliminate all these sources, future cavity designs should make use of the spring washer design appendix A.10 that can be used for static and piezo driven mirrors and eliminate the need for glue without compromising mechanical resonance frequencies. Electrical connector bodies could be machined from MACOR machinable ceramics instead of PEEK. However, while the Kapton used for wire insulation could in principle be replaced by ceramic beads, this would require a lot of space and could make the wiring more prone to electrical shorts.

As long-chain hydrocarbons can fundamentally not be pumped by ion pumps, an additional cleaning step of all vacuum components right before assembly should be added using a plasma cleaner to remove any organic contamination. While plasma cleaning is an established technique in e.g. biology, it has not been used in AMO high-vacuum systems to our knowledge so especially compatibility of atomically smooth, super-polished optical sub-

strates with different types of plasma has to be further investigated. For existing vacuum systems, a plasma asher as used for in-situ chamber cleaning in electron microscopy could be added to the system and the *remote plasma* they generate should be compatible with optics as it uses Oxygen radicals instead of high kinetic energy electrons or ions.

A.2 Motorized Mirror Model

Modern AMO experiments are growing in complexity far beyond the single MOT and bulk gas in an optical dipole trap. Nowadays, quantum gas microscopes, optical tweezer arrays and hybrid cavity or clock experiments have to integrate a large number of optical lattices for trapping, transport or continuous reloading from a reservoir. Achieving and maintaining alignment between different lattices themselves and with cavity modes in different planes is a high dimensional optimization problem that requires interferometric precision. In most cases, the alignment can only be measured at the required precision using the target cloud of atoms themselves that provide measures of Rabi frequencies or map out the coupling to a cavity mode standing wave. This significantly limits the speed of the alignment procedure by the experimental cycle time that can be on the order of seconds up to a minute for experiments using evaporative cooling, and the limited amount of information one can extract per shot. This can make manual alignment of complex experiments an arduous task when a noisy signal only updates every several seconds. Even worse, the alignment can be subjective and non-repeatable as walking several mirrors with a slow, noisy feedback signal is hard.

In our hybrid experiment, one of the hardest alignments is overlapping the vertical transport lattice with the MOT and cavity waist simultaneously where the cloud has the be centered on the cavity modes on the μm scale to not introduce lensing. For static lattices like the ELAT, we can use a cats-eye retro setup where the lens on the bottom of the chamber together with an additional lens in front of the retro mirror form a 4f imaging system that images back the focus in the atom plane independent of incoming position. However, for the

optical conveyor belt lattice, the retro beam needs to be dynamically shifted in frequency by a pair of crossed, double pass AOMs constraining both position and angle of the forward pass beam sensitively. A second set of mirrors after the chamber is needed to exactly undo any movement of the forward pass made to align to the cavity waist and MOT. As this alignment can only be gauged using atoms, the lattice needs to stay sufficiently retro-reflected to hold the cloud requiring four mirrors to be walked simultaneously. Historically, this alignment was performed by hand in very small increments while constantly restoring retro alignment using a photodiode measuring light coupled back into the lattice fiber. Due to the many degrees of freedom, large distance of the mirror pairs from the atom plane, axes coupling from the periscope and many diffraction orders from the doubled passed AOMs the alignment was easily lost and could take a whole day.

As the blue cavity contamination required us to move the location of the atoms longitudinally along the cavity regularly, we decided to improve and speed up this alignment using motorized mirrors and a first order model of the optical system. We started by motorizing the mirrors of the ELAT, that on one hand is easier to align due to the robust 4f retro setup, but at the same time much more sensitive due to the small lattice waist along the cavity axis and resulting high NA of the ELAT beam. The 4f retro-reflector only re-images the small lattice waist up to paraxial order: Any aberrations, especially when the beam isn't centered on all the lenses will destroy or lead to a bi-modal trap for the atoms. Also, the focus of the ELAT forward and backward pass need to overlap exactly in the atom plane, but the vertical absorption imaging path can only be focused on the atoms to the limited resolution of the imaging system. While the alignment of the imaging was initially performed using an intra-cavity dipole trap using resonant 784 nm light in the cavity the ELAT focusing was always hard to reproduce and painful due to optics both on the top of and under the chamber and the foci being coupled. By replacing the key optics, - top focusing telescope, input mirror and retro tip-tilt and piston - on motorized mirror mounts the focusing of the ELAT

can be optimized by measuring the lattice trap frequencies using the atoms themselves from parametric heating. This decouples the lattice alignment from any imaging artifacts such as limited NA, wrong focal plane or chromatic aberrations and enabled to double the trap frequency compared to manual alignment.

In order to perform repeatable, multi-dimensional scans, the motorized mirrors need to be repeatable and free of backlash. While compact, high resolution and passively stable, ultrasonic piezo driven *pico-motors* or equivalent, however, are probabilistic in step size and direction dependent rendering them not useful for systematic scans. We therefore augment commercial off-the-shelf pico motors using absolute magnetic encoders AS5600 [1] to turn them into compact, closed-loop actuators. To keep the design compact, a small radially polarized magnet gets glued to the manual adjustment knob and a 3D printed holder positions a small sensor PCB exactly centered over the rotation axis. As the AS5600 support a $\sim$ 1 mm range of air gaps the encoder can be statically mounted for many full rotations of the fine pitch thread adjuster, but it is important to not exceed the permissible air gap range. A Raspberry Pi SBC then reads out the encoders over I2C and issues commands to the pico-motor controller via TCP/IP implementing a PI controller to close the loop and also keep track of full, integer rotation count.

While this system is already a significant improvement, the experiment framework still controls and scans raw mirror angles which for the optically passively robust ELAT setup is sufficient to optimize lattice overlap and focus. For the VLAT transport lattice, the large waist makes the focusing less sensitive but the double passed AOMs require two pairs of mirrors, one on top of the chamber to ensure overlap with MOT and cavity waist and one pair below the chamber to return the beam to the AOM axis. Moving the atoms in either MOT or cavity plane while maintaining a lattice requires four mirror knobs to be turned simultaneously. As this path references a lot of other alignments we didn't want to replace

1. https://ams-osram.com/products/sensor-solutions/position-sensors/ams-as5600-position-sensor

all the mirrors. Instead, we retrofit existing and stable Thorlabs *Polaris* mounts with closed loop digital servo motors (STS-3032) [2] using custom designed couplers. These have the additional benefit that all axes can be moved synchronously compared to pico-motors that can only move one axis at a time and are also much slower.

The beam position and angle from a single mirror are highly correlated and to control either using a mirror pair they both need to be walked to compensate the undesired tilt or position. This makes the part of the parameter space that we want to actually search a very narrow, (anti-)diagonal ellipse in raw "knob space" $\boldsymbol{\theta} = (\boldsymbol{\theta}^s, \boldsymbol{\theta}^t) \in \mathbb{R}^8$, and we want to employ a model to instead operate in physical beam positions in the plane of interest. If we have a model of the whole optical system $\boldsymbol{y} = f(\boldsymbol{\theta})$ that can predict the ray positions in the MOT and cavity plane as a function of knob angles $\boldsymbol{\theta}$, we can invert it $\boldsymbol{\theta}^* = f^{-1}(\boldsymbol{y}^*)$ to yield the knob angles need to reach a target position.

Previous approaches to this problem include gradient-free optimization (Nelder–Mead, Powell) on a scalar signal which is slow, blind to correlation ellipse structure, and therefore can get easily lost in a sparse, high-dimensional search space. Machine-learning and reinforcement-learning methods are extremely data-hungry, setup-specific, and requiring retraining for new configurations, and full ray-tracing packages are not designed for real-time closed-loop inverse problems. We need an optical model that is easy and fast to invert in real time but still allows treatment of large angles of reflection and lenses in the beam path. While simple ABCD transfer matrices can model the effects of lenses and mirror to first order, they are inherently only suited for elements on a single optical axis and can not treat large displacements and angles that we encounter in the real setup. Extensions such as 3×3 or 5×5 matrices exist [Sie86], these only capture shears and not real rotations. To model the real experiment geometry, one could employ coordinate transform matrices between elements but those will lead to complex final expressions from repeated matrix products. Instead, we

2. https://www.feetechrc.com/6v-45kg-magnetic-code-360-degree-serial-bus-steering-gear.html

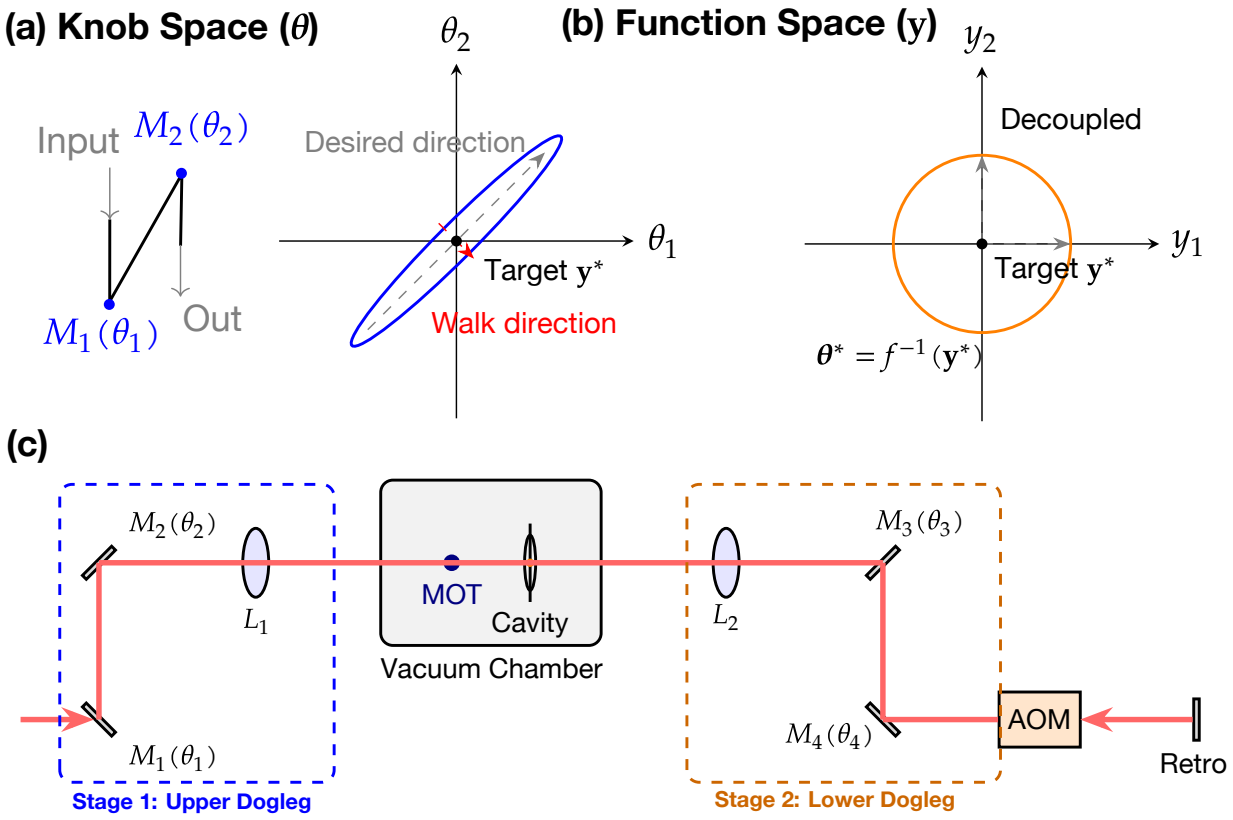


Figure A.3: Alignment search space before and after model-based inversion. (a) **Knob space:** Response of the beam position at the target to perturbations of a two-mirror dogleg. The Jacobian $\mathbf{J}$ has condition number $\kappa \gg 1$, producing a strongly elongated ellipse. To explore different offsets of the output ray requires "beam walking" along the major axis. (b) **Function space** after model inversion: the coordinate transform $\boldsymbol{\theta} \to f^{-1}(\boldsymbol{y})$ maps the ellipse to a circle ($\kappa \approx 1$) that is easy to scan and optimize. (c) **Optical layout** of the retro-reflected transport lattice. The beam exits the input fiber, is steered by the upper dogleg (mirrors M_1, M_2), passes through telescope lens L_1, traverses the vacuum chamber (MOT center and cavity waist separated by $\sim$5 cm), exits through L_2, is steered by the lower dogleg (M_3, M_4), double-passes the AOM and retro-reflects back into the input fiber.

employ a formalism based on the homogeneous coordinates of *projective geometry* [Cor23] that allows for arbitrary angles and displacements in a 2D optical system.

A ray traveling along the line $ax + by + c = 0$ in a transverse plane is represented by the homogeneous vector

$$\boldsymbol{r} = \begin{pmatrix} c \\ a \\ b \end{pmatrix}. \tag{A.1}$$

The slope $m = -a/b$ and intercept $h = -c/b$ recover the usual ABCD state vector $(h, m)^T$, where the additional coordinate makes translations linear.

Any optical element with canonical ray transfer matrix $\mathbf{M}_0$, placed at position (u, v) with orientation φ, has the lab-frame matrix

$$\mathbf{M} = \mathbf{T}^{-1}(u, v)\, \mathbf{R}^{-1}(\varphi)\, \mathbf{M}_0\, \mathbf{R}(\varphi)\, \mathbf{T}(u, v), \tag{A.2}$$

where $\mathbf{T}$ and $\mathbf{R}$ are 3×3 translation and rotation matrices. This is the standard change-of-basis pattern extended from 2×2 ABCD to 3×3 homogeneous coordinates. The system matrix is the ordered product of all element matrices along the beam path, but in an important distinction from ABCD matrices we don't need to insert propagation matrices between the elements.

For a mirror, $\mathbf{M}_0 = \mathrm{diag}(-1,\, 1,\, -1)$; for a thin lens of focal length f,

$$\mathbf{M}_0 = \begin{pmatrix} 1 & 0 & 0 \\ -1/f & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{A.3}$$

Mirror angles enter as orientation $\varphi = \varphi_0 + \theta_i$ in the sandwich transform [Eq. (A.2)] and are the only free parameters in the model where θ_i are the alignment corrections added to the

design angles φ_0.

The formalism inherits the projective duality of homogeneous coordinates, identical to the algebra used in computer graphics which makes it extremely efficient to solve for intersections of rays or a ray through two points

$$\text{Line through two points:} \quad \boldsymbol{r} = \boldsymbol{p}_1 \times \boldsymbol{p}_2, \tag{A.4}$$

$$\text{Intersection of two lines:} \quad \boldsymbol{p} = \boldsymbol{r}_1 \times \boldsymbol{r}_2. \tag{A.5}$$

The same cross-product operation has dual geometric interpretations. Crucially, element positions enter through $\mathbf{T}(u, v)$ and ray intersections are resolved automatically via Eq. (A.5), no explicit propagation distances are needed. The model is assembled directly from roughly measured element positions from the experiment and specified focal lengths for lenses.

The sagittal and tangential planes are treated as two independent 3×3 systems, and this decoupling is exact for setups with parallel or 90° periscopes as our VLAT path and axis-aligned astigmatic elements as the cylindrical telescope in our ELAT. The fully coupled case (arbitrary image rotation) would require a Plücker-coordinate extension to 6D which is still an open question [Cor23]. All equations below are written for one transverse plane $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4) \in \mathbb{R}^4$.

For our lattice beam path, the system matrix is the ordered product of $\sim$20 element matrices: fiber output, mirrors $M_1(\theta_1)$ through $M_4(\theta_4)$, telescope lenses L_1 and L_2, the vacuum chamber (with MOT and cavity positions as fixed points), the AOM, and the retroreflector.

The four-degree-of-freedom inversion $\boldsymbol{\theta}^* = f^{-1}(\boldsymbol{y}^*)$ decomposes into two decoupled two-dimensional problems, solved sequentially in a single forward pass.

The target ray connecting two known spatial points (MOT center $\boldsymbol{p}_{\text{MOT}}$ and cavity waist $\boldsymbol{p}_{\text{cav}}$) is constructed via Eq. (A.4):

$$\boldsymbol{r}_{\text{target}} = \boldsymbol{p}_{\text{cav}} \times \boldsymbol{p}_{\text{MOT}}. \tag{A.6}$$

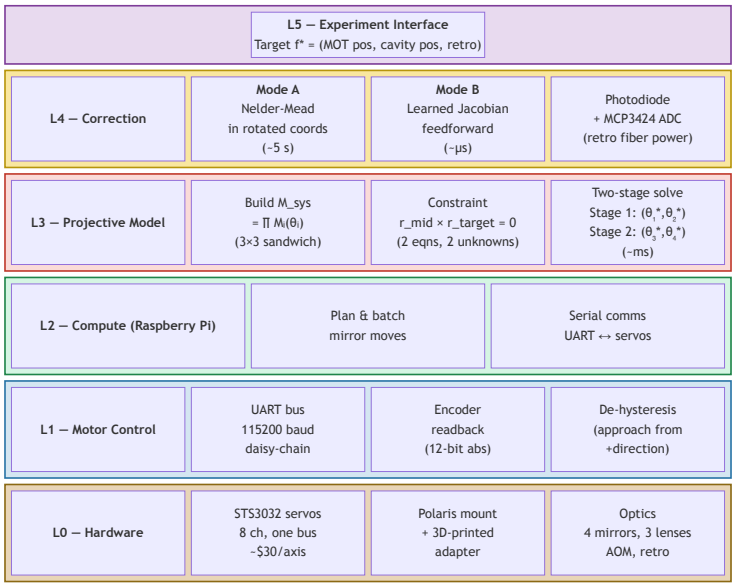


Figure A.4: Layered control architecture (L0–L5). **L0** (Hardware): STS3032 servos on Polaris mounts with 3D-printed adapters drive four mirrors, or Pico motors with encoders. **L1** (Motor Control): UART daisy chain at 115 200 baud with 12-bit absolute encoder readback and uni-directional approach for backlash elimination. **L2** (Compute): Raspberry Pi batches and plans mirror moves, handles serial communication. **L3** (Projective Model): Builds the system matrix $\mathbf{M}_{\text{sys}} = \prod_i \mathbf{M}_i(\theta_i)$ via the 3×3 sandwich transform, solves the collinearity constraint $\boldsymbol{r}_{\text{mid}} \times \boldsymbol{r}_{\text{target}} = \mathbf{0}$ in two sequential stages (∼ms). **L4** (Correction): Mode A applies Nelder–Mead optimization in rotated mirror-pair coordinates (∼5 s); Mode B retrieves empirical corrections via multilinear interpolation of a pre-measured `xarray` table (∼μs); both use retro-coupled fiber power from an MCP3424 ADC as the objective. **L5** (Experiment Interface): Accepts target positions $\boldsymbol{y}^* = (p_{\text{MOT}}, p_{\text{cav}}, \text{retro})$.

The input ray is propagated symbolically through mirrors M_1, M_2 and lens L_1

$$\boldsymbol{r}_{\text{mid}}(\theta_1, \theta_2) = \mathbf{M}_{L_1}\, \mathbf{M}_{M_2}(\theta_2)\, \mathbf{M}_{M_1}(\theta_1)\, \boldsymbol{r}_{\text{in}}. \tag{A.7}$$

and the alignment condition (the propagated ray must coincide with the target ray) is expressed as a collinearity constraint in homogeneous coordinates:

$$\boldsymbol{r}_{\text{mid}}(\theta_1, \theta_2) \times \boldsymbol{r}_{\text{target}} = \mathbf{0}. \tag{A.8}$$

Since the cross product of two three-vectors yields a three-vector with one redundant component (by $\boldsymbol{v} \cdot (\boldsymbol{v} \times \boldsymbol{w}) = 0$), this provides two independent scalar equations in two unknowns (θ_1, θ_2). Equation (A.8) is the mechanism by which the model inversion f^{-1} is evaluated.

The double-pass AOM acts as both an angle filter (Bragg condition) and a spatial filter (finite aperture). Any deviation of the output ray from the design path causes loss of diffraction efficiency and beam walk-off, destroying the retro-coupled signal. Therefore, the output ray must be *invariant* with respect to the design-path reference ray.

The reference ray is obtained by propagating through the unperturbed system (all $\theta_i = 0$):

$$\boldsymbol{r}_{\text{ref}} = \prod_i \mathbf{M}_i(0)\, \boldsymbol{r}_{\text{in}}. \tag{A.9}$$

The output ray, propagated using the Stage 1 solution, is

$$\boldsymbol{r}_{\text{out}}(\theta_3, \theta_4) = \mathbf{M}_{M_4}(\theta_4)\, \mathbf{M}_{M_3}(\theta_3)\, \mathbf{M}_{L_2}\, \boldsymbol{r}_{\text{mid}}(\theta_1^*, \theta_2^*). \tag{A.10}$$

The AOM invariance constraint is

$$\boldsymbol{r}_{\text{out}}(\theta_3, \theta_4) \times \boldsymbol{r}_{\text{ref}} = \mathbf{0}, \tag{A.11}$$

which has the same structure as Eq. (A.8): two equations in two unknowns. The two stages are decoupled and solved sequentially—no iteration is required. They can be solved either symbolically or numerically, however we have found the root finding to be faster than evaluating the symbolic solution.

Algorithm 1 Model inversion: $\boldsymbol{\theta}^* = f^{-1}(\boldsymbol{y}^*)$

Require: Element matrices $\{\mathbf{M}_i(\theta_i)\}$, input ray $\boldsymbol{r}_{\text{in}}$, target points $\boldsymbol{p}_{\text{MOT}}$, $\boldsymbol{p}_{\text{cav}}$

Ensure: Mirror angles $\boldsymbol{\theta}^* = (\theta_1^*, \theta_2^*, \theta_3^*, \theta_4^*)$

— **Stage 1: Forward path** —

$\boldsymbol{r}_{\text{target}} \leftarrow \boldsymbol{p}_{\text{cav}} \times \boldsymbol{p}_{\text{MOT}}$

$\boldsymbol{r}_{\text{mid}}(\theta_1, \theta_2) \leftarrow \mathbf{M}_{L_1} \cdot \mathbf{M}_{M_2}(\theta_2) \cdot \mathbf{M}_{M_1}(\theta_1) \cdot \boldsymbol{r}_{\text{in}}$

Solve $\boldsymbol{r}_{\text{mid}} \times \boldsymbol{r}_{\text{target}} = \mathbf{0}$ for (θ_1^*, θ_2^*)

— **Stage 2: Return path (AOM invariance)** —

$\boldsymbol{r}_{\text{ref}} \leftarrow \prod_i \mathbf{M}_i(0) \cdot \boldsymbol{r}_{\text{in}}$

$\boldsymbol{r}_{\text{out}}(\theta_3, \theta_4) \leftarrow \mathbf{M}_{M_4}(\theta_4) \cdot \mathbf{M}_{M_3}(\theta_3) \cdot \mathbf{M}_{L_2} \cdot \boldsymbol{r}_{\text{mid}}(\theta_1^*, \theta_2^*)$

Solve $\boldsymbol{r}_{\text{out}} \times \boldsymbol{r}_{\text{ref}} = \mathbf{0}$ for (θ_3^*, θ_4^*)

— **Command** —

Send $\boldsymbol{\theta}^*$ to motor controller; await encoder confirmation

The paraxial model has residual errors from aberrations, mechanical imperfections (mount hysteresis, thermal drift), and unmodeled elements (window wedge, cubes etc). In practice, the model places the beam within $\sim$90% of the optimum over $\sim$ 100 μm displacements in the cavity plane, but for very large displacements of the mirrors drops off due to higher-order effects.

We therefore employ two types of correction mechanism that use the power coupled

back into the VLAT fiber as a fast measure of the lattice overlap: A Nelder–Mead simplex optimizer maximizes $P(\boldsymbol{\theta})$ on the lower-dogleg mirrors (M_3, M_4) only, operating in rotated coordinates that decorrelate the mirror-pair response to first order. For the y-direction, the coordinate transform is $(u, v) \mapsto (\theta_3 + \theta_4,\ -\theta_3 + \theta_4)$; for the x-direction, $(u, v) \mapsto (\theta_3 + \theta_4,\ \theta_3 - \theta_4)$. Each evaluation reads the retro-coupled fiber power via the MCP3424 ADC after each servo move and convergence requires ≤ 12 function evaluations per axis with $\sim$5 s total time. This can be used to periodically reoptimize the lattice overlap after thermal drifts or to optimize at a new point that is at the edge of model validity. For multi-dimensional scans it is too slow to optimize every point with individual optimization and we therefore find and interpolate a correction model on top of the first order model from a grid of retro optimizer results. The systematic residual between the model prediction and the true optimum depends smoothly on the target positions $\boldsymbol{y}^*$. We pre-measure this residual on a four-dimensional grid and store it as an `xarray` DataArray $\mathcal{I}$:

$$\delta\boldsymbol{\theta}_{\text{corr}} = \mathcal{I}\Big(p_{\text{cav}}^{(1)},\ p_{\text{cav}}^{(2)},\ p_{\text{MOT}}^{(1)},\ p_{\text{MOT}}^{(2)}\Big)\,, \tag{A.12}$$

where $\mathcal{I}$ is a multilinear interpolant over the grid (6 points per axis, $6^4 = 1296$ nodes, 8 servo corrections per node). At each node, the model solution is commanded and Nelder–Mead refines the alignment, and the measured correction $\delta\boldsymbol{\theta}$ is recorded. At runtime, the correction for an arbitrary target is retrieved by interpolation (with extrapolation beyond grid bounds). Unlike a linear Jacobian, $\mathcal{I}$ captures the full nonlinear residual at the cost of a one-time grid measurement ($\sim$20 min).

This model based lattice optimization has enabled the fine scans of multimode VRS spectra in chapter 5.

A.3 Magnetic Translation Operator

We work in the Symmetric Gauge with the magnetic length set to unity ($\ell_B = 1$). The unnormalized basis functions for the Lowest Landau Level (LLL) are:

$$\tilde{\phi}_m(z) = z^m e^{-|z|^2/4}$$

We compute the normalization constant N_m for the state $|m\rangle$ using the inner product on the plane $\langle f|g\rangle = \int_{\mathbb{C}} \overline{f(z)} g(z)\, d^2z$, yielding:

$$|m\rangle = \frac{1}{\sqrt{\pi 2^{m+1} m!}}\, z^m\, e^{-|z|^2/4} \tag{A.13}$$

A.3.1 The Magnetic Translation Operator

The magnetic translation operator $\hat{T}(\eta)$ corresponding to a spatial shift by a complex number η acts on an LLL wavefunction $\psi(z) = f(z)e^{-|z|^2/4}$ (where $f(z)$ is analytic) as:

$$\hat{T}(\eta)\psi(z) = e^{-|\eta|^2/4} e^{\bar{\eta}z/2} \psi(z - \eta) \tag{A.14}$$

If one naively translates the state by shifting the coordinate $z \to z - \eta$, the Gaussian envelope becomes:

$$e^{-|z-\eta|^2/4} = e^{-|z|^2/4} e^{-|\eta|^2/4} e^{\mathrm{Re}(\bar{\eta}z)/2} = e^{-|z|^2/4} e^{-|\eta|^2/4} e^{(\bar{\eta}z + \eta\bar{z})/4}$$

The $\bar{z}$ term can't be included in the LLL and therefore couples to higher Landau level components, effectively exciting the particle out of the ground state.

The magnetic translation operator includes the specific phase factor $e^{\bar{\eta}z/2}$ to counteract this. When multiplied by the shifted Gaussian, the non-analytic $\bar{z}$ term is exactly cancelled,

leaving only a holomorphic function of z (and constant phase factors):

$$\underbrace{e^{\bar{\eta}z/2}}_{\text{Phase Fix}} \times \underbrace{e^{-|z-\eta|^2/4}}_{\text{Shifted Gaussian}} \propto e^{-|z|^2/4} \times (\text{Holomorphic in } z)$$

We wish to calculate the matrix element $C_{l,\alpha}(\eta)$ representing the overlap of a translated angular momentum state $|\alpha\rangle$ with a target state $|l\rangle$:

$$C_{l,\alpha}(\eta) = \langle l|\hat{T}(\eta)|\alpha\rangle$$

Substituting the states:

$$\hat{T}(\eta)|\alpha\rangle = N_\alpha e^{-|\eta|^2/4} e^{\bar{\eta}z/2} (z-\eta)^\alpha e^{-|z-\eta|^2/4}$$

Standard magnetic translation identities allow us to simplify the Gaussian factors. However, strictly working within the analytic Bargmann space is cleaner. We focus on the transformation of the analytic part (polynomials). The operator acting on the polynomial $P(z) = z^\alpha$ is:

$$\hat{\mathcal{T}}(\eta)[z^\alpha] = e^{-|\eta|^2/4} e^{\bar{\eta}z/2} (z-\eta)^\alpha$$

We expand the two terms in the product using the binomial theorem and the Taylor series for the exponential:

$$(z-\eta)^\alpha = \sum_{k=0}^{\alpha} \binom{\alpha}{k} z^k (-\eta)^{\alpha-k}$$

$$e^{\bar{\eta}z/2} = \sum_{j=0}^{\infty} \frac{1}{j!} \left(\frac{\bar{\eta}}{2}\right)^j z^j$$

Multiplying these two series:

$$\text{Poly}(z) = e^{-|\eta|^2/4} \left[\sum_{k=0}^{\alpha} \binom{\alpha}{k} (-\eta)^{\alpha-k} z^k \right] \left[\sum_{j=0}^{\infty} \frac{1}{j!} \left(\frac{\bar{\eta}}{2}\right)^j z^j \right]$$

We perform a Cauchy product of the series. We introduce a new index $l = k + j$ representing the total power of z. We sum over l from 0 to ∞. For a fixed l, the index k runs from 0 to l. However, we are also constrained by $k \leq \alpha$. Thus, the summation range for k is $0 \leq k \leq \min(l, \alpha)$. Substituting $j = l - k$:

$$\text{Poly}(z) = e^{-|\eta|^2/4} \sum_{l=0}^{\infty} z^l \left[\sum_{k=0}^{\min(l,\alpha)} \binom{\alpha}{k} (-\eta)^{\alpha-k} \frac{1}{(l-k)!} \left(\frac{\bar{\eta}}{2}\right)^{l-k} \right]$$

This gives us the coefficient of the *unnormalized* power z^l. Let's call the term in the brackets $S_{l,\alpha}$.

The expansion in terms of normalized states is:

$$\hat{T}(\eta)|\alpha\rangle = \sum_l C_{l,\alpha} |l\rangle = \sum_l C_{l,\alpha} N_l z^l e^{-|z|^2/4}$$

Comparing this with our derived expression $\hat{T}(\eta)(N_\alpha z^\alpha e^{-|z|^2/4}) = N_\alpha e^{-|z|^2/4} \sum_l S_{l,\alpha} z^l$:

$$C_{l,\alpha} N_l = N_\alpha e^{-|\eta|^2/4} S_{l,\alpha}$$

Solving for the coefficient $C_{l,\alpha}$:

$$C_{l,\alpha} = e^{-|\eta|^2/4} \frac{N_\alpha}{N_l} \sum_{k=0}^{\min(l,\alpha)} \binom{\alpha}{k} (-\eta)^{\alpha-k} \frac{1}{(l-k)!} \left(\frac{\bar{\eta}}{2}\right)^{l-k}$$

Substitute $N_m = (\pi 2^{m+1} m!)^{-1/2}$:

$$\frac{N_\alpha}{N_l} = \sqrt{\frac{\pi 2^{l+1} l!}{\pi 2^{\alpha+1} \alpha!}} = \sqrt{\frac{l!}{\alpha!}} 2^{(l-\alpha)/2} = \sqrt{\frac{l!}{\alpha!}} \left(\frac{1}{2}\right)^{(\alpha-l)/2}$$

Combining the terms, we arrive at the unified formula for the magnetic translation coefficient:

$$C_{l,\alpha}(z_0) = e^{-|z_0|^2/4} \sqrt{\frac{l!}{\alpha!}} \left(\frac{1}{2}\right)^{\frac{\alpha-l}{2}} \sum_{k=0}^{\min(l,\alpha)} \binom{\alpha}{k} (-z_0)^{\alpha-k} \frac{1}{(l-k)!} \left(\frac{\bar{z}_0}{2}\right)^{l-k} \tag{A.15}$$

In the derivation, we focused on the transformation of the polynomials $z^\alpha \to \text{Poly}(z)$ and mostly ignored the Gaussian factor. This is formally rigorous via the **Bargmann-Fock isomorphism**.

The Hilbert space of LLL wavefunctions $L^2(\mathbb{C}, d^2z)$ is isomorphic to the Bargmann space $\mathcal{F}$ of analytic functions. Instead of carrying the ubiquitous $e^{-|z|^2/4}$ factor, we simply work with the analytic function $f(z)$. To compensate for dropping the exponential from the state, we insert a weight into the integration measure:

$$\langle f|g\rangle_{\mathcal{F}} = \int_{\mathbb{C}} \overline{f(z)} g(z) e^{-|z|^2/2} \, d^2z$$

In this space, the magnetic translation operator simplifies to acting purely on the analytic function as $\hat{T}(\eta) f(z) = e^{-|\eta|^2/4} e^{\bar{\eta} z/2} f(z-\eta)$, which matches the "Poly(z)" transformation used in Section 3.

A.4 Cavity mode shift and coupling from Astigmatism

We begin by establishing the equivalence between a particle in a rotating frame and a charged particle in a magnetic field. This allows us to map the cavity eigenmodes onto Landau Levels.

Consider a particle of mass m and charge $q = 1$ moving in a 2D plane under the influence of a vector potential $\mathbf{A}(\mathbf{r})$ and a scalar potential $V_{\text{eff}}(\mathbf{r})$. The Hamiltonian is:

$$H = \frac{(\mathbf{p} - \mathbf{A})^2}{2m} + V_{\text{eff}}(\mathbf{r}) \tag{A.16}$$

We choose the symmetric gauge $\mathbf{A} = \frac{B}{2}(-y, x, 0) = \frac{B}{2}\hat{z} \times \mathbf{r}$. Expanding the kinetic term:

$$\frac{(\mathbf{p} - \mathbf{A})^2}{2m} = \frac{p^2}{2m} - \frac{1}{2m}(\mathbf{p} \cdot \mathbf{A} + \mathbf{A} \cdot \mathbf{p}) + \frac{A^2}{2m} \tag{A.17}$$

$$= \frac{p^2}{2m} - \frac{B}{2m}L_z + \frac{B^2}{8m}r^2 \tag{A.18}$$

where we used $\mathbf{A} \cdot \mathbf{p} = \frac{B}{2}(xp_y - yp_x) = \frac{B}{2}L_z$.

Now, compare this to the Hamiltonian of a neutral particle in a harmonic trap $V_{\text{trap}} = \frac{1}{2}m\omega_0^2 r^2$, viewed from a frame rotating at frequency Ω. The rotation introduces a term $-\Omega L_z$:

$$H_{\text{rot}} = \frac{p^2}{2m} + \frac{1}{2}m\omega_0^2 r^2 - \Omega L_z \tag{A.19}$$

Matching the terms between the Gauge Hamiltonian and the Rotating Hamiltonian:

1. **Coriolis Term:** $-\Omega L_z = -\frac{B}{2m}L_z \implies B_{\text{eff}} = 2m\Omega$.

2. **Effective Potential:** The centrifugal term from the gauge expansion ($\frac{B^2}{8m}r^2 = \frac{1}{2}m\Omega^2 r^2$) must be subtracted from the trap potential to match the net scalar potential V_{eff}.

$$V_{\text{eff}}(r) = \frac{1}{2}m\omega_0^2 r^2 - \frac{1}{2}m\Omega^2 r^2 = \frac{1}{2}m(\omega_0^2 - \Omega^2)r^2 \tag{A.20}$$

Thus, the rotating frame is mathematically equivalent to a system with an effective magnetic field $B_{\text{eff}} = 2m\Omega$ and a reduced scalar potential.

When the rotation frequency matches the trap frequency ($\Omega \to \omega_0$), the effective scalar potential vanishes ($V_{\text{eff}} \to 0$). The force balance is exact: the attractive trap force cancels

the repulsive centrifugal force. The Hamiltonian becomes simply that of a free particle in a magnetic field:

$$H_{\text{LL}} = \frac{(\mathbf{p} - \mathbf{A})^2}{2m} = \hbar\omega_c(n_L + 1/2) \tag{A.21}$$

where the cyclotron frequency is $\omega_c = B_{\text{eff}}/m = 2\Omega \approx 2\omega_0$. The energy spectrum depends **only** on the index n_L, which we identify as the **Landau Level index**.

We describe the states using the circular occupation numbers n_L and n_R.

- n_L: The Landau Level index. Determines the energy $E \approx 2\hbar\omega_0(n_L + 1/2)$.
- $l = n_R - n_L$: The angular momentum.

Since n_R is a number operator ($n_R \geq 0$), we have the constraint:

$$l \geq -n_L \tag{A.22}$$

This naturally organizes the states:

- **Lowest Landau Level ($n_L = 0$):** Must have $l \geq 0$.
- **Higher Landau Levels ($n_L > 0$):** Can support negative angular momentum states (e.g., $l = -1$ exists in $n_L = 1$).

Finally, we can relate this back to the standard Laguerre-Gaussian (LG) radial index p. The total LG mode number is $N = n_R + n_L = 2p + |l|$. Solving for n_L and n_R:

$$n_L = p + \frac{|l| - l}{2}, \quad n_R = p + \frac{|l| + l}{2} \tag{A.23}$$

This confirms that for positive l, the Landau index n_L is simply the radial index p. For negative l, the state is pushed into a higher Landau level $n_L = p + |l|$.

A.5 2. Effective Hamiltonian with Residual Detuning

We utilize the circular creation and annihilation operators:

$$a_R = \frac{1}{\sqrt{2}}(a_x - ia_y) \tag{A.24}$$

$$a_L = \frac{1}{\sqrt{2}}(a_x + ia_y) \tag{A.25}$$

Substituting these into H_{rot} with a small residual detuning $\Delta_l = \omega_0 - \Omega$:

$$H_{\text{rot}}^{(0)} = \hbar\omega_0(2n_L + 1) + \hbar\Delta_l(n_R - n_L) \tag{A.26}$$

The energy depends primarily on n_L (Landau Level index), with a fine structure splitting given by Δ_l.

A.6 Second order astigmatism perturbation

We now expand the perturbation $V = \lambda(x^2 - y^2)$ explicitly. First, express coordinates in terms of a_R, a_L:

$$x = \frac{x_{\text{zpf}}}{\sqrt{2}}(a_R + a_L + a_R^\dagger + a_L^\dagger) \tag{A.27}$$

$$y = \frac{x_{\text{zpf}}}{i\sqrt{2}}(a_R - a_L - a_R^\dagger + a_L^\dagger) \tag{A.28}$$

where $x_{\text{zpf}} = \sqrt{\hbar/2m\omega_0}$.

We calculate $x^2 - y^2$ by summing the squares of the operator groups. Terms involving mixed products like $a_R a_L$ (and hermitian conjugates) cancel out in the subtraction $x^2 - y^2$, as they represent rotationally symmetric breathing modes or simple scaling. The surviving

terms are the quadrupole terms:

$$x^2 - y^2 = x_{\text{zpf}}^2 \left[(a_R^\dagger)^2 + a_R^2 + (a_L^\dagger)^2 + a_L^2 \right] \tag{A.29}$$

We must now determine which of these terms are significant in the Rotating Frame.

- **Terms $(a_L^\dagger)^2$ and a_L^2:** These change the Landau Level index n_L by ± 2. The energy cost is $\approx 4\hbar\omega_0 = 2\hbar\omega_c$ (twice the cyclotron frequency). These are highly off-resonant and effectively average to zero.
- **Terms $(a_R^\dagger)^2$ and a_R^2:** These change the intra-level index n_R but conserve n_L. The energy cost is $2\hbar\Delta_l$. As $\Delta_l \to 0$, these terms become **resonant**.

The effective perturbation Hamiltonian is therefore:

$$V_{\text{eff}} = \hbar\Lambda \left[(a_R^\dagger)^2 + (a_R)^2 \right] \tag{A.30}$$

where $\Lambda = \lambda x_{\text{zpf}}^2/\hbar$. We calculate the coupling between angular momentum states $|l\rangle$ and $|l+2\rangle$ within a fixed Landau Level p (specifically for $l \geq 0$ where $n_L = p$). Recall that $n_R = p + \frac{|l|+l}{2}$. For $l \geq 0$, $n_R = p + l$.

A.6.1 Coupling Calculation

We compute $\langle p, l'|V_{\text{eff}}|p, l\rangle$. The term $(a_R^\dagger)^2$ raises n_R by 2 (so $l \to l+2$).

$$\langle p, l+2|(a_R^\dagger)^2|p, l\rangle = \sqrt{n_R + 1}\sqrt{n_R + 2}\langle p, l+2|p, l+2\rangle \tag{A.31}$$

$$= \sqrt{(p+l+1)(p+l+2)} \tag{A.32}$$

Substituting the prefactor $\hbar\Lambda$:

$$\mathcal{M}_{l\to l+2} = \hbar\Lambda\sqrt{(p+l+1)(p+l+2)} \tag{A.33}$$

For the Lowest Landau Level ($p = 0$):

$$\langle 0, l+2|V|0, l\rangle = \hbar\Lambda\sqrt{(l+1)(l+2)} \tag{A.34}$$

We analyze the effect of this coupling on the spectrum near the degeneracy point ($\Delta_l \to 0$). Consider the subspace of two modes $\{|l\rangle, |l+2\rangle\}$ in the LLL. The Hamiltonian matrix is:

$$H_{2\times 2} = \hbar \begin{pmatrix} \Delta_l l & \Lambda\sqrt{(l+1)(l+2)} \\ \Lambda\sqrt{(l+1)(l+2)} & \Delta_l(l+2) \end{pmatrix} \tag{A.35}$$

Defining the mean energy $\bar{E} = \hbar\Delta_l(l+1)$ and the gap $\delta_{gap} = 2\hbar\Delta_l$:

$$H_{2\times 2} = \bar{E}\mathbb{I} + \hbar \begin{pmatrix} -\Delta_l & g_{\text{eff}} \\ g_{\text{eff}} & \Delta_l \end{pmatrix} \tag{A.36}$$

where $g_{\text{eff}} = \Lambda\sqrt{(l+1)(l+2)}$. The eigenenergies are:

$$E_\pm = \bar{E} \pm \hbar\sqrt{\Delta_l^2 + g_{\text{eff}}^2} \tag{A.37}$$

This result recovers the behavior observed in experiments:

1. **Off-Resonance ($\Delta_l \gg \Lambda$):** The mode mixing is suppressed. The energy shift is quadratic in the astigmatism parameter Λ:

$$\Delta E \approx \frac{\hbar g_{\text{eff}}^2}{2\Delta_l} \propto \frac{\Lambda^2}{\Delta_l} \tag{A.38}$$

2. **At Degeneracy ($\Delta_l \to 0$):** The gap closes, and the modes hybridize completely into symmetric/antisymmetric superpositions. The splitting becomes linear in the

astigmatism parameter:

$$\Delta E = \pm\hbar g_{\text{eff}} = \pm\hbar\Lambda\sqrt{(l+1)(l+2)} \tag{A.39}$$

This linear sensitivity to Λ (and thus to δ^2 if arising from misalignment) explains the high sensitivity of the degenerate cavity.

A.7 Analog Lockbox

Optical cavities are widely used in precision metrology, laser stabilization, and quantum optics. To maintain a cavity on resonance, we typically lock a laser's frequency to the cavity (or vice versa) using a feedback loop. Due to the high dynamic range and high bandwidth of laser and cavity locks, PI controllers using analog electronics are still superior for a lot of applications in the lab. A proportional-integral (PI) controller has the transfer function $H(s) = K_p\left(1 + \frac{\omega_I}{s}\right)$. It provides theoretically infinite gain at DC ($s \to 0$), heavily suppressing low-frequency drift. The PI corner frequency ω_I (the zero) is chosen to be lower than the unity gain crossing but as high as possible without eroding the phase margin at crossover.

This design is based on the original Simon lab lockbox but with a highly optimized critical path for feedback bandwidth up to 10 MHz and low noise performance. The mechanical design employs all PCB mounted connectors and interfaces, enabling mass manufacturing via PCB assembly service and eliminates all error prone Molex crimp connections and manual assembly. Over 30 lockboxes of this new design have been employed across the three experiments in the lab reliably.

To achieve massive low-frequency gain without sacrificing bandwidth, a double-integrator (PII) can be used. A true PII controller is constructed by summing a standard PI controller and a double integrator (II) in parallel: $H(s) = K_p(1 + \frac{\omega_{PI}}{s}) + \frac{K_{II}}{s^2}$, where three opamps

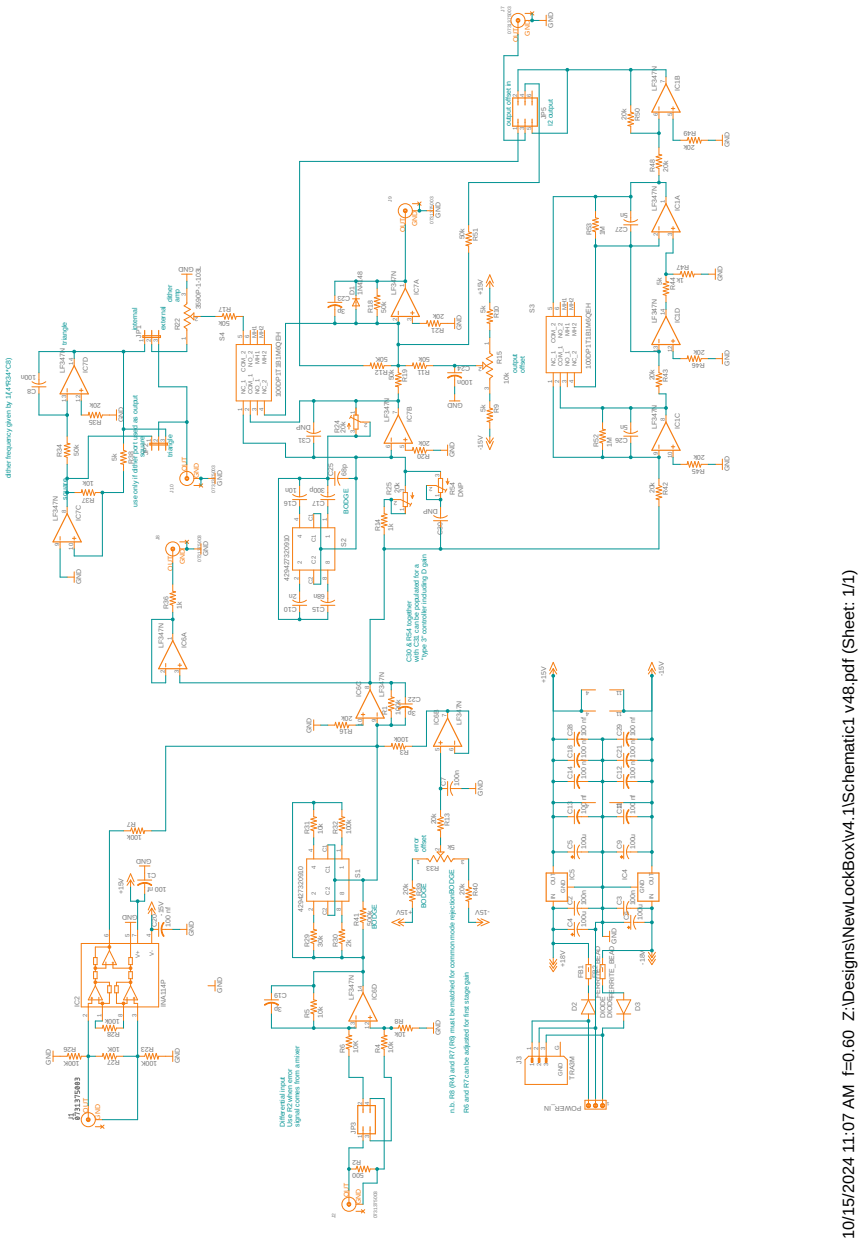


Figure A.5: Analog lockbox design consisting of a PID controller with additional double integrator and flexible front-end and output stage for a variety of lab applications from laser and cavity locking to intensity stabilization. The input signal (PDH signal, transmission, intensity) can be amplified over a large range of gains and offset subtracted to provide a symmetric error signal that is output for monitoring. A single Opamp implements a PI controller with optional D part (components labeled as DNP), and a switch toggles between a dither generator while resetting the integrator and lock while disabling the dither.

are used to match the polarity of the double integrator to the PI controller. A special structure called a *composite* opamp is formed to enable the large gains required for the double integrator with sufficient bandwidth and not limited by the amplifiers slew rate. The final stage combines the PI and II output, the dither generator and an optional output offset input that can be used for measuring transfer functions using a network analyzer.

All component values for the integrator capacitor switch bank are designed to give a wide range of corner frequencies from 100 Hz to 1 MHz.

A.8 Piezo Hysteresis compensation

At their core, piezoelectric materials operate as charge-controlled devices. The macroscopic mechanical strain (and roughly, deflection) is directly proportional to the accumulated electrical charge. The optimal method to completely bypass structural non-linearities and thermal drift is to drive these materials with a strict charge amplifier. However, implementing stable, integrated charge control for positioning close to DC is very hard to implement and not commercially available. Consequently, the vast majority of systems use voltage amplifiers to drive piezos leading to a highly hysteretic and non-linear response curve between piezo voltage and displacement. For our in-vacuum lens alignment this is a huge problem as we need to position one of the lenses with ~ 100 nm precision in two axes. The hysteresis makes it impossible to perform systematic scans without a huge overhead (one could drive the piezo through its full voltage range several times to reset the history before approaching the final point, but this could reduce the piezo lifetime).

Therefore, an accurate phenomenological model is strictly required to implement feed-forward compensation in software. Over the years, numerous mathematical frameworks have been proposed to describe ferroelectric hysteresis. We evaluated several standard models before settling on the Generalized Prandtl-Ishlinskii (GPI) formulation.

The GPI formulation strikes an ideal balance. Unlike the classical Prandtl-Ishlinskii

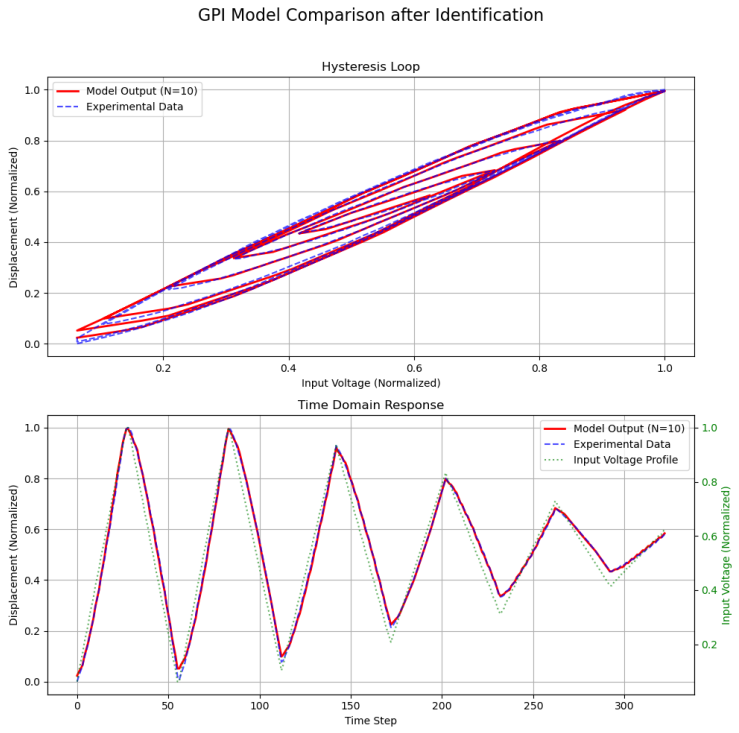
model restricted to symmetric loops, the GPI model substitutes the input with distinct non-linear envelope functions [JRS09]. This guarantees self-consistency and captures highly asymmetric major and minor loops. Crucially, it provides a tractable analytical inverse and is computationally lightweight, parameterizing the density function linearly over N discrete operators [CRZ15, YWFS20]. We can obtain an in-situ measurement of the actual piezo displacement from the frequency shift of cavity modes due to astigmatism appendix A.4 and unwrap the quadratic shape to measure the displacement with high accuracy. The effect of piezo hysteresis and non-linearity is shown in fig. A.6: The piezo is driven with a triangular waveform of diminishing amplitude (green, dashed) and the experimentally measured displacement (blue, dashed) follows with a non-linear response. The output displacement exhibits severe nonlinearity and rate-dependent hysteresis due to internal ferroelectric domain wall friction leading to a complicated hysteresis shape.

From the measured data we can fit a GPI model with a series of play operators and non-linear coefficients providing an excellent fit (red) to the measured data. Most importantly, this online model can be inverted at runtime to calculate the voltage needed to reach a desired displacement. In fig. A.6b, the same triangular scan is shown but with the inverse feed-forward model enabled, collapsing the displacement response to almost a single, straight line with small artifacts at the edges of the actuation range.

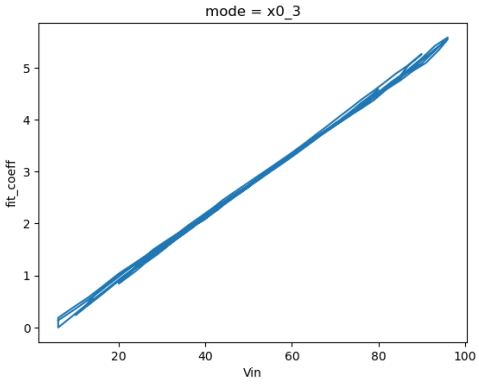
This linearization is performed for both axes of the SmarAct stage in the experiment and finally enabled a systematic scan of the lens alignment at degeneracy which was an arduous, hour-long task before.

A.9 Flexure design for precision optics alignment

Aligning optical cavities with more than just two mirrors quickly becomes challenging due to the explosion of degrees of freedom. For the twisted cavity using intra-cavity lenses, or the cavity array microscope, every mirror and lens has to have the right position and



(a) Hysteresis response of the SmarAct slip-stick stage piezos in scan mode. Driven with a decaying triangular input voltage (dashed green) the measured displacement shows both non-linearity and history dependence. Plot in the time domain (bottom) and voltage-displacement plane (top) show complicated state dependence and non-linearity.



(b) Same hysteresis curve after applying feed-forward correction using the GPI model. The voltage-displacement curve now collapses onto a single line as desired for an ideal actuator.

Figure A.6: Model identification (left) and result (right) of GPI model.

tip/tilt to form a closed cavity path and not introduce aberrations, contributing 5 degrees of freedom. Naively, one might be tempted to install every single optic on a five axis stage to be able to control every degree of freedom, however this is prohibitive due to space constraints in vacuum applications and even in test setups in air the series kinematics of most stages introduce many mechanical resonances due to the long lever arms making the whole cavity shake.

Instead, we started to introduce custom flexures that are designed to constrain most of the degrees of freedom such as the position of a mirror, while the most sensitive degrees of freedom such as tip and tilt angles are close to the design value but can still be fine tuned over a small range. By exploiting geometrical structures such as parallelograms, flexures can

realize translation stages with ranges of several 100 μm and extremely high parallelism. Such a design was used in the original GRIN based CAM design and is shown in fig. A.7. The transverse lens alignment, both in position and angle, is provided by a v-groove that was machined into the structure before the main flexure was created using EDM. At the same time, a scissor-jack design increases the free stroke of a piezo stack from only 30 μm up to 256 μm providing over 8x amplification. This turns the high-force but low-throw piezo stacks into useful devices for optical alignment that have exactly the opposite requirement. For the two GRIN-lens based CAM cavity, the flexure stage kept the transverse alignment constant over its full range of longitudinal translation. A similar structure was also employed in the first in-vacuum CAM to provide three axis translation of the curved endmirror. In order to

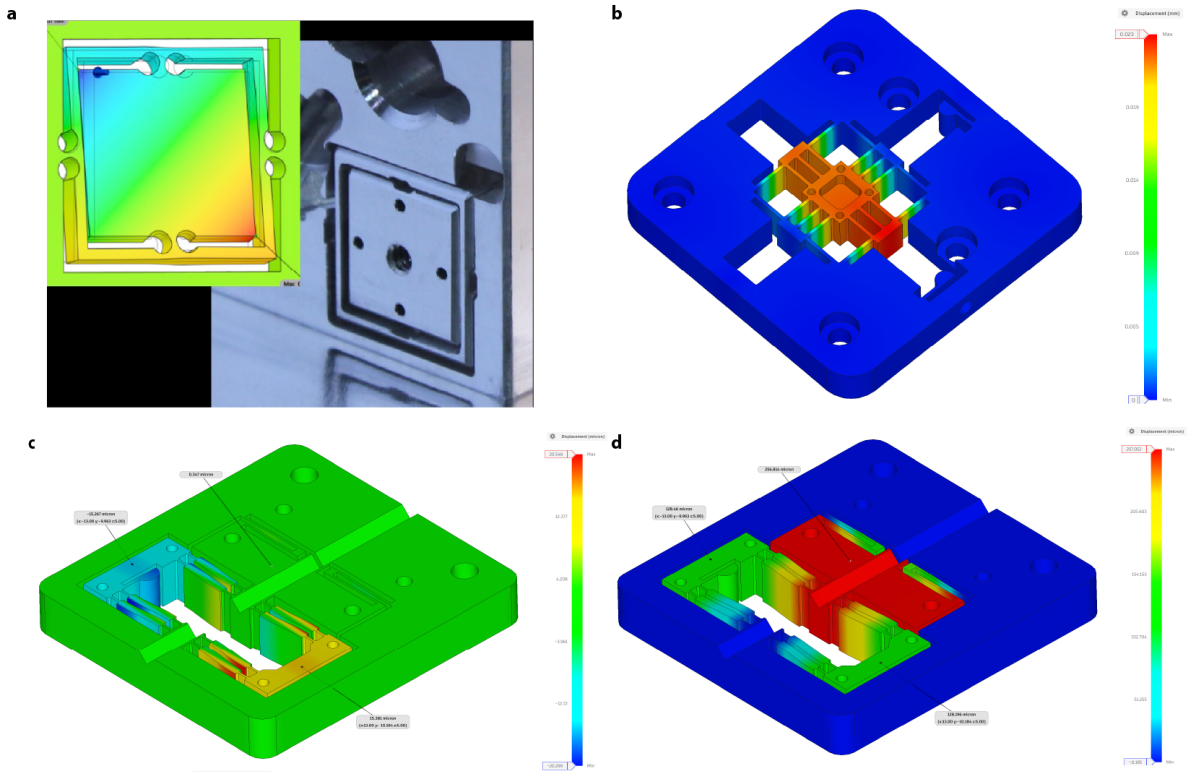


Figure A.7: Flexure designs for optical applications as a (a) tip-tilt stage integrated into the cavity mount structure, (b) xy scanning stage for the cavity AFM and (c) amplified linear translation stage for lens based cavities.

design these flexures, it is advantageous to use a combination of analytical equations and FEM in the CAD software of choice. Such an FEM simulation is shown in fig. A.7 where the load should be specified just under the blocking force of the piezo. The structure itself should be stiff enough to provide a pre-load to the piezo which improves free throw and linearity. To ensure this, the flexure element width should be chosen thick enough such that the deflection under the full piezo force is smaller than its specified free throw. At the same time, the geometry and thickness of the flexing elements will limit the ultimate achievable amplification and need to be chosen thin enough to not impede the desired amplification geometry.

After an initial FEM simulation has been obtained, the thickness (i.e. the width) of the flexing element can be fine tuned by scaling it according to the beam formula

$$\delta x = \frac{Pl^3}{3EI} = \frac{4Pl^3}{Ewh^3} = \frac{4P}{Ew}\left(\frac{l}{h}\right)^3 \tag{A.40}$$

where h is the thickness of the flexing beam in the direction of bending. As the displacement depends cubically on this, analytical scaling of the dimensions can significantly speed up the FEM iteration loop that is still required to check for e.g. maximal strain that is not to exceed for a given material. For the twisted cavity mounting structure, a special tip-tilt flexure (see fig. A.7 (a)) was designed as the in-vacuum translation stage only provides translational degrees of freedom for the lens. Using the flexure machined right into the cavity structure itself, the position of the lens stages stayed constrained exactly to its design position while the exact angle of the lens could be fine adjusted to be normal to the image rotation axis using goniometers attached to the central thread. Once this alignment was completed on the testbench, the small gaps in the flexure could be filled using a low-viscosity glass filled epoxy resin with high dimensional stability maintaining the alignment but making the structure quasi monolithic. In fig. A.7 (b), a custom high-speed translation stage for the sample scanning the cavity AFM is shown that uses parallel kinematics to provide high-bandwidth

actuation along both x- and y-axis.

The small size and parallel kinematics of flexures make them ideal for aligning optical systems such as cavities that need to be fine-tuned after assembly. The small size and stiffness reduces the impact of mechanical resonances significantly compared to large crossed-roller bearings and is also intrinsically compatible with ultra-high vacuum. While prototype stages have successfully been manufactured using water jet machining, the width of the flexing elements needs to be very accurate and uniform to not introduce undesired axes coupling, and can only be achieved by iterating cutting offsets for a particular waterjet and are rarely reproducible. In comparison, electrical discharge machining (EDM) is ideally suited as it is extremely dimensionally accurate and provides near to no cutting force on the delicate flexure structures.

A.10 Spring washer design

In the quest to investigate laser induced contamination (see appendix A.1 for details), we wanted to design a cavity that can be locked using a piezo at high bandwidth but not use any epoxy or glue to not introduce uncontrolled contaminations. The final design is a spring washer that can be used to clamp either a fixed mirror or a piezo-mirror stack. This design was inspired by the spring box design [Sto] used in the cryogenic mm-wave experiment that could not use glue due to cryogenic constraints. Said design used Belleville washers to load the piezo-mirror stack but the high stiffness of the washers requires extreme precision of the mounting structures depth and has to be matched for specific mirror substrates or piezos due to tolerance. Instead, the spring washer approach is designed to operate the individual arms not in the linear but in the plastic deformation regime where the applied force is approximately constant independent of the exact displacement. The hole in the cavity mount still has to be designed for a specific mirror and piezo thickness but affords more flexibility for tolerances on thickness or machining.

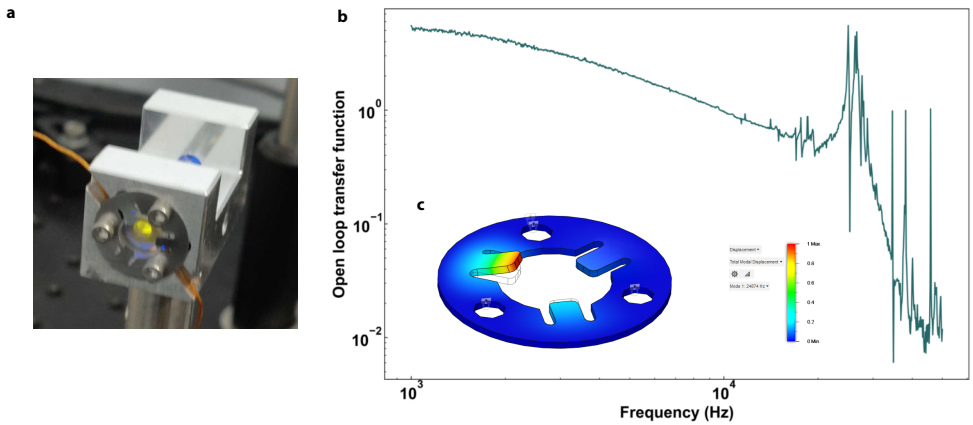


Figure A.8: Spring washer design for high bandwidth lock and vacuum compatibility. The cantilevered spring washer (c) simplifies the cavity structure (a) and eliminates the need of glue. The transfer of the piezo-mirror system from PDH cavity locking shows the first resonance at 25 kHz (b) in agreement with a FEM simulation (c).

Using half-inch mirror and a CTS ring piezo as shown in fig. A.8, the first mechanical resonance occurs at ≈ 25 kHz in agreement with FEM simulations of the first bending mode of the individual arms. This enables high-bandwidth locking and is higher than even most advanced glue based designs as it completely eliminates glue as the softest material in the stack. Furthermore, it allows for the mirror to be readily removed for inspection as was integral for the LIC testbench.

A.11 High bandwidth piezo driver

For many applications in the lab, but predominantly cavity locking, high performance piezo drivers are required.

While some good low-noise commercial drivers exist, they employ switching boost converters and special circuit topologies that make it impossible to ground any of the outputs which makes them incompatible with Rydberg atoms integration. The switching boost converters are also not galvanically isolated and can cause ground loops and high-frequency noise on the input signal chain. Also the usage of large ring piezo stacks for long throw applications or large aperture barycentric mounts as for the cavity array lead to a large ca-

pacitance requiring high driving currents. While historically some in-house class A amplifier designs were explored they suffered from excess power dissipation requiring water cooling and high noise.

After comparing a lot of different designs (Class B, Class AB using high voltage OPAMPS etc) an integrated solution from APEX was chosen: PA164 Based on the Apex PA164 high voltage operational amplifier, this design offers high voltage operation up to 200V or ±100V with output currents up to 4A peak, making it suitable for demanding applications such as high bandwidth cavity locks, precision positioning stages, and atomic force microscopy.

The PA164 can provide output currents up to 4A peak and 1A continuous within its Safe Operating Area (SOA). This is essential for driving piezoelectric loads, where the peak current requirement scales with both frequency and capacitance according to:

$$I_{pk} = V_{pp} \times 2\pi f \times C_{piezo} \tag{A.41}$$

For example, driving a typical 1µF piezo actuator at 1kHz with 100V peak-to-peak requires approximately 628mA peak current. This calculation shows why a high-current amplifier like the PA164 is necessary - even moderate-sized piezos at kHz frequencies demand substantial peak currents.

Therefore make sure to use a power supply that can source enough current! If you want to make sure the piezo doesn't see more current than intended or it is designed for, the PA164 features an integrated current limiter.

The current limit can be programmed using a precision resistor according to: $R_{CL} = 0.465V/I_{limit}$ When driving piezos in feedback applications like cavity locking, the phase response becomes critical. The output impedance of the amplifier combines with the piezo's capacitance to form an additional low-pass filter. This adds phase shift that must be considered in the overall loop stability analysis. The piezo's mechanical resonance can be electronically damped to some extent, provided the amplifier's output impedance is sufficiently

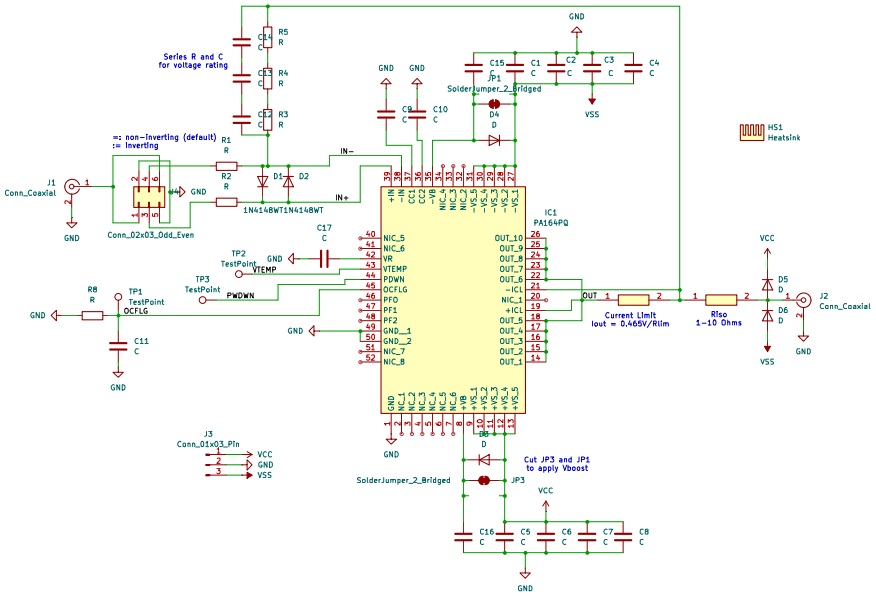


Figure A.9: Schematic of PA164 based high-bandwidth piezo drive. It can be configured for either inverting or non-inverting configuration and provides a configurable output current limit.

low (see https://www.sciencedirect.com/science/article/pii/S0924424721007639, it can make sense to add a small 10-50 Ohm resistor at the output, what the second resistor footprint is for).

The PA164 can be configured in either non-inverting or inverting topologies with gains set by external resistors:

$$G_{\text{non-inv}} = 1 + \frac{R_{\text{feedback}}}{R_{\text{ground}}} \tag{A.42}$$

$$G_{\text{inv}} = -\frac{R_{\text{feedback}}}{R_{\text{input}}} \tag{A.43}$$

The frequency response of the amplifier was measured under various load conditions.

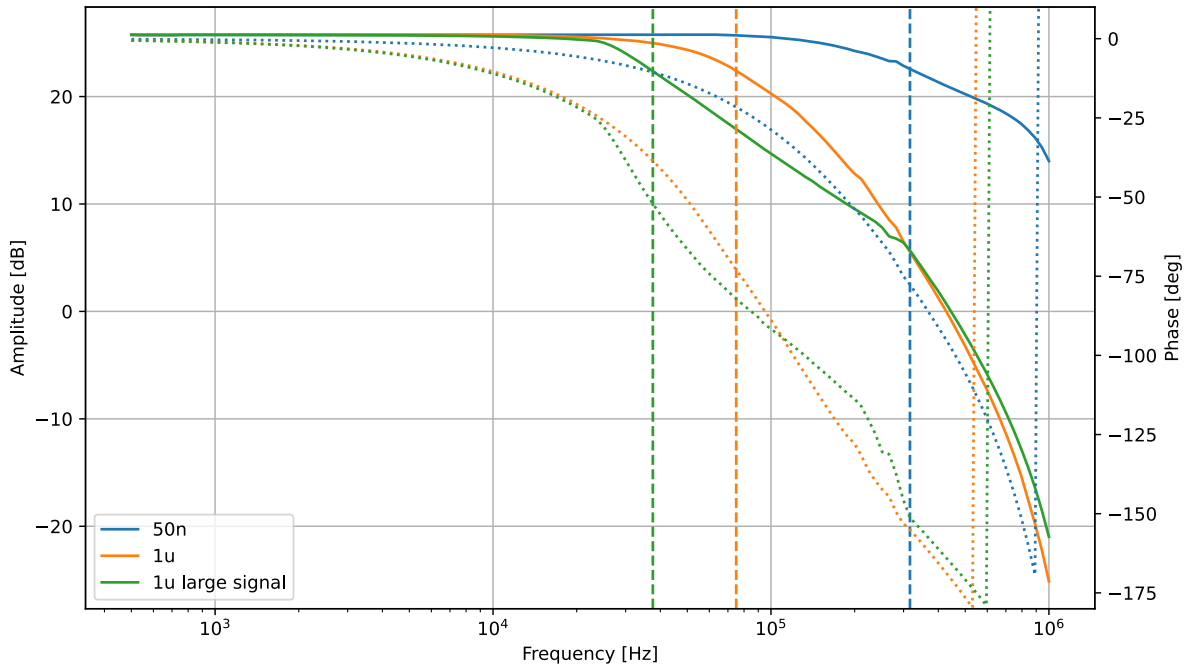


Figure A.10: Amplifier response under various load conditions. A bode plot of output voltage magnitude and phase for different piezo capacitances shows high bandwidths of up to 80 kHz for 1μF piezo capacitance.

Without load, the -3dB bandwidth extends to 200kHz, with 80 degrees of phase lag. When driving a 1µF piezo load, the bandwidth reduces to approximately 75kHz which corresponds to the large signal bandwidth of the amplifier. At that frequency it only adds 75 degrees of phase which leaves enough phase margin for a feedback loop. The green trace shows the current limiting effects of the output stage at high frequencies with large signals.

The amplifier does not add substantial noise to the input signal (20x Gain = 26 dB of voltage gain) and only shows some lines at 60 Hz and harmonics which are presumably power supply ripple.

A.12 Real-time feedback

A key challenge in the hybrid experimental cavity platform is keeping all the different components aligned, in terms of real space alignment but also in frequency space. While the cavities are actively stabilized through the cavity locks, both the science and blue buildup cavity are locked using the fundamental laser wavelength before doubling to not be resonant with the atoms. For the blue cavity, this mostly results in a static frequency offset due to the two wavelengths being reflected at different layers in the Bragg stack, but for the science cavity the relative frequency drift between science and locking light is exacerbated through the large amount of glass in the optical path. The aspheric lenses are made from fused silica but the refractive index changes with temperature of the medium, and this coefficient of change is also different between the two wavelengths. Even for small temperature fluctuations that get conducted into the cavity mount through the flange, the many wavelengths in the glass lead to a substantial phase difference that can cause a frequency drift of 10s of MHz over a few hours.

Furthermore, electric background fields can drift over time from changes in piezo voltages or slow drifts of patch charge distributions changing the field seen by the Rydberg atoms. In order to be able to collect enough data on the narrow dark polariton resonance for $g^{(2)}$ measurements, we need to actively follow both the cavity resonance and the EIT peak by adjusting the 1560 nm and 480 nm frequencies continuously.

As our hybrid experiments operates with several runs per second, we devised the following, network based architecture to store, analyze, and change parameters in the control software in real time with fast TCP based connections.

Internal architecture of the controller is made up of several key parts that are executed on the lab server that also host the experiment database in a python environment. As this pipeline needs to be low latency and robust over tens of thousands of shots, it builds heavily on established technologies and open source software. A MongoDB database is used

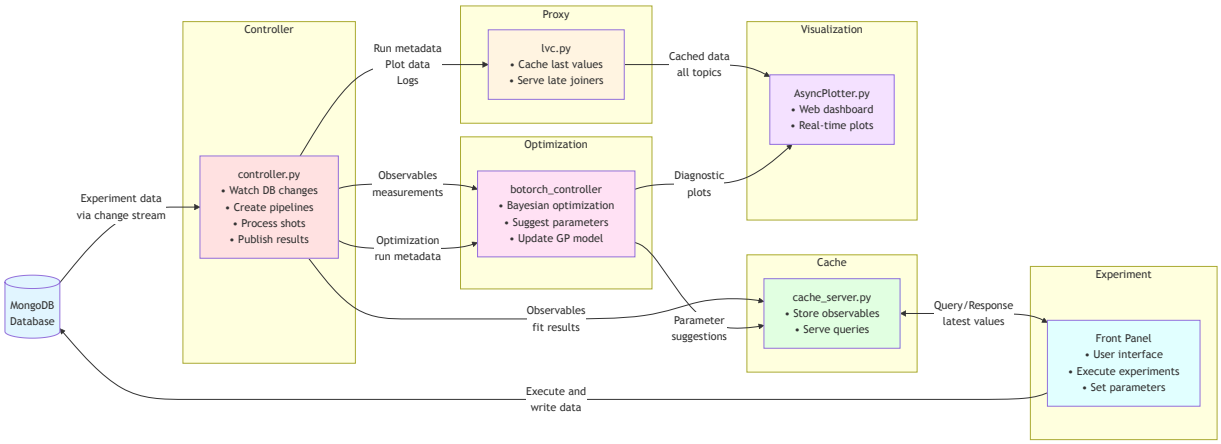


Figure A.11: High-level system architecture showing all 6 components with their key functions and data flow.

for efficient storage of small data packets from the photon counter, timer or cameras. It simultaneously enables change stream processing without polling and provides a persistence layer always committing data to the DB first. The controller creates a device-specific pipeline (Counter, ADC, Camera, EMCCD) that analyzes the data in real time and parses front-panel MVs to interpret it. For example for cavity spectroscopy it can fit Lorentzian, VRS or EIT fitfunctions in real time and forward e.g. center frequencies. All operations are designed using the Streamz library for a graph processing flow and computing mean and variance values in real time as more data is averaged. Multiple output sinks can be defined that feed data to different parts of the system for real time plots in the browser and observables that are sent to the front panel.

This realtime pipeline also enables an additional powerful tool for the experimenter:

Instead of feeding back on the cavity center frequency or other frequency offsets, the system can also be used for a Bayesian optimization workflow: A special optimization controller tries to maximize an observable (could be atom number, VRS splitting g etc.) observed by the analysis controller and send new candidate points to the front panel. First, an initialization is obtained with Sobol sampling of all the enabled parameter MVs within their specified range getting a global picture of the parameter space. From those initial observation a Gaus-

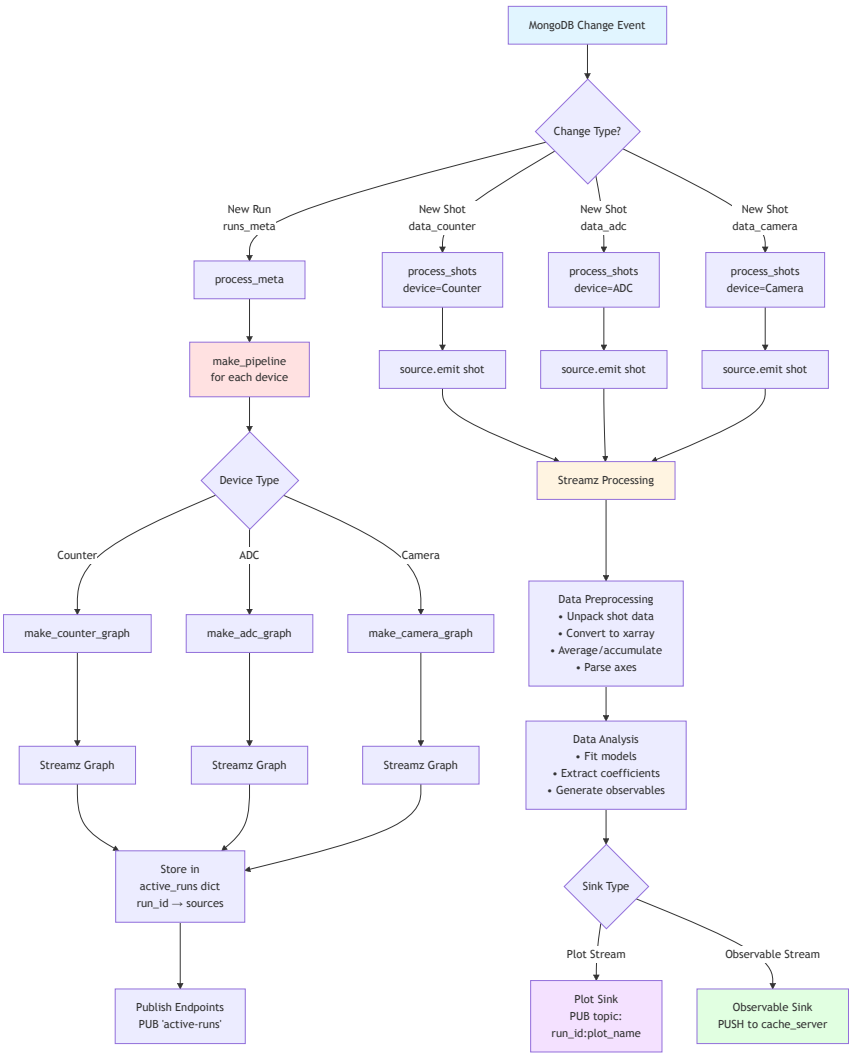


Figure A.12: Internal architecture of the controller showing:

sian process (GP) model is built that is then used to suggest new candidate points while balancing exploitation and exploration.

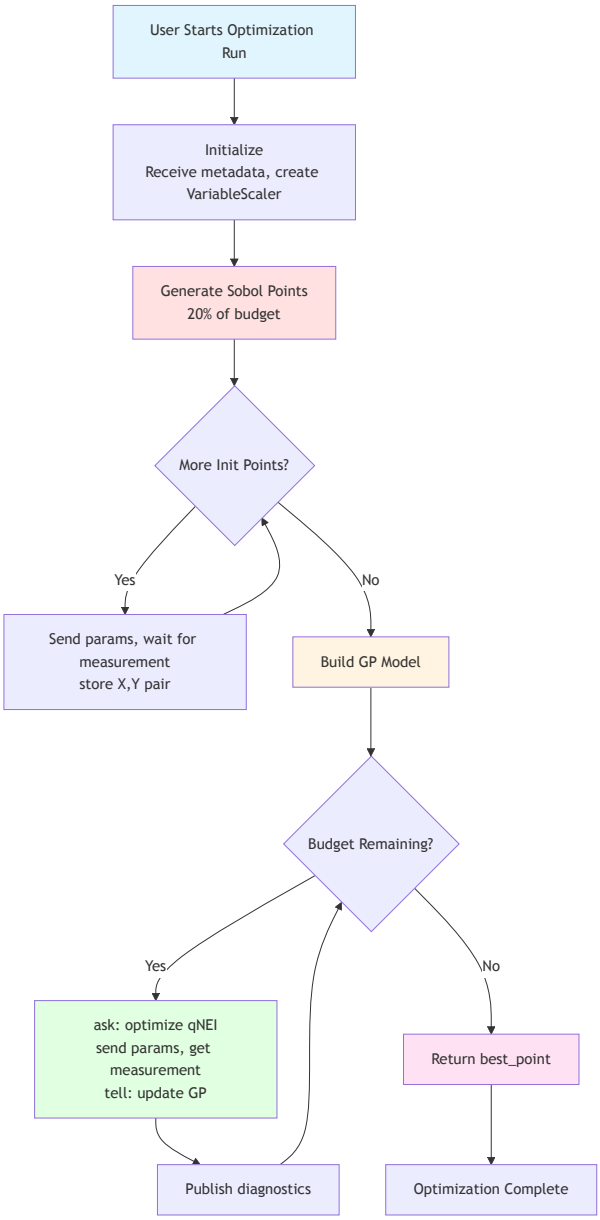


Figure A.13: Bayesian optimization workflow

REFERENCES

[AD17] Miguel A. Alonso and Mark R. Dennis. Ray-optical Poincaré sphere for structured Gaussian beams. *Optica*, 4(4):476–486, April 2017.

[AEE+23] Koji Azuma, Sophia E. Economou, David Elkouss, Paul Hilaire, Liang Jiang, Hoi-Kwong Lo, and Ilan Tzitrin. Quantum repeaters: From quantum networks to the quantum internet. *Reviews of Modern Physics*, 95(4):045006, December 2023.

[AL20] Bernard Ouma Alunda and Yong Joong Lee. Review: Cantilever-Based Sensors for High Speed Atomic Force Microscopy. *Sensors*, 20(17):4784, January 2020.

[Alo14] Miguel A. Alonso. The Connection between Rays and Waves. In Wolfgang Osten, editor, *Fringe 2013*, pages 457–464, Berlin, Heidelberg, 2014. Springer.

[And18] N. V. Andreeva. Atomic force microscopy with interferometric method for detection of the cantilever displacement and its application for low-temperature studies. *Ferroelectrics*, 525(1):178–186, March 2018.

[ANS+25] Lukas Ahlheit, Chris Nill, Daniil Svirskiy, Jan de Haan, Simon Schroers, Wolfgang Alt, Nina Stiesdal, Igor Lesanovsky, and Sebastian Hofferberth. Magic running- and standing-wave optical traps for Rydberg atoms. *Physical Review A*, 111(1):013115, January 2025.

[BBC+24] N. N. Bezuglov, I. I. Beterov, A. Cinins, K. Miculis, V. M. Entin, P. I. Betleni, G. Suliman, V. V. Gromyko, D. B. Tretyakov, E. A. Yakshina, and I. I. Ryabtsev. High-Fidelity Individual Addressing of Single Atoms in Quantum Registers at Three-Photon Laser Excitation of Rydberg States, November 2024.

[BCF+14] P. Bienias, S. Choi, O. Firstenberg, M. F. Maghrebi, M. Gullans, M. D. Lukin, A. V. Gorshkov, and H. P. Büchler. Scattering resonances and bound states for strongly interacting Rydberg polaritons. *Physical Review A*, 90(5):053804, November 2014.

[BCW+12] Justin G. Bohnet, Zilong Chen, Joshua M. Weiner, Dominic Meiser, Murray J. Holland, and James K. Thompson. A steady-state superradiant laser with less than one intracavity photon. *Nature*, 484(7392):78–81, April 2012.

[BGP+09] Waseem S. Bakr, Jonathon I. Gillen, Amy Peng, Simon Fölling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature*, 462:74–77, 2009.

[BJP+23] Claire Baum, Matt Jaffe, Lukas Palm, Aishwarya Kumar, and Jonathan Simon. Optical mode conversion via spatiotemporally modulated atomic susceptibility. *Optics Express*, 31(1):528, January 2023.

[BKB+18] Andrea Bergschneider, Vincent M. Klinkhamer, Jan Hendrik Becher, Ralf Klemt, Gerhard Zürn, Philipp M. Preiss, and Selim Jochim. Spin-resolved single-atom imaging of Li 6 in free space. *Physical Review A*, 97(6):063613, June 2018.

[BL78] R. Bonifacio and L. A. Lugiato. Optical bistability and cooperative effects in resonance fluorescence. *Physical Review A*, 18(3):1129–1144, September 1978.

[Bla01] Eric D. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, January 2001.

[BQG86] G. Binnig, C. F. Quate, and Ch. Gerber. Atomic Force Microscope. *Physical Review Letters*, 56(9):930–933, March 1986.

[BRV+] Stephanie M Bohaichuk, Fabian Ripka, Vijin Venu, Florian Christaller, Chang Liu, Matthias Schmidt, Harald Kübler, and James P Shaffer. A Three-Photon Rydberg Atom-Based Radio Frequency Sensing Scheme with Narrow Linewidth.

[BYS+24] I. I Beterov, E. A Yakshina, G. Suliman, P. I Betleni, A. A Prilutskaya, D. A Skvortsova, T. R Zagirov, D. B Tret'yakov, V. M Entin, N. N Bezuglov, and I. I Ryabtsev. Rabi Oscillations at Three-Photon Laser Excitation of a Single Rubidium Rydberg Atom in an Optical Dipole Trap. *Žurnal èksperimental'noj i teoretičeskoj fiziki*, 166(4):535–547, December 2024.

[CDG24] Claude Cohen-Tannoudji, Jacques Dupont-Roc, and Gilbert Grynberg. *Atom - Photon Interactions: Basic Processes and Applications*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, 2024.

[CFZ+25] Gianvito Chiarella, Tobias Frank, Leart Zuka, Pau Farrera, and Gerhard Rempe. Source of heralded atom-photon entanglement for quantum networking. *Physical Review Letters*, 135:240802, 2025.

[CJS+19] Logan W. Clark, Ningyuan Jia, Nathan Schine, Claire Baum, Alexandros Georgakopoulos, and Jonathan Simon. Interacting Floquet polaritons. *Nature*, 571(7766):532–536, July 2019.

[Cor23] Theodore A. Corcovilos. Beyond the ABCDs: A better matrix method for geometric optics by using homogeneous coordinates. *American Journal of Physics*, 91(6):449–457, June 2023.

[CRZ15] Jinhao Chen, Beibei Ren, and Qing-Chang Zhong. Hysteresis compensation in piezoelectric actuator positioning control based on the uncertainty and disturbance estimator. In *2015 American Control Conference (ACC)*, pages 2537–2542, July 2015.

[CSB+20] Logan W Clark, Nathan Schine, Claire Baum, Ningyuan Jia, and Jonathan Simon. Observation of Laughlin states made of light. *Nature*, 582(June), 2020.

[CWG01] N. R. Cooper, N. K. Wilkin, and J. M F Gunn. Quantum phases of vortices in rotating Bose-Einstein condensates. *Physical Review Letters*, 87(12):120405/1–120405/4, 2001.

[DA19] Mark R Dennis and Miguel A Alonso. Gaussian mode families from systems of rays. *Journal of Physics: Photonics*, 1(2):025003, February 2019.

[DBL+22] Davide Dreon, Alexander Baumgärtner, Xiangliang Li, Simon Hertlein, Tilman Esslinger, and Tobias Donner. Self-oscillating pump in a topological dissipative atom–cavity system. *Nature*, 608(7923):494–498, August 2022.

[Dog19] Nishant Dogra. *Interaction- and Dissipation-Induced Phenomena in a Quantum Gas Coupled to a Cavity*. Doctoral Thesis, ETH Zurich, 2019.

[Fet09] Alexander L. Fetter. Rotating trapped Bose-Einstein condensates. *Reviews of Modern Physics*, 81(2):647–691, 2009.

[GB21] Christian Gross and Waseem S. Bakr. Quantum gas microscopy for single atom and spin detection. *Nature Physics*, 17:1316–1323, 2021.

[Gem07] Nathan Gemelke. Nathan D Gemelke June 2007 t. (June), 2007.

[GFP+20] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret. Stabilization and operation of a Kerr-cat qubit. *Nature*, 584(7820):205–209, August 2020.

[GGB18] A. Grankin, P. Grangier, and E. Brion. Diagrammatic treatment of few-photon scattering from a Rydberg blockaded atomic ensemble in a cavity. *Physical Review A*, 98(5):053860, November 2018.

[GGD+25] Brandon Grinkemeyer, Elmer Guardado-Sanchez, Ivana Dimitrova, Danilo Shchepanovich, G. Eirini Mandopoulou, Johannes Borregaard, Vladan Vuletić, and Mikhail D. Lukin. Error-detected quantum operations with neutral atoms mediated by an optical cavity. *Science*, 387:1301–1305, 2025.

[Gie03] Franz J. Giessibl. Advances in atomic force microscopy. *Reviews of Modern Physics*, 75(3):949–983, July 2003.

[GMP86] S. M. Girvin, A. H. MacDonald, and P. M. Platzman. Magneto-roton theory of collective excitations in the fractional quantum Hall effect. *Physical Review B*, 33(4):2481–2494, February 1986.

[GSC10] Nathan Gemelke, Edina Sarajlic, and Steven Chu. Rotating Few-body Atomic Systems in the Fractional Quantum Hall Regime. pages 1–22, 2010.

[GSS18] Alexandros Georgakopoulos, Ariel Sommer, and Jonathan Simon. Theory of interacting cavity Rydberg polaritons. *Quantum Science and Technology*, 4:014005, 2018.

[GTCH06] Andrew D. Greentree, Charles Tahan, Jared H. Cole, and L. C. L. Hollenberg. Quantum phase transitions of light. *Nature Physics*, 2(12):856–861, December 2006.

[GTW+16] M. J. Gullans, J. D. Thompson, Y. Wang, Q.-Y. Liang, V. Vuletić, M. D. Lukin, and A. V. Gorshkov. Effective Field Theory for Rydberg Polaritons. *Physical Review Letters*, 117(11):113601, September 2016.

[GXF+16] Zhe-Xuan Gong, Minghui Xu, Michael Foss-Feig, James K. Thompson, Ana Maria Rey, Murray Holland, and Alexey V. Gorshkov. Steady-state superradiance with Rydberg polaritons, November 2016.

[Hal] F Duncan M Haldane. Quantum Geometry, Exclusion Statistics, and the Geometry of "Flux Attachment" in 2D Landau levels.

[HK10] M. Z. Hasan and C. L. Kane. Colloquium: Topological insulators. *Reviews of Modern Physics*, 82(4):3045–3067, 2010.

[HKF+93] Fritz Haake, Mikhail I. Kolobov, Claude Fabre, Elisabeth Giacobino, and Serge Reynaud. Superradiant laser. *Physical Review Letters*, 71(7):995–998, August 1993.

[HLZ20] Fei He, Jian Liu, and Ka-Di Zhu. Optomechanical atomic force microscope. *Nanotechnology*, 32(8):085505, December 2020.

[Hof76] Douglas R. Hofstadter. Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields. *Physical Review B*, 14(6):2239–2249, September 1976.

[HPS+06] M. J. Higgins, R. Proksch, J. E. Sader, M. Polcik, S. Mc Endoo, J. P. Cleveland, and S. P. Jarvis. Noninvasive determination of optical lever sensitivity in atomic force microscopy. *Review of Scientific Instruments*, 77(1):013701, January 2006.

[HSB+25] Beili Hu, Josiah Sinclair, Edita Bytyqi, Michelle Chong, Alyssa Rudelis, Joshua Ramette, Zachary Vendeiro, and Vladan Vuletić. Site-selective cavity readout and classical error correction of a 5-bit atomic register. *Physical Review Letters*, 134:120801, 2025.

[HSW+24] Lukas Hartung, Matthias Seubert, Stephan Welte, Emanuele Distante, and Gerhard Rempe. A quantum-network register assembled with optical tweezers in an optical cavity. *Science*, 385:179–183, 2024.

[IHK+25] Alexander Impertro, SeungJung Huh, Simon Karch, Julian F. Wienand, Immanuel Bloch, and Monika Aidelsburger. Strongly interacting Meissner phases in large bosonic flux ladders. *Nature Physics*, 21:895–901, 2025.

[ILSF18] Peter A. Ivanov, Fabian Letscher, Jonathan Simon, and Michael Fleischhauer. Adiabatic flux insertion and growing of Laughlin states of cavity Rydberg polaritons. *Physical Review A*, 98(1):013847, July 2018.

[IMV+20] Denis A. Ikonnikov, Sergey A. Myslivets, Mikhail N. Volochaev, Vasily G. Arkhipkin, and Andrey M. Vyunishev. Two-dimensional Talbot effect of the optical vortices and their spatial evolution. *Scientific Reports*, 10(1):20315, November 2020.

[Jac22] Janusz E. Jacak. Formal derivation of the Laughlin function and its generalization for other topological phases of FQHE. *Scientific Reports*, 12(1):616, January 2022.

[Jai89] J. K. Jain. Composite-fermion approach for the fractional quantum Hall effect. *Physical Review Letters*, 63(2):199–202, 1989.

[JCL23] A. M. Joshua, G. Cheng, and E. V. Lau. Soft matter analysis via atomic force microscopy (AFM): A review. *Applied Surface Science Advances*, 17:100448, October 2023.

[JMM+22] Naijun Jin, Charles A. McLemore, David Mason, James P. Hendrie, Yizhi Luo, Megan L. Kelleher, Prashanta Kharel, Franklyn Quinlan, Scott A. Diddams, and Peter T. Rakich. Micro-fabricated mirrors with finesse exceeding one million. *Optica*, 9(9):965–970, September 2022.

[JRS09] Mohammad Al Janaideh, Subhash Rakheja, and Chun-Yi Su. A generalized Prandtl–Ishlinskii model for characterizing the hysteresis and saturation nonlinearities of smart actuators. *Smart Materials and Structures*, 18(4):045001, April 2009.

[JZG+23] Kaiwen Ji, Qi Zhong, Li Ge, Gregoire Beaudoin, Isabelle Sagnes, Fabrice Raineri, Ramy El-Ganainy, and Alejandro M. Yacomotti. Tracking exceptional points above the lasing threshold. *Nature Communications*, 14(1):8304, December 2023.

[KÇO20] Özgür Karc, Ümit Çelik, and Ahmet Oral. 1 fm/Hz noise level low temperature Fabry-Pérot atomic force/magnetic force microscope operating in 4–300 K temperature range. *Review of Scientific Instruments*, 91(1):013703, January 2020.

[KK18] Peter Kirton and Jonathan Keeling. Superradiant and lasing states in driven-dissipative Dicke models. *New Journal of Physics*, 20(1):015009, January 2018.

[KL09] Alexei Kitaev and Chris Laumann. Topological phases and quantum computation. 2009.

[KML+23] Ronen M. Kroeze, Brendan P. Marsh, Kuan-Yu Lin, Jonathan Keeling, and Benjamin L. Lev. High Cooperativity Using a Confocal-Cavity–QED Microscope. *PRX Quantum*, 4(2):020326, May 2023.

[KRKDT19] Peter Kirton, Mor M. Roses, Jonathan Keeling, and Emanuele G. Dalla Torre. Introduction to the Dicke Model: From Equilibrium to Nonequilibrium, and Vice Versa. *Advanced Quantum Technologies*, 2(1-2):1800043, 2019.

[KTU03] Kenichi Kasamatsu, Makoto Tsubota, and Masahito Ueda. Nonlinear dynamics of vortex lattice formation in a rotating Bose-Einstein condensate. *Physical Review A*, 67(3):033610, March 2003.

[KV08] T. J. Kippenberg and K. J. Vahala. Cavity Optomechanics: Back-Action at the Mesoscale. *Science*, 321(5893):1172–1176, August 2008.

[Lau83] R. B. Laughlin. Anomalous quantum Hall effect: An incompressible quantum fluid with fractionally charged excitations. *Physical Review Letters*, 50(18):1395–1398, 1983.

[LDZ+21] Xiangliang Li, Davide Dreon, Philip Zupancic, Alexander Baumgärtner, Andrea Morales, Wei Zheng, Nigel R. Cooper, Tobias Donner, and Tilman Esslinger. First order phase transition between two centro-symmetric superradiant crystals. *Physical Review Research*, 3(1):L012024, March 2021.

[LHR+24] Philipp Lunt, Paul Hill, Johannes Reiter, Philipp M. Preiss, Maciej Gałka, and Selim Jochim. Realization of a Laughlin state of two rapidly rotating fermions. *Physical Review Letters*, 133:253401, 2024.

[LKK+23] Julian Léonard, Sooshin Kim, Joyce Kwan, Perrin Segura, Fabian Grusdt, Cécile Repellin, Nathan Goldman, and Markus Greiner. Realization of a fractional quantum Hall state with ultracold atoms. *Nature*, 619:495–499, 2023.

[LPJ+25] Bowen Li, Lukas Palm, Marius Jürgensen, Yiming Cady Feng, Markus Greiner, and Jon Simon. Mega-FPS low light camera. *Optics Express*, 33(15):31096–31106, July 2025.

[LWS+10] Qiang Li, Tao Wang, Yikai Su, Min Yan, and Min Qiu. Coupled mode theory analysis of mode-splitting in coupled cavity system. *Optics Express*, 18(8):8367–8382, April 2010.

[Mar16] D. V. Martynov. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Physical Review D*, 93(11), 2016.

[MBL+20] Juan A. Muniz, Diego Barberena, Robert J. Lewis-Swan, Dylan J. Young, Julia R. K. Cline, Ana Maria Rey, and James K. Thompson. Exploring dynamical phase transitions with cold atoms in an optical cavity. *Nature*, 580(7805):602–607, April 2020.

[MCT+08] M. I. Makin, Jared H. Cole, Charles Tahan, Lloyd C. L. Hollenberg, and Andrew D. Greentree. Quantum phase transitions in photonic cavities with two-level systems. *Physical Review A*, 77(5):053819, May 2008.

[MDL+19] Andrea Morales, Davide Dreon, Xiangliang Li, Alexander Baumgärtner, Philip Zupancic, Tobias Donner, and Tilman Esslinger. Two-mode Dicke model from nondegenerate polarization modes. *Physical Review A*, 100(1):013816, July 2019.

[MF07] Alexis G. Morris and David L. Feder. Gaussian potentials facilitate access to quantum hall states in rotating Bose gases. *Physical Review Letters*, 99(24):2–5, 2007.

[MGR+20] Yoichi Miyahara, Harrisonn Griffin, Antoine Roy-Gobeil, Ron Belyansky, Hadallia Bergeron, José Bustamante, and Peter Grutter. Optical excitation of atomic force microscopy cantilever for accurate spectroscopic measurements. *EPJ Techniques and Instrumentation*, 7(1):1–11, December 2020.

[MH10] D. Meiser and M. J. Holland. Steady-state superradiance with alkaline-earth-metal atoms. *Physical Review A*, 81(3):033847, March 2010.

[MSO+19] Ruichao Ma, Brendan Saxberg, Clai Owens, Nelson Leung, Yao Lu, Jonathan Simon, and David I. Schuster. A dissipatively stabilized Mott insulator of photons. *Nature*, 566(7742):51–57, February 2019.

[NA21] Tathagata Nandi and Sri Rama Koti Ainavarapu. Applications of atomic force microscopy in modern biology. *Emerging Topics in Life Sciences*, 5(1):103–111, February 2021.

[NKKN09] Shuhei Nishida, Dai Kobayashi, Hideki Kawakatsu, and Yuki Nishimori. Photothermal excitation of a single-crystalline silicon cantilever for higher vibration modes in liquid. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, 27(2):964–968, April 2009.

[NLC+18] Matthew A. Norcia, Robert J. Lewis-Swan, Julia R. K. Cline, Bihui Zhu, Ana M. Rey, and James K. Thompson. Cavity-mediated collective spin-exchange interactions in a strontium superradiant laser. *Science*, 361(6399):259–262, July 2018.

[Pan56] S. Pancharatnam. Generalized theory of interference, and its applications. *Proceedings of the Indian Academy of Sciences - Section A*, 44(5):247–262, November 1956.

[PdT+22] Marco Piccardo, Michael de Oliveira, Andrea Toma, Vincenzo Aglieri, Andrew Forbes, and Antonio Ambrosio. Vortex laser arrays with topological charge control and self-healing of defects. *Nature Photonics*, 16(5):359–365, May 2022.

[PdvG92] Constant A. J. Putman, Bart G. de Grooth, Niek F. van Hulst, and Jan Greve. A theoretical comparison between interferometric and optical beam deflection technique for the measurement of cantilever displacement in AFM. *Ultramicroscopy*, 42–44:1509–1513, July 1992.

[PGP20] L Palm, F Grusdt, and P M Preiss. Skyrmion ground states of rapidly rotating few-fermion systems. *New Journal of Physics*, 22(8):083037, August 2020.

[PHR22] David Plankensteiner, Christoph Hotter, and Helmut Ritsch. QuantumCumulants.jl: A Julia framework for generalized mean-field equations in open quantum systems. *Quantum*, 6:617, January 2022.

[PKB+24] Felix A. Palm, Joyce Kwan, Brice Bakkali-Hassani, Markus Greiner, Ulrich Schollwöck, Nathan Goldman, and Fabian Grusdt. Growing extended Laughlin states in a quantum gas microscope: A patchwork construction. *Physical Review Research*, 6:013198, 2024.

[PLN+21] Paolo Piergentili, Wenlin Li, Riccardo Natali, David Vitali, and Giovanni Di Giuseppe. Absolute Determination of the Single-Photon Optomechanical Coupling Rate via a Hopf Bifurcation. *Physical Review Applied*, 15(3):034012, March 2021.

[RBT+11] I. I. Ryabtsev, I. I. Beterov, D. B. Tretyakov, V. M. Entin, and E. A. Yakshina. Doppler- and recoil-free laser excitation of Rydberg states via three-photon transitions. *Physical Review A*, 84(5):053409, November 2011.

[RF16] Ben Routley and Andrew Fleming. High sensitivity interferometer for on-axis detection of AFM cantilever deflection. In *2016 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pages 1–5, July 2016.

[RLB+14] Sylvain Ravets, Henning Labuhn, Daniel Barredo, Lucas Béguin, Thierry Lahaye, and Antoine Browaeys. Coherent dipole–dipole coupling between two single Rydberg atoms at an electrically-tuned Förster resonance. *Nature Physics*, 10(12):914–917, December 2014.

[RMT08] Arvind Raman, John Melcher, and Ryan Tung. Cantilever dynamics in atomic force microscopy. *Nano Today*, 3(1-2):20–27, February 2008.

[RNH+12] Stephan Ritter, Christian Nölleke, Carolin Hahn, Andreas Reiserer, Andreas Neuzner, Manuel Uphoff, Martin Mücke, Eden Figueroa, Joerg Bochmann, and Gerhard Rempe. An elementary quantum network of single atoms in optical cavities. *Nature*, 484(7393):195–200, April 2012.

[RRD11] M. Roncaglia, M. Rizzi, and J. Dalibard. From rotating atomic rings to quantum Hall states. *Scientific Reports*, 1(1):43, 2011.

[SBS15] Ariel Sommer, Hans Peter Büchler, and Jonathan Simon. Quantum Crystals and Laughlin Droplets of Cavity Rydberg Polaritons, June 2015.

[SFR+23] Kevin Stitely, Fabian Finger, Rodrigo Rosa-Medina, Francesco Ferri, Tobias Donner, Tilman Esslinger, Scott Parkins, and Bernd Krauskopf. Quantum Fluctuation Dynamics of Dispersive Superradiant Pulses in a Hybrid Light-Matter System, February 2023.

[SGM+04] Benjamin S. Sheard, Malcolm B. Gray, Conor M. Mow-Lowry, David E. McClelland, and Stanley E. Whitcomb. Observation and characterization of an optical spring. *Physical Review A*, 69(5):051801, May 2004.

[Sie86] A. E. Siegman. *Lasers*. 1986.

[SKK+24] Jim Skulte, Phatthamon Kongkhambut, Hans Keßler, Andreas Hemmerich, Ludwig Mathey, and Jayson G. Cosme. Realizing limit cycles in dissipative bosonic systems, January 2024.

[SKS+25] Danial Shadmany, Aishwarya Kumar, Anna Soper, Lukas Palm, Chuan Yin, Henry Ando, Bowen Li, Lavanya Taneja, Matt Jaffe, Schuster David, and Jon Simon. Cavity QED in a high NA resonator. *Science Advances*, 11(9):eads8171, February 2025.

[ŠPAW17] N. Šibalić, J. D. Pritchard, C. S. Adams, and K. J. Weatherill. ARC: An open-source library for calculating properties of alkali Rydberg atoms. *Computer Physics Communications*, 220:319–331, November 2017.

[SRG+16] Nathan Schine, Albert Ryou, Andrey Gromov, Ariel Sommer, and Jonathan Simon. Letter Synthetic Landau levels for photons. *Nature*, 534(7609):671–675, 2016.

[SSDRG11] Nicolas Sangouard, Christoph Simon, Hugues De Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33–80, March 2011.

[SSF20] M. Sánchez-Barquilla, R. E. F. Silva, and J. Feist. Cumulant expansion for the treatment of light-matter interactions in arbitrary material structures. *The Journal of Chemical Physics*, 152(3):034108, January 2020.

[Sto] Mark Stone. *Hybrid Cavity QED with Rydberg Atoms*. PhD thesis.

[SWE+10] Jacob F. Sherson, Christof Weitenberg, Manuel Endres, Marc Cheneau, Immanuel Bloch, and Stefan Kuhr. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature*, 467:68–72, 2010.

[SZL+11] F Serwane, G Zürn, T Lompe, T B Ottenstein, A N Wenz, and S Jochim. Deterministic preparation of a tunable few-fermion system. *Science*, 332(6027):336–338, April 2011.

[Tan] Lavanya Taneja. *Cavity-QED Tools for Quantum Interconnects*. PhD thesis.

[TMA+19] N. Thaicharoen, K. R. Moore, D. A. Anderson, R. C. Powel, E. Peterson, and G. Raithel. Electromagnetically-induced transparency, absorption, and microwave field sensing in a Rb vapor cell with a three-color all-infrared laser system. *Physical Review A*, 100(6):063427, December 2019.

[Ton16] David Tong. Lectures on the Quantum Hall Effect. (January), 2016.

[TSG82] D C Tsui, H L Stormer, and A C Gossard. Two-dimensional magnetotransport in the extreme quantum limit. *Physical Review Letters*, 48(22):1559–1562, 1982.

[TXM+17] D. A. Tieri, Minghui Xu, D. Meiser, J. Cooper, and M. J. Holland. Theory of the crossover from lasing to steady state superradiance, February 2017.

[UC17] R. O. Umucalılar and I. Carusotto. Generation and spectroscopic signatures of a fractional quantum Hall liquid of photons in an incoherently pumped optical cavity. *Physical Review A*, 96(5):053808, November 2017.

[USC21] Rifat Onur Umucalilar, Jonathan Simon, and Iacopo Carusotto. Autonomous stabilization of photonic Laughlin states through angular momentum potentials. *Physical Review A*, 104(2):023704, August 2021.

[UVR+19] Alban Urvoy, Zachary Vendeiro, Joshua Ramette, Albert Adiyatullin, and Vladan Vuletić. Direct Laser Cooling to Bose-Einstein Condensation in a Dipole Trap. *Physical Review Letters*, 122(20):203202, May 2019.

[Vie08] Susanne Viefers. Quantum Hall physics in rotating Bose-Einstein condensates. 2008.

[WLC+24] Can Wang, Feng-Ming Liu, Ming-Cheng Chen, He Chen, Xian-He Zhao, Chong Ying, Zhong-Xia Shang, Jian-Wen Wang, Yong-Heng Huo, Cheng-Zhi Peng, Xiaobo Zhu, Chao-Yang Lu, and Jian-Wei Pan. Realization of fractional quantum Hall state with interacting photons. *Science*, 384:579–584, 2024.

[XZZ19] Xiaoxiao Xue, Xiaoping Zheng, and Bingkun Zhou. Super-efficient temporal solitons in mutually coupled optical cavities. *Nature Photonics*, 13(9):616–622, September 2019.

[YDZJ23] Chen Yang, Chao-Qun Dang, Wu-Le Zhu, and Bing-Feng Ju. High-speed atomic force microscopy in ultra-precision surface machining and measurement: Challenges, solutions and opportunities. *Surface Science and Technology*, 1(1):7, September 2023.

[YHL+23] Zhenjie Yan, Jacquelyn Ho, Yue-Hui Lu, Stuart J. Masson, Ana Asenjo-Garcia, and Dan M. Stamper-Kurn. Superradiant and subradiant cavity scattering by atom arrays. *Physical Review Letters*, 131:253603, 2023.

[YLMS21] Jirayu Yang, Lesli Liu, Jiuwu Mongkolkiattichai, and Peter Schauss. Site-resolved imaging of ultracold fermions in a triangular-lattice quantum gas microscope. *PRX Quantum*, 2:020344, 2021.

[YWFS20] Zhiliang Yu, Yue Wu, Zhiyi Fang, and Hailin Sun. Modeling and compensation of hysteresis in piezoelectric actuators. *Heliyon*, 6(5):e03999, May 2020.

[ZAS+05] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck, and W. Ketterle. Vortices and superfluidity in a strongly interacting Fermi gas. *Nature*, 435(7045):1047–1051, 2005.

[ZDL+19] P. Zupancic, D. Dreon, X. Li, A. Baumgärtner, A. Morales, W. Zheng, N. R. Cooper, T. Esslinger, and T. Donner. P -Band Induced Self-Organization and Dynamics with Repulsively Driven Ultracold Atoms in an Optical Cavity. *Physical Review Letters*, 123(23):233601, December 2019.